

**Alexis Espino-Sánchez
Gabriele Salvatore de Blasio
Carmelo Rubén García
José Carlos Rodríguez
Mariam Quesada-Ojeda
Alexis Quesada-Arencibia (Eds.)**

EUROCAST 2026

Computer Aided Systems Theory

EXTENDED ABSTRACTS

**20th International Conference on Computer Aided Systems Theory
*Las Palmas de Gran Canaria, Spain, February 2026***

**Twentieth International Conference on
COMPUTER AIDED SYSTEMS THEORY**

EUROCAST 2026

Edited by

Alexis Espino-Sánchez
Gabriele Salvatore de Blasio
Carmelo Rubén García
José Carlos Rodríguez-Rodríguez
Mariam Quesada-Ojeda
Alexis Quesada-Arencibia

© 2026 IUCES Universidad de Las Palmas de Gran Canaria

ISBN: 978-84-09-81744-3

REGISTRO: 2026003448

Preface

The Eurocast Conferences are particularly unique among the European Scientific-Technical Congresses because it is one of the few periodic meetings that is promoted and organized exclusively by university and socio-cultural institutions, without the tutelage, direction or funding of associations, professionals or companies. It is currently the oldest of those. It is celebrated every two years. Initially, alternating Las Palmas de G.C. and a university in continental Europe, and since 2001, always in Las Palmas de G.C.

The idea of the first Eurocast was developed in 1988 by Prof. Franz Pichler, from the University of Linz and Prof. Roberto Moreno, who recently passed away in December 2025, at a meeting in Vienna promoted by the past Honorary President, the late Dr. Werner Schimanovich. The first meeting, Eurocast 1989, took place in February of that year, in Las Palmas School of Industrial Engineers, promoted by the Faculty of Informatics of Las Palmas and the Institute of Systems of the University of Linz. The Opening Session took place in the town of Gáldar, February 26th, 1989.

Science, and especially Technology, have moved in an almost vertiginous way, driven by the need and the promotion of consumerism, associated with the change of values that has been printed in the new generations. And Eurocast, within what we understand as a certain freedom, and with prudence, has been adapting the profile of its organization from a meeting of very specific specialists, to a practically multidisciplinary, flexible and changing conference, which in each event try to attract the experts and especially young researchers, facilitating the interaction between them, which is a generator of creativity.

The key to the success of Eurocast for 37 years has been in the quality of the contributions of its participants. This has to be recognized in the first place. They have made possible, with the help of the Springer Verlag publications in Computer Science, the worldwide distribution of the most important effect of Eurocast: that of joining together for many years, scientists and engineers of ages, training, interests and from very different European and non-European institutions. And that they could share their experiences in the design and analysis of systems using the most advanced mathematical methods to make efficient models and algorithms in computers. And this from the socio-economic, biological, medical technologies and sciences and information and communication engineering topics. All in a multidisciplinary atmosphere, which has facilitated the appearance and discussion of new and creative ideas and developments.

In this open multidisciplinary spirit, the 2026 edition consists of 18 major thematic blocks, which sweep a broad spectrum of cutting-edge research in computer and systems sciences and technologies, including systems

theory, applications, pioneers and landmarks, theory and applications of metaheuristic algorithms, mechatronic product development, model-based system design, verification and simulation, applications of signal processing technology, applied data science and engineering for intelligent transportation systems and smart mobility, computer and systems based methods and electronic tools in clinical and academic medicine, systems in industrial robotics, automation and IoT, systems thinking and systems simulation: an interdisciplinary approach to modeling complexity, data science in medical and bio-informatics, modeling, simulation and optimization in production and logistics, "Green AI" and SW-Tools for sustainable energy and materials consumption, stochastic models, statistical methods and applied systems simulations, impact of AI and quantum technologies on cybersecurity, indoor positioning and localization, blackbox systems for generative AI – explainability and applications, cybermedical systems as human-centered AI and digital medicine applications and innovative teaching approaches in computer science and automation: methods, tools and experiences.

Among approximately 200 submissions of papers to participate in the Conference, 166 will be presented and defended by their authors during the four days of scientific sessions. Among them, we will proceed to a final selection, after the Conference, which will be published by Springer Verlag and distributed to the scientific community.

But, again, the kernel of the success of Eurocast lays in the right proposals of subjects for Workshops, their resonance and impact, their diffusion and their strict selection of the many intended contributions, all by the Workshops Chairpersons. They are Eurocast.

Three invited Plenary Conferences are presented: the first, by Prof. Atsushi Ito from Chuo University, known expert in Artificial Intelligence applied to different fields as software learning support and agricultural support. The second, by Prof. Joaquín Torres-Sospedra from University of Valencia, known expert in Neural Networks, Pattern Recognition, Machine Learning, Image Processing, Outdoor robotics, Artificial Intelligence, Indoor Localization&Positioning. And the third, by Prof. Andrzej Gnatowski from Wrocław University of Technology and expert in Quantum Computation.

Special thanks to our hosts, partners and most of all, friends, colleagues of Elder Museum.

The Eurocast 2026 Conference, in line with its 37 years of history, once again offers an opportunity for Canary, Spanish, European and global science and technology for openness and international relations, which will benefit our society. Welcome to **Eurocast 2026**.

Las Palmas de Gran Canaria, February 2026.
The Editors.

eurocast 2026

Twentieth International Conference on Computer Aided Systems Theory

February 23 - 27, 2026

**Museo Elder de la Ciencia y la Tecnología, Las Palmas de Gran Canaria
Canary Islands, Spain**



Supporter Institutions:



Universidad de Las Palmas de Gran Canaria
Las Palmas de Gran Canaria, Spain



Johannes Kepler Universität Linz
Linz, Austria



University of Applied Sciences
Upper Austria



Museo Elder de la Ciencia y la Tecnología
Las Palmas de Gran Canaria, Spain



Fundación Universitaria de Las Palmas
Las Palmas de Gran Canaria, Spain



With the collaboration of

City Council of Gáldar
Convention Bureau



Conference Chairs

Founder and Honorary: Franz Pichler (Linz)

General: Roberto Moreno-Díaz (Las Palmas)

Conference: Alexis Quesada-Arencibia (Las Palmas)

Program: Michael Affenzeller (Hagenberg)

Logistics: Gabriel de Blasio (Las Palmas)



Technical Secretariat

Alexis Espino-Sánchez
Instituto Universitario de Cibernética,
Empresa y Sociedad (IUCES)
Universidad de Las Palmas de Gran Canaria
Campus de Tafira
35017 Las Palmas de Gran Canaria, Spain

Phone: +34-928-457108
E-mail: eurocast@iuctc.ulpgc.es

Roberto Moreno Díaz (1939-2025)

In Memoriam



Professor Moreno Díaz was one of the main driving forces behind the University of Las Palmas de Gran Canaria and a founding figure of computer science studies at this institution. He obtained his degree in Physical Sciences in 1962 and his PhD in 1965 from the University of Madrid. He became a university Professor of Electromagnetism in 1968 and later Professor of Computer Science and Artificial Intelligence.

From 1962 to 1965, he served as an assistant professor of Industrial Physics at the University of Madrid. Between 1965 and 1968, he was a member of the scientific staff at the Charles Stark Draper Laboratory at the Massachusetts Institute of Technology (MIT), under the supervision of Warren S. McCulloch, one of the fathers of cybernetics, and Louis Sutro. During this period, he began designing a robot for the exploration of Mars at the request of NASA. From 1969 to 1979, he was Professor and Head of the Department of Electricity and Electronics at the University of Zaragoza, where he founded a research group focused on neural networks, vision, and computing. In 1979, he returned to Las Palmas, where he established several research groups that are still active today at the ULPGC.

Professor Moreno Díaz was the author or co-author of more than 135 scientific publications in the fields of neurocybernetics, retinal theory, and natural and artificial vision. He supervised 22 doctoral theses on these topics as well as on multimedia technologies.

Throughout his career, he held numerous academic and research leadership positions, including Dean and Vice-Rector at the University of Zaragoza; Director of the Institute of Educational Sciences (ICE); Dean of the Faculty of Computer Science at the ULPGC; Director of the Canary Islands Institute of Technology (ITC); and Director of one of the first research institutes at the ULPGC, the University Institute of Cybernetic Sciences and Technologies (IUCTC).

From 1999 onwards, he was awarded all six levels of research recognition granted by the Spanish Ministry of Education and Science, as well as all six teaching merit levels.

He became a Corresponding Member of the Royal Academy of Exact, Physical and Natural Sciences of Madrid in 1981. In 1985, he received the Canary Islands Research Prize. He was a Founding Member and Vice-President of the Canary Islands Academy of Sciences and received the first ULPGC Plaque for University Excellence. He was named an Honorary Citizen of Gáldar and was awarded the *Can de Plata* for scientific merit by the Cabildo Insular de Gran Canaria. In May 2018, he was appointed Academician of the International Academy of Systems and Cybernetic Sciences (IACSYS), an institution in which only four Spaniards are represented. He was invited to lecture at numerous universities in Spain, Europe, and North America.

From 2009 onwards, he held the position of Professor Emeritus at the University of Las Palmas de Gran Canaria.

As a lecturer, Professor Moreno Díaz had a profound influence on the ULPGC. He was the first lecturer at the university to teach Systems Theory and Artificial Intelligence, and one of the first to teach Electronics. He taught at several schools, including the School of Industrial Engineering and the School of Computer Science. His contribution to doctoral education, through the supervision of PhD theses and the creation of multiple doctoral programmes, was particularly significant.

As a researcher, his pioneering contributions included the theory of logical neural networks, electronic models of neurons and neural networks, retinal models, and computer vision. His highly productive period at MIT under the supervision of Warren McCulloch led to the creation of several research groups at the Universities of Zaragoza, La Laguna, and the ULPGC. He was the driving force behind the oldest international computing conference in Spain, EUROCAST, founded in 1989, and served as its General Chairman. The conference attracted world-renowned scientists such as Franz Pichler, Luigi Ricciardi, Shunsuke Sato, and Paul Cull. In 1995, he organised a conference in honour of his mentor Warren McCulloch, attended by leading figures in cybernetics including Heinz von Foerster, Jerome Lettvin, and Michael Arbib.

Between 1970 and 1974, he was the Spanish representative to the International Federation for Information Processing (IFIP). From 1981 onwards, he served as Chairman of the Board of Directors of the European Southwest Region of the International Congress on Applied Systems Research and Cybernetics (ICASRC). He was a member of the ICASRC Executive Council for Baden-Baden between 1984 and 1986 and, from 1982, a member of the Research Council of the International Systems Institute in San Francisco, California. He also served as General Chairman or Conference Chairman of numerous international conferences and symposia held in Europe and the Americas.

In the area of knowledge transfer, Professor Moreno Díaz made an essential contribution to the creation of the Canary Islands Institute of Technology, where he promoted the establishment of several research centres in fields such as wind energy, medical informatics, and multimedia.

Finally, his commitment to international cooperation deserves special mention, particularly his role in promoting computer science studies at the ULPGC throughout Latin America. This included his involvement in the 18th Latin American Conference on Computer Science (PANEL'92) and the signing of multiple cooperation agreements with universities in Latin America and Europe.

Below, we would like to share his latest extended abstract for Eurocast, a paper he was finalizing to present at his workshop.

Towards a Systems Model of Communication for Consciousness

Roberto Moreno-Díaz

Universidad de Las Palmas de Gran Canaria

Extended Abstract

Keywords: Systems Theory · Consciousness · Communication

Consciousness is a personal experience unique to each individual, yet it can be communicated. We believe that, at present, no scientific theory can fully explain it, as its nature as a private experience prevents consensus among scientists regarding external verification and validation. It may be one of the “*ignorabimus*”, the unreachable, like the absolute zero of thermodynamics.

Despite that, we may propose reasonable models that align with widely accepted properties of consciousness, such as its communicability. Therefore, we suggest that intraspecific communication within a group, conveyed through signals, forms the basis of what we may call “consciousness”. Signal messages will develop into context dependent languages in species with greater brain processing capacity. Consciousness and social life may then be considered consequences of communication.

Restricting the argument to vertebrates, the simplest case would be solitary fishes (e.g., abyssal) with a single mating message. Vertebrates, for example, that care for their offspring can communicate as their offspring develop the ability to engage in contextual communication, thereby increasing the consciousness of both.

Also, consciousness is expected to increase as the ability to handle complex contextual messages increases, which is related to the brain’s complexity, as seen in dolphins. Additionally, it will increase in larger coherent groups, where the complexity of context dependent messages and their possible number are much greater, e.g., in human professional societies.

This interaction is represented within each individual, according to Craik’s “model of the environment for planning” [1], which is updated to include “a model of itself in the environment”, similar to a “mirror reflection” that enables “talking” to itself in natural contextual language and the generation of the concept of the infinite from finite spacetime. This corresponds to the Receptor and Generator of information in a typical communication system. A complete update of Craik’s model essentially requires the inclusion of the Emitter, which formally functions as an inverse Receptor, in addition to Craik’s Effectors or actuators. The total communication process between the individual and its environment is then depicted in the models of both the environment and the individual, using the same signal or context-dependent language.

As a communication model, the Stoics' framework (Von Domarus, McCulloch [2], based on the Lekton) is suitable for the Receptor and part of the Generator. The Stoics' Lekton (to be read) captures messages within their context – a form of contextual memory shared by both the environment model and the self-model. Visually, the Lekton resembles a book, with text and multisensory illustrations that detail the system's conscious history, including its planning throughout life. This serves as the database for current planning and self-reflection. It must be located with its position spread across the cortex, in the secondary sensory areas. It can be read by the current environment model and by the model itself, through “reading” signals.

The Emitter part of the system would be an inverted receptor that receives data from the Generator (the Models) already in the format handled by the Lekton (signal messages or context language) and sends it to the environment for other individuals to receive. Effectors for locomotion and manipulation are omitted here, although proprioceptor signals from them will feed back to the Models. The upper line of a model's diagram will be the Stoics' triadic relation in communication: facts-utterance-lekton.

As can be seen, the role of the Lekton in consciousness is crucial. Genetic material provides the unconscious structures that, for the most part, form the system's core. Lekton and its related areas for receiving and sending messages generate consciousness. Lekton is “learned” throughout the system's lifetime.

All this communication structure is built upon the autonomous, instinctive mechanisms for homeostasis, emergencies, and the selected mode of operation of the vertebrate, which may be functions of the cortex (if it exists), the cerebellum, the hypothalamus, and, mostly, the reticular formation [3].

As a first conclusion, the model points to brain areas involved in communication, language, and/or message processing and generation as the physical locus of consciousness. It also suggests how this work is done.

References

1. Craik, K. *The Nature of Explanation*. Oxford University Press (1943)
2. McCulloch, W.S. XIV Lekton. In *Communication*, edited by Lee Thayer, published by Charles C. Thomas Publisher, Springfield, Illinois, USA (1967)
3. Kilmer, W. *A Global Inhibitory Model for Decision Modes of Attention: Functions of the Nucleus Reticularis Thalami*. In *Brain Processes, Theories and Models*, edited by R. Moreno-Díaz and J. Mira-Mira. The MIT Press (1995)

Table of Contents

Systems Theory, Applications, Pioneers, and Landmarks

From National Records to Local Insights: A Data-Driven Framework for Time-Series.....	30
<i>M. Miró-Julià, R. Alberich, N. Alirio Cruz, A. Tobar Nicolau</i>	
Coordination and Communication Foundations for Agentic AI	32
<i>U.M. Borghoff, P. Bottoni, R. Pareschi</i>	
Quaternionic Observers for Multi-Channel Systems: From Smart Grids to Medical Imaging.	34
<i>M. Sebek</i>	
Analysis of D-Wave Quantum Annealers on the Example of Multi-Depot VRP for Public Transport.....	36
<i>M. Jaroszczuk, M. Uchroński</i>	
The Moon-Hopkins Billing-Machine.....	38
<i>B. Kerschbaumer</i>	

Theory and Applications of Metaheuristic Algorithms

Evaluating the Robustness of AI Image Detectors for Diffusion-Generated Images via Evolutionary Optimization of Prompt Embeddings	40
<i>M. Salvenmoser, M. Affenzeller</i>	
Domain-Informed Machine Learning by Hybridizing Symbolic Regression and Large Language Models	42
<i>D. Nguyen Duy, F. Bachinger, M. Affenzeller</i>	
Physics-Informed Knowledge Validation via LLM-Assisted Shape-Constrained Symbolic Regression.....	44
<i>F. Bachinger, S. Habringer, C. Staudinger, D. Nguyen Duy, M. Affenzeller</i>	

Meta-Modeling Power Inverters for Fast Evaluation of Energy Flow Controllers on Resource Limited Devices.....	46
<i>C. Dalpiaz, K. Preiner, P. Kastner, M. Kargl, F. Eibensteiner, J. Langer</i>	
Visualizing Populations in Multi-Objective Optimization	48
<i>B. Werth, M. Affenzeller</i>	
Shapes un-Constrained: Tracking Phenotypical Behavior in Evolutionary Symbolic Regression.....	50
<i>C. Haider, F. Bachinger, B. Werth, M. Affenzeller</i>	
Improving Genetic Programming-based Symbolic Regression with Exhaustive Methods	52
<i>L. Kammerer, D.J. Bartlett, H. Desmond, P.G. Ferreira, G. Kronberger, S. Winkler</i>	
A Learning Large Neighborhood Search for the Electric Autonomous Dial-A-Ride Problem.....	54
<i>L. Tomandl, M. Bresich, G.R. Raidl, S. Limmer</i>	
Transfer Learning using Template Seeding within Structure-Template Symbolic Regression.....	56
<i>C. Haider, J. Zenisek, P. Fleck, M. Fruhmann, C. Hambrock, M. Affenzeller</i>	
Optimizing Dynamic Dispatching in Elevator Control with a Destination Registration System	58
<i>M. Bresich, G.R. Raidl, S. Limmer</i>	
Timing of Queries in Interactive Job Scheduling.....	60
<i>J. Varga, H. Korinek, G.R. Raidl, T. Rodemann</i>	
Comparative Analysis of GRASP Variants for CVRPTW: A Case Study in Urban Organic Waste Collection.....	62
<i>M. Maarouf Bassaïdi, J.A. Brito Santana, J.A. Moreno Pérez, H. Alonso Afonso</i>	
Superstrings of Uniform-Cardinality Set Systems via the Generalized Traveling Salesperson Problem	64
<i>A. Weissenfels, E. Iurlano, G.R. Raidl</i>	

Function Libraries for Symbolic Regression with Genetic Programming	66
<i>P. Fleck, C. Haider, M. Fruhmann, J. Zenisek</i>	
Accelerators for Metaheuristics	68
<i>C. Smutnicki, Z. Banaszak, G. Bocewicz</i>	
On Solving the Multiple Variable Gapped Longest Common Subsequence Problem	70
<i>M. Djukanović, N. Balaban, C. Blum, A. Kartelj, S. Dzeroski, Z. Zebec</i>	
Improving Sign Language Recognition with Synthetic Oversampling	72
<i>M. Quesada-Ojeda, R. Holmes, A. Ventresque</i>	

Mechatronic Product Development

Electric Motor Cooling System Evaluation Methodology	74
<i>M. Golombek, T. Schlechter</i>	
Back Projection: Exterior Lighting Application for Automotive Industry	76
<i>Y. Balaji, M. Dobrača, H. Haraz, T. Schlechter</i>	
Deep Learning Based Object Recognition Approach for Motorbike Dashboard Software Validation	78
<i>S. Hussein, T. Schlechter</i>	
Digital Twin-Driven Predictive Maintenance for Intelligent Mechatronic Product Development	80
<i>D. Kirchsteiger, M. Jungwirth</i>	
Developing a Structured Database Using MariaDB to Store and Manage Real-Time Process Data Delivered by an MQTT Broker	82
<i>J. Okochi, T. Schlechter</i>	
Development of an Intel RealSense-based ArUco-Marker System for Angle Measurements on a Ski Simulator	84
<i>C. Hofer, F. Hammer, V. Putz</i>	

Sustainability Potential of Emerging Digital Technologies in Modern Regional Waste Management 86
S. Merschak, P. Hehenberger

Development and Validation of a Rotary Impact Drilling System for an Ultra-Lightweight UAV-Deployed Tree Planting Robot Enabling Autonomous Forest Restoration 88
J.L. Voigt, G. Happich, T. Weiser

Regulation-Based Innovation and its Recent Challenges in the Automotive Sector 90
D. Gingo, S. Wundsam, V. Pradeep, A. Strehar, T. Schlechter

Model-Based System Design, Verification and Simulation

The CODES Framework: A Structured Prompting Methodology to Mitigate Cognitive Bias in AI-Assisted Novice Programming 92
J. Luo, A. Ito, A. Otsuka, A. Sasaki

Support Tool for Generative AI-Based Coding 94
A. Otsuka, A. Ito

Classical and Neural Network Modeling of a Heating System in a Smart Building 96
A. Stachno, A. Czemplik

Using Recurrent Neural Networks for Prediction in Dynamic Systems with Support of Artificial Intelligence language models 98
I. Karcz-Duleba, A. Czemplik

Modeling of Thermal Energy Storage Tanks 100
A. Czemplik

Behavior Analysis for Grazing of Cows 102
B. Bayaraa, K. Nogami, N. Tsuruta, M. Hasegawa, R. Fukumori, Y. Nagao, A. Otsuka, M. Onda, Y. Hiramatsu, T. Nakano, A. Ito

Zonal Indoor User Localization in a Room based on Air Quality Measurements using Artificial Neural Networks.....	104
<i>A. Stachno</i>	
Neural Analysis of Thermal Images as the Basis of an Intelligent Presence Sensor in Building Automation	106
<i>A. Stachno, M. Suproniuk, N. Leleniewski</i>	
Optimal Inventory Control Using the Monte Carlo Method	108
<i>M. Olesiak</i>	
Media Representation of Artificial Intelligence and datafication in Smart City campaigns: A case study of Amsterdam and Singapore	110
<i>N. Starzyńska</i>	
Towards Human-Technology Symbiosis in Audience Engagement: A Web- and IoT-Based Continuous Response Digital Interface for Real-Time Musical Tension in Live Concerts	112
<i>M. Kulbacki, A. Porębski, J. Segen, Z. Chaczko, A. Ito, W. Bożejko, A. Gnatowski, R. Klempous, M. Kulbacki</i>	

Applications of Signal Processing Technology

Time-of-Flight based 3D Imaging for Endoscopic Medical Applications	114
<i>D. Doujak, T. Thurner</i>	
Towards Efficient Event-Based Signal Processing Using FIR Filters.....	116
<i>A. Werzi, B. Moser, M. Lunglmayr</i>	
Edge Processing of IMU Data for IoT PAwR Communication	118
<i>T. Kammerhofer, A.M. Stein, T. Thurner</i>	
Envelope Spectrum Analysis of Phonocardiogram Data	120
<i>D. Neussl, T. Kammerhofer, T. Thurner</i>	
Strain Sensor System for Integration in Basalt-Reinforced Concrete Segments	122
<i>L. Heindler, A. Evangelatos, T. Thurner</i>	

A Projection-Based Hybrid Model for Nonlinear Acoustic Echo Cancellation.....	124
<i>A. Nezamdoust, M. Scarpiniti, M. Huemer, D. Comminiello</i>	
Complex-Valued Spline Adaptive Hammerstein Filter	126
<i>B. Plaimer, M. Wagner, O. Lang, M. Huemer</i>	
Time-Domain Estimation of Frequency-Dependent IQ Imbalance for 5G NR-based JCAS Systems	128
<i>A. Meingassner-Lang, M. Tockner, O. Lang, B. Plaimer, M. Wagner, M. Huemer</i>	
An Approximated Lower Test Signal Generator Linearity Bound of the AST Linearity Test.....	130
<i>M. Wagner, O. Lang, M. Huemer</i>	
An Enhanced mm-Wave Tomographic Imaging Platform for Low-Permittivity Materials	132
<i>M. Koller, A. Och, M. Wagner, A. Windner, O. Lang, S. Schuster, S. Scheiblhofer, P.A. Hözl, S. Matzinger, J. Reisinger, M. Huemer</i>	
3D Point Cloud Generation Using mmWave Radar Sensors	134
<i>E. Wollitzer, M. Zeinzinger, M. Kargl, D. Hrubý, F. Eibensteiner, M. Černý, J. Langer</i>	
Cross-Domain Time-Series Analytics for Industrial Monitoring	136
<i>E.D. Ocansey, R.M. Holom, H. Hinterbichler, M. Araz, A. Krenn, V. Dino</i>	

Applied Data Science and Engineering for Intelligent Transportation Systems and Smart Mobility

Vision-Driven Autonomous Traffic Signal Control through Deep Reinforcement Learning	138
<i>D. Caballero-García-Alcaide, D. Yagüe, J.M. Armengol, M.P. Sesmero, A. Sanchis</i>	
Data Extraction and Processing for DIY Drone Forensics	140
<i>A. Attenberger, M. Schmidt, A. Hoyer, L. Hardi</i>	

Investigating Greater Manchester Tram Service Optimization Using Multi-Layer Transport Library Simulation Approach.....	142
<i>K. Han, L.A. Christie</i>	
Bringing People and Luggage to the Airport	144
<i>B. Hu, P. Stern, A. Schaffenberger, T. Hirsch, J. Berndorfer, A. Fallast, S. Lesak</i>	
Evaluating and Improving Data Management in Urban Data Platform Development Research Projects	146
<i>I. Marco-Pérez, F. Larrinaga, A. Pérez, M. Illarramendi, B. Pérez, A.L. Rubio</i>	
CLG-Smooth: Temporal Regularization for Smoother and Safer VLM-RL Driving Policies	148
<i>G. Cordova, H. Iqbal, A. Al-Kaff, B. López Boada, F. García</i>	
Data-Driven Human-AI Co-Simulation and Digital Twin Integration for Smart Urban Mobility.....	150
<i>Y. Liu, A. Aghanouri, C. Olaverri-Monreal</i>	
Improving Highway Construction Site Safety with the Traffic Monitoring and Warning System	152
<i>A. Maletzky, F. Oberhauser, R.M. Holom</i>	
Bridging the Planning-Execution Gap with a Construction Site Revision System	154
<i>F. Oberhauser, A. Maletzky, A. Alshenoudy, B. Sabrowsky-Hirsch, M. Steindl, R.M. Holom, E.D. Ocansey</i>	
Intelligent Parking Management System: An Infrastructure-Based and CARLA Simulation Approach.....	156
<i>B. Pérez, M. Resino, E. Serrano, A. Al-Kaff, F. García</i>	
It's About Time: Measuring Time Error Propagation between V2X Agents.....	158
<i>J. Godoy, A. Gil-Cuadrado, A. Al-Kaff, F. García</i>	
Comparative Evaluation of Real-Time Communication Protocols in a Real Vehicle Teleoperation Platform	160
<i>J. Gómez, A. Al-Kaff, F. García</i>	

Field Test Methodology for V2X Cooperative Perception between Heterogeneous Vehicles	162
<i>M. Resino, J. Godoy, A. Al-Kaff, F. García, J. Espelosín, L. Sanchez, T. Seco</i>	
A Learning-Based Metaheuristic for the Vehicle Routing of Multi-Compartment Silo Trucks.....	164
<i>P. Schulenburg, U. Göhner, S. Gülficici</i>	

Computer and Systems Based Methods and Electronic Tools in Clinical and Academic Medicine

Complex Phasor Representations on the Gaussian Plane to Annotate Arrhythmia Episodes in the ECG Signal.....	166
<i>J. Nikodem, R. Klempous, J. Rozenblit, K. Kluwak, M. Kulbacki, D. Zyśko, D. Jagielski</i>	
Early Detection of Alzheimer's Disease Using Deep Learning Models on MRI Data: A Study Based on ADNI.....	168
<i>B. Suárez Santiago, A. González Hernández, G.S. De Blasio, C.R. García-Rodríguez, A. Quesada-Arencibia</i>	
Framework Development for Remote Health Monitoring during Manned Space Missions	169
<i>M. Kozlovszky, O. Szabó Nagyné, A. Koleszár, P. Kristof, E. Csanák, L. Kovács, V. Némethy</i>	

Systems in Industrial Robotics, Automation and IoT

Rapid Prototyping Workflow for Robotic Gripper Development through CAD-Python Integration.....	171
<i>G. Diaz Peralta, B. Höfig, M. Glück</i>	
Automated Tool Wear Detection for Industry 4.0: Image Segmentation of Photometric Stereo Data	173
<i>T. Jäkel, F. Schirmeier</i>	
SOMP - Stiffness Optimized Motion Planer.....	175
<i>D. Seyser, M. Schebek, L. Pammer</i>	
AI-Driven Defect Detection and Adaptive Control for Winding Process Optimization	177
<i>R. Seliger, U. Göhner</i>	

End-to-End Bill of Process Generation Using Multi-Modal Transformers	179
<i>D. Ulrich, R. Seliger, F. Schirmeier</i>	
Safe Reinforcement Learning using Ideas from Model Predictive Control.....	181
<i>G. Schäfer, J. Rehrl, S. Huber, S. Hirlaender</i>	
Reducing Downtimes in Process Control Systems through the use of Digital Twins	183
<i>M. Taferner, M. Horn, S. Kranzer, H.C. Trenkler, R. Portenkirchner</i>	
Generating Approximately Periodic Multivariate Timeseries with Gaussian Processes	185
<i>E. Reich, S. Huber</i>	
Empirical Optimization of Computer Vision Model for Real-World Part Verification in Manual Order Picking: A Technology Transfer Perspective	187
<i>M.K. Siddiqui, S. Chakraborty, J. Kressel, J. Grinninger</i>	
An Embedded On-Device Training Framework for Neural Networks.....	189
<i>M. Kargl, D. Muttenthaler, F. Eibensteiner, C. Dalpiaz, P. Kastner, J. Langer</i>	
Human-Centred AI for Industrial Maintenance: Prototyping an LLM-Based Knowledge Platform	191
<i>G.A. Zwettler, M. Hanreich, K. Kloiber, J. Schönböck, J. Altmann</i>	
Finetuning LLMs on Synthetic Tasks for Improved STEP Understanding.....	193
<i>F. Lischka, F. Schirmeier, U. Göhner</i>	
Comparing Generalisation in Neural ODEs under Different Regularisation Schemes	195
<i>O. Sassnick, T. Rosenstatter, A. Unterweger, S. Huber</i>	

BLE 5.4 Periodic Advertising with Responses: Eliminating Limitations of Advertisement-Based BLE Communication	197
<i>M.A. Bawiec, A. Duda</i>	

Systems Thinking and Systems Simulation: An Interdisciplinary Approach to Modeling Complexity

What is Systems Theory: Core Concepts for Understanding and Teaching Systems Theory and Cybernetics in Further Education	199
<i>S.N. Groesser</i>	

First Empirical Evidence of Human Decision-Making in an Interactive Water-Food Simulator	201
<i>M. Schaffernicht, S.N. Groesser</i>	

Global Warming: A Systemic Strategy for Weathering Mammoth Complexity	203
<i>M. Schwaninger, S. Ott</i>	

Systems Thinking and Symbols: Reflections on Symbolic Qualities of System Archetypes	205
<i>M. Tilebein, L. Zheng</i>	

Exploring Predictors of Employee Inclusion in Hybrid Teams: A Machine Learning Analysis of Complex Work Dynamics	207
<i>S. Kerschbaumer, S. Boehm, M. Hartner-Tiefenthaler</i>	

Data-Driven Decision Support in Complex Situations- IRIS Integrated Reachback Information System: Case Study: Data-Driven Decision Support and Sector-Based Optimization with IDEA4C	209
<i>S.W. Pickl</i>	

Data Science in Medical and Bio-Informatics

Federated Machine Learning for Risk Prediction in Nursing Care	211
<i>E. Mayrhuber, P. Stampfer, S.P. Kumar Veeranki, S. Winkler</i>	

Assessing ECG Signal Quality via Patient-Specific Embedding Similarity	213
<i>M. Pletz, A. Maletzky, M. Giretzlehner</i>	
Retina-Focused Preprocessing for Enhanced Deep Learning Classification of Multiple Sclerosis from OCT Scans	215
<i>L. Álvarez-Rodríguez, B. Cordón, E. García-Martin, J. Novo, J. de Moura, M. Ortega</i>	
Tracing Clinical Pathways from Unstructured Medical Documents: A Case Study on a Standard Operating Procedure for Chest Pain.....	217
<i>N. Trajkovska, S. Kastenhofer, S. Sandler, S. Bauernfeind, C. Praschl, O. Krauss, R. Hecktor, D. Herzmanek, M. Roiss, M. Haider, M. Rechenmacher, D. Roth</i>	
Feature-Based Clustering of Shockwave Signals for Understanding Variability in Electrohydraulic Extracorporeal Shockwave Therapy.....	219
<i>H. Janout, J. Flatscher, S. Winkler, P. Slezak, C. Slezak</i>	
Towards improving Medical Image Semantic Segmentation using Model Soups.....	221
<i>J. Ortiz de Murua, J. Heras, A. Inés</i>	
A Learning-Embedded Hyper-Heuristic Framework for the Integrated Healthcare Timetabling Competition 2053.....	223
<i>K. Han, T. Papamarkou</i>	
Clinical Data Preparation for Process Mining in Hospitals with Routine Data	225
<i>K. Arthofer, C. Walcher, S. Adlberger, T. Königswieser</i>	
ASA-Specific Detection and Prediction of Complications in Post-Surgery Records of Patients	227
<i>J. Vetter, M. Strobl, G. Halmerbauer, T. Königswieser, S. Winkler</i>	
Estimation of Glycated Hemoglobin from Continuous Glucose Monitoring data by Symbolic Regression	229
<i>J.I. Hidalgo, E. Maqueda, O. Garnica, W. Sun, D. Parra, J.M. Velasco</i>	

Multi-Disease Explainable Deep Learning System for Automated Chest X-Ray Screening.....	231
<i>E. Goyanes, B. Fernández-Moreira, J. de Moura, J. Novo, M. Ortega</i>	

Data Science Approaches to Health Workforce Planning: A Scoping Review and Taxonomy	233
<i>H. Günther, K. Zeume, L. Anderle</i>	

Modeling, Simulation, and Optimization in Production and Logistics

Extending the Piloted-Beam Search with Marginal Survivors.....	235
<i>S. Habringer, F. Holzinger, S. Wagner</i>	

Robust Reaction Strategies for Worker-Flexible Job Shop Scheduling Problems under Uncertainty	237
<i>D. Hutter, T. Steinberger, M. Hellwig</i>	

Model-Based Optimisation of the Energy Efficiency of Machining Processes	239
<i>D. Bernroder, F. Spitzer, J. Giedenbacher, H. Groening</i>	

Closing the Feedback Loop: Adaptive Constraint Generation for Tree-Search Algorithms through Large Language Models.....	241
<i>S. Habringer, S. Leitner, J. Karder</i>	

Towards Less Energy Reliance - A Computational Study on Benchmark Instances for Selected Locations in Austria	243
<i>A. Kinast, D. Falkner, M. Bögl, R. Braune</i>	

An Incremental Algorithm for Optimized Container Transport and Stacking under Dynamic and Uncertain Conditions	245
<i>N. Ralfs, M. Becker</i>	

Real-Time Data Stream Processing for Interpretable Machine Learning	247
<i>P. Neuhauser, F. Holzinger, S. Wagner</i>	

"Green AI" and SW-Tools for Sustainable Energy and Materials Consumption

Towards Improving Substation Load Forecasts in Renewable-Driven Grid Environments.....	249
<i>D. Falkner, F. Oberhauser, F. Ebner, M. Affenzeller</i>	
Modeling of a DELTA-Robot for AI-Based path planning	251
<i>A. Weyermann, T. Weiser, D. Jacob</i>	
Optimizing the Energy Efficiency of a Delta Robot through AI-based Trajectory Planning	253
<i>N. Weber, D. Jacob</i>	
Generic HVAC Model Evaluation for Predictive Maintenance Application.....	255
<i>D. Sengl, E. Paráda, A. Bratukhin, M. Kobelrausch</i>	
Optimization Strategies for Deploying Symbolic Regression Models on Embedded Hardware for Energy Management.....	257
<i>A. Prielinger, C. Dalpiaz, F. Eibenstein, K. Preiner, J. Langer</i>	
Smart Forecasts for Smart Factories	259
<i>S. Arnold, G. Podann, A. Kneißler, K. Zeume, K. Celik</i>	
Sample-Efficient Pareto Front Modeling for Energy-Aware Reinforcement Learning Using Bayesian Optimization.....	261
<i>G. Schäfer, J. Rehrl, S. Huber, S. Hirlaender</i>	
Powder Bed Defect Detection in L-PBF Using YOLO RT-DETR	263
<i>P. Schulenburg, C. Buhl</i>	
Exploring Zero-Shot Foundation Models for Multivariate Time Series Anomaly Detection.....	265
<i>M. Uray, S. Messineo, R. Kwitt, S. Huber</i>	
Operator-in-Loop Validation and Annotation Framework for AI-Driven QA in Metal AM.....	267
<i>F. Waheed, K. Höfig, F. Riß</i>	

AI4Green: Harnessing the Ecological Potential of AI in Industry.....	269
<i>K. Höfig, M. Döller, F. Waheed, F. Riß, M. Scholz, J. Bauer, H. Waclawek, G. Schäfer, S. Huber, B. Heinzl, M. Hellwig, S. Finck</i>	

Stochastic Models, Statistical Methods, and Applied Systems Simulations

Shock Models Driven by a Two-State Markov Modulated Poisson Process	271
<i>N. Giordano, B. Martinucci, P. Paraggio</i>	
Varentropy-Based Analysis of Precipitation Variability.....	273
<i>A. Di Crescenzo, S. Spina, P. Villani</i>	
On the Rejection Sampling Method for Multidimensional Diffusions through a Drift-Based Transformation	275
<i>A. Di Crescenzo, V. Mustaro, S. Spina</i>	
Testing Heteroscedasticity in Diffusion Processes	277
<i>G. Albano</i>	
Some Time-inhomogeneous Pseudo-First-Order Kinetic Models in Random Environment.....	279
<i>V. Giorno, A.G. Nobile</i>	
On Time-Changed Gauss-Markov Processes, First Passage Times, and Simulations	281
<i>L. Caputo, M.F. Carfora, K. Kobayashi, E. Pirozzi</i>	
Fractional Stochastic Processes for Risk Models and Simulations.....	283
<i>E. Pirozzi</i>	

Impact of AI and Quantum Technologies on Cybersecurity

PKI on IIoT Environments: Requirements, Design and Compliance.....	285
<i>I. Pereira-Domínguez, C. Hernández-Goya</i>	

Federated Quantum Machine Learning for Privacy-Preserving Mobile Malware Detection: A Proposal	287
<i>F. Mercaldo, F. Martinelli, M. Petyx, P. Agnello, A. Santone</i>	
Low-Latency Conversational Ranker for Multi-Agent Voice Interaction.....	289
<i>M. Rodriguez-Vega, P. Caballero-Gil</i>	
Detection and Reconstruction of Traffic Signs Under Adversarial Conditions.....	291
<i>A. Barreto-Ramallo, M. Rodriguez-Vega, P. Caballero-Gil</i>	
A PQC-Improved Secure Aggregation Scheme for Federated Learning	293
<i>L. Syne, C. Hernández-Goya, P. Caballero-Gil</i>	
Anomaly Detection in Smart IoT Environments Using Classical and Quantum Machine Learning	295
<i>A. García Carrasco, C. Hernández-Goya, J.M. Molina-Gil</i>	
QUBO-Based Modeling for Biomedical Diagnosis: Application to Breast Cancer	297
<i>S. Diaz-Santos, D. Escanez-Exposito, P. Caballero-Gil, C. Caballero-Gil</i>	
Quantum Key Distribution Simulation in Rust	299
<i>D. Escanez-Exposito, P. Caballero-Gil</i>	
qcrypto: Library for Scientific Development and Testing of Quantum Cryptographic Primitives	301
<i>J. Garcia-Diaz, D. Escanez-Exposito, J. Padilla-Pio, P. Caballero-Gil</i>	
Stochastic Noise Adaptative Quantum Data Locking (SNA-QDL) Protocol.....	303
<i>Y. Cabrera-Aparicio, J. Garcia-Diaz, P. Caballero-Gil</i>	
Indoor Positioning and Localization	
The Impact of BLE channel on RSSI-Based Ranging.....	305
<i>J. Pietrucha, M. Nikodem</i>	

Direction-Finding of IoT devices method combining Amplitude and Phase monopulse comparison functions 307
J.A. López-Pastor, M. Hutar, P. Brida, J.L. Gomez-Tornero

Analysis of Tools for Measuring Wi-Fi Network Performance with an Emphasis on Secure Indoor Localization and Coverage Optimization 309
I. Bridova, J. Machaj, P. Brida

Classical AI Search Algorithms Analysis for Optimized SAR Trajectory 311
R. Aguasca-Colomo, J.M. Molina-Gil, M.A. Jiménez-Suárez, F.J. Navarro-Amador

Blackbox Systems for Generative AI - Explainability and Applications

Detecting Object Consistency Errors in Reference-Conditioned Image Generation 313
A. Maschek, R. Serban, A. Stöckl, D.C. Schedl

Evaluating Agentic Web Behavior on Financial Websites: A Case Study with Raiffeisen Landesbank Oberösterreich 315
J. Nitu, L. Kühhas, M. Ehrengruber, A. Stöckl

Improving LIME with Logistic Surrogate Models 317
S. Dreiseitl

Investigating Memorization and Generalization of LLMs Using Chess as a Testbed 319
A. Stöckl

Guided Code Generation with Large Language Models and Static Code Analysis 321
M. Dalkilic, C. Praschl

Synthetic UAV Imagery for Ecological Monitoring Using Generative Artificial Intelligence 323
C. Praschl, M.T. Karbasch, K. Probst, D.C. Schedl

Explainable and Privacy-Preserving AI for Healthcare: A Local GraphRAG Pipeline for Nursing Care Guidelines 325
O. Krauss, A. Mileva

Combining Traditional Machine Learning and Generative AI to Create a Trustworthy Overview of Austrian Online Shops 327
S. Sandler, O. Krauss

Navigating the Sea of LLM Evaluation: Investigating Bias in Toxicity Benchmarks 329
R. Gugg, S. Niederländer, A. Stöckl, M. Flechl

Explaining GenAI Endurance Training Plans Through Post-Hoc Behavioral Evaluation 331
L. Voggeneder, J. Krösche, M. Kurz, E. Sonnleitner, S. Selinger

Cybermedical Systems as Human-Centered AI and Digital Medicine Applications

xML CardioTwin: A Multimodal, Explainable Cardiac Digital-Twin Pipeline for Personalised Atrial Fibrillation Ablation and Cardioneuroablation 333
M. Kulbacki, J. Segen, Z. Chaczko, J. Nikodem, R. Klempous, E. Szlachcic, J. Rozenblit, P. Skoczynski, D. Jagielski

Immersive Tourniquet Training 335
K. Buchenrieder, U.M. Borghoff

A Holter Smart Patch for Patients with Transient Cardiac Events 337
M. Silaci, K. Buchenrieder

FLOD: a Framework for Explainable Federated Learning Object Detection 339
F. Mercaldo, L. Ciardiello, M. Petyx, P. Agnello, M. Cesarelli, F. Martinelli, A. Santone

Primary and Adjacent Zooming Areas in 3D Medical Image Analysis 341
M. Vincze, A. Benhamida, S. Burian, M. Kozlovszky

Rosetta Stone for Immune Toxicity Intelligence: Explainable Machine Learning Across Immune Checkpoint Inhibitors, Radionuclide Therapies and Antibody Drug Conjugates in Real-World Oncology.....	343
<i>M. Kulbacki, Z. Chaczko, J. Segen, A. Maciejczyk, I. Pawlak, P. Zawadzki, B. Cybulska-Stopa</i>	

Innovative Teaching Approaches in Computer Science and Automation: Methods, Tools, and Experiences

Applying Machine Learning and Automation to Enhance Student Support in Higher Education	345
<i>K. Han, Y. Chen, C. Seager</i>	
Digital Learning Factory As A Workshop Format: Low-Threshold Introduction To Industry 4.0	347
<i>N. Angel, H. Orsolits</i>	
Integrating Makeathons and Problem-Based Learning in MINT Degree Programs: Curriculum Design and Fair Assessment Strategies.....	349
<i>S. Arnold, A. Kneißler, K. Zeume</i>	
Embedded Engineering Experiences: Impact on Learning Efficiency - Smart Packaging Technologies Pilot.....	351
<i>M. Glück, B. Höfig, P. Eichinger</i>	
Understanding Modern Data Architectures Through Visualized NoSQL Cluster Simulations	353
<i>K. Zeume, M. Ayas, D. Stroz</i>	
Design Thinking and Project Based Learning as a Synergetic Approach to Learning	355
<i>H.C. Trenkler, S. Kranzer, M. Horn, M. Taferner, R. Portenkirchner</i>	
Teaching Computational Thinking and Coding in an Extracurricular Context: An Assessment of the Status Quo	357
<i>L. Schachner</i>	

AI Augmented Virtual Student Exchange for 24.7 Global Engineering Education	359
<i>Z. Chaczko, W. Bożejko, I.J. Rudas, A. Ito, R. Klempous, J. Nikodem, J. Rozenblit, M. Kulbacki, J. Segen</i>	
Author Index	361

From National Records to Local Insights: A Data-Driven Framework for Time-Series

Margaret Miró-Julià^[0000-0002-9300-8698], Ricardo Alberich^[0000-0002-0707-8607],
 Nelson Alirio Cruz^[0000-0002-7370-5111], and Adrián Tobar
 Nicolau^[0000-0003-0198-2475]

Departament de Ciències Matemàtiques i Informàtica, Universitat de les Illes Balears,
 07122 Palma, Spain
 margaret.miro@uib.es, r.alberich@uib.es, nelson-aliro.cruz@uib.es,
 adrian.tobar@uib.es

Extended abstract

The Consumer Price Index (CPI) is a statistical indicator that measures the evolution of prices for goods and services consumed by households residing in Spain. The composition of the shopping basket is derived from household consumption patterns, and the relative importance of each good or service in the CPI calculation is determined by its share in total consumption.

The Instituto Nacional de Estadística (INE) provides extensive CPI-related information through data tables and time-series datasets at monthly and annual frequencies. Most of this information is reported at the Autonomous Community level, both in aggregate form and broken down by ECOICOP (European Classification of Individual Consumption by Purpose) groups.

Understanding consumer profiles is a key asset for decision-making processes across economic sectors. The objective of this paper is to examine and evaluate the use of various Statistical and Data Mining techniques to identify which approaches are most suitable for analyzing consumer behavior. Furthermore, we aim to extend this knowledge by integrating methods that enable the translation of Autonomous Community data to more granular geographical units, such as municipalities or sub-regional areas. Specifically, we propose a model for the disaggregation of the CPI time series using local economic indicators as covariates. Within this context, the model incorporates:

Compositional Data Analysis. In Statistics, Compositional Data Analysis (CoDa) [1] deals with vectors of strictly positive components representing proportions, percentages, concentrations or frequencies of a whole. They are constrained by constant sum and widely used in domains such as chemistry, econometrics and food industry. Since the total number of consumers is fixed, compositional data analysis can be applied. The important aspect in compositional data is not the specific value of each component, but the relative ratio of each component.

Benchmarking. Benchmarking systematically compares indicators or outcomes across units (countries, regions, institutions) to identify best practices, gaps and

improvements [2]. In official statistics, benchmarking aligns short-term or regional indicators with more reliable annual or national aggregates while preserving key dynamics.

Anchorage. The anchorage (anchoring point) is the reference level used in a benchmarking or calibration procedure [3]. In empirical disaggregation or reconciliation frameworks, the anchoring aggregate provides the coherence constraint.

Reconciliation of Temporal and Hierarchical Series. Temporal and hierarchical reconciliation techniques adjust forecasts or estimates across multiple aggregation levels to ensure mutual coherence [4]. Given a set of series linked by aggregation constraints (municipalities → provinces → regions → country), independently estimated models may violate additivity. Reconciliation methods modify forecasts to satisfy constraints with minimal distortion.

Data Sources Used as Covariates. The following indicators will be used to disaggregate the time series: Industrial Production Index (monthly indicator of the volume of industrial production); Population (annual municipal register of resident population); Industrial Energy Consumption (survey on energy consumption by industrial firms); Business Demography (monthly data on company creations, dissolutions and capital changes); Household Food Expenditure (annual series of household food consumption); and Labor Force and Activity Rates (quarterly data on economically active population).

Using official INE webpages (<https://www.ine.es>), raw data can be efficiently collected and stored within a database. Such datasets often exhibit a time-series structure, typically consisting of statistical measurements across various time periods. Effective data preprocessing is essential to ensure the reliability of the insights obtained, and selecting an appropriate data transformation method can significantly improve model performance. The integration of statistical and data-mining techniques facilitates the analysis of complex domain-specific problems. In this paper, we propose a comprehensive data preprocessing approach alongside an analytical model that consolidates observations at the Autonomous Community level and disaggregates them to a local level using economic indicators as covariates. The effectiveness of the proposed model will be evaluated, and the corresponding results will be presented and discussed.

References

1. Aitchinson, J. The Statistical Analysis of Compositional Data. *Journal of the Royal Statistical Society. Series B (Methodological)*, **44**(2), 139-177, (1982). <https://doi.org/10.1111/j.2517-6161.1982.tb01195.x>
2. Kütz, Martin: *The Art of Benchmarking: An Analytical Guide to Methods and Applications for IT Management*, Springer Cham, Switzerland, 2022. <https://doi.org/10.1007/978-3-030-90586-6>
3. Furnham, A., and Boo, H. C. A literature review of the anchoring effect. *The Journal of Socio-Economics*, **40** (1): 35-42, (2011). <https://doi.org/10.1016/j.sococ.2010.10.008>
4. Shi, K., Ding, Z., Chen, Z. et al. A multi-scale time series forecasting framework with temporal hierarchical information fusion and reconciliation. *Data Mining Knowledge Discovery* **39**, 32 (2025). <https://doi.org/10.1007/s10618-025-01103-9>

Coordination and Communication Foundations for Agentic AI (Extended Abstract)

Uwe M. Borghoff^{1*}, Paolo Bottoni^{2**}, and Remo Pareschi³

¹ Computer Science, University of the Bundeswehr Munich, Neubiberg, Germany

² Computer Science, Sapienza University of Rome, Rome, Italy

³ Stake Lab, University of Molise, Campobasso, Italy

Coordination and communication have long been central themes in distributed computation and multi-agent systems. We show how these themes take new form in Agentic AI, where intelligent behavior emerges through structured interaction among diverse components such as autonomous agents, humans, and LLMs functioning as collaborators. Classical paradigms in coordination languages and concurrent logic continue to provide durable foundations for these architectures.

LINDA [8] introduced the notion of a *tuple space*—a shared associative memory enabling generative communication through simple primitives (*out*, *in*, *rd*, *eval*). Processes interact by producing and consuming tuples, decoupled in both time and space. The concept persists in modern multi-agent environments and AI orchestration frameworks where asynchronous interaction occurs through shared data layers, blackboards, or vector databases. In contrast, blackboard architectures [7] exemplify centralized coordination. Its metaphor of agents writing and reading from a shared blackboard remains instructive for agentic systems that integrate diverse tools or LLM modules, where global state coordination and interpretability are essential. Linear Logic [9] provided a resource-sensitive semantics later embodied in the Linear Objects language (LO) [3], where objects evolve by consuming and producing facts according to linear rules. Its awareness of resources anticipates constraints central to coordination, such as bounded memory, limited context, and transactional consistency. These ideas supported the development of the Constraint-Based Knowledge Brokers [2], which unified querying across heterogeneous knowledge sources [6] using signed feature constraints [5]. Such architectures prefigure modern federated or multilingual AI systems coordinating partial information across distributed agents. The Concurrent Logical Framework (CLF) [11] extended linear logic with constructs to represent concurrency and temporal independence. It models concurrency at the level of logical inference, treating execution traces as proofs. This framework foreshadows formalisms for verifying communication correctness and ensuring compositional behavior in multi-agent coordination. Concurrent Constraint Programming (CCP) [10] introduced a shared constraint store through which agents interact by *tell* and *ask* operations. The store is monotonic, with constraints that accumulate rather than retract, reflecting incremental knowledge construction. This model

* Corresponding author: uwe.borghoff@unibw.de

** Paolo Bottoni was partly supported by the Italian Ministry of University and Research (MUR) under PRIN grant B87G22000450001 (PINPOINT).

Table 1: Relevance of coordination paradigms to Agentic AI communication.

Model	Coordination	Communication	Relevance to Agentic AI
LINDA	Decentralized shared store	Asynchronous, generative	Loosely coupled collaboration and indirect data sharing
Black-board	Centralized shared state	Controlled, opportunistic	Global state coordination and transparent reasoning
LO	Logical agents using linear logic	Message-driven, resource-aware	Communication as computation and state evolution
CLF	Logical framework for concurrency	Typed, concurrent semantics	Verified interleaving and compositional reasoning
CCP	Constraint-based shared store	Guarded, declarative synchronization	Coordination by entailment and knowledge accumulation

blends logic and concurrency: agents proceed when their preconditions are entailed, enabling asynchronous yet consistent coordination. In modern terms, it corresponds to topic-driven agentic behavior [4], where each contribution refines a shared semantic space and synchronization follows from entailment. Agha's actor model [1] can be seen as a synthesis of these extremes: a decentralized network of agents that communicate asynchronously, with each encapsulating its own state and process, serving as a prototype for modern agents.

Modeling coordination and communication in agentic AI will support defining and comparing strategies for their effective integration into real-world processes.

References

1. Agha, G.: Actors: a model of concurrent computation in distributed systems. MIT press (1986)
2. Andreoli, J., Borghoff, U.M., Pareschi, R.: The constraint-based knowledge broker model: Semantics, implementation and analysis. *Journal of Symbolic Computation* **21**(4), 635–667 (1996). <https://doi.org/10.1006/jsco.1996.0035>
3. Andreoli, J.M., Pareschi, R.: Linear objects: Logical processes with built-in inheritance. *New Generation Computing* **9**, 445–473 (1991)
4. Borghoff, U.M., Bottoni, P., Pareschi, R.: Human-artificial interaction in the age of agentic AI: A system-theoretical approach. *Frontiers in Human Dynamics* **7** (2025)
5. Borghoff, U.M., Pareschi, R., Fontana, F.A., Formato, F.: Constraint-based protocols for distributed problem solving. *Sci. Comput. Program.* **30**, 201–225 (1998)
6. Borghoff, U.M., Schlichter, J.H.: On combining the knowledge of heterogeneous information repositories. *J. Univers. Comput. Sci.* **2**(7), 515–532 (1996)
7. Fennell, R.D., Lesser, V.R.: Parallelism in artificial intelligence problem solving: A case study of hearsay II. *IEEE Trans. Computers* **26**(2), 98–111 (1977)
8. Gelernter, D.: Generative communication in Linda. *ACM Transactions on Programming Languages and Systems (TOPLAS)* **7**(1), 80–112 (1985)
9. Girard, J.: Linear logic. *Theoretical Computer Science* **50**, 1–102 (1987)
10. Saraswat, V.A., Rinard, M.C.: Concurrent constraint programming. In: 17th ACM Symposium on Principles of Programming Languages. pp. 232–245. ACM (1990)
11. Watkins, K., et al.: A concurrent logical framework: The propositional fragment. In: TYPES '03. pp. 355–377. LNCS **3085**, Springer (2003)

Quaternionic Observers for Multi-Channel Systems: From Smart Grids to Medical Imaging.

Michael Sebek^[0000-0003-0927-2988]

Czech Technical University in Prague, Faculty of Electrical Engineering,
Department of Control Engineering, Prague, Czech Republic
michael.sebek@fel.cvut.cz

Abstract. Building on recent advances in quaternionic observer design, we show how to deploy the methods across industries. Representing four related signals as a single quaternion enables consistent estimation across coupled channels. We illustrate the approach with practical case studies in power systems, wireless and optical communications, imaging and video, industrial sensing and structural monitoring, medical imaging, and quantum control.

Keywords: Quaternions · Observer design · Multi-channel estimation

A quaternion is a *hypercomplex* number $q = a + bi + cj + dk \in \mathbb{H}$ with $a, b, c, d \in \mathbb{R}$ and $i^2 = j^2 = k^2 = ijk = -1$. Like \mathbb{R} and \mathbb{C} , \mathbb{H} is a division algebra, but multiplication is noncommutative; many determinant-based shortcuts (characteristic and minimal polynomials, Cayley–Hamilton) do not carry over verbatim.

Beyond rigid-body attitude and graphics, \mathbb{H} conveniently packs *four* coupled real channels with fixed cross-axis mixing, e.g., polarization parameters (dual-polarization optical and wireless links), the two orthogonal components of the Clarke transform plus a neutral and zero-sequence channel (three-phase power), tri-axial signals with a bias (IMUs and vibration monitoring), RGB with a slowly varying illumination term (vision and remote sensing), small-tip coherences (MRI and NMR), short-window coherences (quantum control and spin systems), biopotential pairs (EEG in clinical and BCI recordings), and multi-sensor fusion (robotics and IoT).

Additional uses include quad-polarization radar and SAR returns in Earth observation, B-format ambisonics in spatial audio and XR, Bayer RGGB quads in computational photography, multi-wavelength PPG and fNIRS in digital health, 3-axis IMU data with a thermal and bias channel in drones and wearables, LiDAR range–intensity–ambient–confidence stacks in autonomous driving, and tri-axial magnetometry with an auxiliary calibration channel in geophysics.

Quaternions also surface in AI: quaternion layers for RGB, IQ and ambisonics, quaternion attention for multimodal fusion, unit- \mathbb{H} priors for 3D pose, SLAM and NeRF, and \mathbb{S}^3 generative models for robotics and XR.

While other representations (e.g., \mathbb{R}^4 vectors) exist, we adopt the hypercomplex view and work directly with hypercomplex states and operators. This makes

left/right order explicit, avoids embedding artifacts, and keeps pure-quaternion factors and similarity classes visible for analysis and design.

Working natively in \mathbb{H} keeps the four channels phase- and gain-coherent and allows simple cross-channel couplings otherwise block-structured in \mathbb{R}^4 . This motivates quaternionic LTI observers that mirror the complex case while respecting right-eigenvalue similarity classes and adjoint duality [6]. We use the right-module convention and right-eigenvalue (similarity) classes throughout.

Following [7], we consider a SISO continuous-time quaternionic state-space model $\dot{x} = Ax + Bu$, $y = Cx + Du$, where all matrices and signals live in \mathbb{H} . A full-order quaternionic observer $\dot{\hat{x}} = A\hat{x} + Bu + L(y - C\hat{x})$, yields error dynamics $\dot{x} - \dot{\hat{x}} = (A - LC)(x - \hat{x})$, see [7]. We choose a quaternion vector L so that the right spectrum (right-eigenvalue similarity classes) of $A - LC$ lies in the open left half-plane; the discrete-time case is analogous.

The paper's *primary contribution* is the *practical design and use* of quaternionic observers, demonstrated through *case studies* drawn from multiple fields. Representative cases to be presented at the conference include:

Three-phase grids: Grid voltages and currents decompose into two orthogonal components plus a zero sequence; packing $(\alpha, \beta, 0, \text{aux})$ in \mathbb{H} enables coherent PLL and phasor tracking and robust frequency estimation under mild unbalance [1].

Per-pixel color video: Each pixel's RGB evolves smoothly; a damped second-order temporal model per channel, packed as (R, G, B, aux) supports denoising and deblurring with fixed cross-channel mixing [2].

Small-tip MRI and NMR voxel: For small flip angles, magnetization is nearly linear; linearized Bloch dynamics over short windows with quaternion multipliers capture phase and gain for estimation and sequence design [4].

Two-spin qubit network: Two weakly coupled spins behave as coupled oscillators; model reduction to a second-order LTI with small rotating-frame offsets yields a compact observer target [5].

Spacecraft thermal panel: A lumped two-node model (front and back panel temperatures) captures conduction between layers and radiative exchange; linearization about an operating point gives an LTI plant. Bundling $(T_{\text{front}}, T_{\text{back}}, q_{\text{in}}, \text{aux})$ in $y \in \mathbb{H}$ enables coherent estimation of temperatures and heat input with one observer [3].

1. Akagi, H., Watanabe, E.H., Aredes, M.: *Instantaneous Power Theory and Applications* (2nd ed.). Wiley-IEEE (2017).
2. Ell, T.A., Sangwine, S.J.: *Quaternion Fourier Transforms for Signal and Image Processing*. Wiley (2013).
3. D. G. Gilmore (ed.), *Spacecraft Thermal Control Handbook, Volume I: Fundamental Technologies*, 2nd ed., The Aerospace Press, El Segundo, CA, 2002.
4. Haacke, E.M., Brown, R.W., Thompson, M.R., Venkatesan, R.: *Magnetic Resonance Imaging* (2nd ed.). Wiley (2014).
5. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information* (10th Anniversary ed.). CUP (2010).
6. Rodman, L.: *Topics in Quaternion Linear Algebra*. Princeton (2014).
7. Sebek, M: Observer Design over Hypercomplex Quaternions. *Submitted to the ECC 2026*, Reykjavik, Iceland. DOI: 10.48550/arXiv.2510.15598

Analysis of D-Wave Quantum Annealers on the example of multi-depot VRP for public transport

Michał Jaroszczuk¹[0009–0009–8617–889X] and Mariusz Uchroński^{1,2}[0000–0002–9185–1841]

¹ Wrocław University of Science and Technology
Wyb. Wyspiańskiego 27, 50-370 Wrocław, Poland
michal.jaroszczuk@pwr.edu.pl

² Wrocław Centre for Networking and Supercomputing
Wyb. Wyspiańskiego 27, 50-370 Wrocław, Poland
mariusz.uchronski@pwr.edu.pl

Extended Abstract

Well-organized public transport is the basis for the proper functioning of large urban centers. Managing a good technical condition of vehicles and infrastructure can be a challenge for highly developed transport networks. Typically, the vehicle moves between two types of location: depots and terminal loops. The depots serve as overnight parking with service facilities, whereas in the terminal loops, drivers can take a break or swap with another employee at the end of the shift.

Each terminal loop must have appropriate rest and refreshment facilities, required by the law, which must be maintained in proper condition. This is typically provided by cleaning crews, whose starting point is the depot, where they begin their work by collecting the appropriate equipment and then travel between loops to perform their duties. The service time at a given station and the travel time between them are usually fixed, as crews do not work during peak traffic hours. The order of loop visits is predominantly once established, used daily and only changed in exceptional circumstances.

Optimizing routes between terminal loops, assigned to the nearest depots, can be modeled as Multi-Depot Capacitated Vehicle Routing Problem (MDCVRP), where the vehicle capacity corresponds to the maximum number of services that crew can handle. Optimization can be performed using traditional computers and algorithms, but due to the development of quantum technologies, it is worth examining whether the problem can be solved with QPU. There are several studies, such as [1], [2], which present the possibility of practical application of quantum annealers in VRP optimization, but do not provide specific QUBO models. The review also found works on quantum optimization in application to public transport: [3], [4], [5], but none of them concern service crews.

Currently one of the most promising quantum solutions which can be used for discrete optimization purposes, according to [6], are D-Wave quantum annealers. It is possible to use a hybrid approach called Leap, but if the calculations have to be run only natively, it can be chosen from two quantum processors: Advantage in two versions 4.1 and 6.4 and the latest Advantage 2 QPU [7]. However,

regardless of the processor used, the principle of operation of the quantum annealer remains the same. By changes of transverse field it is necessary to move from the initial Hamiltonian H_0 to the final Ising Hamiltonian H_P , whose minimal energy configuration corresponds to the best solution of the given problem [8]. The problem to be solved must be modeled using QUBO (Quadratic Unconstrained Binary Optimization), which requires presenting constraints in the objective function [9]. QUBO decision variables represent then a real structure of qubits in quantum processor architecture.

In the paper, the QUBO formulation of a mentioned MDCVRP problem is given with a description of the model properties. A practical appliance is shown, as the experiments are carried out on real dataset from Wroclaw public transport provider. The research focuses on two aspects. Firstly, it is needed to check the results obtained to show that Quantum Annealing technology may be used to solve VRP domain problems that have a practical appliance. Secondly, the analysis of D-Wave QPUs is made to find out the differences and to check if a new Advantage 2 unit may lead to better results, compared to previous solutions.

References

- [1] S. Yarkoni et al. "Quantum annealing for industry applications: introduction and review". In: *Reports on Progress in Physics* 85 (2022), p. 104001. DOI: 10.1088/1361-6633/ac8c54.
- [2] Shriyank Somvanshi et al. "Quantum Computing in Transportation Engineering: A Survey". In: (2025).
- [3] S. Yarkoni et al. "Quantum Shuttle: traffic navigation with Quantum computing". In: *APEQS* (2020), pp. 22–30. DOI: 10.1145/3412451.3428500.
- [4] S. Yu and T. Nabil. "Applying the Hubbard-Stratonovich Transformation to Solve Scheduling Problems Under Inequality Constraints With Quantum Annealing". In: *Front. Phys.* 9 (2021). DOI: 10.3389/fphy.2021.730685.
- [5] Ravi Jagirdar. "Quantum Leap for Mobility: Revolutionizing Transportation with Quantum Computing". In: *Journal of Research in Engineering and Computer Sciences* 2 (2024), pp. 41–48. DOI: 10.63002/jrecks.24.617.
- [6] E. Osaba, E. Villar-Rodriguez, and I. Oregi. "A Systematic Literature Review of Quantum Computing for Routing Problems". In: *IEEE Access* 10 (2022), pp. 55805–55817. DOI: 10.1109/ACCESS.2022.3177790.
- [7] D-Wave Company. *Advantage 2 Quantum Computer*. https://www.dwavequantum.com/media/3ikjqkyw/d-wave-advantage2-data-sheet-final-5_16_25.pdf. Accessed: 2025-07-26. 2025.
- [8] S. Feld et al. "A Hybrid Solution Method for the Capacitated Vehicle Routing Problem Using a Quantum Annealer". In: *Frontiers in ICT* 6 (2019). DOI: <https://doi.org/10.3389/fict.2019.00013>.
- [9] Ramkumar Harikrishnakumar et al. "A Quantum Annealing Approach for Dynamic Multi-Depot Capacitated Vehicle Routing Problem". In: *ArXiv* (2020). DOI: <http://dx.doi.org/10.48550/arXiv.2005.12478>.

The Moon-Hopkins Billing-Machine

Berthold Kerschbaumer^[1]

^[1] School of Informatics, Communications and Media - Campus Hagenberg
 University of Applied Sciences Upper Austria
 Softwarepark 11, 4232 Hagenberg, Austria
 Berthold.kerschbaumer@fh-hagenberg.at

Keywords: billing machine, computing and writing machine.

Extended Abstract

At the beginning of the 20th century, the Moon-Hopkins computing and writing machine was the first successful attempt to bring a billing machine onto the market. In this context, a billing machine is understood to be a machine that combines a (three/four-species) calculating machine with a typewriter. In comparison, accounting machines use counters to add up columns or rows, whereas billing machines can also perform multiplications and store them in different registers. The challenge was to combine both machines with their full functionality and to output the calculated values as well as the texts on a common paper carrier. The solution to this task resulted in a machine of high technical complexity and was reflected in a corresponding price. As the market for its use was also limited, it was always a niche product that remained unrivalled for decades. In addition, patent disputes delayed its rapid market launch.

Standard literature [1] cites Hubert Hopkins (1859–1930) from St. Louis, Missouri, as the (sole) inventor, although his older brother, Williams Wallace Hopkins (1850–1916), was also granted several patents. The machine was ready for production as early as 1902, and a prototype was presented to the public in St. Louis in October.[2]

The first patent for a combined writing and calculating machine (US1049093A) was filed by Hubert Hopkins on 9 June 1902 but was not granted until 31 December 1912. The visible-writing, type-bar machine described in the patent is very similar to the Oliver typewriter with its unique U-shaped type-bars. The “Addograph Manufacturing Company” in Missouri is named as the submitting company. In the next patent filed (US1039130) on 24 January 1903, changes were made primarily to the input unit of the calculating machine; the writing unit remained virtually unchanged. The patent (US844519) dated 12 May 1904 introduces significant changes to the typewriter part, which represent a technological step backwards with the use of a blind-writing, upstrike machine. One possible explanation for this could be the expiry of patent protection for this class of machines. The patent was filed by the “Moon Hopkins Manufacturing Company” of St. Louis, Missouri, and William W. Hopkins, the older brother of Hubert Hopkins. The changes made to the calculating mechanism in the patent (US1133029)

laid the foundation for the basic design of the following patents. In Austria, too, a patent application (AT37748) was filed on 17 August 1907 for a combined writing and calculating machine in which „beim Ausschreiben von Rechnungen, Aufstellen von Tabellen,... auf dem im Papierschlitten befindlichen Papierbogen außer den Zahlen auch die erforderlichen Worte, Phrasen ... abgedruckt werden können.“ (‘when writing out invoices, drawing up tables, ... the necessary words, phrases ... can be printed on the sheet of paper in the paper carriage in addition to the numbers.’)

In 1907 the “Moon-Hopkins Billing Machine Company” has been incorporated at St. Louis to manufacture and sell typewriters, adding machines and other mechanical apparatus for performing mathematical calculations and recording same. Hubert and William Hopkins were the major shareholders.[3] Until 1912, the machine was discussed and presented in trade journals [4,5], but since the granting of the patent in the USA was delayed until 1912 (in Austria, for example, the patent was granted on 1 January 1909), no advertising took place and only a few machines are likely to have been produced. The reason for this could be a patent dispute with the Burroughs Adding Machine Company. In 1902, the Hopkins brothers (50% share) and financier James Lewis Dalton founded the “Dalton Adding Machine Company” to produce a ten-key adding machine. Just one year later, Hubert Hopkins sold his shares, which then passed into the ownership of the Burroughs Adding Machine Co. in 1905. As this also transferred the patent rights to essential parts of the calculating machine, the above-mentioned patent dispute arose, which was only resolved in 1921 through the sale to Burroughs.[6] From then on, the machine was marketed under the name Burroughs-Moon-Hopkins and was used in businesses until the 1960s.[7]

This article aims to highlight the pioneering work of the Hopkins brothers, who developed the first usable billing machine and reassesses their contribution to early office automation technologies, which laid the groundwork for later integrated calculating-writing systems.

References

1. Martin, E.: Die Rechenmaschinen und ihre Entwicklungsgeschichte. 1st edn. Johannes Mayer, Pappenheim (1925)
2. Martin, E.: Die Schreibmaschine und ihre Entwicklungsgeschichte. 5th edn. Johannes Mayer, Pappenheim (1934)
3. Typewriter Topics. Vol. V, No.3, Mar. 1907, The Typewriter Publishing Co., New York, (1907)
4. Schreibmaschinen-Zeitung, Jul. 15.1909, p.217, Hamburg (1909)
5. Büro-Bedarf-Rundschau, Nov. 15.1911, p.769, Berlin (1911)
6. Business Equipment Topics. Vol. XLIX, No.1, Sept. 1921, The Business Equipment Publishing Co., New York, (1921)
7. Frank, H.: Von der Burroughs-Moon-Hopkins zum Unisys-Aquanta. Historische Bürowelt **52**, 4–8 (1998)

Evaluating the Robustness of AI Image Detectors for Diffusion-Generated Images via Evolutionary Optimization of Prompt Embeddings

Marcel Salvenmoser¹[0009-0001-9772-718X] and Michael Affenzeller¹[0000-0001-5692-5940]

University of Applied Sciences Upper Austria, Austria
 {marcel.salvenmoser,michael.affenzeller}@fh-ooe.at
<https://heal.heuristiclab.com/team>

Keywords: Evolutionary Computing · Diffusion Models · Adversarial Image Generation · Prompt Embeddings · AI Image Detection · Robustness Evaluation

Extended abstract

Diffusion Models have become a popular method for generating synthesized images through text conditioning [3]. As they reach photorealistic quality, distinguishing between human-made and AI-generated content has become increasingly difficult. This motivates the development of AI-image detectors, but their robustness across different diffusion architectures and reliability remains largely unclear. To address this, we evaluate open AI-image detectors through automated adversarial optimization using *Evolutionary Algorithms (EAs)*. Our method operates directly in the *prompt embedding space* of a Diffusion Model rather than on textual inputs. Each embedding represents a latent vector encoding the semantic conditioning for image generation. The objective is a robustness evaluation, not the development of tools that intentionally evade detection.

Method. For image generation, we use *SDXL-Turbo* [6] for its high sampling speed and strong visual quality. Its text-conditioning tensors are treated as high-dimensional search vectors. An EA applies mutation, crossover, and selection to these embeddings (*genotype*), whose decoded images (*phenotype*) are scored by an AI-image detector. This work extends prior single-objective studies on aesthetic optimization in embedding space [5], building on the principles of interpolation and local search in diffusion conditioning [4]. A population of embeddings is initialized randomly and evolved through variation and selection. Two optimization tasks are considered:

1. **Unconstrained search:** find any embeddings that produce images classified as non-AI, exploring regions of the embedding space that detectors misclassify.
2. **Prompt-constrained search:** enforce a given topic or style while minimizing detector confidence, to test whether semantically valid images can remain undetected.

Preliminary results. Using the public *Organika/sdxl-detector* [2] model from Hugging Face, we initialized a population of 100 embeddings whose generated images were initially detected as AI with over 90%+ probability. After approximately 25 generations, the EA discovered individuals with detection probabilities as low as 1%. In prompt-constrained runs, convergence was slower but still achieved around 1% detection probability after roughly 100 generations. Both experiments show that, for this detector, it is possible to reduce detection confidence solely by optimizing prompt embeddings.

Ongoing work. We extend this study to additional open detectors with accompanying papers, for example: DIRE (Diffusion Reconstruction Error) [7], which measures the reconstruction residual between an image and its diffusion-based inversion and SSP (A Single Simple Patch) [1], which focuses on high-frequency patch fingerprint. These detectors capture complementary mechanisms and reconstruction-based, patch-level allowing comparison of how different architectures respond to embedding perturbations. Furthermore, this also opens up the possibility of multi-objective-optimization by using multiple detectors on a single image result.

Contributions. This paper contributes to understanding and improving the robustness of AI-image detectors. It demonstrates that evolutionary search in the prompt embedding space can systematically expose weaknesses of existing detectors and provide a reproducible benchmark for robustness evaluation in diffusion-based image generation.

References

1. Chen, J., Yao, J., Niu, L.: A single simple patch is all you need for ai-generated image detection (2024), <https://arxiv.org/abs/2402.01123>
2. Organika/sdxl-detector · Hugging Face, <https://huggingface.co/Organika/sdxl-detector>
3. Po, R., Wang, Y., Golyanik, V., Aberman, K., Barron, J., Bermano, A.H., Chan, E.R., Dekel, T., Holynski, A., Kanazawa, A., Liu, C.K., Liu, L., Mildenhall, B., Nießner, M., Ommer, B., Theobalt, C., Wonka, P., Wetzstein, G.: State of the Art on Diffusion Models for Visual Computing . <https://doi.org/10.48550/arxiv.2310.07204>
4. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical Text-Conditional Image Generation with CLIP Latents . <https://doi.org/10.48550/arxiv.2204.06125>
5. Salvenmoser, M., Affenzeller, M.: Evolving the embedding space of diffusion models in the field of visual arts. In: Machado, P., Johnson, C., Santos, I. (eds.) Artificial Intelligence in Music, Sound, Art and Design. pp. 402–416. Springer Nature Switzerland, Cham (2025)
6. Sauer, A., Lorenz, D., Blattmann, A., Rombach, R.: Adversarial Diffusion Distillation . <https://doi.org/10.48550/arxiv.2311.17042>
7. Wang, Z., Bao, J., Zhou, W., Wang, W., Hu, H., Chen, H., Li, H.: Dire for diffusion-generated image detection (2023), <https://arxiv.org/abs/2303.09295>

Domain-informed Machine Learning by Hybridizing Symbolic Regression and Large Language Models

Du Nguyen Duy¹[0009-0000-3422-1918],
 Florian Bachinger²[0000-0002-5146-1750], and
 Michael Affenzeller²[0000-0001-5692-5940]

¹ Software Competence Center Hagenberg, Austria
 du.nguyen.duy@scch.at

² University of Applied Sciences Upper Austria, Austria
 {michael.affenzeller, florian.bachinger}@fh-ooe.at

The demand for interpretable and trustworthy machine learning models is rising in industrial and scientific applications. Stakeholders in fields such as chemical engineering, finance, and physics need models that not only predict accurately but also align with established domain knowledge and provide human-understandable insights. Symbolic Regression (SR) is a promising technique that directly addresses this by searching the space of mathematical expressions to find an analytical formula $f(\mathbf{x})$ that best fits a given dataset [1, 2].

However, a fundamental limitation of standard SR is its purely data-driven nature. SR algorithms, whether based on Genetic Programming (GP) or neural networks, optimize a fitness function typically based on error metrics like Mean Squared Error (MSE) or R^2 [3]. This process can “overfit” to the data, discovering spurious correlations or overly complex models that are semantically nonsensical. For example, an SR model for a chemical process might find an equation suggesting that reaction yield decreases with catalyst concentration, a finding that violates fundamental chemical principles and would be immediately rejected by an expert.

Prior attempts to inject domain knowledge into SR have focused on formally encoded constraints. A notable example is shape-constrained symbolic regression [4], which enforces shape-properties like monotonicity or curvature (convexity or concavity) on the candidate functions. While useful, this method has a significant bottleneck: it requires an expert to manually translate their implicit domain knowledge into a rigid, formal mathematical structure (e.g., monotonic increasing, as defining that a function’s derivative is non-negative $\frac{\partial f}{\partial x} \geq 0$). This is often impractical, as the vast majority of expert knowledge is unstructured, existing in textbooks, research papers, and natural language discourse.

In this work, we propose a novel framework, **Domain-informed Symbolic Regression (DiSR)**, to bridge this gap. We leverage Large Language Models (LLMs) [5, 6], which encode extensive scientific and domain-specific principles from their training on massive text and code corpora, in two distinct phases.

- First, during the SR search, the LLM acts as a “domain knowledge agent”. We define two objectives: (1) **Accuracy** (\mathcal{L}_{fit}), the traditional data-fit met-

ric (e.g., MSE), and (2) **Plausibility** (S_{LLM}), a “domain sensibility” score from 0.0 (nonsensical) to 1.0 (plausible), generated by an LLM. The LLM is prompted with the candidate formula and the semantic context (e.g., “You are an expert in thermodynamics. x_1 is temperature (K), y is reaction yield (%). Evaluate the following equation’s physical plausibility on a scale from 0.0 to 1.0”). The final fitness \mathcal{F} for a candidate is a weighted sum:

$$\mathcal{F} = (1 - w) \cdot \mathcal{L}_{fit} + w \cdot (1 - S_{LLM}),$$

where w is a hyperparameter balancing data-fit against domain knowledge.

- The second phase begins after the SR search yields a set of high-performing, plausible models. DiSR uses the LLM as a “scientific interpreter” to perform automated knowledge extraction on the final symbolic expression. The LLM is prompted to analyze the formula and describe the relationships it implies in natural language (e.g., “the model suggests that yield y increases linearly with temperature x_1 but has a saturating, inverse-square relationship with pressure x_2 ”). Furthermore, we task the LLM with validating these extracted principles against its vast training corpus to determine if they are: (1) **Established** (supported by existing literature), (2) **Novel** (a potentially new discovery), or (3) **Contradictory** (challenging existing theory).

This bi-directional framework positions DiSR as more than just a model-fitting tool. It is an end-to-end knowledge discovery system. By first integrating unstructured domain knowledge into the search, we ensure the discovery of robust and trustworthy models. By then using the LLM to extract and validate new insights from these models, we create a mechanism for discovering human-understandable, scientifically-grounded knowledge that can, in turn, inform future research. It also addresses the frequent claim of *interpretability* and *explainability* in standard SR by using LLMs to convey both these attributes in natural language to the human operator. We intend to evaluate our DiSR framework by comparing it to standard SR in terms of runtime performance, accuracy and *closeness* to the known underlying set of benchmark equations.

References

1. Koza, J. Genetic programming as a means for programming computers by natural selection. *Statistics And Computing*. **4**, 87-112 (1994)
2. Affenzeller, M., Wagner, S., Winkler, S. & Beham, A. Genetic algorithms and genetic programming: modern concepts and practical applications. (Taylor & Francis, 2009)
3. Kronberger, G., Burlacu, B., Kommenda, M., Winkler, S. & Affenzeller, M. Symbolic regression. (Chapman, 2024)
4. Kronberger, G., França, F., Burlacu, B., Haider, C. & Kommenda, M. Shape-constrained symbolic regression—improving extrapolation with prior knowledge. *Evolutionary Computation*. **30**, 75-98 (2022)
5. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, L. & Polosukhin, I. Attention is all you need. *Advances In Neural Information Processing Systems*. **30** (2017)
6. Naveed, H., Khan, A., Qiu, S., Saqib, M., Anwar, S., Usman, M., Akhtar, N., Barnes, N. & Mian, A. A comprehensive overview of large language models. *ACM Transactions On Intelligent Systems And Technology*. **16**, 1-72 (2025)

Physics-Informed Knowledge Validation via LLM-Assisted Shape-Constrained Symbolic Regression

Florian Bachinger¹, Stefan Habringer¹, Clemens Staudinger², Du Nguyen Duy³, and Michael Affenzeller¹

¹ HEAL, University of Applied Sciences Upper Austria

² voestalpine AG, Linz, Austria

³ SCCH, Hagenberg, Austria

Symbolic regression (SR) searches for analytic expressions $f(x)$ that accurately describe studied phenomena in a given problem domain. A major challenge of SR is to find models that not only fit the data well, but also conform to expert knowledge about the underlying phenomena. Incorporating such knowledge in the form of shape-constraints (SCs) has been shown to improve the reliability and generalization capabilities of SR, especially, when data is limited [2]. One example for SCs is monotonicity, which enforces that the function $f(x)$ is monotonically non-decreasing $\frac{\partial f}{\partial x_i} \geq 0$ over a given input variable x_i (cf. Figure 1).




Fig. 1: Constrained models (blue) are constrained to be monotonically non-decreasing $\frac{\partial f}{\partial x_i} \geq 0$ and exhibit more plausible extrapolation behavior.

Such knowledge is often available from first principles or from domain experts. However, these constraints must be formally defined and validated before they can be incorporated into shape-constrained SR, a process that is both time-consuming and prone to errors, and overall difficult for domain experts who are not familiar with formal mathematical definitions. In this work, we aim to tackle this challenge, by proposing a novel approach for knowledge extraction [1] and knowledge validation that supports domain experts by automatically extracting, validating, and explaining expert knowledge based on only observational data.

The experimental setup is designed as visualized in Figure 2: We sample observational data from known equations of the AI Feynman benchmark set [3], which allows us to infer ground-truth constraints by symbolically deriving the

equations. The knowledge extraction approach [1] operates solely on the observational data, without access to the original equations, variable descriptions, or names. In contrast, the LLM is provided only with the broad application domain (e.g., physics, mechanics), but not the equation itself. This setup enables us to evaluate both the constraint extraction process and the LLM's validation capabilities against known ground truth. As a last step, a domain expert assesses the LLM's output for correctness and interpretability.




Fig. 2: The dashed box indicates the core focus of this work. A LLM is provided with a description of the problem domain, the names and descriptions of the variables, and presented with a set of constraints extracted from observational data [1]. These constraints might be erroneous and are in general hard to interpret for domain experts. The LLM validates the constraints and tries to provide a natural language description of each.

We evaluate this approach across three key dimensions: First, we assess the LLM's ability to distinguish valid from invalid constraints, where the set of known, valid constraints is provided as a reference. Second, we examine the accuracy and interpretability of the LLM's natural language explanations for each constraint. Finally, we aim to log the effort spent by the domain expert in assessing the discovered constraints, this is the amount of valuable human resource time our approach may be saving. By combining automated constraint extraction with LLM-assisted validation and explanation, this work demonstrates a promising path toward making shape-constrained SR more accessible to domain experts and progressing the field of physics-informed machine learning and explainable AI.

References

1. Bachinger, F., Haider, C., Zenisek, J., de França, F.O., Affenzeller, M.: Automated inference of domain knowledge in scientific machine learning. In: Quesada-Arencibia, A., Affenzeller, M., Moreno-Díaz, R. (eds.) Computer Aided Systems Theory – EU-ROCAST 2024. pp. 122–130. Springer Nature Switzerland, Cham (2025)
2. Haider, C., de França, F.O., Burlacu, B., Kronberger, G.: Using shape constraints for improving symbolic regression models (2021). <https://doi.org/10.48550/ARXIV.2107.09458>, <https://arxiv.org/abs/2107.09458>
3. Udrescu, S.M., Tegmark, M.: Ai feynman: A physics-inspired method for symbolic regression. *Science Advances* **6**(16), eaay2631 (2020). <https://doi.org/10.1126/sciadv.aay2631>, <https://www.science.org/doi/abs/10.1126/sciadv.aay2631>

Meta-modeling Power Inverters for Fast Evaluation of Energy Flow Controllers on Resource Limited Devices

Christoph Dalpiaz¹, Kathrin Preiner², Phillip Kastner¹, Michael Kargl¹,
Florian Eibensteiner¹, and Josef Langer¹

¹ Josef Ressel Centre for AI on Resource Limited Devices

University of Applied Sciences Upper Austria, Hagenberg, Austria

{christoph.dalpiaz, phillip.kastner, michael.kargl, florian.eibensteiner,
josef.langer}@fh-hagenberg.at

² Fronius International GmbH, Thalheim, Austria

preiner.kathrin@fronius.com

Abstract. In recent publications, energy flow controllers have been trained using simulation based optimization with a complex inverter simulation model, which leads to rather long training times.

In this paper, the training and results of a computationally efficient symbolic regression meta-model for the power inverter is presented, which can be applied for faster training and adaptation of energy flow controllers. The results show that the performance of energy flow controllers can be assessed correctly using the meta-model in almost all simulated scenarios. Moreover, the speedup of the controller evaluation using the meta-model results in approximately a factor 24 compared to the evaluation using the complex inverter simulation model.

Keywords: Meta-model · Power Inverter · Energy Management System · Genetic Programming · Symbolic Regression

1 Introduction

In order to tackle global warming, the usage of renewable energies, like photovoltaic (PV) energy together with batteries for storing energy in times with low demand has been greatly expanded in residential households. For households with a PV system and a battery, it is desired to apply an energy management system (EMS) in order to control the energy flows between the PV system, the battery and the grid. In a recent publication by Kefer et al. [2], energy flow controllers (EFCs) have been realized using symbolic regression models determined by genetic programming, where the objective is to minimize the energy costs of the household. In order to provide a near-optimal solution in a wide range of scenarios, it is desired to adapt the symbolic regression models during runtime. Moreover, the computational complexity of the adaptation should be as low as possible. The complexity can be reduced by applying a simplified meta-model [3] instead of a complex inverter simulation model for the EFC evaluation, which is described in the next section.

2 Methods

In Fig. 1, the applied structure for meta-modeling the inverter using separate grid power and battery loss power models is depicted. The training data for model training has been generated by simulating a large number of scenarios. The models are trained using symbolic regression, where an implementation of the Offspring Selection Genetic Algorithm (OSGA) [1] has been applied.




Fig. 1. Structure for training and implementation of meta-models.

3 Results

In order to determine if the trained meta-models can be successfully applied for the evaluation of EFCs, it has to be shown that the following condition is satisfied for all possible pairs of energy flow controllers E_1, E_2

$$C_D\{E_1\} < C_D\{E_2\} \iff C_M\{E_1\} < C_M\{E_2\} \quad (1)$$

where $C_D\{E_i\}$ denotes the cost from applying E_i for EFC on the detailed simulation model and $C_M\{E_i\}$ denotes the cost from applying E_i for EFC on the meta-model. The simulation results on an example household with a set of 21 EFCs show, that the above condition is satisfied in almost all cases. Moreover the speedup for the evaluation using the meta-model results in approximately a factor 24 compared to the detailed simulation model.

Acknowledgments. The financial support by the Austrian Federal Ministry of Labour and Economy, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged.

References

1. Affenzeller, M., Wagner, S.: Offspring selection: A new self-adaptive selection scheme for genetic algorithms. In: Ribeiro, B., Albrecht, R.F., Dobnikar, A., Pearson, D.W., Steele, N.C. (eds.) *Adaptive and Natural Computing Algorithms*. pp. 218–221. Springer Vienna, Vienna (2005)
2. Kefer, K., Hanghofer, R., Kefer, P., Stöger, M., Hofer, B., Affenzeller, M., Winkler, S.: Simulation-based optimization of residential energy flows using white box modeling by genetic programming. *Energy and Buildings* **258**, 111829 (2022)
3. Lidberg, S.: Evaluating Fast and Efficient Modeling Methods for Simulation-Based Optimization. Ph.D. thesis, University of Skövde (2021)

Visualizing Populations in Multi-Objective Optimization *

Bernhard Werth^{1,2[0000-0002-7830-3138]} and Michael Affenzeller^{1,2[0000-0001-5692-5940]}

¹ Heuristic and Evolutionary Algorithms Laboratory
University of Applied Sciences Upper Austria, 4232 Hagenberg, Austria

² Institute for Symbolic Artificial Intelligence
Johannes Kepler University, 4040 Linz, Austria

bernhard.werth@fh-hagenberg.at

Understanding the dynamics of multi-objective evolutionary algorithms (MOEAs) remains challenging, particularly in high-dimensional search spaces where objective trade-offs interact with multimodal landscapes. Established performance indicators such as hypervolume or static Pareto front plots provide outcome summaries, but they offer little insight into why algorithms progress, stagnate, or diverge. For explainability in evolutionary computation—critical for both practitioners and researchers—more fine-grained perspectives are needed that connect algorithm behavior with the structural properties of the search space.

We aim to make the evolving state of MOEAs more interpretable by combining two complementary ideas: First, we visualize the local correlation of objectives in the neighborhood of candidate solutions, which gives an indication of the modality of the problem [3]. Positive correlations signal regions where simultaneous improvements are likely, while negative correlations reveal where objectives are locally in conflict. This local information acts as a proxy for the evolvability of different regions in the search space, highlighting where mutation or variation operators are more or less promising.

Second, we employ embeddings that project the current population into a low-dimensional representation (e.g., via Principle Component Analysis (PCA) or t-SNE [2]) while still allowing samples drawn from the embedding to be mapped back into the original search space. This dual property makes it possible to enrich each population snapshot with additional structural information from around the population, rather than being restricted to the solutions the algorithm currently maintains. The resulting views show not only where the population is, but also how its immediate surroundings are structured in terms of objective correlations and potential improvements.

Figure 1 illustrates the approach for an NSGA-II [1] run on a 20-dimensional two-objective benchmark problem in later stages of convergence. On the left, the population is displayed in the embedded search space, where thicker dots

* This research has been conducted within the scope of the Human Centered Artificial Intelligence (HCAI) project, funded by the Austrian Science Fund (FWF) [DFH 23-N].




Fig. 1. Example visualization of an NSGA-II run, showing population embedding (left) and objective space (right).

(colored by objective 0 in order to allow visual identification of the corresponding objective values in on the right side) denote the candidate solutions. These are overlaid on a background of sample points: the embedded space was filled with a grid of samples, which were then back-transformed, evaluated, and colored according to the correlation between objectives (green indicating positive correlation, red indicating negative correlation). On the right, the same population is shown in the objective space. The combined visualization reveals that the algorithm has reached a relatively homogeneous region: while fringe solutions can still improve through variation, the center of the population mainly benefits from improved spacing once it matches the qualities of its structural background.

Viewed across the full runtime, such visualizations enable a microscopic, per-run analysis of MOEA dynamics. Early experiments on two-objective benchmark functions with real-valued encodings show that the method captures distinct phases of algorithm behavior like global exploration, focusing on promising subspaces, temporary stagnation in local optima, and final refinement of Pareto fronts. These dynamics are difficult to discern from aggregate performance indicators alone.

The contribution of this work is a visualization technique that enriches standard objective- and search-space views with local correlation and evolvability information. By systematically supporting detailed per-run analysis, the framework complements established metrics and provides a new tool for reconstructing and explaining the trajectory of evolutionary search.

References

1. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE transactions on evolutionary computation* **6**(2), 182–197 (2002)
2. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
3. Tušar, T., Cork, J.N.: Using local correlation between objectives to detect problem modality. In: International Conference on the Applications of Evolutionary Computation (Part of EvoStar). pp. 527–542. Springer (2025)

Shapes un-Constrained: Tracking Phenotypical Behavior in Evolutionary Symbolic Regression

Christian Haider¹[0000-0002-3039-8844], Florian Bachinger^{1,2}[0000-0002-5146-1750], Bernhard Werth^{1,2}[0000-0002-7830-3138], and Michael Affenzeller^{1,2}[0000-0001-5692-5940]

¹ Heuristic and Evolutionary Algorithms Laboratory
University of Applied Sciences Upper Austria, 4232 Hagenberg, Austria

² Institute for Symbolic Artificial Intelligence
Johannes Kepler University, 4040 Linz, Austria

`christian.haider@fh-hagenberg.at`

In symbolic regression, explainability efforts have traditionally focused on the analysis of individual solutions. Techniques such as symbolic simplification or node importance measures aim to improve the interpretability of candidate models. In addition, methods from the broader field of explainable artificial intelligence (XAI) can often be applied to symbolic models as well [2]. In contrast, explainability at the level of the search process itself has received comparatively little attention. Standard indicators, including symbol or variable frequency statistics and population quality curves, provide only a coarse, often reduced, view of how evolutionary search unfolds [1]. These measures capture “what” occurs during optimization, but rarely shed light on the functional structures that emerge and disappear as the search progresses.

In this work, we propose to study algorithm dynamics through the occurrence and evolution of shape properties in symbolic regression populations. Shape properties describe broad functional behaviors—such as monotonicity, curvature (convexity, concavity), non-negativity, or constancy, that of a model. Typically, such properties are considered in the context of shape-constrained regression, where models are required to adhere to predefined desired properties [4]. Instead, rather than enforcing constraints, we ask when and how such properties emerge naturally during search, how prevalent they are throughout the population, and how their presence evolves over time. Identifying when and how such shape-properties emerge or disappear might provide us with more information, useful in e.g. the development of approach to enforce of desired properties.

To this end, we adopt a recently proposed method for annotating benchmark problems with shape constraints [3]. The approach samples the intended input domain of a candidate model and evaluates both the function and its symbolic derivatives. From this, the set of “empirically adhered-to” shapes can be extracted. By organizing the input space with decision trees, the granularity of detected shapes can be controlled: for example, one might obtain a description such as “the model is constant for $x_1 > 0$ and $x_2 > 3.14$, but unconstrained otherwise.” When applied to every individual in every generation, this procedure produces a timeline of structural properties present in evolving populations.

This type of representation allows us to move beyond generic quality metrics and directly observe when populations enter phases of shape dominance—for instance, when monotonicity in one variable dominates the population for several generations, before being replaced by more specialized or irregular behaviors. Such analysis may be particularly informative in cases where the true generating function is known, as they allow us to ask whether correct shape properties were discovered, overlooked, or only temporarily present during the run.

We expect that this type of analysis will help to characterize typical patterns of exploration, exploitation, and specialization in evolutionary symbolic regression. In particular, it provides an avenue for investigating possible connections between the stability of shape properties and the generalization ability of evolved models. More broadly, the approach offers a new perspective on algorithm-level explainability in evolutionary computation, complementing existing measures such as symbol frequencies, allele analysis, or diversity indicators.




Fig. 1. First draft for a possible visualization of our idea.

References

1. Affenzeller, M., Winkler, S.M., Kronberger, G., Kommenda, M., Burlacu, B., Wagner, S.: Gaining Deeper Insights in Symbolic Regression, pp. 175–190. Springer New York, New York, NY (2014). https://doi.org/10.1007/978-1-4939-0375-7_10
2. Aldeia, G.S.I., de França, F.O.: Interpretability in symbolic regression: a benchmark of explanatory methods using the feynman data set. *Genetic Programming and Evolvable Machines* **23**(3), 309–349 (May 2022). <https://doi.org/10.1007/s10710-022-09435-x>
3. Bachinger, F., Werth, B., Zenisek, J., Haider, C., Olivetti de França, F.: Scrbenchmark: A benchmarking library for shape-constrained regression. In: *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. pp. 2505–2513 (2025)
4. Kronberger, G., de França, F.O., Burlacu, B., Haider, C., Kommenda, M.: Shape-constrained symbolic regression—improving extrapolation with prior knowledge. *Evolutionary computation* **30**(1), 75–98 (2022)

Improving Genetic Programming-based Symbolic Regression with Exhaustive Methods

Lukas Kammerer^{1,2[0000-0001-8236-4294]}, Deaglan J. Bartlett^{3[0000-0001-9426-7723]}, Harry Desmond^{4[0000-0003-0685-9791]}, Pedro G. Ferreira^{3[0000-0002-3021-2851]}, Gabriel Kronberger^{1[0000-0002-3012-3189]}, and Stephan Winkler^{1,2[0000-0002-5196-4294]}

¹ Heuristic and Evolutionary Algorithms Laboratory,
University of Applied Sciences Upper Austria, Hagenberg, Austria

² Department of Computer Science, Johannes Kepler University, Linz, Austria
³ Astrophysics, University of Oxford, Oxford, UK

⁴ Institute of Cosmology & Gravitation, University of Portsmouth, Portsmouth, UK
lukas.kammerer@fh-hagenberg.at

Introduction and Motivation: Symbolic regression (SR) is a machine learning task that aims to identify mathematical closed-form expressions of arbitrary structure that capture dependencies in given data [4]. The representation of models as human-readable and interpretable mathematical formulas distinguishes SR algorithms from other machine learning methods.

The most common algorithm for solving SR problems is Genetic Programming (GP), which iteratively improves a randomly initialized population through mutation and crossover operations. GP-based methods are capable of finding accurate models for high-dimensional and highly nonlinear problems. However, despite their promise of interpretability, GP-based SR methods suffer from bloat and tend to produce overly complex models with limited interpretability [5]. Recent non-evolutionary approaches [1, 3, 6] address this drawback by aiming to create both accurate and compact expressions. However, their applicability is limited to problems with few variables, low noise levels, or limited nonlinear dependencies.

A notable deterministic SR algorithm is Exhaustive Symbolic Regression (ESR) [1], which performs a brute-force search over all algebraically unique model structures up to a certain length. It uses a variant of Description Length (DL) as a loss function to find models that are both short and accurate and that generalize well to unseen data. The DL accounts for both the error and the structural and parametric complexity of a model without requiring additional configuration for balancing these properties. ESR has been shown to be effective at finding short, accurate models for real-world problems, particularly in physics. Its drawback is the limited maximum model length, as the number of all possible models grows exponentially with the allowed model length.

Aims: This work combines the advantages of ESR and scalable GP-based methods to find models that are both short and accurate, with lengths beyond those reachable by ESR alone. For a given dataset, we first use ESR to generate a set of small, sufficiently accurate models. We then improve these models heuris-

tically with GP-based SR methods to obtain more accurate but still concise models outside of ESR's search space.

Methods: We link both approaches via population initialization: instead of randomly initializing a population of models, as is usually done in GP, we use a set of ESR models with the lowest DL as the initial population. Given that models produced by ESR are syntactically diverse and contain only meaningful building blocks, we hypothesize that such an informed start will lead to shorter final models compared with GP runs with randomly initialized populations.

An advantage of using ESR as the baseline algorithm is that it produces models that usually generalize well. Therefore, we expect our GP-based extension to provide well-generalizing models as any crossover and mutation step uses well-fitting and meaningful sub-models. Since ESR does not require any hyper-parameters, our approach does not add further complexity an overall modelling methodology.

We use the ESR implementation from [1] and *Operon* [2] as the GP-based SR implementation. We choose Operon because it is among the most accurate SR algorithms in recent benchmarks [5]. Although it also tends to create some of the longest models, we investigate whether this effect is mitigated by our approach. Operon also includes an implementation of NSGA-2, enabling multi-objective search and the generation of a Pareto front over model length and accuracy.

We will apply our methods to both synthetic problems from the SRBench benchmark [5] and to real-world problems in physics, as many problems in this domain are suitable for ESR and thus for further enhancement. We compare plain ESR and Operon with randomly initialized populations against our ESR-initialized Operon configuration, highlighting improvements in model length and accuracy.

References

1. Bartlett, D.J., Desmond, H., Ferreira, P.G.: Exhaustive symbolic regression. *IEEE Transactions on Evolutionary Computation* **28**(4), 950–964 (2023)
2. Burlacu, B., Kronberger, G., Kommenda, M.: Operon c++ an efficient genetic programming framework for symbolic regression. In: *Proceedings of the 2020 Genetic and Evolutionary Computation Conference Companion*. pp. 1562–1570 (2020)
3. Kammerer, L., Kronberger, G., Burlacu, B., Winkler, S.M., Kommenda, M., Affenzeller, M.: Symbolic regression by exhaustive search: reducing the search space using syntactical constraints and efficient semantic structure deduplication. In: *Genetic Programming Theory and Practice XVII*, pp. 79–99. Springer (2020)
4. Koza, J.R.: Genetic programming as a means for programming computers by natural selection. *Statistics and computing* **4**(2), 87–112 (1994)
5. La Cava, W., Burlacu, B., Virgolin, M., Kommenda, M., Orzechowski, P., de França, F.O., Jin, Y., Moore, J.H.: Contemporary symbolic regression methods and their relative performance. *Advances in Neural Information Processing Systems* **2021**(DB1), 1 (2021)
6. Udrescu, S.M., Tegmark, M.: Ai feynman: A physics-inspired method for symbolic regression. *Science Advances* **6**(16), eaay2631 (2020)

A Learning Large Neighborhood Search for the Electric Autonomous Dial-A-Ride Problem

Laurenz Tomandl¹, Maria Bresich¹, Günther R. Raidl¹, and Steffen Limmer²

¹ Institute of Logic and Computation, TU Wien, Vienna, Austria
 {ltomandl, mbresich, raidl}@ac.tuwien.ac.at

² Honda Research Institute Europe GmbH, 63073 Offenbach/Main, Germany
 steffen.limmer@honda-ri.de

The Dial-A-Ride Problem (DARP), first introduced by Cordeau and Laporte [4], involves determining a set of minimum-cost vehicle routes that satisfy user transportation requests. In the DARP, each user must be transported from a pickup to a drop-off location within a predefined time window. Additionally, users are subject to a maximum ride time constraint that limits how long they may remain in a vehicle. The DARP allows multiple passengers to share a vehicle up to its capacity, a characteristic that differentiates it from other routing problems and makes it more complex in practice. Furthermore, routes are constrained in duration to account for driver break requirements.

Several extensions of the DARP have been proposed. Among them, the Electric Autonomous Dial-A-Ride Problem (EADARP), introduced by Bongiovanni et al. [1], considers fleets of self-driving electric vehicles. This variant relaxes certain constraints, such as maximum route duration (since autonomous vehicles do not require driver breaks), but introduces new challenges related to energy management. In particular, vehicles must visit charging stations, a new class of service locations, to replenish their batteries as needed. The cost of a route is defined as a weighted sum of user excess ride time, the excess time a user spends on board of a vehicle over their direct route travel time, and total routing distance.

We build upon the approach of Bresich et al. [3], which is a Large Neighborhood Search (LNS) with a novel charging management strategy. Their method inserts charging stops on the fly only after the stop order of a route has been determined and uses a fragment-wise route representation, where route fragments are subsets of stops. This representation simplifies the computation of features such as excess ride time and improves efficiency. Their approach achieves state-of-the-art results for instances of various sizes. Later, Bresich et al. [2] extended this method to also handle very large instances via an aggressive heatmap-based pruning scheme. Other variants of the EADARP, such as the dynamic version, have been addressed using a genetic programming hyperheuristic [6] and a reinforcement learning supported LNS [8].

In this work, we further explore improvements to solving the EADARP by introducing a Learning LNS (LLNS) strategy. In contrast to previous studies that focused on sparsification or hyperparameter tuning via deep learning, our approach learns a function that evaluates how promising a route is for inclusion in the destroy set of the LNS. Specifically, we employ a route-wise destroy operation that removes entire routes from the solution. Each route is scored by a

neural network and based on this information selected to be destroyed in a probabilistic manner. The neural network model evaluates routes using features such as route duration, the number of covered orders, average user detour, energy consumption margin, time-window slack, and how long the route was already part of the incumbent solution, and when it was last selected to be destroyed. We employ a supervised learning strategy for training the neural network offline on representative problem instances. The target is based on the improvement in objective value after removing a route and repairing the partial solution.

To the best of our knowledge, this is a first approach to select routes to destroy in a more educated way by means of a neural network trained by supervised learning for a DARP variant. The approach can adapt to changing incumbent solutions and helps focusing the search, potentially improving convergence speed and solution quality compared to heuristic or random destroy strategies.

Detailed tuning and computational experiments are still in progress, but related learning-based destroy mechanisms have shown strong results in other optimization problems [5]. We expect to achieve improvements with minimal additional computational effort, as the number of routes even for very large instances is relatively small and the feature extraction is efficient. Results of an extensive evaluation on benchmark instances from [3, 7], comparing LLNS against state-of-the-art LNS variants in terms of solution quality, runtime, and robustness will be presented at the conference.

References

1. Bongiovanni, C., Kaspi, M., Geroliminis, N.: The electric autonomous dial-a-ride problem. *Transportation Research Part B* **122**, 436–456 (2019)
2. Bresich, M., Raidl, G.R., Coleman, C., Welke, P., Limmer, S.: Search space reduction through machine learning for the electric autonomous dial-a-ride problem. In: Springer Lecture Note in Computer Science. LOD '25 (in press)
3. Bresich, M., Raidl, G.R., Limmer, S.: Letting a large neighborhood search for an electric dial-a-ride problem fly: On-the-fly charging station insertion. In: Proceedings of the Genetic and Evolutionary Computation Conference. pp. 142–150. GECCO '24, ACM (2024)
4. Cordeau, J.F., Laporte, G.: A tabu search heuristic for the static multi-vehicle dial-a-ride problem. *Transportation Research Part B* **37**(6), 579–594 (2003)
5. Hottung, A., Wong-Chung, P., Tierney, K.: Learning destroy operators for vehicle routing problems with deep neural networks. *Symposium on Transportation Analysis Conference* (2025)
6. Huang, W., Mei, Y., Raidl, G., Zhang, F., Tomandl, L., Limmer, S., Zhang, M., Rodemann, T.: Genetic programming hyper-heuristic for the dynamic electric dial-a-ride problem. In: IEEE Congress on Evolutionary Computation. IEEE (2025)
7. Limmer, S.: Bilevel large neighborhood search for the electric autonomous dial-a-ride problem. *Transportation Research Interdisciplinary Perspectives* **21**, 100876 (2023)
8. Tomandl, L., Bresich, M., Raidl, G.R., Mei, Y., Limmer, S., Rodemann, T.: A reinforcement learning guided large neighborhood search for the dynamic electric autonomous dial-a-ride problem. In: Springer Lecture Note in Computer Science. LOD '25 (in press)

Transfer Learning using Template Seeding within Structure-Template Symbolic Regression

Christian Haider¹, Jan Zenisek¹, Philipp Fleck¹, Maximilian Fruhmann²,
Carina Hambrock³, and Michael Affenzeller¹

¹ Heuristic and Evolutionary Algorithms Laboratory
University of Applied Sciences Upper Austria, 4232 Hagenberg, Austria

² K1-MET GmbH
Stahlstraße 14, 4020 Linz, Austria
³ voestalpine Stahl GmbH
voestalpine-Straße 3, 4020 Linz, Austria

christian.haider@fh-hagenberg.at

Introduction: Symbolic regression (SR) is a machine learning method that maps real-valued inputs (features) to real-valued outputs (targets). This applies for most regression tasks, however, unlike other regression approaches, SR does not require a predefined model structure. Instead, it simultaneously searches for both optimal parameters and functional form [1]. SR typically uses genetic programming (GP) to evolve mathematical expressions, making it well suited for exploring large model spaces. The resulting models are closed-form analytical expressions, enabling interpretability. This makes SR interesting for integrating knowledge into the modeling process. In particular, when implemented with tree-based GP, SR represents candidate models as symbolic expression trees, allowing to incorporate domain knowledge via structural or shape constraints [2].

One approach to integrating prior knowledge is using structure template regression (STR), which is an extension to conventional symbolic regression. STR combines data-based modeling with physical-based modeling. This is achieved by directly integrating first-principle expressions into the model training process.

Method: The core idea is to learn reusable structural building blocks, the so-called templates, during GP runs, and transfer them to related problems as informed priors. This transformation will be performed in three steps. *In this paper, we focus on the first part: template extraction and mining. Steps 2 and 3 will be presented in a follow-up paper.*

1) Template mining on initial training task: From SR modeling results, we extract frequently occurring and high-performing sub-expressions. We normalize them into templates by (i) replacing constants with free parameters, (ii) abstracting variables into typed placeholders, and (iii) canonicalizing algebraic form. Each template is stored with metadata (domain, units, ...).

2) Library curation: We rank templates by frequency and cross-validated performance, filter near-duplicates, and keep a balanced set spanning low- to mid-complexity templates.

3) Transfer via template seeding: For a new target task (inductive transfer learning), we instantiate a fraction of the initial GP population with templates assigned to target variables. We add template-aware operators that preserve

substructures during crossover/mutation while allowing coefficient adaptation and variable remapping. The remainder of the population is drawn from standard primitives to maintain diversity.




Fig. 1. Structure-template transfer from communicating vessels problem (hydrodynamics domain) to heat transfer problem (thermodynamic domain). In a first step we extract possible structures, to setup the library, to then seed these templates into the GP population.

Toy problem (Figure 1): Transfer a hydrodynamics system of communicating vessels to a thermally coupled two-mass system. In the source domain, transfer between tanks is captured by $\frac{dh_1}{dt} = \sqrt{h_1 - h_2}$ (Torricelli-like flow), which means that the water flows from higher levels to lower levels. In the target domain, heat transfer follows $\frac{dT_1}{dt} = -c(T_1 - T_2)$, this relation indicates that heat flows from hot systems to colder ones. Despite different physics, both share a difference-driven transfer motif of the captured templates.

This strategy biases search toward physically plausible forms, reduces the effective search space, accelerates convergence, and preserves interpretability, while still allowing the algorithm to learn task-specific interactions.

References

1. Koza, J.R.: Genetic programming as a means for programming computers by natural selection. *Statistics and computing* **4**(2), 87–112 (1994)
2. Kronberger, G., de França, F.O., Burlacu, B., Haider, C., Kommenda, M.: Shape-constrained symbolic regression—improving extrapolation with prior knowledge. *Evolutionary computation* **30**(1), 75–98 (2022)

Optimizing Dynamic Dispatching in Elevator Control with a Destination Registration System*

Maria Bresich¹✉, Günther R. Raidl¹, and Steffen Limmer²

¹ Institute of Logic and Computation, TU Wien, Vienna, Austria
 {mbresich, raidl}@ac.tuwien.ac.at

² Honda Research Institute Europe GmbH, 63073 Offenbach/Main, Germany
 steffen.limmer@honda-ri.de

Nowadays, high-rise buildings and skyscrapers have become an integral part of major towns and shape each cityscape in a distinctive way. As the urban population keeps growing, vertical urbanization has been on the rise and will be required to satisfy the housing needs and demands of the population. This also leads to an urgent need for efficient vertical transportation within buildings, for which multi-elevator systems are a widespread solution. Such systems often consist of multiple elevator cars with each car having its own shaft, and elevator group control systems (EGCSs) are designed to optimize the passenger transportation by means of dispatching algorithms for the elevators. Traditionally, there are hall buttons on each floor for calling elevators, and the destination floor is only selected within an elevator car. We focus on a different approach and useful enhancement, i.e., destination registration systems (DRSs) which provide a terminal on each floor where passengers can indicate their desired destination already before entering an elevator, thus providing more information to the system earlier. This allows for better dispatching decisions and increased efficiency but also requires more advanced algorithms for assigning passengers to elevators.

In this work, we tackle the assignment and dispatching problem for an EGCS with a destination registration system (EGCS-DRS). Our considered system also allows for specifying variable group sizes in case multiple people arrive together which further increases the practical complexity of this NP-hard problem. The goal is to minimize the average waiting time per individual from requesting an elevator until entering it. Literature on this topic is sparse and due to the lack of a standardized formulation, a variety of slightly different problem variants have been considered. An early work on EGCS-DRSs is by Beielstein et al. [1] who propose evolution strategies in combination with neural networks (NNs) to judge the system's performance. Exact approaches such as a branch-and-bound algorithm and a mixed integer linear programming (MILP) formulation were introduced by Tanaka et al. [3] and Ruokokoski et al. [2] respectively.

We first address the static (offline) problem where all orders are known in advance and formalize it as a MILP before approaching the dynamic (online) variant in the same way. The latter is more realistic and relevant in practice as orders only become known over time when passengers arrive. In both cases, the movement of elevators over time is modeled as network flows. The offline

* This project is financially supported by Honda Research Institute Europe GmbH.

problem is solved once and can be used as a baseline to determine globally optimal solutions while for the online problem, the MILP is resolved at each decision point when a new order arrives or an elevator reaches its next stop. As such exact approaches often struggle to scale to larger instances, we also propose a fast greedy algorithm based on a simple heuristic for the dynamic problem. It determines fast pickups for each elevator taking into account its position and the of waiting passengers, and also the current choices for other elevators. At each decision point, the algorithm iteratively determines only the next best action for each elevator which includes whether to wait or move (and where to) and which passengers to pick up next. A similar approach is used in our third proposed method for the online problem: a reinforcement learning (RL) algorithm for predicting for each elevator the best next destination. We model the problem as a Markov decision process with terminal as well as immediate rewards to enable a faster learning process. Observations of a system's state are split into sub-observations for each valid elevator-action combination to get a more compact representation. A proximal policy optimization algorithm is used for learning the policy, and the NN model consists of multiple blocks of multi layer perceptrons and aggregation steps for each elevator-action pair. As the final output is then a single elevator-action of highest probability, the NN is repeatedly applied in an autoregressive way with updates of the state in between to consider already made decisions until an action for each elevator is determined. These actions are enhanced by a greedy heuristic for selecting the passengers to be picked up.

Our approaches are evaluated in a computational study on different elevator system sizes, two scenario lengths, and various traffic patterns inspired by real-world use cases. Results show that the offline MILP approach is only feasible for small instances whereas the online MILP provides high quality solutions also for larger instances but nevertheless requires too much time to be used in real-time applications. In terms of computation time, the greedy algorithm is most efficient while also achieving reasonable solutions for most instances. The performance of the RL approach varies across the different traffic patterns and elevator systems, but it is in general also able to find good solutions even for large instances.

Acknowledgments. The implementation of the algorithms as well as all the empirical evaluation was done by Jonas Kompauer as part of his master's thesis done at TU Wien, Austria, in 2025. We want to thank him for his excellent work.

References

1. Beielstein, T., Ewald, C., Markon, S.: Optimal elevator group control by evolution strategies. In: Genetic and Evolutionary Computation Conference 2003 – Part II. LNCS, vol. 2724, pp. 1963–1974. Springer (2003)
2. Ruokokoski, M., Ehtamo, H., Pardalos, P.M.: Elevator dispatching problem: a mixed integer linear programming formulation and polyhedral results. *Journal of Combinatorial Optimization* **29**(4), 750–780 (2015)
3. Tanaka, S., Uraguchi, Y., Araki, M.: Dynamic optimization of the operation of single-car elevator systems with destination hall call registration: Part II. The solution algorithm. *European Journal of Operational Research* **167**(2), 574–587 (2005)

Timing of Queries in Interactive Job Scheduling

Johannes Varga¹, Harald Korinek¹, Günther R. Raidl¹, and Tobias Rodeman²

¹ Institute of Logic and Computation, TU Wien, Vienna, Austria

{jvarga,raidl}@ac.tuwien.ac.at, haraldkor@protonmail.com

² Honda Research Institute Europe, Offenbach, Germany

tobias.rodemann@honda-ri.de

Motivation and Problem. We consider the Interactive Job Scheduling Problem in which human users, e.g. patients in radiation therapy [2], regularly require access to a limited and high-demand resource, e.g. linear accelerators for the treatment of cancer. To avoid no-shows and frustration of the users, it is important to consider the preferences of the users, in particular the times in which they are available to use the resource. Usually, it is not practical for the users to specify their full availabilities in all detail. Instead, we assume that users initially propose only few timeslots in which they are available to access the resource and then a scheduling system interacts with the users to find out more about the relevant parts of their preferences. During interaction, the scheduling system proposes alternative timeslots to the users, who accept or reject them depending on their availability times. Previous work [1] does this in multiple interaction rounds, and in each round, the scheduling system sends out a small number of queries and awaits all replies before starting the next round. We refer to this as *blocking approach*. In practice, some queries will be answered immediately, while others remain unanswered for a longer time. Since the scheduling system waits for all of them, each round takes up a significant amount of time, reducing the total amount of queries that can be sent in a fixed timeframe. In this work, we explore different strategies for the timing of interaction in the Interactive Job Scheduling Problem to achieve better schedules with the same time budget.

Approaches. We consider two models for the durations users require for their reply after having received queries: One with normally distributed, the other with exponentially distributed durations. We develop two interaction approaches to reduce the overall interaction time that is required to achieve a schedule of good quality, see also Algorithms 1 and 2. In the first approach, the *responsive approach*, the scheduling system starts by computing and sending out one query per user. Whenever a reply is received, another query is computed and sent until the time budget is reached. A disadvantage of this approach is that a reoptimization is necessary for each reply and that a lot of queries are sent in total. Therefore, we refine it to the more advanced *hybrid approach*. Here, we also start with one query per user. Once the first reply arrives, we wait for a fixed amount of time t^{wait} . Then, we compute and send out one new query for each reply that was received during waiting and continue the cycle by waiting again for the next reply. This is repeated until the time budget is exhausted.

Algorithm 1: Responsive approach

```

1 Send  $b$  queries ;
2 while  $Time\ remaining$  do
3   Wait for reply ;
4   Send one query ;
5 Compute final schedule ;

```

Algorithm 2: Hybrid approach

```

1 Send  $b$  queries ;
2 while  $Time\ remaining$  do
3   Wait for reply ;
4   Wait for  $t^{\text{wait}}$  time units ;
5   Send one query for each received reply ;
6 Compute final schedule ;

```




Fig. 1. Mean objective over the number of received replies.

Results. To evaluate the approaches, we extend the simulation of the interaction between scheduling system and users from Varga et al. [1] with the reply-time characteristics of the users. Each approach is given a time budget of 600 time units and we evaluate the payoff between the number of queries and final schedule quality in form of operational costs. We evaluate the approaches on 30 instances with two machines, twelve users and 48 jobs and take the mean over all instances for the results. User delays for answering queries follow a normal distribution with mean 100 and standard deviation 50 and an exponential distribution with scale parameter 100, respectively. Figure 1 shows the development of the mean objective value over the number of received replies to the queries for both user timing models. All approaches have a similar objective development over the number of queries. But while the blocking approach is only able to send approximately 50 and 35 queries, respectively, the responsive approach is able to send more than 80 queries and therefore reaches a better objective value. The hybrid approach reduces the number of query computations from ≥ 60 (responsive) to around 20 ($t^{\text{wait}} = 20$), also translating into time savings in practice.

References

1. Varga, J., Raidl, G.R., Rönnberg, E., Rodemann, T.: Scheduling jobs using queries to interactively learn human availability times. *Computers & Operations Research* **167**, 106648 (2024)
2. Vieira, B., Demirtas, D., van de Kamer, J.B., Hans, E.W., Jongste, W., van Harten, W.: Radiotherapy treatment scheduling: Implementing operations research into clinical practice. *Plos one* **16**(2), e0247428 (2021)

Comparative Analysis of GRASP Variants for CVRPTW: A Case Study in Urban Organic Waste Collection

Mustapha Maarouf Bassaidi^{1[0000-0003-3657-5669]}, Julio A. Brito Santana^{1,2[0000-0003-3857-4707]}, José A. Moreno Pérez^{1,2[0000-0001-9506-5197]}, and Helí Alonso Afonso^{1[0009-0004-0839-1875]}

¹ Instituto Universitario de Desarrollo Regional, Universidad de La Laguna, SPAIN
 {mmaarouf, jbrito, jamoreno, halonsoa}@ull.edu.es

² Departamento de Ingeniería Informática y de Sistemas, Universidad de La Laguna

The management of urban solid waste is one of the most significant challenges that contemporary cities face, with organic waste constituting 40% to 60% of the entire municipal solid waste stream. The need for prompt collection of organic waste arises from its natural properties: rapid decomposition, offensive smells, and public health risks. This immediate need imposes strict time window constraints on the classical vehicle routing problem, resulting in a much more intricate optimisation scenario known as the Capacitated Vehicle Routing Problem with Time Windows (CVRPTW). Although this problem is well-documented in the optimisation literature, achieving effective solutions for real-world cases continues to present computational difficulties. The CVRPTW characteristics in practical contexts should be solved by using metaheuristic methods. Among these methods, GRASP (Greedy Randomized Adaptive Search Procedure) has been shown to be particularly effective for this class of problems. GRASP integrates a greedy randomised construction phase with a local search step, usually included in an iterative multistart framework [4]. This two-stage design enables GRASP to strike a careful balance between exploring diverse regions of the solution space and intensively exploiting promising areas, a trade-off that is crucial for good performance on complex optimisation problems.

The effectiveness of the basic GRASP framework has led to the development of numerous variants and hybrid schemes, each aimed at improving particular aspects of its performance. Reactive GRASP dynamically tunes its parameters based on the quality of the solutions obtained during the search [1]. GRASP with Path Relinking enhances the method by exploring trajectories that connect elite solutions in the solution space [8]. Hybridisation of GRASP and Variable Neighborhood Search (VNS) takes advantage of multiple neighbourhood structures to achieve a more thorough local exploration [5, 2]. GRASP with Tabu Search introduces memory-based mechanisms that discourage revisiting recently explored regions, thereby helping the algorithm escape local optima [7]. Several multistart GRASP variants exploit parallelism to investigate several regions of the solution space simultaneously [6].

We develop a comprehensive comparative analysis of several GRASP variants using the urban organic waste collection routing problem as a controlled case study. The selection of this particular problem is motivated by both its practical relevance and its algorithmic complexity. From a practical point of view, efficient organic waste collection directly impacts municipal operating costs, service quality, and environmental outcomes through the reduction of greenhouse gas emissions. From an algorithmic perspective, the temporal constraints inherent in organic waste collection create a demanding evaluation environment for testing the intensification and diversification capabilities of each

GRASP variant. Our evaluation uses real-world instances from Tenerife (Spain) comprising approximately 300 nodes with mixed time window tightness. Performance is evaluated using standard metrics: total travel distance, time window violations, vehicle utilisation, and computational time. This makes it an interesting problem for a systematic comparison, as success or failure of each variant in managing these temporal constraints will provide valuable insight into their relative strengths and weaknesses.

The research represents a dedicated effort to bridge the traditional gap between theoretical metaheuristic development and practical real-world problem solving by systematically evaluating multiple GRASP variants on a single, well-defined practical problem. The methodology and findings contribute not only to the specific domain of waste collection optimisation but also offer broader insights applicable to a wide range of vehicle routing problems and scheduling applications where GRASP has previously been applied [3]. In addition, this work aims to tune a practical tool to assist local administrators in adopting optimised collection systems that reduce operational costs and greenhouse gas emissions while simultaneously improving service quality and organic waste valorisation.

Acknowledgments. This study is funded by Cabildo de Tenerife, Talentum - *Antonio González* Programme.

References

1. Boudia, M., Louly, M.A., Prins, C.: A reactive GRASP and path relinking for a combined production-distribution problem. *Computers & Operations Research* **34**(11), 3402–3419 (2007). <https://doi.org/10.1016/j.cor.2006.02.005>
2. Bruglieri, M., Ferone, D., Festa, P., et al.: A GRASP with penalty objective function for the green vehicle routing problem with private capacitated stations. *Computers & Operations Research* **143**, 105770 (2022). <https://doi.org/10.1016/j.cor.2022.105770>
3. Corbea-Pérez, A., Brito, J., Moreno-Pérez, J.: Optimisation of the dynamic waste collection. In: Pereira, A., et al. (eds.) *Optimization, Learning Algorithms and Applications*. OL2A 2024, Communications in Computer and Information Science, vol. 2281, chap. 11. Springer (2024). https://doi.org/10.1007/978-3-031-77432-4_11
4. Ferone, D., Festa, P., Resende, M.: Grasp. In: Martí, R., Pardalos, P., Resende, M. (eds.) *Handbook of Heuristics*, chap. 23. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-00385-0_23
5. González-Navasa, C., Pérez, J., Brito, J.: A hybrid metaheuristic for a tourist route recommender. In: Quesada-Arencibia, A., Affenzeller, M., Moreno-Díaz, R. (eds.) *Computer Aided Systems Theory – EUROCAST 2024*, LNCS, vol. 15172, chap. 21. Springer (2025). https://doi.org/10.1007/978-3-031-82949-9_21
6. Martí, R., Lozano, J., Mendiburu, A., Hernando, L.: Multi-start methods. In: Martí, R., Pardalos, P., Resende, M. (eds.) *Handbook of Heuristics*, chap. 1. Springer, Cham (2025). https://doi.org/10.1007/978-3-032-00385-0_1
7. Póvoa, C.L.R., Roboredo, M.C., Velasco, A.S., Pessoa, A.A., Paes, F.G.: A hybrid grasp and tabu-search heuristic and an exact method for a variant of the multi-compartment vehicle routing problem. *Expert Systems with Applications* **259**, 125319 (2025). <https://doi.org/https://doi.org/10.1016/j.eswa.2024.125319>
8. Usberti, F.L., França, P.M., França, A.L.M.: GRASP with evolutionary path-relinking for the capacitated arc routing problem. *Computers & Operations Research* **40**(12), 3206–3217 (2013). <https://doi.org/10.1016/j.cor.2011.10.014>

Superstrings of Uniform-Cardinality Set Systems via the Generalized Traveling Salesperson Problem

Alexander Weissenfels, Enrico Iurlano[✉], and Günther R. Raidl

Algorithms and Complexity Group, TU Wien, Austria
 [firstName].[lastName]@tuwien.ac.at

In this work we address computational approaches to the *shortest P_n^k -cover* problem, a variant of the *Shortest Superstring Problem* (SSP) [1] by Lipski Jr. [7]. Given the collection of all cardinality- k subsets of a cardinality- n alphabet, one seeks here a shortest string over this alphabet such that, for each subset in the collection, an arbitrary arrangement of its elements appears as contiguous substring of the superstring; we call such a string a P_n^k -cover. In this sense, the string 1234515241352 covers all ten 3-subsets of the alphabet $\{1, \dots, 5\}$; e.g., at the eighth position the occurring substring 241 is responsible for the coverage of $\{1, 2, 4\}$. For the classical SSP no rearrangements are allowed, instead a pre-specified collection of *strings* (rather than sets) over a finite alphabet is given, and a shortest string has to be found which contains all of them as contiguous substrings [1]. The algorithmic study and the computation of short so-called superstrings is motivated by applications in data compression [9] and computational biology [5]. In the latter field, one frequently faces the problem to reconstruct DNA sequences from a collection of fragments (oligonucleotides) which stem from experimentally obtained samples.

Related work. The P_n^k -cover problem has strong connections to the one of finding so-called *universal cycle coverings*, in which the strings are cyclic, i.e., the first element of the superstring is the successor of the last one. Recently, in [4], a necessary divisibility-condition on n and k for the existence of so-called *universal cycles* (where each k -subset is covered by one unique occurrence of a substring of length k) has been shown to be an asymptotically sufficient condition, too. This settled a longstanding conjecture of Chung, Diaconis, and Graham [2] formulated in 1989. Apart from selected values for n and k satisfying the particular divisibility condition, for which this conjecture is valid and yields optimal P_n^k -covers, the optima's structure is not fully understood so far. An upper bound of magnitude $\binom{n}{k} + O(n^{\lfloor k/2 \rfloor})$ has been obtained by a combinatorial construction [8]. Further bounds taking into account the magnitude of k as function of n can be found in [3] for the cyclic problem version. If we more generally demand coverage of any set X in a given subset P of the alphabet's power set, determining the existence of such a length- m superstring becomes NP-complete [6]. Computational approaches for the problem (also due to Lipski Jr. [7]) of superstrings of the entire power set of a finite alphabet are presented in [11], where also some selected results on P_n^k -covers are given for up to $n \leq 7$.

Our contribution. Pursuing an approach relying on the maximum overlap shared by suffixes and prefixes of strings as a distance metric, we design

particular instances of the Generalized Traveling Salesperson Problem (GTSP) whose feasible solutions, i.e., generalized Hamiltonian cycles, correspond to feasible superstrings for the P_n^k -cover problem. The vertex clusters correspond to the cardinality- k subsets and their elements to all $k!$ strings representing the subsets. For tractable, smaller values $k \leq n \leq 15$ we approach these derived GTSP instances by GLNS [10], an established metaheuristic software for the GTSP being a hybrid of Large Neighborhood Search and Simulated Annealing. The pursued approach, furthermore, allows to tackle a more general interpolation of the SSP and the P_n^k -cover problem by banning custom strings from selected clusters. With the use of intense computational resources, we thereby get insights into the size of (near-)optimal solutions and their structure.

For $k = 3$ and $n \in \{9, 12\}$ as well as $(n, k) = (8, 4)$ we are able to identify, to the best of our knowledge, so-far unknown minimum-length solutions; for other (n, k) -values a small optimality gap remains open, which we believe to be reducible in future by tailored (meta)heuristics that further strengthen the primal bounds. For some claims of optimality, we fall back on a—compared to the literature on P_n^k -covers—slightly strengthened purposefully derived dual bound

$$\beta_n^k := k + (n - k) \left\lceil \frac{\binom{n}{k}}{n} \right\rceil + \sum_{i=1}^k \left\lceil \frac{\binom{n}{k}}{n} - \frac{i}{k} \right\rceil.$$

References

1. Armen, C., Stein, C.: Short superstrings and the structure of overlapping strings. *Journal of Computational Biology* **2**(2), 307–332 (1995)
2. Chung, F.R.K., Diaconis, P., Graham, R.L.: Universal cycles for combinatorial structures. *Discrete Mathematics* **110**(1-3), 43–59 (1992)
3. Dębski, M., Lonc, Z.: Universal cycle packings and coverings for k -subsets of an n -set. *Graphs and Combinatorics* **32**(6), 2323–2337 (2016)
4. Glock, S., Joos, F., Kühn, D., Osthus, D.: Euler tours in hypergraphs. *Combinatorica* **40**(5), 679–690 (2020)
5. Kasprzak, M.: On the link between DNA sequencing and graph theory. *Computational Methods in Science and Technology* **10**(1), 39–47 (2004)
6. Kou, L.T.: Polynomial complete consecutive information retrieval problems. *SIAM Journal on Computing* **6**(1), 67–75 (1977)
7. Lipski Jr., W.: On strings containing all subsets as substrings. *Discrete Mathematics* **21**(3), 253–259 (1978)
8. Lonc, Z., Traczyk, T., Truszczyński, M.: Optimal f -graphs for the family of all k -subsets of an n -set. In: Ghosh, S.P., Kambayashi, Y., Lipski Jr., W. (eds.) *Data base file organization*. pp. 247–270. Academic Press (1983)
9. Mayne, A., James, E.B.: Information compression by factorising common strings. *The Computer Journal* **18**(2), 157–160 (1975)
10. Smith, S.L., Imeson, F.: GLNS: An effective large neighborhood search heuristic for the generalized traveling salesman problem. *Computers & Operations Research* **87**, 1–19 (2017)
11. Tabatabai, P.: On sequences covering subsets of a finite set. Master’s thesis, TU Graz (2018)

Function Libraries for Symbolic Regression with Genetic Programming

Philipp Fleck¹, Christian Haider¹, Maximilian Fruhmann², and Jan Zenisek¹

¹ University of Applied Sciences Upper Austria
Softwarepark 11, 4232 Hagenberg, Austria

{philipp.fleck, christian.haider, jan.zenisek}@fh-hagenberg.at

² K1-MET GmbH
Stahlstraße 14, 4020 Linz, Austria
maximilian.fruhmann@k1-met.com

1 Introduction

The task of Symbolic Regression (SR) is to find mathematical expressions that fit a given dataset [5]. Unlike linear regression, SR must identify both the structure of the equation and its numeric coefficients. Genetic Programming (GP) is the most prominent algorithm for this task, evolving expression trees [4]. Recent advances under the umbrella terms “physics-informed machine learning”, “knowledge-based modelling” or “hybrid modelling” have extended SR and other machine learning algorithms to incorporate prior knowledge about the modelled system [2, 3]. In SR for example, such knowledge can be integrated through shape constraints on model behavior [1].

Many engineering, physical and chemical domains have established equations describing relations between variables. To leverage this knowledge in SR, we introduce a *function library* concept containing predefined functions or parameterized templates, as illustrated in Figure 1. Domain experts curate these libraries with established equations, recurring sub-expressions or domain-specific functions.




Fig. 1. Schematic of the function library concept. A library of predefined functions is made available to the GP algorithm during evolution.

The function library serves as a knowledge repository for GP, accessible to genetic operators during evolution. In mutation, rather than replacing a randomly selected subtree with a newly generated one, we can insert a subtree from

the library. Likewise, during crossover we may insert a library subtree instead of swapping between two individuals. Figure 2 shows an operator using a library function with a placeholder.




Fig. 2. Example operator acting like crossover but inserting a library function. The child root is taken from an individual and a placeholder in the library subtree is filled by the remaining parent subtree.

This work makes two main contributions:

- First, we formalize the function library concept for GP-based SR, detailing function encoding and integration with GP through new genetic operators.
- Second, we conduct experimental analysis on benchmark problems using curated function libraries containing both useful and non-useful functions. We then compare GP performance with and without function libraries across different operator variants. For example, we evaluate whether GP should be allowed to modify library functions or whether they should remain immutable within evolved trees.

References

1. Haider, C., de França, F.O., Burlacu, B., Kronberger, G.: Using shape constraints for improving symbolic regression models. arXiv preprint arXiv:2107.09458 (2021)
2. Karniadakis, G.E., Kevrekidis, I.G., Lu, L., Perdikaris, P., Wang, S., Yang, L.: Physics-informed machine learning. *Nature Reviews Physics* **3**(6), 422–440 (2021)
3. Keren, L.S., Liberzon, A., Lazebnik, T.: A computational framework for physics-informed symbolic regression with straightforward integration of domain knowledge. *Scientific Reports* **13**(1), 1249 (2023)
4. Koza, J.R.: Genetic programming as a means for programming computers by natural selection. *Statistics and computing* **4**(2), 87–112 (1994)
5. Kronberger, G., Burlacu, B., Kommenda, M., Winkler, S.M., Affenzeller, M.: Symbolic regression. Chapman and Hall/CRC (2024)

Accelerators for metaheuristics

Czeslaw Smutnicki¹[0000-0003-4640-5364], Zbigniew
 Banaszak²[0000- 0001-7219-3903], and Grzegorz
 Bocewicz²[0000-0002-5181-2872]

¹ Wroclaw University of Science and Technology, Poland

² Koszalin University of Technology, Poland

*Czeslaw.Smutnicki@pwr.edu.pl, Zbigniew.Banaszak@tu.koszalin.pl,
 Grzegorz.Bocewicz@tu.koszalin.pl*

Abstract. We provide a short overview of our own methods and methodologies used to design accelerators. There are algorithms or their components dedicated to speed up metaheuristics. They were based chiefly on some properties of graphs. The proposed approaches have an impact on the advantageous properties of the algorithms in numerical tests.

Keywords: optimization · algorithm · graphs · accelerator · complexity

The accelerator means a special procedure (software) designed to reduce the computational cost of an algorithm or a component of the algorithm. It speeds up existing or newly designed algorithms. Notice that recent trends in optimization focus mainly on metaheuristics; see, for example, numerous references to evolutionary approaches in [4]. Accelerators are especially expected in this area.

We provide a short overview of existing approaches and discuss some of the methods and methodologies used to design accelerators. They were based chiefly on graphs and their properties. In many cases, these approaches lead to the “success” of an algorithm. The advantages from the use of an accelerator can be evaluated in several contexts: (1) reduction of the theoretical computational complexity, (2) aggregation and decomposition of calculations, (3) applying parallel computing, and (4) proper implementation of data structures and algorithm.

The first method recognized by us in the literature refers to the efficient implementation of the NEH algorithm from [6] used for the flow-shop scheduling problem. The original NEH algorithm with computational complexity $O(n^3m)$ was converted to the new version, which runs in $O(n^2m)$ time [12]. For instances with $n = 100$ tasks, the expected acceleration is approximately 50-100 times.

More significant advantages have been obtained for local search metaheuristics, which analyze the entire neighborhood, calculating the value of the goal function for close solutions [4]. The first among the accelerators proposed by us refers to the tabu search (TS) approach designed for the flow-shop scheduling problem [7]. Due to critical path properties, as well as an aggregation and decomposition of the calculations, the analysis of the entire INS neighborhood can

be performed in at most $O(n^2m)$ and $O(nm^2)$ time (for basic and advanced implementations, respectively), instead of the original $O(n^3m)$. Several accelerators for various types of neighborhood (API, NPI, INS) were described in [10].

Another, more advanced idea of the accelerator has been proposed in our paper [9]. For the job-shop scheduling problem, h solutions in the neighborhood can be checked in $O(h \cdot n \lg n)$ time instead of originally $O(h \cdot nm)$ time. Graph properties have also been used to accelerate neighborhood verification in TS for more complex scheduling problems [8]; the acceleration is up to 600 \times .

There are several other research directions associated with accelerators. The first is represented by [5]. He shows how to analyze the neighborhood of the exponential size in polynomial time. The second idea refers to parallel calculations [3]; it speeds up calculations, however, does not need to bring an accelerator. The third replaces the various optimization models in the cyclic scheduling by an iterative procedure using graphs [1, 2, 11]. Besides fast calculation of the goal function value, one can have interest in fast verification of some auxiliary elements, e.g. move status. The list of known accelerators is quite long. In each case considered, the verification of results should be performed on common benchmarks (see for example [13]), in adequately wide computer tests.

References

1. Banaszak Z., Bocewicz G., Smutnicki C., Graph-based feasibility screening for metaheuristic optimization under time uncertainty, *Journal of Intelligent & Fuzzy Systems*, 47(3), 5127–5140, 2024.
2. Bożejko W., Bocewicz G. (eds), Modelling Performance Analysis of Cyclic Systems, Studies of Systems, Decision and Control 241, Springer, 2020.
3. Bożejko, W. A new class of parallel scheduling algorithms. Oficyna Wydawnicza Politechniki Wrocławskiej, 2010.
4. Campelo F., EC-bestiarie <https://github.com/fcampelo/EC-Bestiarie>.
5. Congram, R. K., Potts, C. N., and van de Velde, S. L., An iterated dynasearch algorithm for the single-machine total weighted tardiness scheduling problem. *INFORMS Journal on Computing*, 14(1), 52-67, 2002.
6. Nawaz M., Enscore Jr. E.E., Ham I., A heuristic algorithm for the m -machine, n -job flow-shop sequencing problem. *OMEGA International Journal of Management Science*, 11, 91–95, 1983.
7. Nowicki E., Smutnicki C., A fast taboo search algorithm for the flow shop. *European Journal of Operational Research*, 91, 160–175, 1996.
8. Nowicki, E., Smutnicki, C. The flow shop with parallel machines: A taboo search approach. *European Journal of Operational Research*, 106(2-3), 226-253, 1998.
9. Nowicki, E., Smutnicki C., An advanced taboo search algorithm for the job shop problem. *Journal of Scheduling*, 8, 145-159, 2005.
10. Smutnicki C., Algorytmy szeregowania, EXIT, 2002.
11. Bożejko W., Pempera J., Smutnicki C., Wodecki M. (eds) Zaawansowane modele i algorytmy optymalizacji w systemach cyklicznych, EXIT, 2017.
12. Taillard E., Some efficient heuristic methods for flow shop sequencing. *European Journal of Operational Research*, 47, 64–74, 1990.
13. Taillard E., Benchmarks for basic scheduling problems, *European Journal of Operational Research*, 64, 278–285, 1993.

On Solving the Multiple Variable Gapped Longest Common Subsequence Problem

Marko Djukanović^{1,2,6}, Nikola Balaban², Christian Blum³, Aleksandar Kartelj⁴, Sašo Džeroski⁵, and Žiga Zebec⁶

¹ University of Nova Gorica, Nova Gorica, Slovenia
 marko.djukanovic@ung.si

² Faculty of Natural Sciences and Mathematics, University of Banja Luka, Banja Luka, Bosnia and Herzegovina
 nikola.balaban@student.pmf.unibl.org, marko.djukanovic@pmf.unibl.org

³ Artificial Intelligence Research Institute (IIIA-CSIC), Barcelona, Spain
 christian.blum@iiia.csic.es

⁴ Faculty of Mathematics, University of Belgrade, Belgrade, Serbia
 kartelj@matf.rs

⁵ Jožef Stefan Institute, Ljubljana, Slovenia
 saso.dzeroski@ijs.si

⁶ Institute of Information Sciences (IZUM), Maribor, Slovenia
 ziga.zebec@izum.si

The *Longest Common Subsequence Problem* (LCSP) [1] is a well-known combinatorial optimization problem with numerous applications in computational biology, particularly in the analysis of molecular sequences. Given an arbitrary set $S = \{s_1, \dots, s_m\}$ of input sequences over an alphabet Σ , the goal is to determine the longest possible subsequence common to all $s_i \in S$. In recent decades, many practically motivated variants of the LCSP have been proposed by considering additional structural or biological constraints such as the constrained, arc-preserving, repetition-free version, etc. This work is focused on solving the *Variable Gapped LCSP* (VGLCSP), introduced in [3] and theoretically examined in [2]. This problem extends the standard LCSP by imposing flexible distance limits (“gaps”) between consecutive symbols in the resulting subsequence, allowing these gaps to vary along the input sequences. Formally, for each input sequence $s_i \in S$, a gap function $G_{s_i}: \{1, \dots, |s_i|\} \mapsto \mathbb{N}$ defines the maximum distance allowed between consecutive characters of a subsequence. That is, if the characters of a subsequence s appear at the positions $i_i^1 < \dots < i_i^{|s|}$ in s_i , we say that the gap constraint is satisfied if $i_i^x - i_i^{x-1} \leq G_{s_i}[i_i^x] + 1$ for all $x = 2, \dots, |s|$. A common subsequence s is considered feasible if the corresponding gap constraints hold across all sequences in S . The VGLCSP represents a flexible and realistic model of sequence alignment, which is particularly useful for comparing DNA or protein sequences with variable structural distances. Although exact dynamic programming approaches exist for the two-sequence case ($m = 2$), their computational complexity becomes prohibitive for larger m . To the best of our knowledge, no effective exact or approximate approaches currently exist for the generalized VGLCSP, a clearly NP-hard problem.

Building on the well-established state-space representation of the LCSP [1], we propose a general *state graph* formulation for the VGLCSP. Each node rep-

resents a partial feasible subsequence, characterized by the vector of current positions in each input sequence and the subsequence length. Specifically, a state $v = (\mathbf{p}^{L,v}, l^v)$ is induced by a partial solution s^v if: (i) $p_i^{L,v} - 1$ is the smallest index such that s^v is a subsequence of the prefix string $s_i[1] \cdots s_i[p_i^{L,v} - 1]$; (ii) $l^v = |s^v|$; and (iii) all the gaps G_{s_i} are preserved with s^v . A directed arc $\alpha = (v_1 v_2)$, labeled $lett(\alpha) = a \in \Sigma$, exists between states v_1 and v_2 if $l^{v_2} = l^{v_1} + 1$ and $s^{v_2} = s^{v_1} \cdot a$. To expand a state v , all letters that occur in each suffix of s_i defined by respective positions $p_i^{L,v}$ are considered. For each such letter a , the minimal positions $p_i^{L,a} \geq p_i^{L,v}$ of its appearances are identified fulfilling $p_i^{L,a} - p_i^{L,v} \leq G_{s_i}(p_i^{L,a})$. The resulting child node, if not suboptimal, is given by $w = (p_i^{L,a} + 1, l^v + 1)$. The root node $r := ((1, \dots, 1), 0)$ represents the empty subsequence, while goal states correspond to non-extendable feasible subsequences. Unlike the standard LCSP, the VGLCSP admits multiple (possibly exponentially many) root nodes. For example, in sequences $s_1 = \text{ATGGAAAA}$ and $s_2 = \text{ATCCAAAA}$ with $G_{s_1} = G_{s_2} = 1$, the state $((1, 1), 0)$ cannot reach any state with the position vector $(5, 5)$ via transitions, missing optimal subsequence AAA. This motivates the methodology of *Iterative Multi-Source Beam Search* (IMSBS). BS is a popular meta-heuristic approach that performs in a breadth-first-search manner, controlled by the size of beam β and a heuristic h . In the IMSBS, BS is iteratively executed with a set of promising root nodes as an initial beam, selected from a dynamically maintained global pool of root nodes, controlling diversification of the approach. Each iteration performs a complete BS, guided by a heuristic h based on letter frequencies, minimal residual substring lengths, or probability-weighted estimates.

Preliminary experiments on synthetically generated instances demonstrate that IMSBS consistently outperforms the single-root BS baseline initialized at $r = (\mathbf{1}, 0)$. The results reveal clear benefits of multi-source exploration and highlight differences between heuristic strategies in terms of efficiency and solution quality. These findings open promising directions for designing more advanced hybrid and learning-assisted approaches to tackle the generalized VGLCSP.

Acknowledgments. This publication is co-funded by the European Union's Horizon Europe research and innovation program under the Marie Skłodowska-Curie COFUND Postdoctoral Programme grant agreement No.101081355– SMASH and by the Republic of Slovenia and the European Union from the European Regional Development Fund. Co-funded by European Union. Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or European Research Executive Agency (REA). Neither the European Union nor the REA can be held responsible for them.

References

1. Finding longest common subsequences: New anytime A* search results. *Applied Soft Computing* **95**, 106499 (2020)
2. Adamson, D., Kosche, M., Koß, T., Manea, F., Siemer, S.: Longest common subsequence with gap constraints. In: *International Conference on Combinatorics on Words*. pp. 60–76. Springer (2023)
3. Penga, Y.H., Yangb, C.B.: The longest common subsequence problem with variable gapped constraints. In: *Proceedings of the 28th Workshop on Combinatorial Mathematics and Computation Theory*, Penghu, Taiwan. pp. 17–23 (2011)

Improving Sign Language Recognition with Synthetic Oversampling

Mariam Quesada-Ojeda[✉], Ruth Holmes[✉], and Anthony Ventresque[✉]

School of Computer Science and Statistics, Trinity College Dublin {quesadam, ruth.holmes, anthony.ventresque}@tcd.ie

Sign Language (SL) is the primary mode of communication for many Deaf and Hard-of-Hearing individuals, with the World Federation of the Deaf (WFD) estimating that more than 70 million people belong to Deaf signing communities worldwide [3]. SL is not a single universal language, but rather a collective term for the over 200 distinct languages [3] using the visual-gestural modality of communication.

Sign Language Recognition (SLR) is a field of Computational Linguistics that aims to identify and extract linguistic features and patterns from visual data of SL in use. Naturally, recent work increasingly draws on Machine Learning (ML) techniques that require large-scale datasets to detect these underlying patterns effectively. However, SL datasets are often scarce and limited in both size and scope – an issue that is amplified in low-resource settings such as Irish Sign Language (ISL) [5]. Furthermore, SL datasets often present severely imbalanced sign frequencies [6], a pattern typical of natural language data but especially impactful to SLR. This pronounced imbalance means that many key vocabulary items occur only rarely, if at all, constraining the capability of SLR models to learn and generalise across the long tail of low-frequency sign categories.

Our approach is motivated by two key challenges in SLR: learning from scarce corpora and achieving robust recognition under heavily skewed class distributions. While resampling techniques are widely used in ML in general [4], they are not straightforward to apply to SLR. Undersampling discards already limited data, while naive oversampling such as Random Oversampling increases minority frequency but adds no variation and risks overfitting.

Data augmentation offers a more principled alternative. SMOTE [2] and its variants generate synthetic minority samples in feature space, enriching class diversity. However, despite its success in other areas, SMOTE has seen limited application in SLR, and existing work uses it only as an auxiliary component within larger pipelines [1,8]. Furthermore, to the best of our knowledge no SLR work has positioned class imbalance as the primary research focus. Therefore, our work directly addresses this gap via the first architecture-agnostic, systematic evaluation of multiple SMOTE variants, including decision-boundary reinforcement via Borderline-SMOTE [4] and support vector guided synthesis using SVM-SMOTE [7], alongside Random Oversampling as a comparative baseline.

Our experiments use keypoint-based representations of the Irish Sign Language Hand-Shape (ISL-HS) dataset. Leveraging its relatively balanced class distribution, we first establish an empirical best-case scenario or *Upper Baseline* (UB) which represents an idealised reference point for model performance. We

then develop a *Lower Baseline* (LB) scenario by undersampling class frequencies according to Zipf's Law [9] to represent a realistic benchmark of performance under conditions of severe class imbalance. Finally, to this LB data, we apply each of our chosen oversampling algorithms to generate multiple synthetically *re-balanced* scenarios. Each scenario is evaluated across six standard classifiers to assess their effect on recognition performance.

Overall, the experiments show that oversampling techniques consistently improve performance compared to the imbalanced LB, with an average increase of 33.85% in macro-F1 score. Although no single technique emerges as a clear standout across all classification models, all *do* show an encouraging trend of strong and reliable performance. These findings confirm that oversampling is an effective approach for mitigating class imbalance and improving recognition of minority hand shapes in SLR.

References

1. Alani, A.A., Cosma, G.: Arsl-cnn: a convolutional neural network for arabic sign language gesture recognition. *Indonesian journal of electrical engineering and computer science* **22**(2), 1096–1107 (2021)
2. Chawla, N.V., Bowyer, K.W., Hall, L.O., Kegelmeyer, W.P.: Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research* **16**, 321–357 (2002)
3. of the Deaf, W.F.: Wfd – frequently asked questions. <https://wfdeaf.org/contact/faqs/> (nd), last accessed: 2025-12-01
4. Han, H., Wang, W.Y., Mao, B.H.: Borderline-smote: a new over-sampling method in imbalanced data sets learning. In: *International conference on intelligent computing*. pp. 878–887. Springer (2005)
5. Holmes, R., Rushe, E., De Coster, M., Bonnaerens, M., Satoh, S., Sugimoto, A., Ventresque, A.: From scarcity to understanding: Transfer learning for the extremely low resource irish sign language. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2008–2017 (2023)
6. Kimchi, I., Wolters, L., Stamp, R., Arnon, I.: Evidence of zipfian distributions in three sign languages. *Gesture* **22**(2), 154–188 (2023)
7. Nguyen, H.M., Cooper, E.W., Kamei, K.: Borderline over-sampling for imbalanced data classification. *International Journal of Knowledge Engineering and Soft Data Paradigms* **3**(1), 4–21 (2011)
8. Passos, W.L., Araujo, G.M., Gois, J.N., de Lima, A.A.: A gait energy image-based system for brazilian sign language recognition. *IEEE Transactions on Circuits and Systems I: Regular Papers* **68**(11), 4761–4771 (2021)
9. Piantadosi, S.T.: Zipf's word frequency law in natural language: A critical review and future directions. *Psychonomic bulletin & review* **21**, 1112–1130 (2014)

Electric Motor Cooling System Evaluation Methodology

Moritz Golombek¹ and Thomas Schlechter²,

¹ University of Applied Sciences Upper Austria, Wels, Austria
 moritz.golombek@students.fh-wels.at

² University of Applied Sciences Upper Austria, Wels, Austria
 thomas.schlechter@ieee.org

Abstract. Keywords: Electric Motors · Cooling Systems · Development Process.

Electric Vehicles (EVs) have already become a common sight on public roads, but in the competitive context of racing they have not yet gained a similar status. A global push for a higher recognition in sustainability Key Performance Indicators (KPIs) motivates race series to change this [1]. This use case subjects the powertrain to a high frequency of peak load acceleration events with minimal cool-down periods in between. This results in a demand for highly effective thermal management systems in the powertrain.

Electric motors convert electric energy input into mechanical power output through the interplay between the magnetic fields of the stator and rotor. This conversion underlies a varying efficiency resulting in thermal losses. Therefore, two fundamental classification parameters for electric motors are their peak and continuous output power. While the peak power is defined by the electromagnetic design, the continuous power is primarily defined by the cooling system's capability of extracting the thermal losses and keeping the electric motor within its operational temperature range. At the same time vehicle-level efforts to maximize dynamic KPIs impose dimensional constraints. Cooling systems for high performance EVs, therefore, must balance maximum heat extraction capabilities with minimal associated packaging and weight to accomplish the aforementioned goals [2]. To address this, the perception of thermal management must shift from an add-on, considered at a later development stage, to an integral part of the vehicle development process. This article describes a multidisciplinary characterization approach for electric motor cooling systems to support this.

While requirements for a cooling system are highly dependent on the interfacing components, the development can be parallelized by calibrating boundary conditions at an early stage and therefore shaping cooling requirements. The basis for this evaluation is built by localizing and quantifying explicit causes for thermal losses present in an electromagnetic design. Losses are attributed to ohmic resistance and alternating current effects in windings and a combination of hysteresis and eddy current losses in magnets and cores. The loss structure largely depends on the operating speed and phase current. The quantitatively

biggest losses are localized in the windings and therefore they are focused for the subsequent cooling efforts [3]. An impact analysis of elevated temperatures on the output performance correlates a continuous decrease of the electric motor efficiency to increasing component temperatures and lists causes for irreversible damages when magnet or conductor limits are exceeded.

After the significance of the thermal management is clarified, thermodynamic principles are reviewed to explain the relevance of factors including the heat transfer coefficient and the thermal conductivity to the cooling performance. Because they are difficult to estimate mathematically, especially in cooling methods using liquid coolant mediums [4], an integration of experimental testing with thermal simulation models is proposed to create an accurate model that can be used for the development of the cooling system. The goal being an incremental broadening of the test scope to isolate validation parameters. First, the loss structure of an electric motor model is compared to a prototype in several electromagnetic KPIs to compensate for manufacturing-related uncertainties. A Dry Rise Test (DRT) is then conducted where the thermal inertia in the electric motor is monitored to identify thermal resistances and masses between the mechanical components of the system. This is done without any coolant medium present. After the simulation model is calibrated to the output the Wet Rise Test (WRT) is conducted with the cooling system filled and fully operational. The delta between the DRT and WRT allows the contribution of the liquid coolant to be isolated and evaluated. The output is in turn used to hone the simulation model, allowing optimized matching of the cooling system components to the performance demands while minimizing the weight and size penalty for the vehicle.

The presented methodology allows the parallel development of electric motor cooling systems by combining loss localization and quantification in electromagnetic designs with an iterative validation loop between thermal simulation and experimental testing. This enables the integration of the thermal management system into the vehicle's design from the beginning, allowing for optimized performance in the competitive context of racing EVs.

References

1. A review of the disclosed sustainability credentials of 91 global two and four wheeled motorsport championships. Tech. Rep. 4, Enovation Consulting Ltd (Mar 2025)
2. Wang, X., Yan, Y., Li, Y.: Study on high-speed electric motor cooling with oil spray. *e-Prime - Advances in Electrical Engineering, Electronics and Energy* **4**, 100170 (Jun 2023). <https://doi.org/10.1016/j.prime.2023.100170>
3. Selvin Raj, J.A.P., Asirvatham, L.G., Angeline, A.A., Manova, S., Rakshith, B.L., Bose, J.R., Mahian, O., Wongwises, S.: Thermal management strategies and power ratings of electric vehicle motors. *Renewable and Sustainable Energy Reviews* **189**, 113874 (Jan 2024). <https://doi.org/10.1016/j.rser.2023.113874>
4. Ghahfarokhi, P.S., Podgornovs, A., Kallaste, A., Cardoso, A.J.M., Belahcen, A., Vaimann, T.: The Oil Spray Cooling System of Automotive Traction Motors: The State of the Art. *IEEE Transactions on Transportation Electrification* **9**(1), 428–451 (Mar 2023). <https://doi.org/10.1109/TTE.2022.3189596>

Back Projection: Exterior Lighting Application for Automotive Industry

Yugesh Balaji¹, Medina Dobrača¹, Hassan Haraz¹ and Thomas Schlechter²

¹ University of Applied Sciences Upper Austria, 4600 Wels, Austria

{yugesh.balaji, medina.dobraca, hassan.haraz}@students.fh-wels.at

² University of Applied Sciences Upper Austria, 4600 Wels, Austria

thomas.schlechter@ieee.org

Keywords: Rear-Projection Systems · Automotive Lighting · Digital Light Processing · Vehicle-to-Everything Communication · Exterior Lighting Design

The transition to electric vehicles has introduced new spatial and design possibilities within the automotive sector, particularly with regard to front-end architecture. This paper presents a study of a compact rear-projection system designed to display dynamic content through the vehicle grille. Based on an initial ZKW Digital Art Screen (DAS) prototype that demonstrated high brightness and image quality, but was impractical to scale up due to its size and the complexity of its optical alignment, the aim of this work was to reduce the system volume by at least 70% while maintaining a wide field of view (150°) and a brightness of at least 10,000 nits.

The project was conceived as a realistic automotive-grade solution capable of operating reliably under environmental conditions ranging from -40°C to 120°C, as well as prolonged exposure to ultraviolet radiation and vibration. Regulatory requirements introduced significant constraints that directly shaped the system design. These regulations include the maximum 70 mm distance from the main daytime running light and restrictions on illuminated content while the vehicle is in motion. Such conditions influenced both material and layout decisions, as well as the definition of possible use cases. Consequently, the project focused on non-textual, symbol-based projections and static or low-speed scenarios, such as parked vehicles or low-speed autonomous operation. From a human-centered perspective, incorporating lighting-based external human-machine interfaces (eHMIs) significantly improves the communication of vehicle intent to pedestrians. Empirical studies have shown that the congruent use of light-based eHMI and vehicle behavior increases pedestrians' trust in and understanding of autonomous or highly automated vehicles, essential to ADAS system safety [1].

To achieve these objectives, multiple light-emission and projection surface technologies were systematically evaluated, focusing on optical efficiency, thermal performance, integration depth and regulatory compliance. In collaboration with ZKW, iterative trade-off analyses were conducted to enable a structured comparison of projection architectures, light sources and screen materials. The eval-

uation incorporated Digital Light Processing (DLP) modules [7], Liquid Crystal on Silicon systems [2], laser scanning solutions and Light Emitting Plasma sources [4], as well as screen technologies such as Optical Diffusive Projection Films [6] and Switchable Polymer-Displayed Liquid Crystal Smart Films [5]. Recent work has shown that use of diffractive optical elements and engineered exit-pupil/diffuser assemblies can allow designers to reach wide fields of view while keeping optical stack depth and volume to a minimum [3].

Two viable concept solutions were developed. The first solution employs a high-power DLP projection system combined with phosphor-converted white LEDs and a polycarbonate projection surface. The second solution uses a mirrorless configuration in which a centralized light engine delivers illumination via fiber-optic routing to the projection surface. Both systems were analyzed in terms of optical output and thermal management under realistic automotive boundary conditions. The results of this project demonstrate that compact rear-projection systems are technically feasible and capable of transforming exterior vehicle lighting into an intelligent interface for communication, safety, personalization, and potential Vehicle-to-Everything interaction. This paper highlights the potential of projection technologies to simplify system design while broadening the functional capabilities of next-generation vehicles.

References

1. Dey, D., Matviienko, A., Berger, M., Pfleging, B., Martens, M., Terken, J.: Communicating the intention of an automated vehicle to pedestrians: The contributions of ehmi and vehicle behavior. *it - Information Technology* **63**(2), 123–141 (2021), <https://doi.org/10.1515/itit-2020-0025>
2. Henkel AG & Co. KGaA: Creating the automotive displays of the future today (2021), <https://www.henkel.com/press-and-media/press-releases-and-kits/2021-07-05-creating-the-automotive-displays-of-the-future-today-1252124>, last accessed: 2025-06-15
3. Khan, M.S., Saeed, W.M., Roth, B., Lachmayer, R.: Diffractive optics based automotive lighting system. *Advanced Optical Technologies* **10**(1), 49–57 (2021), <https://doi.org/10.1515/aot-2020-0055>
4. Lighting Research Center: Plasma lighting systems (2013), <https://www.lrc.rpi.edu/programs/NLPIP/lightingAnswers/pdf/print/PlasmaLightingSystems.pdf>, last accessed: 2025-06-25
5. Pro Display: Switchable Smart Film, <https://prodisplay.com/smart-glass/switchable-smart-film/>, last accessed: 2025-06-16
6. Pro Display: Pro Diffusion Rear Projection Film (2022), <https://prodisplay.com/projection-screens/rear-projection/rear-projection-film/>, last accessed: 2025-06-16
7. Texas Instruments: DLP Technology Overview, https://www.ti.com/dlp-chip/overview.html?keyMatch=DLP&tisearch=universal_search, last accessed: 2025-06-24

Deep Learning Based Object Recognition Approach for Motorbike Dashboard Software Validation

Seifeldien Hussein¹ and Thomas Schlechter²

¹ University of Applied Sciences Upper Austria, 4600 Wels, Austria
 seifeldienhussein@gmail.com

² University of Applied Sciences Upper Austria, 4600 Wels, Austria
 thomas.schlechter@ieee.org

Keywords: Deep Learning · Object Detection · Software Validation.

With the current shift in the motorbike industry toward the direction of software defined vehicle, dashboard software that provides the rider with a graphical user interface to monitor and control the bike has gained a huge attention from original equipment manufacturers in the past years to continuously provide a seamless and safe ride experience. As a result, the dashboard market is expected to grow from 10.71\$ billion to 18.64\$ billion in the time frame from 2021 to 2028 with a compound annual growth rate of 8.2% [1]. Therefore, ensuring its reliability and functionality has become a significant challenge in the automotive industry as manual validation is time consuming, expensive and error prone.

This study explores the application of deep learning based object detection techniques to automate the dashboard software validation process. To implement the proposed testing environment, a camera system is integrated with a hardware in the loop setup. The system captures and processes dashboard images using a deep learning model, which detects key elements like check engine light and compares their existence and location to reference values for validation. This approach significantly enhances testing accuracy, speed and repeatability.

The study begins by reviewing traditional object detection methods and their limitations. It then examines modern deep learning based techniques, with a focus on one-stage and two-stage detectors. One-stage detectors are found to offer a balance between speed and accuracy which makes them suitable for dashboard validation [2]. Additionally, the study highlights the importance of camera system selection. It was found that high resolution and high dynamic range sensors improve detection performance [3] [4]. Other factors such as dataset preparation, environment control and software architecture are also discussed in details.

The proposed methodology is structured around three milestones. It covers the implementation of the test setup, the training of the detection model and the final system integration. A modular and scalable standalone application was developed to handle image capturing, detection and result comparison. Additionally, it was successfully integrated with a hardware in the loop setup for comprehensive testing.

The experimental results were evaluated over a dataset composed of camera images and Photoshop exports which were artificially augmented to simulate four different operating conditions including varying camera angles to simulate different testing setups, increasing temperature which results in sensor noise [5], a combined scenario with both factors and a normal operating condition scenario. The results demonstrated an overall average accuracy of 99% and precision of 100% in detecting dashboard elements across all testing scenarios.

The initial findings of this study demonstrate that deep learning based object detection, combined with optimized camera system and hardware in the loop setup, presents a reliable and scalable solution for dashboard software validation with a very satisfactory accuracy. Given that the test dataset was composed partially of Photoshop exports and is not sufficient to draw a complete conclusion, future research may focus on expanding the training and testing dataset. Additionally, adding new validation layers to test colors and read texts using optical character recognition. Finally, enabling real time processing to test the correctness of the animations could be part of the future research.

References

1. A. X. Qun Pang, J. Jansen Lin, C. H. Ong, et al.: Deep learning based layout recognition approach for HMI software validation. In: 2024 IEEE Conference on Artificial Intelligence (CAI), pp. 429–437 (2024). <https://doi.org/10.1109/CAI59869.2024.00085>
2. J. Kaur and W. Singh: Tools, techniques, datasets and application areas for object detection in an image: A review. *Multimedia Tools and Applications* **81**(27), 38297–38351 (2022). <https://doi.org/10.1007/s11042-022-13153-y>
3. V. Thambawita, I. Strümke, S. A. Hicks, P. Halvorsen, S. Parasa, and M. A. Riegler: Impact of image resolution on deep learning performance in endoscopy image classification: An experimental study using a large dataset of endoscopic images. *Diagnostics* **11**(12), 2183 (2021). <https://doi.org/10.3390/diagnostics11122183>
4. R. Mukherjee, M. Bessa, P. Melo-Pinto, and A. Chalmers: AObject detection under challenging lighting conditions using high dynamic range imagery. *IEEE Access* **9**, 77771–77783 (2021). <https://doi.org/10.1109/ACCESS.2021.3082293>
5. Cognex In-Sight 9000 Series Vision System Manual, <https://www.support.cognex.com/docs>, last accessed 2025/07/09

Digital Twin-Driven Predictive Maintenance for Intelligent Mechatronic Product Development

Danut Kirchsteiger ^{1[0009-0005-5331-1077]} and Mario Jungwirth ^{1[0000-0002-1652-8745]}

¹ University of Applied Sciences Upper Austria, Wels Campus, 4600 Wels, Austria
mario.jungwirth@fh-wels.at

Abstract. This study presents a digital twin-based approach for predictive maintenance in mechatronic product development. A high-fidelity virtual model of an electromechanical drivetrain - comprising a DC motor, shaft, bearings, and an imbalance disk - was developed using Simscape Multibody to replicate a laboratory demonstrator. The model enables the generation of synthetic sensor data under varied operating conditions, addressing the challenge of acquiring labelled fault data from physical systems. Validation through time- and frequency-domain analysis confirmed the model's accuracy, with vibration amplitude deviations below three percent. A machine learning classifier trained exclusively on synthetic data successfully detected imbalance conditions in the real system. These findings demonstrate the potential of digital twins to accelerate prototyping, support intelligent monitoring, and reduce development costs. Future work will expand fault coverage and sensor integration to further enhance predictive capabilities.

Keywords: Digital Twin, Mechatronic Product Development, Virtual Prototyping, Predictive Maintenance

1 Introduction

In the evolving landscape of mechatronic product development, digital twins are increasingly recognised as key enablers for accelerating design, validation, and system intelligence [1]. This paper presents the development of a digital twin for an electromechanical drivetrain shown in Fig. 1, aimed at supporting predictive maintenance strategies during early product development phases [2]. The model, built in Simscape Multi-body, replicates a laboratory demonstrator comprising a DC motor, shaft, bearings, and an imbalance disk. It enables the generation of synthetic sensor data under controlled and varied operating conditions, addressing the challenge of acquiring extensive labelled fault data from physical systems. This approach facilitates the integration of machine learning-based fault detection into the design process, reducing reliance on costly and time-consuming experimental setups [3].

2 Methodology

To ensure the fidelity of the digital twin, a comprehensive validation was performed using time-domain and frequency-domain analyses [2]. The dominant frequency peak, a critical indicator of imbalance in rotating systems, was used to calibrate the model. The simulated and measured signals showed excellent agreement, with vibration amplitude deviations below three percent. This validated environment was then used to generate synthetic datasets for training a machine learning classifier in MATLAB, capable of distinguishing between balanced and unbalanced states. Notably, the model trained exclusively on synthetic data successfully identified real im-balance conditions in the physical system, demonstrating the effectiveness of simulation-based data generation for predictive maintenance applications.




Fig. 1. Electrical drivetrain demonstrator [2]

3 Conclusion

This study demonstrates how digital twins can serve as both virtual prototypes and intelligent data sources within mechatronic product development. By enabling scalable and cost-efficient fault detection, the approach supports faster prototyping, iterative design refinement, and the deployment of predictive maintenance features across the product lifecycle. The integration of simulation and machine learning enhances system reliability and reduces development overhead [1]. Future work will extend the model to additional fault types and sensor modalities, further strengthening its role in intelligent product design and lifecycle management.

References

1. M. Saraswat, et.al, Congress on Intelligent Systems, Springer, Wiesbaden (2021)
2. D. Kirchsteiger, Simulationsgestützte Unwuchtdetektion in elektrischen Maschinen auf Basis synthetischer Daten und maschinellen Lernens, Masterthesis, Wels (2025)
3. G. Peter, S. Christian, G. Andreas und T. Stefan, Automatisierung von Geschäftsprozessen im Maschinen- und Anlagenbau – Fallstudie zu Predictive Maintenance in IoT-Best Practices, pp. 195–206, Springer Nature, Wiesbaden (2021)

Developing a Structured Database Using MariaDB to Store and Manage Real-Time Process Data Delivered by an MQTT Broker

James Okochi¹ Thomas Schlechter¹

University of Applied Sciences Upper Austria, 4600 Wels, Austria

james.okochi@students.fh-wels.at, thomas.schlechter@ieee.org

Keywords: Industrial Engineering · Learning Factory · IIoT · Digital Twin

The future of industry and education currently faces significant indicators of change. Industry seeks new talents, while the upcoming talents foster new means of worklife and as well attitude towards the means of education. Traditional teaching concepts seem deprecated. Both industry and academia are currently undergoing a transformative process towards novel product development and educational mantras. In this work we hit multiple flies with one strike. First, we focus on the actual development of a Learning Factory, which facilitates the education of the next generation of industrial engineering students. Second, we include in this development the use of artificial intelligence (AI) to showcase the product development approach of the future. Third, we spotlight the result of the Learning Factory development in the context of the educational targets within the LEONARDO Erasmus+ Project, which focuses on industrial 5.0 higher educational students enablement.

To turn more concrete, this article explores the transformative potential of Learning Factories in mechatronic systems development. Learning Factories offer a dynamic, collaborative environment that bridges the gap between academia and industry, creating a mutually beneficial ecosystem. [1, 3] The LEONARDO project aims to develop innovative teaching methods, materials and tools for human-centric industrial engineering and management education leveraging on an industry 5.0 replica of a brewing system. Brewing as a process can be considered as highly complex, while brewing as a procedure serves as a "sexy vehicle" for appealing students' interest in industry 5.0 applications and human-centric production. The brewing process is and will increasingly be more automated and highly supervised. [5] For the latter, modern implementations of sensors such as electronic nose, electronic tongue, and infrared spectroscopy are required to be installed on the brewing equipment. To efficiently use the sensor outputs, the produced signals need to be merged locally and processed adequately, researched and investigated deeply by the authors up-front with the results to be summarized. [4] Furthermore, to enable the physical bridging of various involved institutions across Europe, connecting the relevant sites virtually presents another technological challenge. Adequate IoT equipment needs to be selected and included in the whole setup as well.

In the given context, this paper presents the development of a structured MariaDB database for storing and managing real-time process data delivered via an MQTT broker. [2] The implementation targeted industrial applications and was demonstrated using a brewery process simulation. Two Python-based ingestion scripts - one manually written and the other AI generated - were deployed under identical 30-minute continuous simulation conditions to compare throughput, reliability, resource usage, and maintainability. The human-written script ingested 10,820 messages (6.01msg/s) with only 0.0164% data loss, achieving low CPU usage (13.5%) and RAM consumption (78.3MB), alongside minimal runtime errors (3 total). The AI-generated script ingested 9,305 messages (5.18msg/s) with higher data loss (0.1541%), greater CPU (19.2%) and RAM usage (106.7MB), and significantly more runtime errors (37 total). Code evaluation by professional developers yielded higher scores for the human-written approach (average = 4.75) compared to the AI-generated one (average = 2.5). Findings show that while AI-assisted development accelerates implementation, it requires refinement to match the stability, efficiency, and data integrity of manually crafted solutions. The results underscore the value of human expertise in building robust, resource-efficient, and fault-tolerant ingestion systems for real-time industrial data environments.



This work is supported by the LEONARDO project, funded by the European Commission under the Erasmus+ program KA-220 Cooperation Partnerships for Higher Education - No. 2023-1-IT02-KA220-HED-000164699. The views expressed belong to the author(s) alone and do not necessarily reflect the views of the European Union or the Erasmus+ National Agency INDIRE. Neither the European Union nor the granting administration can be held responsible for them.

References

1. Abele, E., Chryssolouris, G., Sihn, W., Metternich, J., ElMaraghy, H., Seliger, G., Sivard, G., ElMaraghy, W., Hummel, V., Tisch, M., Seifermann, S.: Learning factories for future oriented research and education in manufacturing. *CIRP Annals* **66**(2), pp. 803–826 (2017)
2. Bevywise: MQTT Manufacturing Solutions for Industry 4.0 Efficiency, <https://www.bevywise.com/mqtt-usecases/manufacturing-solutions.html>
3. Hulla, M., Hammer, M., Karre, H., Ramsauer, C.: A case study based digitalization training for learning factories. *Procedia Manufacturing* **31**, pp. 169–174 (2019), research. Experience. Education. 9th Conference on Learning Factories 2019 (CLF 2019), Braunschweig, Germany
4. Schlechter, T.: Impact of AI on the Brewing Industry: A Comprehensive Summary. *BrewingScience* **77**(3-4), pp. 18–32 (2024), <https://brewingscience.de/index.php/brewingscience/article/view/197>
5. Yu, W.: Design of real-time temperature monitoring system in liquor brewing process. *Advances in Engineering Innovation* **16**(3), pp. 44–48 (2025). <https://doi.org/10.54254/2977-3903/2025.21915>

Development of an Intel RealSense-based ArUco-Marker System for Angle Measurements on a Ski Simulator

Christina Hofer¹, Florian Hammer¹, and Veronika Putz¹

Linz Center of Mechatronics GmbH, 4040 Linz, Austria

christina.hofer@lcm.at

<https://www.lcm.at/>

Keywords: Marker Detection · Human-Machine-Interaction · Computer Vision

The precise detection and localization of ArUco markers using depth cameras represent a central research topic in mobile robotics, augmented reality and motion capture systems. Previous studies have demonstrated that the combination of RGB and depth data significantly enhances the accuracy of marker localization, particularly in dynamic environments [1, 2]. Open-source implementations and ROS-based systems illustrate the practical applicability of this approach in mobile robots and autonomous platforms [3]. Furthermore, complementary studies on 3D marker placement and graph-based mapping techniques provide valuable contributions to improving spatial orientation and navigation [4]. In addition, research on cyber-physical mobile arm gesture recognition using ultrasound and motion data highlights the integration of multimodal sensing for advanced interaction paradigms [5].

This study focuses on a low-cost motion capture system for a ski boot test stand designed to measure the correct positioning, key angles of the boot and test participant during skiing. The system enables precise angle measurements under both static and highly dynamic conditions and represents a novel implementation of real-time kinematic tracking during a simulated skiing motion. By integrating ArUco-based pose estimation with depth sensing using an Intel RealSense D435i camera, the system acquires spatial coordinates at multiple points and processes them through a dedicated algorithm to compute dynamic angular and kinematic parameters with high precision.

To evaluate the experimental setup, key criteria such as camera position, viewing distance, field of view and marker placement were first defined and validated through preliminary calibration tests. The camera was centrally positioned at an optimal distance to ensure that all markers were clearly visible within the field of view. For initial tests, printed ArUco markers of various sizes (2–6 cm) were first mounted on a cardboard sheet attached to the lower extremities. Based on preliminary measurements, a comprehensive test trial was conducted with a subject positioned on the ski simulator wearing ski boots. Before initiating motion sequences simulating ski turns, ArUco markers were placed at predefined anatomical and mechanical locations to ensure accurate tracking. One marker served as a reference on the rigid frame of the ski simulator, two were mounted

on the ski bindings to monitor relative motion and orientation, one on each ski boot for measuring foot dynamics, and another below each knees for lower limb movement. The trials also assessed environmental factors such as lighting, exposure time, and background influences.

The evaluation of the measurement results revealed that, during dynamic recordings, stable and continuous marker detection was achievable with marker sizes ranging between 3 to 5 cm, when the camera was positioned at a distance of approximately 1 to 1.2 meters. This configuration demonstrated high reliability, providing stable motion tracking and allowing marker localization with sub-centimeter accuracy, up to a few millimeters. Based on these findings, an experimental setup was established that demonstrated both practical applicability and robustness, forming a dependable foundation for calculating joint angles following a dynamic run using the Ski Simulator. Furthermore, it was observed that marker sizes below 3 cm significantly reduced the detection rate in the given setup, especially during rapid movements. Conversely, larger markers improved visibility but introduced potential constraints on movement and increased the likelihood of visual artifacts. Therefore, the selected configuration offers an optimal balance between precision and freedom of motion.

Acknowledgements This work has been supported partly by the COMET-K2 Center of the Linz Center of Mechatronics (LCM) funded by the Austrian federal government and the federal state of Upper Austria and supported partly funded by the COMET project DiMo-NEXT, which is supported by the Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology (BMK), the Federal Ministry for Labour and Economy (BMAW), and the provinces of Salzburg, Upper Austria, and Tyrol within the framework of COMET – Competence Centres for Excellent Technologies. COMET is managed by the Austrian Research Promotion Agency (FFG).

References

1. Minervini, A. et al.: Enhanced Localization of ArUco Markers for Autonomous Robotics: A Comparative Study. In: Secchi, C., Marconi, L. (eds) European Robotics Forum 2024. ERF 2024. Springer Proceedings in Advanced Robotics, vol 32. Springer, Cham. (2024) https://doi.org/10.1007/978-3-031-76424-0_27
2. Zhang, S. et al.: An Improved LED Aruco-Marker Detection Method for Event Camera. In: Chen, H., Fan, P., Wang, L. (eds) Communications, Networking, and Information Systems. CNIS 2023. Communications in Computer and Information Science, vol 1839. Springer, Singapore. (2023)
3. Sampathkrishna, A.: ArUco Marker based localization and Node graph approach to mapping. (2022)
4. Oščádal, P., Heczko, D., Vysocký, A., Mlotek, J., Novák, P., Virgala, I., Sukop, M., Bobovský, Z.: Improved Pose Estimation of Aruco Tags Using a Novel 3D Placement Strategy. Sensors, 20(17), 4825. (2020) <https://doi.org/10.3390/s20174825>
5. Putz, V., Mayer, J., Fenzl, H., Schmidt, R., Pichler-Scheder, M., Kastl, C.: Cyber-physical mobile arm gesture recognition using ultrasound and motion data. 2020 IEEE Conference on Industrial Cyberphysical Systems (ICPS). (2020)

Sustainability Potential of Emerging Digital Technologies in Modern Regional Waste Management

Simon Merschak¹ [0000-0001-8903-6146] and Peter Hehenberger¹ [0000-0001-5104-6525]

¹ Research Group Smart Mechatronics Engineering,
University of Applied Sciences Upper Austria, Wels Campus, Wels, Austria
simon.merschak@fh-wels.at

Keywords: Digital Waste Management, Sustainability, IoT, AI.

Emerging technologies such as Artificial Intelligence (AI), Digital Twins, Digital Platforms, and the Internet of Things (IoT) are increasingly being implemented across various industries and application domains. These developments can have both beneficial and adverse effects on the sustainability performance of the respective systems. In this context, sustainability is not limited to its environmental dimension. Rather, a comprehensive assessment that equally considers environmental, economic, and social aspects is required.

An example of the increasing adoption of novel digital technologies is found in waste management. Despite its importance for achieving circular economy objectives, the waste management sector is often regarded as conservative, and digitalization is still in its early stages [1]. Many facilities continue to rely on manual operations and long-established processes, which slows down the introduction of digital solutions. Although pilot projects and research initiatives are experimenting with AI-based sorting systems, IoT-enabled container tracking, or digital platforms for secondary material exchange, their large-scale implementation is still limited. Moreover, while these technologies hold potential to improve operational efficiency, transparency, and resource recovery, their actual impact on sustainability remains uncertain. It is often not clear whether such digital applications lead to measurable environmental improvements or how they influence the social and economic dimensions of waste management systems.

Building on these considerations, this paper investigates possible applications and the potential for improvement arising from the application of emerging digital technologies in regional waste management facilities. In Austria, the planned construction of the first fully automated and AI-supported waste management facility in the city of Wels (Upper Austria) represents a pioneering step toward the digitalization and automation of municipal waste management. This facility is intended to serve as a pilot project for demonstrating how advanced digital technologies can support higher material recovery rates and operational efficiency. To evaluate the effects of this transformation, a sustainability assessment was carried out, focusing on selected indicators. The analysis aims to identify whether and to what extent digitalized operations contribute to improved sustainability performance compared to conventional waste

management facilities. Relevant emerging technologies which are analyzed are Internet of Things (IoT), Artificial Intelligence (AI) and Digital Platforms.

The integration of IoT technologies in regional waste management facilities enables data-driven optimization and improves operational efficiency. Mobile and sensor-based access control systems automate user authentication, reducing administrative effort while enhancing accessibility outside typical operating hours. Continuous remote monitoring of critical components supports predictive maintenance through real-time performance tracking and automated anomaly alerts. Intelligent sensors combined with automated planning enable demand-oriented container emptying, projected to reduce transport-related CO₂eq emissions by 15–20% [2]. Furthermore, extended opening hours supported by automation are expected to increase collected recyclables by 5–15%, emphasizing that temporal availability of facilities can be more effective than geographic proximity [3]. Altogether, coordinated digital systems can optimize resource use, lower emissions, and strengthen the sustainability of regional waste management operations.

The application of AI in regional waste management centers can advance automation, safety, and user engagement, thereby enhancing operational efficiency and sustainability. LLM-based chatbots can be used to deliver contextual disposal guidance via natural language, promoting compliance and awareness. AI-powered thermal imaging enables early fire detection and automated emergency responses. AI-based optical systems can identify incorrect disposal to provide targeted feedback and improve sorting accuracy. The new automated waste management facility in the city of Wels targets a minimum material purity of 85% and a mis-throw rate below 5%. Literature and case studies indicate these objectives are feasible with AI-supported sorting technologies, especially for plastics, which yield substantial emission reductions [4].

The implementation of digital platforms in regional waste management facilities fosters circular, user-centered systems by enabling interaction, knowledge exchange, and local value retention. An integrated digital marketplace allows contactless, secure exchanges of underused products, extending lifespans and retaining materials locally. A dedicated e-learning platform guides users on facility use and promotes proper waste practices. At the societal level, sharing economy models within regional waste management facilities can enhance community engagement, optimize resource flows, and improve access to products for economically disadvantaged groups, advancing social equity.

References

1. Berg, H., Sebestyén, J., Bendix, P., Le Blevennec, K., Vrancken, K.: Digital waste management. European Environment Agency (2020)
2. Smart waste disposal - Saubermacher AG, <https://www.greentech.at/en/smart-waste-disposal/>, last accessed 2025/11/13
3. Abfallbericht 2024 des Bezirksabfallverbandes Schärding, https://www.umweltprofis.at/fileadmin/user_upload/redakteure/sd/Zahlen_und_Fakten/Statistik/Abfalldatenbericht_2024_HP.pdf, last accessed 2025/11/13
4. Edo, M., Granström, L., Spelhaug, C., Potter, C.: Advanced sorting technologies in the waste sector - Case studies compilation. IEA Bioenergy (2024)

Development and Validation of a Rotary Impact Drilling System for an Ultra-Lightweight UAV-Deployed Tree Planting Robot Enabling Autonomous Forest Restoration

Johann-Lukas Voigt^{1[0009-0000-9294-3759]}, Georg Happich^{1[0009-0001-4975-3703]}
and Tobias Weiser^{1[0009-0005-2396-654X]}

¹ Hochschule Kempten, Institut für KI und Robotik, 87616 Marktoberdorf, Germany
johannlukas.voigt@hs-kempten.de

Abstract. Climate change and forest degradation have created an urgent reforestation crisis in Germany, with current estimates indicating over 525,000 [1] hectares requiring restoration. Drone-based seed dispersal systems exist commercially but suffer from challenges such as low germination rate [2] and random seedling distribution. Container seedling planting can achieve success rates exceeding 90% [3] but requires complex soil preparation and precise planting methodology. The DraAuf project (Drohnen-gestützte automatisierte Aufforstung) [4], funded by the German Federal Ministry for Environment as part of the KI-Leuchttürme initiative, addresses this challenge by developing an autonomous UAV-deployed planting robot. This paper systematically investigates impact drilling strategies optimized for the extreme constraints of aerial deployment.

The project is led by the Institute for Applied AI and Robotics (IKR) at Kempten University of Applied Sciences in collaboration with Munich University of Applied Sciences and Wood-in-Vision GmbH. The complete system architecture comprises three components: a DJI Matrice M350RTK exploration drone for LiDAR mapping and AI-based planting site identification, a DJI Flycart 30 transport drone with 40kg payload capacity, and an autonomous planting robot weighing maximum 25kg including 15-30 seedlings.

The transport drone has a substantial hover power consumption of 8.2kW. Each second of hover flight consumes more energy than an entire drilling operation, making drilling speed the key factor in system efficiency. For conventional rotary drilling with 25kg total system weight, the required stabilization moment limits downforce to approximately 100N. In contrast, rotary impact drilling theoretically permits downforce approaching the full gravitational force of the robot due to the significantly reduced continuous reaction moments and allows for higher contact forces and feed speeds at the same system weight while maintaining robot stability. We therefore develop a testing method to verify and optimize a drilling strategy for DraAuf.

Our experimental methodology employs a purpose-built outdoor drilling test stand designed for field deployment across varied terrain. The stand features a 1m linear rail with closed-loop motor control. The main part of the measurement system is a 6-axis ME-Systems force sensor [6] calibrated to 300N nominal force in all spatial directions. The deliberately non-rigid bearing simulates the elastic properties expected in the final lightweight robot frame. Drill rotation speed is

monitored by an external encoder, while a laser reference system tracks lateral deflection from the intended drilling axis.

The experimental program is split into two phases. Phase 1 optimizes a necessary pilgrim-step drilling strategy—periodic retraction during drilling essential for maintaining cylindrical hole geometry. Variables include but are not limited to retraction frequency and retraction distance. Phase 2 conducts comparative analysis between the optimized rotary impact approach and conventional auger drilling across parameters including rotation/impact frequency, drill diameter and downforce ranging from 100N to 250N.

Data acquisition enables continuous recording of 6D force profiles (peak and RMS values), rotation speed and torque profiles, feed rate and displacement, total drilling time to 20cm target depth, and calculated energy consumption. Additionally, we measure lateral axis deviation and evaluate borehole quality

Field testing locations will vary in soil compaction (measured via AG-BS-70 penetrometer), moisture content (PCE-SMM1 sensor), and slope angle. Special attention focuses on sites with root networks and stones to validate system robustness in conditions representing actual reforestation sites.

With targeted drilling times under 10 seconds and complete planting cycles under 30 seconds, the system could enable planting up to 800 seedlings daily in previously inaccessible terrain. Results will directly impact structural dimensioning of the robot frame, optimal drilling parameter selection for adaptive control strategies, and correlation models between drilling characteristics and planting success probability.

The paper will include detailed results of the outlined methodology, integration challenges between the drilling system and the complete robotic platform and an outlook on the overall DraAuf Projekt development.

Keywords: Reforestation, Automation, Robotics, Drones, Soil-Drilling

References

1. BMLEH: Waldzustandserhebung 2024 ist ein erneuter Weckruf. Pressemitteilung Nr. 47/2025. Bundesministerium für Landwirtschaft, Ernährung und Heimat (2025)
2. Castro, J., Morales-Rueda, F., Alcaraz-Segura, D., Tabik, S.: Automated precise seeding with drones and artificial intelligence: a workflow. *Restoration Ecology* 32(5), e14164 (2024)
3. Pinto, J.R., Marshall, J.D., Dumroese, R.K., Davis, A.S., Cobos, D.R.: Establishment and growth of container seedlings for reforestation: A function of stocktype and edaphic conditions. *Forest Ecology and Management* 261(11), 1876–1884 (2011)
4. Hochschule Kempten: Drohnen und KI pflanzen die Wälder von morgen - Forschungsprojekt „DraAuf“ als KI-Leuchtturmpunkt ausgezeichnet, <https://www.hs-kempten.de>, last accessed 2025/09/30
5. BaSt-Ing GmbH: DrillFast – Der erste Schlagschrauber-Erdbohrer, <https://bast-ing.de/Drill-Fast-Erdbohrer/91111>, last accessed 2024/10/31
6. ME-Meßsysteme GmbH: K6D150 6-Achsen Kraftsensor Datenblatt. ME-Meßsysteme, Hennigsdorf (2024)

Regulation-Based Innovation and its Recent Challenges in the Automotive Sector

Daniel Gingo¹, Simon Wundsam¹, Venkat Pradeep¹, Ambrož Strehar¹, and Thomas Schlechter²

¹ University of Applied Sciences Upper Austria, 4600 Wels, Austria

{Daniel.Gingo, Simon.Wundsam, Venkat.Pradeep, Ambroz.Strehar}@students.fh-wels.at

² University of Applied Sciences Upper Austria, 4600 Wels, Austria

thomas.schlechter@ieee.org

Keywords: Electric Vehicle Regulation · Policy-Driven Innovation · Supply Chain Resilience · Industrial Strategy · NEV Transition · Zero-Emission Mobility.

The automotive industry across Europe is experiencing a notable decline in New Energy Vehicle (NEV) sales, with Germany registering the sharpest drop. This trend coincides with the rapid ascent of China's automotive sector, which benefits from a coherent and centralized NEV strategy. Unlike in Germany, where the direction of the policy often changes with political cycles, China's approach anchored in industrial planning, streamlined bureaucracy and targeted financial incentives has created a stable environment for private investment in the development of NEVs. This paper discusses how regulatory strategies across EU countries, particularly Germany, and similarly structured systems such as the United States (US) can draw lessons from the success of China and Norway to accelerate the transition to NEVs and protect their domestic automotive market [1].

Norway, although not a major automotive producer, illustrates how consistent consumer incentives can be successfully implemented through effective policy making [2]. China's rapid adoption of NEVs is driven by a centrally planned strategy and its dominance in battery production [3]. In contrast, fragmented policy implementation and abrupt changes, such as Germany's sudden withdrawal of subsidies, undermine market confidence. The US, with its dual federal-state regulatory framework, produces uneven adoption of NEVs in states, with California leading due to its stringent Zero Emission Vehicle (ZEV) mandate [4].

The case of Germany highlights the consequences of delayed strategic action [5]. The Dieselgate scandal triggered regulatory tightening and a shift toward electrification, yet insufficient investment in batteries and software has left German automakers trailing vertically integrated rivals such as BYD [6].

On a broader scale, the automotive industry in these regions, excluding China, is obliged to face systemic challenges such as regulatory instability in NEVs, vulnerable raw material supply chains and lagging charging infrastructure [7] [8].

Drawing from these observations [9], the article proposes a theoretical 'super hypothesis': countries with an early, coherent, and sustainable policy coupled with industrial strategy (mainly domestic battery production) and targeted financial incentives have higher adoption of NEVs and maintain strong domestic OEM market share. In contrast, lacking these elements increases the risk of foreign OEM dominance, even if overall NEV uptake grows. When tested across multiple regions, only China fulfilled all the criteria. The success criteria was defined as having at least 30 percent NEV market share within 10 years of policy implementation and more than 24 percent of that market share being dominated by domestic OEMs. China's winning strategy is rooted in its ability to balance policy side together with industrial (mainly battery production) and financial incentives. Future research will apply advanced modeling techniques to explore the complex interdependencies among these factors.

References

1. European Union, "Regulation (EU) 2023/1804 on the deployment of alternative fuels infrastructure, 2023." Accessed: July 18, 2025. [Online]. Available: <https://eur-lex.europa.eu/eli/reg/2023/1804/oj/eng>
2. International Energy Agency, "Norway's electric vehicle incentives." Accessed: June 21, 2025. [Online]. Available: <https://www.iea.org/policies/17809-norways-electric-vehicle-incentives>
3. C. Su and B. Yan, "Policy strategies of China's NEV sector," *Global Transitions Proceedings*. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2949720524000298>
4. United States Congress, "Public Law 117-169 Inflation Reduction Act of 2022." Accessed: July 15, 2025. [Online]. Available: <https://www.govinfo.gov/content/pkg/PLAW-117publ169/pdf/PLAW-117publ169.pdf>
5. L. Zhang, Y. Lin, S. Zhou, and Q. Qin, "Evolution and trends of new energy vehicle policy research hotspots: an analysis based on a Citespace knowledge graph," *Environmental Science and Pollution Research*, vol. 31, no. 5, pp. 7728–7750, Jan. 2024, doi: 10.1007/s11356-023-31485-6.
6. Á. F. Kossuth, "The Transformation of the European Car Industry: Electrification and Its Challenges," *Európai Tükör*, vol. 27, no. 2, pp. 123–141, 2024, doi: 10.32559/et.2024.2.6.
7. G. Bridge and E. Faigen, "Towards the lithium-ion battery production network: Thinking beyond mineral supply chains," *Energy Research & Social Science*, vol. 89, p. 102659, July 2022, doi: 10.1016/j.erss.2022.102659.
8. International Energy Agency, "Global EV Outlook 2025." IEA, May 2025. Accessed: July 21, 2025. [Online]. Available: <https://www.iea.org/reports/global-ev-outlook-2025>
9. X. Zhao and X. Li, "Market-incentive environmental regulation and business performance of new energy vehicle enterprises: empirical evidence from China's dual-credit policy," *Environment, Development and Sustainability*, Dec. 2024, doi: 10.1007/s10668-024-05834-5.

The CODES Framework: A Structured Prompting Methodology to Mitigate Cognitive Bias in AI-Assisted Novice Programming

Jinshan Luo¹, Atsushi Ito¹, Ami Otsuka², Akira Sasaki³

¹ Chuo University, Hachioji, Tokyo, Japan

² Hundreds LLC, Shibuya, Tokyo Japan

³ GClue Inc., Aizuwakamatsu, Fukushima, Japan

Email: kingsamlawsan@gmail.com, atc.00s@g.chuo-u.ac.jp,
amiotsuka.se@gmail.com, akira@gclue.jp

Abstract: Effective use of AI has become a key aspect of digital literacy in programming learning, yet novices often struggle due to cognitive biases that lead to inefficient prompting and defective code generation. Based on observations from programming education, this study identifies typical ineffective prompting patterns and interprets them through cognitive psychology. To address the mismatch between novices' mental models and LLMs' functional mechanisms, we propose the CODES Framework, a structured prompting method. A controlled experiment comparing novice-style and CODES-enhanced prompts demonstrates improvements in code correctness and completeness. These results provide a theoretically grounded and pedagogically applicable model for improving human–AI collaboration in programming education.

Keywords: AI-Assisted Programming; Prompt Engineering; Cognitive Biases.

1 The novice-AI interaction gap

Novice programmers, lacking structured problem-solving schemas, construct prompts based on cognitive biases [1]. To investigate this, our observational study of novice programming students using ChatGPT identified a taxonomy of four recurrent ineffective patterns: (1) Phantom Querying, characterized by an avoidance of diagnostic effort and a lack of context; (2) Copy-Paste Prompting, indicating cognitive overload and offloading the entire diagnostic process; (3) Wishful Querying, demonstrating a failure in problem decomposition; and (4) Abandonment & Reset Regulation, reflecting learned helplessness and a failure of metacognitive strategy. These patterns, rooted in cognitive theory, directly conflict with the technical nature of LLMs. LLMs are probabilistic, autoregressive models, not logical reasoners [2]. Ambiguous, bias-driven prompts force the LLM into a high-entropy state with minimal guidance, leading to error propagation and flawed code.

2 The CODES framework and validation

To address this misalignment, we propose the CODES Framework (Context & Role, Objective, Delimiters & Data, Expectations & Constraints, Scaffolding Instruction). The framework systematically structures the prompting process. Specifically, [C] and [D] provide explicit grounding that mitigates flawed mental models and reduces the occurrence of Phantom Queries [3] [4]; [O] and [E] enforce problem decomposition and the specification of edge cases, counteracting Wishful Querying and automation bias [5] [6]; and [S] models metacognitive planning to prevent Abandonment & Reset behaviors by allowing for logical validation prior to code generation [7].

We conducted a controlled experiment to validate the framework's effectiveness. Two programming tasks were designed: Task 1, a simple game development task requiring the generation of a functional shooting game using Python; and Task 2, a more complex task involving expanded concurrent logic based on Task 1. Each task was tested with a baseline prompt and a CODES-Enhanced prompt, executed 100 times each in isolated GPT-4o sessions. Three performance metrics were evaluated: Execution Rate, runnable scripts per 100 trials; Component Completeness Rate, the proportion of runnable scripts containing all required components; and Functional Completeness Rate (FCR), the proportion of functionally correct scripts among those with complete components. The results clearly demonstrate the CODES framework's superiority. The baseline prompts achieved FCRs of 19% in Task 1 and 16% in Task 2, whereas the CODES-enhanced prompts achieved 33% and 78%, respectively. Analysis further revealed that baseline failures primarily stemmed from the LLM's misinterpretation of ambiguous language, while the explicit structure introduced by the CODES framework effectively eliminated such ambiguity and ensured consistent, semantically accurate code generation.

References

1. J. Sweller, "Cognitive load theory," in *Psychology of Learning and Motivation*, vol. 55, pp. 37–76, Academic Press, 2011
2. J. Liu, et al., "Is your code generated by ChatGPT really correct? Rigorous evaluation of large language models for code generation," *Advances in Neural Information Processing Systems*, vol. 36, pp. 21558–21572, 2023
3. M. Zheng, et al., "Is 'A Helpful Assistant' the best role for large language models? A systematic evaluation of social roles in system prompts," *arXiv preprint arXiv:2311.10054*, 2023, Available: <https://arxiv.org/abs/2311.10054>
4. T. Brown, et al., "Language models are few-shot learners," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020
5. D. Zhou, et al., "Least-to-most prompting enables complex reasoning in large language models," *arXiv preprint arXiv:2205.10625*, 2022
 L. Ouyang, et al., "Training language models to follow instructions with human feedback," *Advances in Neural Information Processing Systems*, vol. 35, pp. 27730–27744, 2022
 J. Wei, et al., "Chain-of-thought prompting elicits reasoning in large language models," *Advances in Neural Information Processing Systems*, vol. 35, pp. 24824–24837, 2022

Support Tool for Generative AI-Based Coding

Ami Otsuka¹, Atsushi Ito²

¹ Hundreds LLC, Shibuya, Tokyo, Japan

² Chuo University, Hachioji, Tokyo, Japan

Email: amiotsuka.se@gmail.com, atc.00s@g.chuo-u.ac.jp

Keywords: Generative AI; Coding; Prompt; Support Tool.

1 Spreading Coding using Generative AI

The application of generative AI to software development has been expanding rapidly. According to Reference [1], the use of GitHub Copilot resulted in a 6.5% increase in project-level productivity, approximately 5.5% improvement at the individual level, and a 5.4% increase in contributor count. Reference [2] reports that developers using AI completed tasks 55.8% faster than those who did not. Reference [3] estimates that the North American market size for generative AI tools for programming was US\$265.9 million in 2022, increased to US\$341.3 million in 2023, and is projected to reach US\$2,833.9 million in 2030, representing a Compound Annual Growth Rate (CAGR) of 35.3%.




Fig. 1 Market size of AI for software development in North America [3]

2 Problems of Coding using Generative AI

A major issue in software development using generative AI—particularly AI specialized for programming, such as Codex—is the difficulty of preserving work logs. GitHub stores code version histories but does not store prompts or error messages.

When using generative AI, substantial changes may occur. In such cases, developers often need to review differences from earlier code versions and the instructions previously given to the AI. However, in practice, this isn't easy.

To address this issue, we have been investigating mechanisms to preserve and utilize work histories in environments like Codex. The following discussion focuses on Codex as a representative example.

3 Designing the Support Tool

As shown in Figure 2, there are two approaches to storing work histories. Case 1 involves modifying Codex itself to record logs. However, because Codex's internal code is not publicly available, this approach is impractical.

A more feasible approach is Case 2, where a proxy is placed between the user and Codex to collect logs. This requires no modification to Codex itself, making implementation easier. However, it introduces a drawback: the proxy must constantly adapt whenever the Codex API specification changes.

For these reasons, we decided to implement a Case 2-type proxy. This allows the proxy to be run locally, improving processing speed and offering greater flexibility in log handling. In the future, the proxy may incorporate its own AI to provide coding advice.

Further details will be presented at Eurocast 2026 in February.




Fig.2 Two design patterns

References

- [1] Fangchen Song, Ashish Agarwal, Wen Wen, “The Impact of Generative AI on Collaborative Open-Source Software Development: Evidence from GitHub Copilot”, <https://doi.org/10.48550/arXiv.2410.02091>, last accessed 2025/11/15
- [2] Sida Peng, Eirini Kalliamvakou, Peter Cihon, Mert Demirer, “The Impact of AI on Developer Productivity: Evidence from GitHub Copilot”, <https://doi.org/10.48550/arXiv.2302.06590>, last accessed 2025/11/15
- [3] “Generative AI in Software Development Lifecycle Market Size”, https://www.for-tunebusinessinsights.com/generative-ai-in-software-development-lifecycle-market-109041?utm_source=chatgpt.com”, last accessed 2025/11/15

Classical and Neural Network Modeling of a Heating System in a Smart Building

Andrzej Stachno¹[0000-0002-1678-8037] and Anna Czemplik²[0000-0002-0979-7322]

¹ Wroclaw University of Science and Technology, Wroclaw, Poland
 andrzej.stachno@pwr.edu.pl

² Wroclaw University of Science and Technology, Wroclaw, Poland
 anna.czemplik@pwr.edu.pl

Abstract. This study investigates a thermal object with electric under-floor heating, and its mathematical representations. Two types of models are considered: a classical, physics-based model derived from differential equations and data-driven models based on artificial neural networks. The objective is to analyze and compare the correspondence between the physical system and the models by analyzing measured time series data, with a specific focus on indoor and outdoor temperature dynamics.

Keywords: Differential equation model · Neural network-based models.

1 Characteristics of the Thermal Object

The object under study is a single open-plan room with a floor area of 42.9 m². The heating cables are distributed within the floor structure and embedded in the screed layer. The heating system is divided into two independently controlled zones. Measurement and control of the facility are provided by a building automation system based on standard components of the KNX protocol [1]. The room is equipped with temperature sensors typical for smart buildings, which are mounted on the ceiling, in a low-level wall outlet near the floor, and within the light switch assembly by the door. Additionally, multiple extra temperature sensors, along with occupancy detection and people counting sensors, were installed to expand the scope of the analysis. Experimental data were collected during the heating system's start-up and shutdown cycles, as well as during periods of normal operation. The analysis of data from individual sensors will allow for the verification of assumptions regarding the expected temperature distribution within the room. Consequently, this will enable the optimization of the temperature measurement strategy, for instance, by determining the optimal number and placement of sensors required to obtain representative readings.

2 Classical Differential Equation Model (DEM)

The classical model of the thermal object is formulated as a system of ordinary differential equations (ODEs). This model, hereafter referred to as the Differential Equation Model (DEM), is based on heat balance equations for the primary

elements with thermal capacitance, namely the walls, floor and indoor air volume. The values of the model's parameters are derived from the geometric dimensions of the room and the thermophysical properties of the building materials. The model primarily enables the calculation of the temperature inside a room. The inputs to the model include the outdoor temperature, the supplied heating power, and the internal heat gains from persons. The simulation results obtained from the DEM are validated against the experimental data collected from the physical thermal object. The DEM is, by design, a simplified representation and is therefore not expected to be perfectly accurate. A preliminary analysis will be conducted to quantify its accuracy relative to the physical system.

3 Artificial Neural Network-based Model (ANNM)

Models based on artificial neural networks (ANNs) are developed using various architectures specialized for time series analysis [2, 3]. These models were initially employed to optimize the number and placement of sensors. This was achieved by comparing the readings from the standard KNX system sensors with those from supplementary sensors installed for the study. Furthermore, the ANNs were utilized for indoor temperature prediction. In this application, the networks operate on three input variables: outdoor temperature, heating power, and the number of persons. Their primary objective is to predict the indoor room temperature. The ANN-based models were trained using two distinct approaches: a) a purely data-driven method using only experimental data, b) a hybrid method incorporating synthetic data generated by simulations of DEM, which exemplifies the concept of Physics-Informed Neural Networks (PINN) [4].

4 Summary

A key objective of this study is to determine the feasibility of substituting direct indoor temperature measurements with predictions from an artificial neural network model. We also examine the potential advantages of a hybrid modeling strategy that combines the classical differential equation models with neural network architectures.

References

1. KNX Association: KNX Advanced Course Documentation – Januar 2019 ISBN-13: 979-8635617861
2. Stanley N., Eugene H.: An Introduction to Econophysics. Correlations and Complexity in Finance, PWN, 2001.
3. Peitgen H.O., Jurgens H., Saupe D.: Introduction to Fractals and Chaos, PWN, 2002. Technical Analysis of the Financial Markets, WIG PRESS, 1999.
4. Raissi M., Perdikaris P., Karniadakis G.E.; Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, Journal of Computational Physics, vol. 378, pp. 686-707 (2019)

Using recurrent neural networks for prediction in dynamic systems with support of artificial intelligence language models

Iwona Karcz-Duleba^[0000-0003-4946-2225] and Anna Czemplik^[0000-0002-0979-7322]

Department of Control Systems and Mechatronics
 Faculty of Information and Communication Technology
 Wroclaw University of Science and Technology
 iwona.karcz-duleba@pwr.edu.pl, anna.czemplik@pwr.edu.pl

Keywords: dynamical systems, recurrent neural networks, physics informed neural networks, artificial intelligence language models, learning support

Introduction. Models that characterize dynamics of various systems described by differential equations allow for a prediction of a behaviour of complex systems, from weather to cancer development. In most cases, solving differential equations analytically is impossible, and a computational cost of numerical calculations is very high. The popularity of recurrent neural networks (RNNs) in many fields and an availability of various application tools have led to attempts to apply them to solve or to predict differential equations. Two main approaches are distinguished: Physics Informed Neural Networks [1] and Neural Operator Learning [2]. Physics Informed Neural Networks (PINNs) facilitate network learning by incorporating physical knowledge about the system. PINNs are successfully used, among others, for piecewise function approximation, theory of functional connections, multiple sets of irregular geometries, inverse computations.

Characteristics of the research subject. To discover the advantages of deep neural networks, they have been used in simple models of dynamic thermal objects. The aim is to answer the following question: Are RNNs suitable for studying dynamic models, particularly for forecasting systems (objects) with different dynamics? Forecasting can be used, among other things, to optimize a control of inertial systems with large time constants. Examples of such systems include heating systems with high thermal capacity. To find the answer, artificial intelligence models are used to train RNNs in Matlab and to support their application with analyzed dynamic models.

As the constructed RNNs will be used in both research and teaching, it is also being examined whether AI can take on the role of a teacher and teach a layperson (student) a completely new subject. Our previous research [3] concerned the role of the teacher in an AI-driven world. The study examined whether AI could pass exams taken by undergraduate students in engineering and technical sciences. We concluded that AIs can facilitate learning, but their use should be cautious and critical.

In this study, we adopted the perspective of a student who wants to learn with the help of AI. The insights gained will help us to use language models in teaching and to support students during their independent learning with AI.

Methodology. Various language models are exploited to learn how to use RNN in Matlab software. The most popular models Chat GPT 5, Gemini, and Cloud (Opus and Sonnet), were selected for this purpose. Two researchers who are unfamiliar with the subject undertook the learning independently. The AI is asked comprehensive questions or sequences of questions in English and Polish [4]. The quality of the teaching is assessed using questions such as: Did I receive a properly functioning program? After learning with AI, am I fully aware of what I am doing and how? Can I say that I have learned the subject and can use these networks in various tasks?

Additionally, to help students choose an appropriate AI model, they were compared using the LM Arena comparison tool, as well as with Matlab learning support. LM Arena [5] is a free platform for comparing and testing artificial intelligence (AI) models, including large language models (LLMs), created by scientists at UC Berkeley.

Conclusions. During the first attempts to use AI, almost none of the LLM models tested generated a well-functioning program. Various errors were encountered (ranging from basic syntax errors to data inconsistencies) and different responses were obtained in different sessions. Without prior knowledge of the data structure and the various network variants, it was impossible to improve the generated programs. Ultimately, however, trained networks were obtained that provided single- or multi-step predictions for all LLM models (Gemini, GPT, CloudOpus, CloudSonnet), but only after applying an appropriate sequence of queries specifying key parameters. Therefore, using AI to learn completely new topics is not straightforward. Without prior knowledge (which can be provided, for example, by a lecturer), the process could be long and error-ridden, so AI should be approached with great caution. Based on the programs generated by AI, neural models were prepared, enabling for the prediction studies to be conducted in simple dynamic models of thermal objects.

References

1. M. Raissi, P. Perdikaris, G. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations.", *Journal of Computational Physics* 378: 686–707 (2019)
2. N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, A. Anandkumar, "Neural operator: Learning maps between function spaces", *Journal of Machine Learning Research.* 24: 1–97 (2021).
3. A. Czemplik, I. Karcz-Duleba, "ChatGPT Language Model as Support for Teaching and Self-learning in the Field of Engineering and Technical Sciences", *Computer Aided Systems Theory – EUROCAST 2024: Part I*, pp.376 – 387, Springer (2024)
4. Y. Kim, J. Russell, M. Karpinska, M. Iyyer, "One ruler to measure them all: Benchmarking multilingual long-context language models", *arXiv:2503.01996v3* (2025)
5. LM Arena platform <https://larena.ai/>

Modeling of Thermal Energy Storage Tanks

Anna Czemplik^[0000-0002-0979-7322]

Wroclaw University of Science and Technology, Wroclaw, Poland
 anna.czemplik@pwr.edu.pl

Abstract. Modern building heating systems that use energy from the environment usually include a thermal energy storage tank (heat buffer, accumulation tank). The buffer allows energy to be stored in the form of heated water. The aim of this work is to develop a tank model that can be used in simulation studies of heating systems. The model will be used to prepare simulations of various heating system variants and to conduct research that will enable the design and testing of different control methods. The model is dedicated to students and engineers in the field of automation and control.

Keywords: Thermal Energy Storage Tank · Heating System.

1 Characteristics of the plant

This presentation outlines a model of a thermal energy storage (TES) tank, integral to heating systems powered by intermittent sources like heat pumps and solar collectors. Sources that harness environmental energy are characterized by a high variability in performance and a lack of coherence with the heat demand. Specifically, during periods of peak demand, these sources typically exhibit their lowest output. The buffer tank serves to store excess thermal energy generated during periods of high availability or low cost, allowing for its dispatch on-demand to the heating system. It is therefore possible to separate the task of energy production from current demand.

A key feature of heat storage tanks is the stratification of water temperature, a phenomenon driven by density differences, where the water is coldest at the bottom and hottest at the top. The establishment and preservation of this temperature profile is facilitated by specific internal design features and by the strategic placement of external connection ports.

Manufacturers offer various buffer tank designs tailored for specific applications. These designs differ in the number and arrangement of inlet and outlet ports, which connect the tank to multiple heat sources and heat consumers. In a basic configuration, hot water from the heat source is introduced through ports in the upper section of the tank, while cooler water is drawn from the bottom ports to be reheated. Similarly, the heating system draws the hottest water from the top of the tank and returns the cooled water to the bottom. This principle ensures that both the charging (heating) and discharging (cooling) processes occur in a stratified manner, preserving the temperature layers. A more advanced

design incorporates an internal heat exchanger coil, through which a heat transfer fluid from the source circulates. The same principles extend to bivalent tanks designed to integrate multiple sources and loads, where high-temperature circuits are connected at higher elevations and low-temperature circuits at lower ones. The temperature stratification reduces the frequency of charging cycles and provides higher available water temperatures for the connected heating systems.

2 Characteristics of the model

A highly accurate model of the tank would necessitate solving complex partial differential equations (PDEs) derived from fluid mechanics [1, 2]. The proposed model is based on the discretization of the tank into several layers and the modeling of forced convection between them using ordinary differential equations (ODEs). These layers can be classified into three types: the top and bottom layers, which interface with the external circuits, and the intermediate layers. The number of intermediate layers selected establishes a trade-off between model fidelity and computational complexity. A buffer tank may also incorporate an auxiliary electric heater, typically located in its lower section. The model of such a tank must be extended to include natural (free) convection, a phenomenon in which warmer, less dense water rises spontaneously to the upper region without mixing with the cooler layers below.

3 Application

The developed tank model has been applied to investigate a heating system that comprises a boiler, the storage tank, and a heating installation. The analysis was conducted within the Matlab/Simulink environment. The models feature a modular design based on a dedicated block library, which enables the construction of various design variants of the tank model. The simulation results demonstrate the tank's behavior during its charging and discharging phases, in addition to its performance under typical operating conditions.

References

1. Chen J., Zhang Y. and others: Thermal performance analysis of a thermocline storage tank with integrated annular distributors, *Frontiers in Energy Research* (2023), doi.org/10.3389/fenrg.2023.1190760,
2. Szczęśniak A., Milewski J. and others: Determination of Thermocline Heat Transfer Coefficient by Using CFD Simulation, *Energies* (2023), dx.doi.org/10.3390/en16073150

Behavior Analysis for Grazing of Cows

Batchimeg Bayaraa¹, Kazuya Nogami¹, Naoya Tsuruta¹, Madoka Hasegawa¹,
 Rika Fukumori², Yoshikazu Nagao¹, Ami Otsuka³,
 Mei Onda⁴, Yuko Hiramatsu⁴, Tomoko Nakano⁴, Atsushi Ito⁴

¹ Utsunomiya University, Utsunomiya, Tochigi, Japan

² Rakuno Gakuen University, Ebetsu, Hokkaido, Japan

³ Hundreds LLC, Shibuya, Tokyo, Japan

⁴ Chuo University, Hachioji, Tokyo, Japan

Email: cheegii15@yahoo.com, 25nogami@icl.is.utsunomiya-u.ac.jp, naoya@is.utsunomiya-u.ac.jp, madoka@a.utsunomiya-u.ac.jp, fukumori@rakuno.ac.jp, ynagao@cc.utsunomiya-u.ac.jp, amiotsuka.se@gmail.com, a23.xc7j@g.chuo-u.ac.jp, susana_y@tamacc.chuo-u.ac.jp, tnakano001a@g.chuo-u.ac.jp, atc.00s@g.chuo-u.ac.jp

Keywords: Grazing; dairy farming; Cow's Behavior; Animal Welfare; IoT; Deep Learning.

1 Grazing and Animal Welfare

Rising feed costs account for approximately half of all operating expenses and are an important aspect of continuing to run a dairy farm [1]. In Japan, cows are usually kept in barns. Most of the concentrate feed is imported from foreign countries. Therefore, it isn't easy to secure stable supplies due to ocean freight costs, fluctuating exchange rates, and other factors.

In-house rearing, restricting cows' behavior, not only causes stress, falls due to weak legs and feet, and the risk of lower milk production, but also makes it easy to neglect proper management and maintenance of excreta disposal, which can harm cow health and milk quality. In contrast, grazing not only reduces stress by allowing unrestricted behavior but also promotes leg and foot growth and even reduces the risk of accidents at night [2]. Grazing is also ideal for reduced feed costs and animal welfare. Eating roughage is expected to reduce feed costs and ensure stable rearing. As mentioned before, a rearing system that relies on imports could be difficult if the supply chain becomes distorted. On the other hand, roughage contains high levels of vitamin A, which can improve milk quality, and a shift to a rearing system based on roughage is considered to have sufficient merit [3]. In addition, from the perspective of 'animal welfare' in the SDGs, which has been promoted in recent years. Pasture, where animals are raised in the wild, is preferable to barn rearing, where animals are kept in a small space. Japan has received an E rating in the Animal Protection Index (API) published by the World Animal Protection Council (WAP) [4]. A shift to grazing would stabilize the dairy farming management system.

2 Behavioral Classification

To measure the effectiveness of grazing, such as how much a cow eats grass and walks, we are developing an IoT system using an accelerometer attached to the cow's body. We are now developing the following technologies.

(1) Find the most suitable position of the body for measuring the cow's behavior.

We are now testing two positions: neck belt and leg.

(2) Develop a learning model to analyze the cow's behavior. This learning model classifies three patterns: eating, walking, and stopping.

(3) Develop an effective test data collection tool for the cow's behavior.

In this paper, we mainly mention the test data collection tool for the cow's behavior. Figure 1 displays the flow of behavior classification. The problem with this flow is that it takes a long time to add labels at the third step. To increase the accuracy of the learning model, we need enough data on cows' behavior. We think we need to introduce AI to get labels. So, we are now introducing a pose estimation tool [5]. Using this tool, we can automatically extract the labels for the cow's behavior from the video. Figure 2 shows the accuracy of the extracted behavior label. The result shows that this approach is practical.




Fig.1 Work flow

	precision	recall	f -score
Eating	0.98	1.00	0.99
Walking	0.97	0.86	0.91
Stopping	0.83	0.42	0.56

Fig.2 Accuracy of automatically extracted data

References

- [1] Central Dairy Council: Survey on the State of Dairy Farm Management in Japan, <https://www.dairy.co.jp/20230317.pdf>(in Japanese), last accessed 2025/11/15
- [2] Ministry of Agriculture, Forestry and Fisheries: Grazing Room, <https://www.maff.go.jp/j/chikusan/sinko/shiryo/houboku/houboku.html>(in Japanese),, last accessed 2025/11/15
- [3] Feed Division, Livestock Production Bureau, Ministry of Agriculture, Forestry and Fisheries : The situation regarding feed, <https://www.maff.go.jp/j/chikusan/sinko/lin/1siryo/attach/pdf/index-970.pdf>(in Japanese),, last accessed 2025/11/15
- [4] World Animal Protection : Animal Protection Index, <https://api.worldanimalprotection.org/>, last accessed 2025/11/15
- [5] MMPoseContributors, “GitHub - open-mmlab_mmpose_OpenMMLab Pose Estimation Toolbox and Benchmark.,”GitHub, <https://github.com/openmmlab/mmpose>, 2020. , last accessed 2025/11/15

Zonal indoor user localization in a room based on air quality measurements using artificial neural networks

Andrzej Stachno¹[0000-0002-1678-8037]

¹ Wroclaw University of Science and Technology, Wroclaw, Poland
andrzej.stachno@pwr.edu.pl

Abstract. This work presents a system for approximate user localization within four zones of a room, based solely on existing air quality sensors of a building automation system. A classical barycenter method and a multilayer perceptron (MLP) neural network are employed to classify the zone based on deviations in CO₂, temperature, humidity, and pressure. Experiments demonstrate clearly superior performance of the neural network approach, despite limitations related to sensor inertia and the presence of only a single user.

Keywords: Smart Home, User localization, Environmental sensors, Artificial neural networks.

1 Introduction and research objective

In modern building automation installations, the KNX [1] system is commonly used to monitor environmental conditions and control user comfort. Standard air quality sensors (CO₂, temperature, humidity) are mainly used to assess environmental conditions and to control ventilation and heating. In this work, the possibility of using such sensors also for approximate localization of a person in a room is investigated, without installing any additional infrastructure.

A concept of a zonal user localization system was developed, based on four sensors placed in the corners of a rectangular room. The room was divided into four zones, and the goal was to assign the person to one of these zones under the assumption of no mechanical ventilation, as well as to investigate the system behavior in the presence of multiple occupants.

2 Methodology and system implementation

The first step was the calibration of an “empty” room, in which baseline values of CO₂, temperature, and humidity were determined for each sensor. Subsequent calculations were performed on deviations from these baseline values, which reduced the impact of drift and background changes. On this basis, a simple presence detection algorithm was implemented, relying on an increase in the average CO₂ concentration and the gradient between sensors [2].

Two approaches were used to determine the active zone. In the classical approach, a virtual “emission center” was computed: each sensor was assigned a weight dependent on the CO₂ increase (with temperature and humidity taken into account), after which the barycenter was calculated and assigned to the appropriate quadrant. However, analysis showed that the nonlinear distribution of CO₂ and temperature limited the accuracy of this method, especially near zone boundaries.

In the second approach, artificial neural networks were used. A multilayer MLP network was implemented, whose input was a feature vector composed of deviations in CO₂, temperature, humidity, and pressure from the four sensors, as well as simple time-related features [3]. The output was a label corresponding to one of five classes: four zones and no person present. The network was trained on data recorded in a real room, in which the user stayed successively in each zone. The computation layer was implemented on an external computer with a KNX interface, and the classification results (zone index) were sent back to the system and used, among other things, for controlling zone-based lighting.

3 Experimental results

The experiments conducted showed that, under the assumed conditions (no ventilation, single user), the neural network achieved clearly higher zone assignment accuracy than the barycenter method. The largest errors occurred near zone boundaries and for short residence times in a newly occupied zone, which was related to the inertia of the CO₂ sensors.

4 Summary

Despite limitations such as response delays and reduced accuracy in the presence of multiple users, the experiment confirmed that existing KNX environmental sensors can be used for coarse zonal user localization with neural networks, without the need for dedicated localization systems. This solution enables the extension of the functionality of conventional building automation installations with simple localization features at minimal cost in additional infrastructure.

References

1. KNX Association: KNX Advanced Course Documentation January 2019 ISBN-13: 979-8635617861
2. Stanley N., Eugene H.: An Introduction to Econophysics. Correlations and Complexity in Finance, PWN, 2001.
3. Peitgen H.O., Jurgens H., Saupe D.: Introduction to Fractals and Chaos, PWN, 2002. Technical Analysis of the Financial Markets, WIG PRESS, 1999.

Neural Analysis of Thermal Images as the Basis of an Intelligent Presence Sensor in Building Automation

Andrzej Stachno^{1[0000-0002-1678-8037]}, Marek Suproniuk^{2[0000-0003-3723-6116]}, Nikodem Lelewelski³

¹ Wroclaw University of Science and Technology, Wroclaw, Poland
andrzej.stachno@pwr.edu.pl

² Military University of Technology in Warsaw, Warszawa, Poland
marek.suproniuk@wat.edu.pl

³ Wroclaw University of Science and Technology, Wroclaw, Poland
263503@student.pwr.edu.pl

Abstract. This work describes an intelligent presence sensor for building automation, based on a low-resolution thermal imaging camera and convolutional neural networks. A dedicated dataset for training and testing the people-detection model was prepared from recorded sequences of thermal images of occupants in various scenarios. The trained network was implemented on an embedded platform operating in real time, generating presence, people-count and zone-occupancy signals. Experiments conducted in near-real conditions demonstrated high detection performance and a clear advantage over classical PIR sensors, especially in darkness and in scenarios with limited user movement.

Keywords: Presence sensor, Thermal imaging camera, Convolutional neural networks, Thermal image analysis

1 Introduction and motivation

This paper presents a completed concept of a presence sensor for building automation, based on a low-resolution thermal imaging camera and neural image analysis. The starting point was the limitation of classical PIR and optical sensors, which failed under low illumination, limited user movement and complex scenes with multiple occupants and significant disturbances. The goal was to develop a device that reliably detects people under all lighting conditions and provides not only a presence signal, but also information on people count and zone occupancy.

2 Measurement methodology and implementation

The prototype consisted of a thermal imaging camera, an embedded processing unit and a communication interface for integration with the building control system. For training purposes, thermal image sequences were recorded and manually annotated for various person configurations, postures and background conditions. Preprocessed images were

fed to a convolutional neural network adapted to low resolution, which enabled automatic feature extraction and robustness to disturbances. Detection results were converted into presence, people-count and zone-occupancy signals made available to the supervisory control system.

3 Experimental results

Experiments conducted in conditions close to real operation (different lighting levels, varying numbers of people, various postures and partial occlusions) confirmed high silhouette-detection performance. The neural model operated stably despite the limited image resolution, while temporal filtering and decision thresholds significantly reduced false alarms. The device generated reliable presence and people-count signals, outperforming traditional PIR sensors, particularly in darkness and scenarios with limited user movement.

4 Conclusion

The developed and tested concept of a thermal presence sensor with neural analysis confirmed the usefulness of combining vision techniques and machine learning in building automation. High detection performance was achieved across diverse conditions while maintaining real-time operation and moderate hardware requirements. It was shown that even a low-resolution camera, supported by a convolutional neural network, can form an effective basis for an intelligent presence sensor and a promising evolution of classical PIR-based solutions.

References

1. Cristinel Popescu: Neural Networks: Convolutional Neural Networks, Kindle edition, Amazon, 2022. Kindle edition, Amazon, 2022.
2. Robert A. Kosinski: Sztuczne sieci neuronowe, Wydawnictwo Naukowe PWN, 2018.
3. R. Tadeusiewicz, M. Flasinski: Rozpoznawanie obrazów, PWN, 1991.

Optimal Inventory Control Using the Monte Carlo Method

Marek Olesiak 0009-0006-9973-6449

*Faculty of Engineering, University of Applied Sciences, Walcz,
Poland*

Contact: marek.olesiak@pwsz.eu

Extended Abstract

This contribution addresses warehouse management as an optimal inventory-control problem under uncertainty. We formulate a single-state mathematical model where the warehouse level represents both on-hand inventory (positive values) and backlog (negative values). The cost functional aggregates holding and shortage penalties with ordering and transportation costs. Transportation is modeled by fixed-capacity lots (e.g., truck, railcar, pallet); the shipping cost does not depend on the actual load factor, while orders are placed in integer lots with a user-defined rounding rule. Demand is random but follows a known probabilistic profile. In this work we treat inventory management as an optimal control problem in which the system state is described by a single aggregated variable (R): positive values correspond to on-hand inventory, negative values to backlog (shortage). This representation simultaneously “resets” the positive part (physical storage) and preserves information on shortage, which incurs higher costs than regular replenishment. The objective function sums holding and shortage costs with logistics components: a fixed ordering cost (K_t) and a per lot transportation cost (k_{jedn}) with a fixed lot capacity (d_{jedn}). Orders are placed in an integer number of lots with a rounding threshold $s_{prog} \in [0, 1]$ (when the fractional part exceeds s_{prog} , we add one extra lot). Daily demand is modeled by a random distribution with given parameters; in the experiments we use the normal distribution with mean μ_p and standard deviation σ_p . The simulation horizon is $n=365$ days and the lead time is $L=1$. We evaluate strategies with Monte Carlo on $tm=600$ independent annual demand scenarios. We use CRN (common random numbers): for every parameter setting we compute the cost on the same stream of scenarios, which guarantees fair comparisons. Besides the mean cost (k) we report estimation uncertainty: the standard error $SE = \sigma / \sqrt{tm}$ and the relative uncertainty index $KE = 3SE/k$. This “estimate + uncertainty” pair lets us separate meaningful differences from sampling noise and report results with an intuitive interpretation ($\pm 3SE \approx 99.7\%$ of the probability mass). We compare five authorial strategies (one stepwise and four “smooth”: linear, parabolic, inverted parabolic, exponential) with standard policies (S, S), (S) and (Q, R). In the authorial strategies we control the threshold S_{kr} and the order levels x_1 (for $0 < s \leq S_{kr}$) and x_2 (for $s \leq 0$); in smooth profiles the order size increases as the deficit deepens according to the chosen functional form. In standard policies we tune S, s, Q, R according to their classical definitions. Parameters are optimized with a compact-domain grid search, and—if needed—refined by a simple evolutionary heuristic (local mutations with improvement acceptance), always on the same set of CRN scenarios. In the experimental study we consider three cost-logistics regimes (A/B/G) and three demand variability levels $\sigma_p \in [50, 100, 200]$. The regimes differ in service level, lot size, K_t and k_{jedn} (e.g., A: 98% service, lot 1000; B: 97.5%, lot 1200; G: small lots, low K_t). We calibrate k_m and k_n following common practice (e.g., $k_m = [rv]/365$; $k_n \approx k_m \alpha / (1 - \alpha)$, where α is the cycle service level) and then keep them fixed across all comparisons to isolate the effect of the policy itself. Results: (i) the objective landscape is irregular and nonconvex (terraces, thresholds and jumps introduced by integer lots and the rounding threshold s_{prog}), which weakens gradient methods; a simple grid plus mutations is stable and reproducible; (ii) short rules (1-3 thresholds/levels) suffice—adding more levels rarely yields material savings; (iii) with large lots

and high K_t , threshold type rules (stepwise, (s, S)) are very competitive; with small lots and low K_t , smooth profiles (e.g., linear) gain advantage; (iv) differences between the best strategies are

often small (fractions to a few percent), hence reporting KE is necessary for sound interpretation. We illustrate this in tables ($op=50/100/200$, regimes A/B/G) with optimal parameters, k_{min} , KE and a simple sensitivity measure to parameter perturbations. The role of Monte Carlo is twofold: it is an objective evaluator of the cost function (the same scenario stream for all comparisons) and a tool for quantifying uncertainty (SE, KE) and choosing the sample size. We show that $tm \approx 600$ is a reasonable compromise between compute time and stability (typical KE at the hundredths of a percent), enabling reliable comparisons even between strategies with very close costs. We recommend confirming final results on an independent set of scenarios (a safeguard against overfitting to a single random draw). Our contributions are fourfold. First, we propose a single state (yet informative) representation that naturally couples' stock and backlog. Second, we use a coherent cost accounting based on geometry of areas at daily transitions $R0 \rightarrow R1$. Third, we adopt CRN + KE as a reporting standard (estimate + uncertainty), which enables comparisons robust to sampling variance. Fourth, we show that simple stepwise/smooth rules (1-3 levels/thresholds), properly calibrated, achieve near optimum performance across a wide range of cost and demand regimes. The methodology is easy to implement in Excel/VBA and Python. Limitations and future work: (i) we used the normal distribution in experiments—applications should consider empirical samples/bootstrapping; (ii) we analyze a single item and fixed L ; (iii) transportation cost is modeled as per lot without freight discounts—extensions with discounts, consolidation and time windows are natural. In parallel we plan to generalize multistep policies ($m > 3$) with $O(1)$ index computation without loops, which proved very efficient in preliminary tests. Practical takeaway: if logistics enforces integer lots and K_t is material, start from a two level stepwise rule (Sk_r, x_1, x_2, s_{prog}) calibrated by a grid on CRN scenarios; for small lots/low K_t consider the linear profile. In all cases report k and KE and validate on an independent scenario sample.

Keywords: inventory control; Monte Carlo; stepwise and smooth policies; grid optimization; estimation uncertainty; CRN; ordering and transportation cost; lead time.

References

- [1] Metropolis, N., Ulam, S.: The Monte Carlo Method. *J. Amer. Stat. Assoc.* 44(247) (1949) 335–341.
- [2] Arrow, K.J., Harris, T., Marschak, J.: Optimal Inventory Policy. *Econometrica* 19(3) (1951) 250–272.
- [3] Silver, E.A., Pyke, D.F., Thomas, D.J.: *Inventory and Production Management in Supply Chains*, 5th ed. CRC Press (2021).
- [4] Ax:säter, S.: *Inventory Control*, 4th ed. Springer (2022).
- [5] Zipkin, P.H.: *Foundations of Inventory Management*, 2nd ed. McGraw-Hill (2023).
- [6] Rubinstein, R.Y., Kroese, D.P.: *Simulation and the Monte Carlo Method*, 4th ed. Wiley (2024).
- [7] Gaspar-Cunha, A., Vieira, A.: Evolutionary algorithms for multiobjective inventory optimization. *Comput. Ind. Eng.* 174 (2023) 108836.
- [8] Chopra, S., Meindl, P.: *Supply Chain Management*, 8th ed. Pearson (2024).
- [9] Simchi-Levi, D., Chen, X., Bramel, J.: *The Logic of Logistics*, 4th ed. Springer (2022).
- [10] Bijvank, M., van der Heijden, M.C.: Inventory control with advanced demand information. *Eur. J. Oper. Res.* 302(1) (2022) 83–97.
- [11] Yuce, B.: *Stochastic Modelling of Inventory Systems*. Springer (2023).

Media Representation of Artificial Intelligence and datafication in Smart City campaigns: A case study of Amsterdam and Singapore

Natasza Starzyńska [0009-0004-1027-6358]

University of Amsterdam, Faculty of Humanities
 BG1, Turfdraagsterpad 9 1090 GN Amsterdam, The Netherlands
 nataszastarzynska@gmail.com

Abstract. This paper explores how artificial intelligence and datafication are represented in the media communication of two leading smart cities, Amsterdam and Singapore. Although these cities differ in context and governance, both position technology as central to urban progress. The study examines the narratives presented on their official media platforms, Amsterdam Smart City and Smart Nation Singapore, to understand how the idea of the smart city connects digitalization with sustainability and to explore how the same concept of the smart city can be portrayed in different ways.

Keywords: Smart Cities, Artificial Intelligence, Datafication, Sustainable Development, Green Energy, Media Representation, Platform Governance, Techno- Solutionism, Surveillance Capitalism, Amsterdam, Singapore, Algorithmic Ethics

1 Extended Abstract

The concept of the smart city has become one of the defining frameworks of modern urban development, combining digital transformation with sustainability goals. While smart city initiatives promise efficiency, they also introduce new forms of surveillance and data control. Within the diverse models of intelligent cities six key dimensions economy, governance, mobility, environment, people, and living reflect how technological systems are integrated with social participation to create sustainable urban ecosystems [1]. This study investigates how artificial intelligence, datafication, and renewable energy are represented through the official media channels of Amsterdam Smart City and Smart Nation Singapore. Additionally, it follows a comparative methodology combining qualitative discourse analysis of visual and textual materials. It focuses on how these narratives construct legitimacy for data-driven governance and sustainable innovation. The research is based on theories of representation [3], surveillance capitalism [4], and techno-solutionism [2]. These perspectives make it possible to examine how language and imagery in smart city platforms transform technological systems into cultural and moral symbols of progress. In Amsterdam, artificial intelligence and

renewable energy are presented as instruments of inclusive innovation. Projects such as ATELIER Positive Energy Districts and GEMINI Mobility as a Commons highlight sustainability. The visual materials often emphasize diversity and community involvement, reinforcing the city's image as democratic. Tools like the Algorithm Register and the Data Ethics Decision Aid (DEDA) further promote transparency [5]. Yet, beneath this participatory narrative, there remains a strong element of techno-solutionism, where social and environmental challenges are reframed as technical problems, hiding the deeper power structures of data governance [2]. On the other hand, in Singapore, the smart city narrative links artificial intelligence and energy efficiency to national unity and discipline. The Smart Nation 2.0 and National AI Strategy 2.0 programs present technological advancement as essential to economic stability and public trust. Renewable energy projects and smart grid solutions are represented as collective achievements contributing to social harmony. Through slogans like "A Smart Nation we can trust," [6] surveillance and algorithmic control are reimagined as protective and helpful. For example, an app like Healthy 365 is promoting civic tools that encourage responsible behavior, normalizing the idea that constant data monitoring is part of good citizenship. Comparing these two cases reveals contrasting ideological models of sustainable urbanism. Amsterdam promotes an image of ethical participation and openness, while Singapore focuses on order, and national identity. Both, however, rely on the same logic of technological inevitability and data optimization. The smart city, in both versions, becomes not only a technical infrastructure but also a symbolic project that redefines social relations and the meaning of progress. The study also highlights the potential of applying optimization methods such as data-based resource allocation or predictive analytics to understand how cities balance social inclusion and technological efficiency. Optimization, in this context, becomes not only a technical approach but also a social one guiding the equitable distribution of technological benefits and minimizing exclusion in digital participation.

The paper concludes that representations of artificial intelligence and datafication play a key role in shaping the cultural politics of urban sustainability. Through selective storytelling, Amsterdam and Singapore transform digital infrastructures into markers of moral and environmental advancement, often masking inequalities in control and visibility.

References

1. Smart cities, smart governance: the role of communication and social participation in building cities of the future. 2025. Scientific Papers of Silesian University of Technology Organization and Management Series 2025 (227).
2. Milan, S. Techno-solutionism and the Standard Human in the Making of the COVID-19 Pandemic. *Social Media + Society*, 7(3), 1–13 (2021).
3. Hall, S. *Representation: Cultural Representations and Signifying Practices*. Sage, London (1997).
4. Zuboff, S. *The Age of Surveillance Capitalism*. PublicAffairs, New York (2019).
5. Amsterdam Smart City. <https://amsterdamsmartcity.com>
6. Smart Nation Singapore. <https://www.smartnation.gov.sg>

Towards Human–Technology Symbiosis in Audience Engagement: A Web- and IoT-based Continuous Response Digital Interface for Real-Time Musical Tension in Live Concerts

Michał Kulbacki^{1,3}[0000-0001-9113-962X], Adam Porebski¹[0009-0001-4214-0226], Jakub Segen^{2,3}[0000-0002-8697-6025], Zenon Chaczko^{3,5}[0000-0002-2816-7510], Atsushi Ito⁶[0000-0002-9686-4019], Wojciech Bożejko⁴[0000-0002-1868-8603], Andrzej Gnatowski⁴[0000-0002-9185-1841], Ryszard Klempous⁴[0000-0003-4889-0930], and Marek Kulbacki^{2,3}[0000-0003-4609-106X]

¹ The Karol Lipinski Academy of Music, Wrocław, Poland

² Polish-Japanese Academy of Information Technology, Warsaw, Poland

³ DIVE IN AI, Wrocław, Poland

⁴ Wrocław University of Science and Technology, Wrocław, Poland

⁵ School of Elect. and Data Engineering, University of Technology Sydney, Australia

⁶ Chuo University, Tokyo, Japan

Extended Abstract

Continuous Response Digital Interface (CRDI) paradigms have established the validity and reliability of real-time listener ratings of musical tension and affective response [3, 4, 9]. Yet most implementations remain confined to bespoke laboratory setups with small samples and off-line analyses that poorly reflect real concert ecologies or contemporary digital infrastructures [6]. Time-series approaches such as pre-whitened cross-correlation and ARIMA-based modelling have shown how serial dependencies between musical structure and continuous tension trajectories can be quantified [1, 2, 5, 8], and recent work on contemporary music has demonstrated systematic listener agreement in perceived segmentation and absorption in live concert settings [7, 10]. However, these methods have rarely been integrated into scalable, cloud-native systems and, to our knowledge, have not been used to measure composer–listener alignment in real time at concert scale. We present a web- and IoT-based CRDI platform that operationalises Human–Technology Symbiosis in live music perception research, understood here as tightly coupled human and digital components that co-produce and adapt a shared representation of musical experience. The system embeds industrial full-stack, MQTT-based data streaming into concert practice. A browser-based mobile interface and optional dial devices synchronously capture continuous tension ratings at 20–50 Hz, time-locked to digital audio and concert metadata, and stream them to a Google Cloud backend for on-line and off-line analysis across venues and sessions. Three public concerts of contemporary music (15 new works) have been recorded to date (N = 108 audience members, all with pre- and post-concert questionnaires), yielding multi-participant

temporal profiles of perceived tension. For a subset of works, composers also provided continuous tension traces while listening to their own pieces in performance, enabling direct analysis of composer–audience alignment that has not been reported in previous CRDI studies [1, 2]. The analysis pipeline instantiates established practice from empirical musicology and time-series statistics: step-wise resampling, per-participant z -scoring, first-order differencing for stationarity, ARIMA-based pre-whitening on composer traces, and cross-correlation of residuals to estimate lags and effect sizes, complemented by rolling inter-subject correlations and a directional agreement index that quantifies local convergence between composer and audience trajectories [1, 2, 8]. Beyond concert deployments, we have implemented a cloud-based experimental sandbox that synchronises with smartphones and supports individual and group CRDI sessions on user-specific recordings, enabling reproducible tension experiments outside the lab. Related work on collectively playable wearable music highlights the potential of networked devices in performance contexts, but does not provide a scalable infrastructure for continuous perception tracking across audiences and works [11]. Conceptually, our platform treats the coupled human–device ensemble as a symbiotic sensor of collective musical experience and as reusable infrastructure for future AI-based audio and time-series models, providing a transferable template for continuous, real-time measurement of aesthetic experience in the wild [9].

References

1. Dean, R.T., Bailes, F.: Time series analysis as a method to examine acoustical influences on real-time perception of music. *EMR* **5**(4), 152–175 (2010)
2. Farbood, M.M.: A parametric, temporal model of musical tension. *Music Perception* **29**(4), 387–428 (2012)
3. Geringer, J.M., Madsen, C.K., Gregory, D.: A fifteen-year history of the continuous response digital interface: Issues relating to reliability and validity. *Bulletin of the Council for Research in Music Education* **160**, 1–15 (2004)
4. Madsen, C.K., Fredrickson, W.E.: The experience of musical tension: A replication of Nielsen’s research using the continuous response digital interface (CRDI). *Journal of Music Therapy* **30**(1), 46–63 (1993)
5. McAdams, S., Vines, B.W., Vieillard, S., Smith, B.K., Reynolds, R.: Influences of large-scale musical form on continuous ratings in response to a contemporary piece. *Music Perception* **22**(2), 297–350 (2004)
6. Nagel, F., et al.: EMuJoy: Software for continuous measurement of perceived emotions in music. *Behavior Research Methods* **39**(2), 283–290 (2007)
7. Phillips, M.A., et al.: What determines the perception of segmentation in contemporary music? *Frontiers in Psychology* **11**, 1001 (2020)
8. Schubert, E.: Modeling perceived emotion with continuous musical features. *Music Perception* **21**(4), 561–585 (2004)
9. Schubert, E., McPherson, G.E.: The perception of emotion in music. Oxford University Press (2006)
10. Swarbrick, D., et al.: Audience musical absorption: Exploring attention and affect in the live concert setting. *Music & Science* **7** (2024)
11. Thorn, S.D., et al.: Collectively playable wearable music: Practice-situated participatory investigations. *Frontiers in Psychology* **13**, 804816 (2022)

Time-of-Flight based 3D Imaging for Endoscopic Medical Applications

Dominic Doujak^{1[0009-0000-4636-0800]} and
Thomas Thurner^{1[0000-0003-3783-1193]}

Chair of Automation and Measurement, Technical University of Leoben,
A-8700 Leoben, Austria. thomas.thurner@unileoben.ac.at

Abstract. Existing imaging systems for clinical endoscopic treatments either provide only 2D information and therefore lack true 3D perception, or offer 3D capabilities but suffer from limited resolution, insufficient short-range accuracy, and substantial integration complexity. This study presents the development and evaluation of a proof-of-concept high-resolution, ultra-compact time-of-flight (ToF) 3D imaging system designed for minimally invasive endoscopic procedures. The system employs ultra-high-frequency modulation in an indirect ToF camera module combined with optically enhanced illumination and imaging components tailored for short-range in-vivo geometries, enabling high resolution in both lateral and depth dimensions. The resulting high-density 3D point clouds at video frame rates improve usability and performance for real-time 3D registration and perception in endoscopic imaging environments.

Keywords: Endoscopic imaging · 3D imaging · Time-of-Flight · ToF

Contribution

Three-dimensional (3D) imaging offers significant advantages over traditional 2D methods in diagnosis and treatment. By capturing spatial relationships and depth, 3D modalities, such as CT, MRI, and 3D ultrasound, provide a more accurate representation of anatomical structures and increase diagnostic confidence [1]. In minimally invasive surgery, conventional 2D video lacks depth information, forcing surgeons to rely on indirect indicators such as shadows or relative size [2]. This increases cognitive load and can alter hand–eye coordination; notably, 97% of errors in laparoscopic gallbladder procedures are linked to depth misperception in 2D views. Stereoscopic 3D laparoscopes improve depth perception, reduce fatigue, and improve precision [2], underscoring the clinical need for true depth-resolving imaging in minimally invasive procedures.

This contribution describes the proof-of-concept (PoC) for a time-of-flight (ToF) based 3D endoscopic imaging approach designed for true depth-resolved visualization in confined anatomical cavities. The approach utilizes an indirect ToF camera with flood illumination by modulated near-infrared light, featuring

a high-performance VGA-resolution 3D imaging sensor of type IRS2976C (Infineon Technologies). Because the off-the-shelf camera module targets mid-range sensing from approximately 20 cm to 5 m with a $68 \times 55^\circ$ field of view, the existing lens system and illumination optics are not suited for the requirements of imaging small body cavities.

Our PoC system for ToF-based endoscopic 3D imaging is illustrated in Figure 1, with the following enhancements to the baseline module: First, the imaging and illumination optics are extended by additional lenses and mirrors to enable 3D imaging within small cavities; second, the modulation frequency of the ToF camera is increased to $f_{mod} = 300$ MHz to improve depth accuracy to the millimeter range.




Fig. 1. Endoscopic 3DI ToF imaging system overview (a). The VGA-resolution 3DI ToF module (b) is enhanced by integration into a proof-of-concept optical setup for 3D imaging within small cavities (c). Results from a 3D staircase object (d) demonstrate promising performance for 3D-perception within the limited imaging volume.

The obtained results demonstrate decent imaging performance for both lateral and depth resolution in small volume cavities (Figure 1 d,e). Integrating ToF 3D sensor technology with a custom-designed optical pathway proves to be a viable strategy for in vivo 3D sensing.

References

1. Fu, B., Damer N.: Biometric Recognition in 3D Medical Images: A Survey. In: IEEE Access, vol. 11, pp. 125601-125615, 2023. <https://doi.org/10.1109/ACCESS.2023.3331118>.
2. Sinha, R.Y., Raje, S.R., Rao, G.A.: Three-dimensional laparoscopy: Principles and practice. In: Journal of Minimal Access Surgery, vol. 13(3), p. 165, 2017. <https://doi.org/10.4103/0972-9941.181761>.

Towards Efficient Event-Based Signal Processing Using FIR Filters

Anna Werzi¹[0009-0006-1336-6660], Bernhard Moser^{1,2}[0000-0001-8373-7523], and Michael Lunglmayr¹[0000-0002-4014-9681]

¹ Institute of Signal Processing, Johannes Kepler University, Linz, Austria

{anna.werzi,bernhard.moser,michael.lunglmayr}@jku.at

² Software Competence Center Hagenberg GmbH, Hagenberg, Austria

The concept of event-based sampling addresses oversampling during periods of low signal activity by adapting to local signal properties. An example is threshold-based sampling (TBS), which samples the signal when its amplitude changes by a specified threshold [1]. These samples are represented by a spike train, which encodes the samples with their position in time and the direction of the relative amplitude change, depicted by the output of the first block in Fig. 1. While several event-based analog-to-digital-converters (ADCs) are proposed in literature [2], there is only limited work on signal processing algorithms exploiting the time-sparsity of the spikes.

This work describes a filtering technique operating on and resulting in spike sequences using common FIR filters to remove noise-induced spikes and potentially changing the positions of the remaining spikes. An iterative algorithm will be described, allowing the application of standard FIR filters on non-uniform sampling sequences, resulting in a computational complexity that scales with the input spike density. To reduce some spikes in advance, a moving average (MA) filter is applied on the non-uniform sequences as a preprocessing stage. This filter has a less optimal frequency response compared to the subsequent equiripple filter, but does not require multiplications, which can improve the overall computational complexity. Furthermore, the proposed algorithm facilitates a piecewise polynomial representation of the impulse response, which results in reduced coefficient storage requirements compared to existing concepts, where the number of coefficients scales with the (typically highly upsampled) time resolution of the spikes [3].




Fig. 1. Proposed event-based filtering chain to remove noise-induced spikes.

The potential of the proposed filtering concept depicted in Fig. 1 is demonstrated on a speech command processing task. The results on a 5.2s sequence

of three randomly selected commands superimposed with a high-pitch tone and noise are listed in Tab. 1. The audio sequence is sampled and processed uniformly and non-uniformly. The non-uniform approach yields a signal-to-noise-ratio-gain (SNRG) of 6.73 dB while compressing the samples to 8.52 %. This is achieved requiring only 34 % of the multiplications and 97 % of the additions compared to the uniform FIR evaluation at a fixed rate of 24 kHz.

Table 1. Filter performance and complexity for $T = 5.2$ seconds of speech commands.

	uniform	non-uniform w/o MA	non-uniform w/ MA
$N_{\text{out}}/N_{\text{in}}$	1	0.1574	0.0852
SNRG (dB)	4.00	5.23	6.73
Mult.	9.83e6	2.45e7	3.35e6
Add.	9.71e6	3.27e7	9.45e6

This emphasizes the capabilities of event-based sampling and processing due to its combined data compression and noise and interference suppression. Such efficiency gains can be directly reflected as complexity reductions in subsequent Spiking Neural Networks (SNNs), e.g. for speech command classification [4], in addition to the expected accuracy gains caused by noise suppression. Depending on the application, one can expect reduced computational effort, storage requirements and/or communication traffic, due to the proposed filtering concept.

Acknowledgement. This work is supported by the “University SAL Labs” initiative of Silicon Austria Labs (SAL) and its Austrian partner universities for applied fundamental research for electronic based systems, by Austrian ministries BMK, BMDW, the State of Upper-Austria in the frame of SCCH and the COMET-K2 “Center for Symbiotic Mechatronics” of the Linz Center of Mechatronics (LCM), funded by the Austrian federal government and the federal state of Upper Austria.

References

- [1] M. Miskowicz. “Send-On-Delta Concept: An Event-Based Data Reporting Strategy”. In: *Sensors* 6 (2006), pp. 49–63.
- [2] Y. Zhao and Y. Lian. “Event-Driven Circuits and Systems: A Promising Low Power Technique for Intelligent Sensors in AIoT Era”. In: *IEEE Transactions on Circuits and Systems II: Express Briefs* 69.7 (2022), pp. 3122–3128.
- [3] K. Han, Y. Wei, and X. Ma. “An efficient non-uniform filtering method for level-crossing sampling”. In: *2016 IEEE International Conference on Digital Signal Processing (DSP)*. 2016, pp. 623–626.
- [4] A. Castagnetti, A. Pegatoquet, and B. Miramond. “Trainable quantization for Speedy Spiking Neural Networks”. In: *Frontiers in Neuroscience* Volume 17 - 2023 (2023).

Edge Processing of IMU Data for IoT PAwR Communication

Thomas Kammerhofer^{1[0009-0007-3878-7741]}, Anna-Maria Stein^{1[0009-0002-2270-8723]}, and Thomas Thurner^{1[0000-0003-3783-1193]}

Chair of Automation and Measurement, Technical University of Leoben, Austria
 thomas.kammerhofer@unileoben.ac.at

Abstract. This paper presents a compact, battery-powered system for vibration-based anomaly detection that combines ultra-low-power edge processing with Bluetooth Low Energy (BLE) 5.4. An LSM6DSOX IMU performs continuous low-power activity detection and wakes a CYW20829 microcontroller only when abnormal vibrations are suspected. The MCU retrieves buffered samples, extracts lightweight spectral features, and reports a 1-byte health state via PAwR, enabling efficient, connectionless communication. This event-driven design minimizes radio use and supports multi-year operation from a coin-cell battery, making the system suitable for unobtrusive retrofitting in industrial environments. Utilizing this minimalist, edge-processing architecture, energy-intensive data transmission can be minimized, resulting in a significant reduction of communication overhead and enabling retrofitting in industrial settings.

Keywords: Edge Processing · IoT · PAwR

1 Introduction

Condition monitoring and anomaly detection are critical for predictive maintenance of industrial and structural systems. Accelerometer-based sensing, especially for seismology and vibration monitoring, has grown rapidly in recent years [1]. Traditional wireless sensor networks often rely on continuous data streaming, which demands significant energy and limits battery life. Advances in low-power communication and embedded processing show that moving computation closer to the data source can significantly improve efficiency [2]. This work presents a battery-powered embedded system for vibration-based anomaly detection that performs onboard signal processing. Bluetooth LE 5.4 adds the PAwR feature, enabling efficient broadcast-style communication between a central device and many peripherals. Unlike traditional BLE connections, PAwR provides time-synchronized, connectionless exchanges within assigned subevents, allowing hundreds of nodes to report status updates with minimal overhead.

2 System Design

The system pairs a CYW20829 BLE 5.4 microcontroller with an LSM6DSOX six-axis IMU featuring a hardware FIFO and embedded finite-state machine. In its low-power mode with activity detection, the IMU consumes just 4.4 μ A. Moreover, it autonomously handles vibration data without MCU involvement. The microcontroller remains in deep sleep until the IMU triggers an interrupt.

The microcontroller remains in deep sleep until a wakeup is triggered by the IMU. Under normal vibration levels, the IMU autonomously manages its FIFO, requiring no MCU activity. Under normal vibration, the IMU manages its FIFO internally. When an anomaly is suspected, it switches to high-performance mode (sample rate of e.g., 1000 Hz). After a wakeup, the MCU retrieves the buffered samples and performs lightweight feature extraction, comparing results to calibrated thresholds to classify the event. Firmware operation follows a state-based flow: deep sleep, interrupt handling, and periodic health reporting. The node operates as a PAwR peripheral, transmitting a 1-byte health-state packet. If no wakeup occurs within a set interval (e.g., 1 min) low-power timer wakes the MCU to send a “0” (normal); confirmed anomalies trigger a “1”. This event-driven communication strategy minimizes radio use, enabling more than two years of autonomous operation from a 220 mA coin-cell battery. The compact, self-contained module can be retrofitted to existing machinery without affecting operation or requiring certification.

3 Evaluation and Conclusion

Initial tests validated the sensing, processing, and communication pipeline. The sensor node was mounted on an electromechanical shaker capable of producing sinusoidal vibrations, with normal operation defined as a stable 50 Hz signal. The embedded algorithm reliably detected abnormal conditions. Measurements showed that the PAwR status updated immediately after a change in the vibration signal (see Fig. 1). The binary anomaly flag can be expanded into a multi-state indicator by adding a lightweight filter bank to the edge-processing stage. After an interrupt, buffered IMU samples are passed through several band-pass filters, and the short-time power in each band is computed to form a compact spectral-power vector. The resulting state is encoded as an integer, preserving minimal communication overhead while providing richer vibration information.




Fig. 1. PAwR messages depending on recorded vibration signals

References

1. Antonino D'Alessandro, Salvatore Scudero, and Giovanni Vitale. A Review of the Capacitive MEMS for Seismology. *Sensors*, 19(14):3093, January 2019. Publisher: Multidisciplinary Digital Publishing Institute.
2. Brian Kelly. The Impact of Edge Computing on Real-Time Data Processing. *International Journal of Computing and Engineering*, 5:44–58, July 2024.

Envelope Spectrum Analysis of Phonocardiogram Data

David Neuss¹[0009-0005-3065-9177], Thomas Kammerhofer²[0009-0007-3878-7741], and Thomas Thurner¹[0000-0003-3783-1193]

Chair of Automation and Measurement, Technical University Leoben, Austria,
 david.neussl@unileoben.ac.at

Abstract. In this paper we propose a method for the investigation of heart sounds by extracting and analysing the envelope and envelope spectrum. Based on analysis of the envelope spectrum, a set of parameters is derived, that can help in classifying phonocardiographic (PCG) data of different heart diseases. Efficient signal processing techniques enable the deployment on low-power microcontrollers.

Keywords: Heart Sounds · Envelope Spectrum · Signal Processing

1 Introduction

For an accessible application to assist in auscultation, such a device must be portable, fast, and affordable. Automatic heart sound classification is a well-researched topic, however often relying on complex machine learning algorithms [1] that require significant compute power.

To bridge this gap and provide a low-threshold assistance system, this project investigates acoustic heart emissions using efficient signal processing techniques. Features describing the shape of the envelope signal (ENV) and the harmonic structure of the envelope spectrum (ES) are extracted to provide clearly interpretable results. These parameters can help differentiate between healthy and non-healthy hearts and may support the diagnosis of underlying heart diseases.

2 Analysis and Parameter Extraction

Typically, the envelope spectrum is used to detect localized faults in rotating machinery, specifically bearings, which produce short, high-frequency fault signals. Due to the similarity between bearing fault signals and the PCG envelope, a comparable approach to bearing fault analysis [3] is proposed.

To obtain the envelope signal, the PCG data is bandpass filtered (50th order FIR, $50 \text{ Hz} < f < 200 \text{ Hz}$), the absolute value calculated and then passed through a lowpass filter (50th order FIR, $f_c = 50 \text{ Hz}$). A 20.000 point Fourier transform yields the envelope spectrum, where the heart rate is detected as the maximum value within the expected window $80\text{BPM} \pm 40\text{BPM}$. Figure 1 illustrates the results of these steps for a healthy heart (NORM) and two diseased hearts (Late Aortic Stenosis LAS and Mitral Stenosis MR). The source data was taken from [2], played back via a simulator setup and rerecorded using a digital stethoscope.

In the calculated envelope signal, the determined crest factor (CRST) quantifies peak height and sharpness. In the ES, the average amplitude of the first 9 harmonics (average harmonic ratio, AHR) and the average amplitude between harmonics (inter-harmonic level ratio, IHL) describe the structure of heart rate multiples, both normalized by the ES amplitude at the heart rate as shown in figure 2.




Fig. 1. Time signals (1st row), spectrograms (2nd row), envelope signals (3rd row) and envelope spectra (4th row) of healthy and pathological heart sounds

3 Results & Outlook

We see that there are notable differences in the discussed parameters between a healthy and non-healthy heart as well as between the presented disease cases.

The novelty of this paper is the analysis of the PCG by investigating the envelope and envelope spectrum, relying exclusively on very efficient signal processing methods. Future investigations will focus on further testing and the derivation of additional parameters for improved classification.




Fig. 2. Comparison of signal parameters CRST, AHR and IHL for two diseases, normalized to the parameter value of the healthy heart.

References

1. Chen, W., Sun, Q., Chen, X., Xie, G., Wu, H., Xu, C.: Deep learning methods for heart sounds classification: A systematic review. *Entropy* **23**(6), 667 (2021)
2. Kammerhofer, T., et al.: Phonocardiographic noise suppression for low-cost ble enhanced digital stethoscopes. In: 2025 IEEE Medical Measurements & Applications (MeMeA). pp. 1–6. IEEE (2025)
3. Randall, R.B., Antoni, J.: Rolling element bearing diagnostics—a tutorial. *Mechanical systems and signal processing* **25**(2), 485–520 (2011)

Strain Sensor System for Integration in Basalt-Reinforced Concrete Segments

Lukas Heindler¹[0000-0003-3951-8838], Alexandros Evangelatos¹[0009-0000-7300-1686], and Thomas Thurner¹[000-0003-3783-1193]

Technical University Leoben, Leoben 8700, Austria
 lukas.heindler@unileoben.ac.at
 http://www.unileoben.ac.at

Abstract. Modern construction aims for continuous structural health monitoring to recognize possible failures. State-of-the-art commercial measurement systems do not fit the requirements due to cost, form factor, and power requirements. We propose a custom tailored measurement system that tackles exactly these points in basalt reinforced climate-friendly concrete tunnel segments. In this work we focus on the signal chain optimizations and noise minimization. The proposed measurement system delivers a noise floor with a standard deviation of $1.6 \mu\text{V}$, a 25-fold improvement in comparison to previous prototypes. These improvements enable further integration in real tunnel segments.

Keywords: Signal chain · structural health monitoring · analog front-end · tunnel lining segments.

1 Introduction

Current climate protection initiatives are pushing modern construction towards sustainable materials and smart buildings that facilitate continuous structural health monitoring. This demand has led to the development of cost-effective, distributed, and easily integrated sensor networks for cutting-edge projects. These built-in sensor-arrays require minimal power and offer sufficient precision for most applications. They allow for wireless data transfer and are thus ideally suited for remote monitoring. We introduce sensor nodes that allow to measure strain and derived parameters like concrete crack widths in climate-friendly, non-steel-reinforced structural lining elements in modern tunnels, where currently only a few high-precision measurement systems are installed [1][2].

This work aims to minimize the noise floor of the measurement system's output signal by optimizing the signal chain including bridge analog front-end, the anti-aliasing filter (AAF) and the following digital filtering of the converted voltage signals from currently non-integrated sensors. Lab-scale bending beam tests are then conducted as a precursor to full-scale tunnel segment measurements, considering the significant production and handling efforts involved.




Fig. 1. a) Picture of the developed measurement system. b) Block diagram of the measurement system for quarter or half bridge configuration.

2 Measurements and test set-up

The measurement system, with its block-diagram shown for quarter and half-bridge configurations, as shown in Figure 1, is designed around the automotive certified 24-bit $\Delta\Sigma$ analog-to-digital converter (ADC) Analog ADS1261. It features several programmable options, including a programmable gain amplifier (PGA), an AAF, and multiple digital filters. In order to minimize sources of interference and influencing mechanisms of the sensors integrated into the fibre-reinforced basalt concrete samples, preliminary tests are being conducted using non-integrated sensor wire loops. It can be assumed that parasitic electromagnetic interference (EMI) in these loops are similar to those found in integrated sensor wire loops. Therefore, the noise optimizations should apply to both cases.

3 Results, discussion and future work

In our preliminary tests, we achieved a noise standard deviation of $1.6 \mu\text{V}$, which corresponds to a strain of $1.87 \mu\text{m}/\text{m}$. Compared to previous work in [3], this represents a 25-fold improvement, but highlights opportunities for enhancing other aspects of the measurement system. These improvements warrant further efforts to construct full-size test specimens reinforced with woven basalt and to conduct measurements using fully integrated equipment. In future work, the primary focus will be on the system's electrical characteristics, including power consumption, long-term stability, and radio frequency (RF) capabilities.

References

1. Kromoser, B., Reichenbach, S., Stoiber, N., Preinstorfer, P., Huber, T.: Potentiale von nichtmetallischer Bewehrung im Infrastruktur-Betonbau NIMETBEW, Endbericht (VIF 2019), Vienna Austria (2020).
2. Monsberger, C., Buchmayer, F., Winkler, M., Cordes, T.: Distributed fiber optic sensing in civil structural health monitoring at the next level – Realization of a comprehensive sensing network along the Brenner Base Tunnel, DOI: 10.3217/978-3-99161-057-1-004.
3. Evangelatos, A., Heindler, L., Galler, R., Thurner, T.: Integrated Sensor Technology for Basalt-Reinforced Segmental Lining Elements, DOI: 10.3217/978-3-99161-057-1-100.

A Projection-Based Hybrid Model for Nonlinear Acoustic Echo Cancellation

Alireza Nezamdoust^{*}[0000-0003-4999-263X],
 Michele Scarpiniti[†] [0000-0002-3164-6256], Mario Huemer^{*}[0000-0002-3164-7232],
 and Danilo Comminiello[†][0000-0003-4067-4504]

[†] Dept. Information Eng., Electr. and Telecom., Sapienza University of Rome, Italy

^{*}Institute of Signal Processing, Johannes Kepler University Linz, Austria
 alireza.nezamdoust@jku.at

The Nonlinear Acoustic Echo Cancellation (NAEC) remains a fundamental and critical challenge in audio signal processing, particularly because performance is strongly degraded by inherent nonlinear distortions arising from loudspeakers and complex room acoustics. Conventional linear adaptive filters, such as the Normalized Least Mean Squares (NLMS) algorithm [6], are fundamentally inadequate for accurately modeling and mitigating the effects of these nonlinearities. To address this limitation, nonlinear adaptive filters have been extensively employed due to their significant modeling versatility [4]. To this end, a common and effective implementation uses Linear-in-the-Parameters (LIP) models, which are typically adapted via time-domain filtering. However, while powerful, this methodology presents a considerable challenge for real-time NAEC applications, which often necessitate extremely long filters to model room impulse responses [5]. The primary limitation in such scenarios is the substantial computational burden, as the cost of processing scales directly with the filter length [1].

This paper presents a comprehensive study of advanced nonlinear adaptive filtering techniques, specifically focusing on the Volterra [3] and the Adaptive Exponential Functional Link (AEFL) [5] models, and introduces a novel framework to enhance their computational efficiency and performance using Projection-based Decomposition (PD) [2]. The proposed methodology integrates PD to simultaneously split the echo path into early and late components, and projects high-dimensional feature vectors onto lower-dimensional subspaces using dedicated projection matrices. This dimensionality reduction yields the Volterra-PD and AEFL-PD models, which significantly reduce computational complexity while maintaining high adaptability. As confirmed in Fig. 1, simulation results conducted by using colored noise and a sigmoid nonlinearity, demonstrate that the PD-enhanced models achieve superior Echo Return Loss Enhancement (ERLE) compared to the conventional models. The achieved balance of enhanced performance and reduced complexity confirms the ability of the PD-based models for efficient, real-time deployment in demanding applications such as teleconferencing systems and smart devices.

In the full paper, we will provide a detailed mathematical formulation of the PD framework, an analysis of computational complexity reductions, and ex-




Fig. 1. Comparative analysis of the performance in terms of ERLE provided by NLMS, AEFL, Volterra, AEFL-PD, and Volterra-PD models when applied to a colored input signal subjected to a sigmoidal nonlinearity.

tended simulation results in various input and noise conditions to further validate the robustness and applicability of the proposed models.

Acknowledgments This work is supported by the COMET-K2 “Center for Symbiotic Mechatronics” of the Linz Center of Mechatronics (LCM), funded by the Austrian federal government and the federal state of Upper Austria.

References

1. Comminiello, D., Scarpiniti, M., Azpícueta-Ruiz, L.A., Arenas-García, J., Uncini, A.: Functional link adaptive filters for nonlinear acoustic echo cancellation. *IEEE Transactions on Audio, Speech, and Language Processing* **21**(7), 1502–1512 (2013)
2. Galántai, A.: Projectors and projection methods, vol. 6. Springer Science & Business Media (2003)
3. Guérin, A., Faucon, G., Le Bouquin-Jeannès, R.: Nonlinear acoustic echo cancellation based on Volterra filters. *IEEE Transactions on Speech and Audio Processing* **11**(6), 672–683 (2004)
4. Nezamdoust, A., Huemer, M., Uncini, A., Comminiello, D.: Efficient functional link adaptive filters based on nearest Kronecker product decomposition. In: ICASSP 2024 – 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 886–890. IEEE (2024)
5. Patel, V., Gandhi, V., Heda, S., George, N.V.: Design of adaptive exponential functional link network-based nonlinear filters. *IEEE Transactions on Circuits and Systems I: Regular Papers* **63**(9), 1434–1442 (2016)
6. Uncini, A., et al.: Fundamentals of adaptive signal processing. Springer (2015)

Complex-Valued Spline Adaptive Hammerstein Filter

Bernhard Plaimer^[0009-0007-6377-5726], Matthias Wagner^[0000-0003-1031-2983],
 Oliver Lang^[0000-0002-6612-0692], and Mario Huemer^[0000-0002-3164-7232]

Institute of Signal Processing, Johannes Kepler University, Linz, Austria,
 {bernhard.plaimer, matthias.wagner, oliver.lang, mario.huemer}@jku.at

In various technical applications, the task is to identify nonlinear systems. In this work, we consider a class of nonlinear systems called Hammerstein systems [8], which consist of a static memoryless nonlinearity followed by a linear time-invariant (LTI) subsystem. Such systems occur in different applications, two examples are the modeling of radio-frequency digital-to-analog-converters [9] and the use of cascaded models for the cancellation of air-induced passive intermodulation [3]. For the identification of these systems, spline adaptive filters (SAFs) and, in particular, the Hammerstein SAF (HSAF), have been introduced, where spline interpolation is used to approximate the nonlinear block of the Hammerstein system [7].

In many cases, Hammerstein systems are complex-valued (CV) and process CV signals, e.g., nonlinear wireless communication transmitters modeled in the CV baseband [1, 2, 5]. In such situations, the goal is to approximate a static, CV, memoryless, nonlinear function $f(x_k)$ using splines, with x_k as the input signal of the Hammerstein system, and with k being the discrete time index. If this function can be expressed as $f(x_k) = f_{\text{Re}}(\text{Re}[x_k]) + j f_{\text{Im}}(\text{Im}[x_k])$, then real-valued (RV) HSAFs may be used for identifying $f_{\text{Re}}(\cdot)$ and $f_{\text{Im}}(\cdot)$, respectively. While in [1, 2, 5] CV Hammerstein systems are identified using only RV HSAFs, most practical examples involve more general CV functions. Hence, similar to the CV Wiener SAF (WSAF) presented in [4], CV HSAFs may be required.

In the following, the contributions of the full paper are briefly summarized:

- Similar to [4], this work utilizes two-dimensional rational tensor-product B-splines to derive a novel CV spline-based Hammerstein model. In the full paper, it will be shown that the input–output relation of this CV spline-based Hammerstein model can be expressed in bilinear form as

$$y_k = \mathbf{h}^H \mathbf{X}_k \mathbf{g}, \quad (1)$$

where \mathbf{h} models the impulse response of the LTI subsystem of the unknown Hammerstein system, and \mathbf{g} denotes the CV knots of the spline-function, which should approximate $f(x_k)$. The product $\mathbf{X}_k \mathbf{g}$ corresponds to the evaluation of the CV splines, where \mathbf{X}_k depends on the input signal x_k .

- The goal of CV HSAFs is to adapt the vectors \mathbf{g} and \mathbf{h} to approximate $f(x_k)$ and the subsequent LTI system. This can be achieved by using CV bilinear (BL) filters, such as the CV BL normalized least mean squares (CBNLMS) filter, introduced in [6]. The combination of the model in (1) with the update equations of the CBNLMS filter results in a novel CV HSAF.




Fig. 1. (a) Block diagram of a CV nonlinear communication transmitter in combination with a CV linear channel \mathbf{h} . (b) Convergence behavior of HSAFs.

– Finally, this CV HSAF is used to identify a nonlinear CV communication transmitter with analog system non-idealities, such as I/Q imbalance and a nonlinear power amplifier (PA), as shown in Fig. 1a. Using the instantaneous squared error (ISE) as the performance measure, preliminary results demonstrate significantly better convergence behavior of this CV HSAF filter compared to that obtained with two RV HSAFs, as shown in Fig. 1b.

References

1. Campo, P.P., Anttila, L., Korpi, D., Valkama, M.: Cascaded Spline-Based Models for Complex Nonlinear Systems: Methods and Applications. *IEEE Trans. Signal Process.* **69**, 370–384 (2021)
2. Campo, P.P., Korpi, D., Anttila, L., Valkama, M.: Nonlinear Digital Cancellation in Full-Duplex Devices using Spline-Based Hammerstein Model. In: *IEEE Globecom Workshops*. pp. 1–7 (2018)
3. Lampu, V., Anttila, L., Turunen, M., Fleischer, M., Hellmann, J., Valkama, M.: Cancellation of Air-Induced Passive Intermodulation in FDD MIMO Systems: Low-Complexity Cascade Model and Measurements. In: *IEEE/MTT-S Int. Microw. Symp.* pp. 33–36 (2023)
4. Liu, C., Zhao, H.: A 2D-LUT Scheme Design for Complex-Valued Spline Adaptive Filter. *IEEE Trans. Circuits Syst. II, Exp. Briefs* **70**(8), 3154–3158 (2023)
5. Paireder, T., Motz, C., Huemer, M.: Spline-Based Adaptive Cancellation of Even-Order Intermodulation Distortions in LTE-A/5G RF Transceivers. *IEEE Trans. Veh. Technol.* **70**(6), 5817–5832 (2021)
6. Plaimer, B., Wagner, M., Lang, O., Huemer, M.: Optimum and Adaptive Complex-Valued Bilinear Filters. *arXiv preprint: 2505.09215* (2025)
7. Scarpiniti, M., Comminiello, D., Parisi, R., Uncini, A.: Hammerstein uniform cubic spline adaptive filters: Learning and convergence properties. *Signal Process.* **100**, 112–123 (2014)
8. Schoukens, M., Tiels, K.: Identification of block-oriented nonlinear systems starting from linear approximations: A survey. *Automatica* **85**, 272–292 (2017)
9. Trampitsch, S., Kalcher, M., Enzinger, H., Gruber, D., Lunglmayr, M., Huemer, M.: A Circuit-Inspired Digital Predistortion of Supply Network Effects for Capacitive RF-DACs. *IEEE Trans. Microw. Theory Techn.* **69**(1), 271–283 (2021)

Time-Domain Estimation of Frequency-Dependent IQ Imbalance for 5G NR-based JCAS Systems

Andreas Meingassner-Lang^[0000-0002-4296-9157], Moritz Tockner^[0009-0001-0404-2959], Oliver Lang^[0000-0002-6612-0692], Bernhard Plaimer^[0009-0007-6377-5726], Matthias Wagner^[0000-0003-1031-2983], and Mario Huemer^[0000-0002-3164-7232]

Institute of Signal Processing, Johannes Kepler University, Linz, Austria
 {andreas.meingassner-lang, moritz.tockner, oliver.lang, bernhard.plaimer, matthias.wagner, mario.huemer}@jku.at

Extended Abstract

The 5G New Radio (NR) standard provides a foundation for joint communications and sensing (JCAS) by enabling the reuse of communication signals for sensing applications [1–3, 6, 7]. However, precise up- and down-conversion is required, as any imbalance between the in-phase and quadrature-phase (IQ) components can introduce significant signal distortions causing an increased noise floor and ghost targets in the range-Doppler map (RDM).

Although many studies have proposed compensation techniques for IQ imbalance in communication systems, only a few have addressed its impact on the sensing performance of JCAS systems, especially when considering imbalances in both, the transmitter and receiver paths. Existing approaches focus on robust waveform design to mitigate IQ imbalance effects [4]. However, such methods are not applicable in the context of 5G NR, where the waveform is standardized, and the uplink data is random. To the best of our knowledge, there are currently no published solutions that address these specific constraints.

In our previous work [7], we introduced a method to mitigate the effects of frequency-independent (FID) IQ imbalance in the RDM. In this contribution, we extend our approach to jointly estimate and compensate for frequency-dependent (FD) IQ imbalance at the receiver and FID imbalance at the transmitter, following the modelling in [4, 5]. Our method is based on estimating the impulse response of the FD IQ imbalance, which allows for efficient IQ imbalance compensation using only a limited number of filter taps. In the full paper, we will also provide a comparison with a frequency-domain estimation approach.

In Fig. 1a, the effect of IQ imbalance is exemplarily illustrated by covering weaker objects in the surrounding. In Fig. 1b, the improvement achieved by applying the proposed compensation method can be observed. The weaker objects can now be detected by a simple peak search. Simulations, which will be detailed in the full paper, demonstrate that the proposed compensation strategy significantly reduces the impact of IQ imbalance in the RDM, thereby improving the

sensitivity and reliability of 5G NR-based sensing systems. A detailed performance analysis as well as the announced comparison with a frequency-domain estimation approach will be given in the full paper.




Fig. 1: IQ imbalance effects on the RDM (a) without and (b) with IQ imbalance compensation and illustrated in dB.

References

1. 3GPP: NR; Physical channels and modulation. Technical Specification (TS) 38.211, 3rd Generation Partnership Project (3GPP) (2025), <http://www.3gpp.org/DynaReport/38211.htm>
2. Barneto, C.B., Liyanaarachchi, S.D., Heino, M., Riihonen, T., Valkama, M.: Full Duplex Radio/Radar Technology: The Enabler for Advanced Joint Communication and Sensing. *IEEE Wireless Commun.* **28**(1), 82–88 (2021). <https://doi.org/10.1109/MWC.001.2000220>
3. Hakobyan, G., Yang, B.: A Novel Intercarrier-Interference Free Signal Processing Scheme for OFDM Radar. *IEEE Trans. Veh. Technol.* **67**(6), 5158–5167 (2017). <https://doi.org/10.1109/TVT.2017.2723868>
4. Lang, O., Hofbauer, C., Tockner, M., Feger, R., Wagner, T., Huemer, M.: OFDM-Based Waveforms for Joint Sensing and Communications Robust to Frequency Selective IQ Imbalance. *IEEE Trans. Veh. Technol.* pp. 1–15 (2024). <https://doi.org/10.1109/TVT.2024.3463802>
5. Lopez-Estraviz, E., Rore, S.D., Horlin, F., Der Perre, L.V.: Optimal Training Sequences for Joint Channel and Frequency-Dependent IQ Imbalance Estimation in OFDM-based Receivers. In: *IEEE Int. Conf. Commun.* vol. 10, pp. 4595–4600 (2006). <https://doi.org/10.1109/ICC.2006.255364>
6. Meingassner, A., Lang, O., Tockner, M., Hofstadler, M., Lindorfer, G., Lunglmayr, M., Huemer, M.: Near Range, Leakage Aware Object Localization Based on 5G NR-Compliant JCAS. In: *Asilomar Conf. Signals, Syst. Comput.* pp. 32–38 (2024). <https://doi.org/10.1109/IEEECONF60004.2024.10942687>
7. Meingassner, A., Lang, O., Tockner, M., Plaimer, B., Wagner, M., Lindorfer, G., Hofstadler, M., Huemer, M.: Tx and Rx IQ Imbalance Compensation for JCAS in 5G NR. *arXiv preprint arXiv:2504.10272* (2025)

An Approximated Lower Test Signal Generator Linearity Bound of the AST Linearity Test

Matthias Wagner¹, Oliver Lang¹, and Mario Huemer¹

Institute of Signal Processing, Johannes Kepler University, Linz, Austria
 matthias.wagner@jku.at

Nonlinear RF devices, e.g., power amplifiers, are usually characterized by their intercept points, such as the third-order intercept point (IP3). The IP3 is typically measured via two-tone tests, which heavily rely on the spectral purity of the test signal, necessitating highly specialized test equipment, making it a challenging task when performed with on-chip equipment [1]. The issue of on-chip IP3 testing has recently been addressed by the artificial slow-time (AST) linearity test, which supports test signals from low-quality on-chip test signal generators (TSGs) [2]. To do so, the AST linearity test couples the same imprecise test signal $x(t)$ into the receiver M times. Each time, the test signal is analogously modified according to so-called slow-time amplitude modulation (STAM) values $\cos\left(\frac{2\pi}{M}n\right)$. [2]

This work discusses the effects of a low TSG IP3 on the AST linearity test. Furthermore, to support the implementation of the AST linearity test, an approximated lower TSG linearity bound is presented. Similar to [2], consider the radar receiver illustrated in Fig. 1 as the system under test (SUT), and the AST linearity test signal $x_T^{\text{mod}}(n, t) = \cos\left(\frac{2\pi}{M}n\right)x(t)$, where the measurement index n is called slow-time index, and where t is the time variable within a single measurement referred to as fast-time. The nonlinear low noise amplifier (LNA) output signal is modeled as

$$y(t) = \alpha_1 (h_1(t) * x_T^{\text{mod}}(n, t)) + \alpha_3 (h_3(t) * x_T^{\text{mod}}(n, t))^3, \quad (1)$$

with α_1 the linear gain of the LNA, α_3 the third-order coefficient, $\alpha_1\alpha_3 < 0$, and the two linear subsystems $h_1(t)$ and $h_3(t)$, which enable the modeling of a frequency dependent IP3. Then, as shown in [2], the analog-to-digital converter (ADC) output signal may be written as

$$\begin{aligned} y_{\text{ADC}}[n, k] \approx & [(\alpha_1 x_1(t) x_{\text{LO}}(t)) * h_{\text{LP}}(t)]|_{t=kT_s} \cos\left(\frac{2\pi}{M}n\right) + \\ & \left[\left(\frac{1}{4} \alpha_3 x_3^3(t) x_{\text{LO}}(t) \right) * h_{\text{LP}}(t) \right] \Big|_{t=kT_s} \cos\left(3 \cdot \frac{2\pi}{M}n\right), \end{aligned} \quad (2)$$

with $x_1(t) = h_1(t) * x(t)$, $x_3(t) = h_3(t) * x(t)$, $h_{\text{LP}}(t)$ the impulse response of a low-pass filter, $x_{\text{LO}}(t)$ the local oscillator (LO) signal, T_s the sampling period of the ADC, and k the discrete fast-time index. The two summands in (2) can be separated by an FFT across the slow-time dimension, which separates the linear part of the SUT from its third-order term. Nevertheless, the third-order term




Fig. 1: Block diagram of radar receiver with TSG.




Fig. 2: Third-order intermodulation product voltage levels.

still depends on the on-chip generated, imprecise two-tone signal $x(t)$. The top graph in Fig. 2 illustrates the voltage level of the third-intermodulation product $A_{ADC,3}|_{dB}$ as a function of the input tone power level $P|_{dB}$, whereby the blue line is obtained with an ideal TSG, the orange line considers a TSG input third-order intercept point (IIP3) of -10 dBm, and the green line is obtained with a TSG IIP3 of 0 dBm. As can be seen, for a sufficiently high input tone power level (>-26 dBm) and a sufficiently high TSG IIP3 (green line), the voltage level of the third-order intermodulation product matches the ideal value. However, for lower TSG IIP3 values (orange line), the determined voltage level deviates from its ideal value for higher input power levels.

In the full paper, it will be shown that the TSG's IIP3 must fulfill

$$A_{IIP3}|_{dB} > A|_{dB} - 10 \log \left(1 - 10^{\frac{\epsilon_{P3}}{20}} \right) + 11.8 \text{ dB} \quad (3)$$

in order to guarantee the functionality of the AST linearity test. Here, $A_{IIP3}|_{dB}$ is the TSG IIP3, $A|_{dB}$ is the tone amplitude of the two-tone input signal, and ϵ_{P3} is the difference between the third-order intermodulation voltage level obtained by the AST linearity test and the ideal one. Given $\epsilon_{P3} = -1$ dB, the bottom graph in Fig. 2 highlights when the inequality holds true, marked by the indicator variable $\epsilon_b = 1$. This result establishes an approximated lower linearity bound for the TSG, supporting the design process when implementing the AST linearity test.

Acknowledgments. This work has been supported by the "LCM - K2 Center for Symbiotic Mechatronics" within the framework of the Austrian COMET-K2 program.

References

1. Ahmad, S., Azizi, K., Zadeh, I.E., Dabrowski, J.: Two-tone PLL for on-chip IP3 test. In: IEEE Int. Symp. Circuits Syst. pp. 3549–3552. IEEE (2010)
2. Wagner, M., Lang, O., Dorrer, S., Ghafi, E.K., Schwarz, A., Huemer, M.: Accurate On-Chip Linearity Monitoring With Low-Quality Test Signal Generation. In: IEEE Int. Symp. Circuits Syst. pp. 1–5. IEEE (2023)

An Enhanced mm-Wave Tomographic Imaging Platform for Low-Permittivity Materials

Manuel Koller^{1,2}, Andreas Och², Matthias Wagner¹, Andreas Windner², Oliver Lang^{1,4}, Stefan Schuster³, Stefan Scheiblhofer³, Patrick A. Hölzl², Stefan Matzinger², Johann Reisinger³, and Mario Huemer¹

¹ Institute of Signal Processing, Johannes Kepler University, Linz, Austria

² Infineon Technologies Austria AG, Linz, Austria

³ voestalpine Stahl GmbH, Linz, Austria

⁴ Christian Doppler Laboratory for Steel Industry Signal Processing and Machine Learning, Linz, Austria

Tomographic microwave imaging can be applied in a multitude of applications in both medical and industrial fields, including cancer detection, nondestructive testing, and material characterization. The vast majority of classical tomography approaches utilize magnitude-based evaluation techniques, which show poor performance for low-permittivity materials due to their small dielectric contrast. A previously proposed concept addresses this limitation by exploiting time-of-flight (ToF) variations based on carrier-phase changes in coherent radar signals, rather than magnitude attenuation [2].

The ToF difference between a transmit–receive antenna pair separated by a distance d for signals propagating through two non-magnetic materials with relative permittivities ϵ_{r1} and ϵ_{r2} , can be approximated as $\Delta\tau \approx \frac{d}{c_0}(\sqrt{\epsilon_{r2}} - \sqrt{\epsilon_{r1}})$, where τ_1, τ_2 denote the respective ToFs [2]. In coherent radar systems, these ToF variations manifest as phase differences $\Delta\Phi$ according to $\Delta\tau = \frac{\Delta\Phi}{2\pi f_c}$, where f_c denotes the carrier frequency. By measuring $\Delta\tau$, or equivalently $\Delta\Phi$, across multiple transmit–receive antenna pairs, spatially dependent permittivity changes can be derived, forming the basis for tomographic reconstruction.

The feasibility of this concept for tomographic reconstruction of phantoms with $\epsilon_r \approx 1$ was previously demonstrated using two 76–81 GHz frequency-modulated continuous-wave (FMCW) radar sensors operating in a bi-static configuration. The sensors had a single transmit channel and two receive channels and were placed around a motorized rotation stage [2], enabling measurements from multiple angles. To further improve spatial coverage, one sensor and the common LO-distribution were manually repositioned after each acquisition, improving reconstruction quality at the expense of longer measurement times.

In this work we present an enhanced system that overcomes the limitations of the previous setup. The system, illustrated in Fig. 1, consists of five radar sensors, each featuring two transmit and four receive channels equipped with sectoral horn antennas. The sensors are arranged around a motorized rotation stage and synchronized via a fixed common LO-distribution for multi-static operation. This increases the effective number of transmit–receive antenna pairs while reducing measurement time from several hours to approximately three minutes. Total variation (TV)-regularized tomographic reconstructions [1] with

corresponding ground truths for two different phantoms are illustrated in Fig. 2. The accurate recovery of edges and internal features demonstrates that the system provides superior spatial resolution and phase stability for mm-wave tomographic imaging. Furthermore, it establishes a robust platform for subsequent methodological and algorithmic advancements.

In the full paper we will outline the theory behind ToF-based mm-wave low-permittivity tomography, describe the new system architecture in detail, and quantitatively compare reconstructions and ground truths for different phantoms.




Fig. 1: The enhanced tomography system. Primary and secondary sensors (1, 2), RX antennas (3), TX antennas (4), foam phantoms (5), motorized rotation stage (6), LO-synchronization link (7).




Fig. 2: Ground truths and tomographic reconstructions of XPS foam phantoms.

Acknowledgments. This work has been supported by the COMET-K2 Center of the Linz Center of Mechatronics (LCM) funded by the Austrian federal government and the federal state of Upper Austria.

The financial support by the Austrian Federal Ministry of Economy, Energy and Tourism, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged.

References

1. Hansen, P., Jørgensen, J., Lionheart, W. (eds.): *Computed Tomography: Algorithms, Insight, and Just Enough Theory*. Society for Industrial and Applied Mathematics (2021)
2. Och, A., Hölzl, P.A., Schuster, S., Scheiblhofer, S., Zankl, D., Pathuri-Bhuvana, V., Weigel, R.: High-resolution millimeter-wave tomography system for nondestructive testing of low-permittivity materials. *IEEE Transactions on Microwave Theory and Techniques* **69**(1), 1105–1113 (2021). <https://doi.org/10.1109/TMTT.2020.3030662>

3D Point Cloud Generation Using mmWave Radar Sensors

Elias Wollitzer¹[0009-0000-9024-4727], Martina Zeinzinger¹[0009-0009-5471-9129], Michael Kargl¹[0009-0009-7753-3335], Dan Hrubař²[0009-0009-7109-9734], Florian Eibensteiner¹[0000-0002-5311-4082], Martin Černý²[0000-0002-8893-2587], and Josef Langer¹

¹ ¹ University of Applied Sciences Upper Austria, Hagenberg, Austria
 {elias.wollitzer, martina.zeinzinger, michael.kargl,
 florian.eibensteiner, josef.langer}@fh-hagenberg.at

² ² VSB – Technical University of Ostrava, Faculty of Electrical Engineering and
 Computer Science, Czech Republic
 {dan.hrubař, martin.černý}@vsb.cz

Abstract. We present an evaluation of 3D point cloud generation using compact mmWave radar sensors with different antenna configurations and compare them for the task of human presence and fall detection. Unlike optical systems such as LiDAR or cameras, radar enables privacy-preserving, robust and temperature-independent object detection under challenging environmental conditions such as darkness, steam, or dust. This work focuses on the FMCW radar sensor BGT60TR13C (Infineon) and compares it to the IWR6843 (Texas Instruments) to assess the impact of antenna configuration on point cloud accuracy and resolution. Results from 19 controlled posture scenarios show that the BGT60TR13C produces sparse short-range point clouds, whereas the IWR6843 provides higher density and angular resolution. These findings underscore the strong impact of antenna design on point cloud performance.

Keywords: FMCW · point cloud · fall detection

1 Introduction

Compact millimeter-wave (mmWave) radar systems have become increasingly attractive for person detection in healthcare and industrial applications. Unlike camera- or LiDAR-based methods, radar sensors operate reliably under poor visibility and provide privacy-preserving environment sensing. A key factor influencing point cloud density and angular resolution is the antenna configuration. The Texas Instruments IWR6843, with its 3 transmit (TX) and 4 receive (RX) antenna array, provides high angular resolution and is widely used for 3D point cloud generation in studies such as [1], achieving 93% fall-detection accuracy. In contrast, the Infineon BGT60TR13C employs a minimal 1-TX/3-RX setup sufficient only for basic angle estimation, and such reduced antenna configurations remain largely unexplored in existing research. This disparity raises the

question to what extent antenna configuration influences 3D point cloud density and accuracy in human presence and fall detection. This work investigates this question by comparing both sensors in terms of 3D point cloud generation capability.

2 Concept

Both evaluated systems are based on Frequency-Modulated Continuous-Wave (FMCW) radar operating in the 60 GHz mmWave band and capture raw intermediate frequency data for digital signal processing. A complete signal processing chain was implemented for the BGT60TR13C, including range and angle estimation, clutter suppression, and digital beamforming for improved angular resolution. The processed data are transformed into Cartesian coordinates to generate a 3D point cloud. The complete processing chain is illustrated in Figure 1 and described in detail in [2]. The IWR6843 served as the reference system using its integrated point cloud generation.




Fig. 1. Radar signal processing chain for 3D point cloud generation [2].

3 Results

Nineteen test scenarios were conducted to measure human postures under controlled conditions and to assess point cloud quality. The results show that point cloud generation with the minimal antenna setup of the BGT60TR13C is feasible. The radar produced sparse but consistent 3D point clouds at short ranges in simple environments, averaging 2-8 points per scene, while the IWR6843 yielded 25–95 points. Due to its limited antenna configuration with 1 TX and 3 RX, the BGT60TR13C exhibits reduced angular resolution and range. These findings highlight the strong influence of antenna design on point cloud density and quality in mmWave radar systems.

References

1. Dang, X., Jin, P., Hao, Z., Ke, W., Deng, H., Wang, L.: Human Movement Recognition Based on 3D Point Cloud Spatiotemporal Information from Millimeter-Wave Radar. *Sensors* **23**(9), 9430 (2023). <https://doi.org/10.3390/s23239430>
2. Xing, M., Zhang, G., Geng, X., Lin, Y.-J.: Comprehensive mPoint: A Method for 3D Point Cloud Generation of Human Bodies Utilizing FMCW MIMO mm-Wave Radar. *Sensors* **21**(19), 6455 (2021). <https://doi.org/10.3390/s21196455>

Cross-Domain Time-Series Analytics for Industrial Monitoring

Evans Doe Ocansey¹[0009-0007-4281-2567] Roxana-Maria Holom¹[0000-0002-1321-0475], Hannes Hinterbichler²[0000-0002-9374-1099], Michael Araz³, Andreas Krenn⁴, and Velic Dino⁴.

¹ RISC Software GmbH, Softwarepark 32a, 4232 Hagenberg, Austria
 evansdoe.ocansey@risc-software.at, roxana.holom@risc-software.at

² Fronius International GmbH, Pettenbach, Austria
 hinterbichler.hannes@fronius.com

³ BENTELER Automobiltechnik GmbH, Paderborn, Germany
 michael.araz@benteler.com

⁴ Voestalpine Signaling Austria GmbH, 8740 Zwettlweg, Austria
 krenn.andreas@voestalpine.com, dino.velic@voestalpine.com

Background and Motivation. The digitalization of industrial processes and transport infrastructure relies heavily on the intelligent analysis of sensor data. This work presents a comparative study of time-series analysis frameworks applied to two distinct European projects: **metaFacturing** (manufacturing) and **IAM4RAIL** (railway logistics). While the domains differ, both face the challenge of extracting actionable insights from high-frequency, heterogeneous sensor streams.

In the context of **metaFacturing** (Data and Metadata for Advanced Digitalization of Manufacturing Industrial Lines), a primary driver of cost is the rework originating from out-of-specification components. A critical bottleneck is the quality of weld seams. To mitigate defects, it is essential to equip operators with AI-based decision support systems. Our hypothesis is that latent patterns in welding process data correlate with seam quality. The objective is twofold: (1) Predict weld seam quality to enable early defect detection, and (2) Identify process variables responsible for defects to guide corrective actions.

In parallel, the **IAM4RAIL** project addresses the need for automated monitoring of railway infrastructure. Efficient asset management and maintenance require precise knowledge of rolling stock usage. Currently, relying solely on scheduled timetables is insufficient due to operational deviations. Operational train parameters serve as critical inputs for monitoring algorithms and data clustering; therefore, they must be estimated directly from time-series data. The objectives for this use case are: (1) To classify the specific type of train passing the crossing, and (2) To accurately estimate the train's speed.

By validating our methodology on these divergent use cases, this research demonstrates a generalized approach to signal-based AI that enhances efficiency in both material production and transport logistics.

Method. For the manufacturing use case, the dataset consists of high-frequency welding process data sampled at 10 kHz. For the railway use case, the dataset comprises of high frequency time-series sensor data. These measurements are captured specifically at either switch and crossing sections of the track, utilizing multi-sensor nodes. **Feature Engineering and Modeling:** We employ a

comprehensive feature extraction pipeline, though specific implementations vary by domain: (1) **metaFacturing (Welding)**: We compare two paradigms. The first is domain-specific signal processing, computing the Power Spectral Density (PSD) to capture frequency-domain characteristics. The second utilizes the Python library, `tsfresh` [1], to systematically generate a high-dimensional feature matrix. (2) **IAM4RAIL (Railway)**: We rely exclusively on the algorithmic feature extraction via `tsfresh` [1], omitting the PSD step. This generates a robust feature set capable of characterizing the temporal signatures of different rolling stock types as they traverse the crossing and sleeper sections. Machine learning models are trained on these feature matrices. For welding, these are binary classifiers optimized for balanced accuracy to detect defects. For the railway case, we train multi-class classifiers for train type prediction and regression models for speed estimation. **Explainable AI (XAI)**: To ensure reliability and actionable feedback, we apply XAI techniques to both domains. A model-agnostic approach calculates SHAP (SHapley Additive exPlanations) values to quantify feature importance. This is crucial for identifying which process parameters cause welding defects or which signal characteristics (e.g., lateral acceleration or longitudinal rotation) define a specific train type. Additionally, we explore model-intrinsic insights using an attention-based deep Multiple Instance Learning model [2].

Conclusion and Outlook. This work establishes a cross-domain AI framework capable of handling complex time-series data. In the manufacturing domain, the models move beyond binary classification to provide diagnostic insights into defect formation. In the railway domain, the methodology successfully classifies train types and estimates velocities using sensor data from the switch and crossing infrastructure. Future efforts will focus on benchmarking these models on held-out production datasets and deploying them in live environments—both on the factory floor and at rail crossings—to validate real-time performance.

Acknowledgments. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them. The work described in this paper is supported by the **metaFacturing** project (GA 101091635) and the **IAM4RAIL** project (GA 101101966), which have received funding under the Horizon Europe programme.

References

1. Christ, M., Braun, N., Neuffer, J., Kempa-Liehr, A.W.: Time series feature extraction on basis of scalable hypothesis tests (tsfresh – a python package). *Neurocomputing* **307**, 72–77 (2018). <https://doi.org/doi.org/10.1016/j.neucom.2018.03.067>
2. Ilse, M., Tomczak, J., Welling, M.: Attention-based deep multiple instance learning. In: Dy, J., Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 80, pp. 2127–2136. PMLR (10–15 Jul 2018), <https://proceedings.mlr.press/v80/ilse18a.html>

Vision-Driven Autonomous Traffic Signal Control through Deep Reinforcement Learning

Diego Caballero-García-Alcaide^[0009-0005-4451-9642], David Yagüe^[0000-0001-9679-5589], José María Armingol^[0000-0002-3353-9956], M. Paz Sesmero^[0000-0001-9473-6809], and Araceli Sanchis^[0000-0002-1429-4092]

Universidad Carlos III de Madrid, Av. de la Universidad 30, 28911 Leganés, Spain

Keywords: Traffic Congestion · Traffic Signal Control · Deep Reinforcement Learning · Intelligent Infrastructures.

1 Introduction

Urban traffic congestion, largely exacerbated by fixed-time traffic signal controllers, remains a critical challenge for modern cities. Traditional systems lack the adaptability to respond to dynamic traffic conditions, leading to unnecessary delays and increased emissions. To address this, we propose an intelligent traffic management system driven by a Deep Reinforcement Learning (DRL) agent.

The main innovation of our work is an agent that learns to control traffic lights using only raw visual input. By processing photorealistic images of an intersection, our approach eliminates the dependency on costly physical sensor infrastructures, paving the way for more scalable solutions.

2 Methodology

Our system is built on a high-fidelity co-simulation platform where CARLA renders the photorealistic environment from which the agent perceives its state, and SUMO manages the underlying traffic dynamics and vehicle behavior. The agent's decision-making core is a Deep Q-Network (DQN) [1], a DRL algorithm ideal for handling high-dimensional inputs like images. The key components of our DRL framework were defined as follows:

- **State Space:** 100x100 pixel grayscale images providing a complete view of the intersection.
- **Action Space:** A discrete set of actions, each corresponding to activating a specific traffic signal phase.
- **Reward Function:** The negative of the total accumulated vehicle waiting time during a control step.

* This work was supported by MCIN/AEI/10.13039/501100011033/ and FEDER UE, under Grant PID2022-140554OB-C32, and by the Comunidad de Madrid under the project SEGVAUTO 5 Gen-CM (TEC-2024/ECO-277).

3 Experimentation and Results

The agent's efficacy was benchmarked against a traditional fixed-time controller in different traffic scenarios and conditions.

1. **Performance in the training intersection:** In its native training environment, the agent demonstrated superior performance. It reduced average vehicle waiting times by up to 22.68% in real traffic conditions and by 12.74% under heavy saturation, where it also increased vehicle throughput by 4.81%. This confirms its ability to create more efficient and adaptive traffic flow.
2. **Generalization to a new intersection:** The agent's ability to generalize was tested by deploying it in a new intersection with a different topology. A direct transfer of the trained model was unsuccessful, and a fine-tuning approach yielded suboptimal results. However, retraining the agent from scratch in the new environment proved highly effective, achieving a waiting time reduction of up to 23.76%. This indicates that while visual feature extraction may be partially transferable, the control policy is highly dependent on the specific intersection geometry.
3. **Multi-agent scenario:** A two-agent system controlling adjacent intersections was implemented to investigate emergent synergistic effects. The results showed minimal impact on the first intersection, but substantial improvement at the second due to the first agent's presence. With both agents active, the second intersection saw its average reward improve by nearly 60% and average waiting times fall by up to 25.19% compared to when it operated alone. This demonstrates a clear positive synergy, where intelligent control at one node creates a more orderly traffic flow that enables greater optimization at adjacent nodes.

4 Conclusions

This work validates that a vision-based DRL agent is a feasible and highly effective solution for intelligent traffic signal control, capable of significantly outperforming traditional methods without relying on massive sensor infrastructures.

Future research will focus on developing more advanced multi-agent coordination strategies and exploring techniques to improve knowledge transfer, thereby reducing the need for retraining and enhancing the scalability of the system.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D.: Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015). <https://doi.org/10.1038/nature14236>

Data Extraction and Processing for DIY Drone Forensics

Andreas Attenberger, Martha Schmidt, Alexander Hoyer, and Lukas Hardi

Research Unit Digital Forensics, Central Office for Information Technology in the Security Sector (ZITiS), 81677 München, Germany

poststelle@zitis.bund.de

<https://www.zitis.bund.de>

Extended Abstract

With the increasing processing power of embedded systems as well as hardware capabilities, the construction of do-it-yourself (DIY) drones with powerful flight and control capacities is possible with commercial off-the-shelf components. The usage of DIY drones presents several advantages to end users including flexible use of lower cost components for customized application scenarios as well as the employment of open source control software. However, features as the possible circumvention of no-fly-zones by way of control software manipulation presents novel challenges to law enforcement agencies. With further evolution of the DIY drone market, it is crucial for digital forensic practitioners to dispose of the necessary information for swift forensic extraction of relevant drone data during investigations as well as for fast response in emergency situations. In this paper we identify the main components carrying drone and user data for subsequent extraction and forensic processing. Challenges and future research questions in digital forensics that arise by the custom hardware and software format of these devices are given.

Hardware components of DIY drones typically consist of the drone frame, battery and power distribution board, radio receiver, electronic speed controllers, motors and propellers, flight controllers as well as the remote controller for operating the drone. Additional components can be various sensors, cameras, video transmitter or air units, first person view (FPV) goggles and a dropping mechanism for load transportation. Furthermore, it is sometimes possible to connect a smartphone for additional control features [4, 3]. Relevant digital forensic data is generally confined to the flight controller as well as the remote controller, with some media data potentially left on the air unit. Additionally, meta data might be present on the receiver including configuration data, which could be relevant in select cases.

For forensic extraction we advocate the prioritized analysis of the flight controller for flight data including GPS coordinates as well as the remote control and air unit for potential video or photo recordings. While it can be challenging to access the individual physical memory content for these components in case of damage or only residual internally stored data, this can be solved by regular

chip-off procedures. Additionally, in many cases, both the flight controller as well as the air unit are equipped with separate digital memory cards for easy data access. Any access to further small scale information like meta data is mandated by the course of the forensic investigation. While the aforementioned process for data extraction does not pose specific challenges, it is necessary to individually analyze different components due to the highly customized and flexible configuration possibilities of DIY drones. Once extracted, data can be further analyzed to yield information like flight logs for further re-construction, for example of flight paths. While the available autopilot software modules employ different formats including encryption for storage of this meta data, open source tools exist for processing select file types [2]. This also allows the creation of customized log processing tools.

Future challenges include the highly dynamic market development with some traditional drone manufacturing companies contributing with their own individual DIY drone components as well as the high amount of customization which precludes the development of general drone forensic tools for quick extraction. Besides hardware storage, we expect to observe the advent of cloud storage in DIY drone forensics as a future source of forensic information to be considered. Current digital forensic research should focus on identifying similarities between the individual control components with a focus on autopilot software solutions. Subsequently, more powerful data extraction, processing and visualization tools can be developed. Machine learning should be considered for identifying individual drone operators in the scope of criminal investigations. While this approach has successfully been applied in automotive forensics [1], it needs to be verified in general drone forensics. Furthermore, future research should consider the impact of artificially intelligent algorithms introduced both in flight operation of individual drones as well as drone swarms. Forensic artefacts that allow identification of swarm members as well as swarm operators need to be determined.

References

1. Dološ, K., Meyer, C., Attenberger, A., Steinberger, J.: Driver identification using in-vehicle digital data in the forensic context of a hit and run accident. *Forensic Science International: Digital Investigation* **35**, 301090 (2020)
2. Hoyer, A., Lohre, K., Schmidt, M., Hardi, L., Attenberger, A.: Uncovering hidden traces of open source uav autopilots. *European Academy of Forensic Science* (2025)
3. Klier, S., Baier, H.: Beware of the rabbit hole - A digital forensic case study of DIY drones. In: Iwaya, L.H., Kamm, L., Martucci, L.A., Pulls, T. (eds.) *Secure IT Systems - 29th Nordic Conference, NordSec 2024 Karlstad, Sweden, November 6-7, 2024 Proceedings*. Lecture Notes in Computer Science, vol. 15396, pp. 325–344. Springer (2024). <https://doi.org/10.1007/978-3-031-79007-2-17>
4. Schmidt, M., Hoyer, A., Lohre, K., Hardi, L., Attenberger, A.: Flying evidence: First insights into the forensic investigation of diy drones. *European Academy of Forensic Science* (2025)

Investigating Greater Manchester Tram Service Optimization Using Multi-Layer Transport Library Simulation Approach

Kate Han¹  and Lee A. Christie² 

¹ Salford Business School, University of Salford, Manchester, UK

² School of Computing, Engineering and Technology, RGU, Aberdeen, Scotland, UK

Workshop: Applied Data Science and Engineering for Intelligent Transportation Systems and Smart Mobility

Keywords: Public transport · Tram network optimization · Multi-layer transport simulation · Evolutionary algorithms · Routing heuristics

Extended Abstract

Urban mobility across Greater Manchester is undergoing strategic transformation through Transport for Greater Manchester (TfGM) and the Bee Network initiatives, which aim to provide an integrated, accessible, and low-emission public transport system. However, tram connectivity in parts of Salford remains limited, restricting the potential for modal shift from private vehicles and equitable access to employment, education, and essential services. Recent plans to expand the Metrolink to Salford's Rugby area highlight the need for evidence-based assessment of potential network extensions [6]. Additionally, prior work on sustainable transport behaviour and personal carbon credits in Greater Manchester [5] underscores the broader environmental importance of improving public transport accessibility. These factors motivate the present study to explore potential tram network expansions in Salford and the wider Greater Manchester area.

This research will investigate candidate scenarios, including the addition of new tram stops and extensions along existing corridors such as the Blue, Pink, and Yellow lines. The study aims to: (i) identify high-impact stop-line configurations, (ii) explore potential improvements in accessibility, travel time, and network coverage, and (iii) provide evidence-based guidance for operational feasibility under varying demand conditions.

A multi-layer transport simulation library developed at RGU has been employed to simulate the multi modal transportation network in Greater Manchester Area. This framework integrates GTFS timetable data, OpenStreetMap (OSM) infrastructure, and polygon-based geographic boundaries to construct a spatially accurate representation of the tram system and its interactions with other transport modes. Evolutionary algorithms and routing heuristics have been applied to explore optimal stop placements and line extensions while respecting operational constraints.

Experimental scenarios are generated and assessed using metrics, both travel times and accessibility metrics. Key outcomes of interest include potential reductions in travel times, increased reachable population, and insights into network performance improvements. This work aims to provide a structured framework for planning tram network enhancements in Greater Manchester. It is expected to contribute to the sustainability principles of the public transport design core values.

By combining multi-layer transport simulation with evolutionary optimisation, this study proposes a flexible, data-driven methodology for evaluating future tram network expansions. The framework is intended to support urban planners in designing cost-effective, high-impact interventions and is generalisable to other urban areas and transport modes, providing a foundation for sustainable public transport planning.

Acknowledgements

The authors would like to acknowledge the support of Salford Business School computing resources and the RGU Transport Simulation Laboratory for both usage of the transport library and expert guidance.

References

1. Han, K., Christie, L.A., Zăvoianu, A.-C., McCall, J.A.W.: Optimising Public Transport Through the Integration of Micro and Macro-level Simulations. In: International Conference on Computer Aided Systems Theory, pp. 107–121. Springer (2024)
2. Han, K., Christie, L.A., Zăvoianu, A.-C., McCall, J.A.W.: Exploring Representations for Optimizing Connected Autonomous Vehicle Routes in Multi-Modal Transport Networks Using Evolutionary Algorithms. *IEEE Transactions on Intelligent Transportation Systems* (2024)
3. Han, K., Christie, L.A., Zăvoianu, A.-C., McCall, J.: Optimising the introduction of connected and autonomous vehicles in a public transport system using macro-level mobility simulations and evolutionary algorithms. In: Proceedings of the Genetic and Evolutionary Computation Conference Companion, pp. 315–316 (2021)
4. Han, K., Christie, L.A., Zăvoianu, A.-C., McCall, J.: On discovering optimal trade-offs when introducing new routes in existing multi-modal public transport systems. In: International Conference on Computer Aided Systems Theory, pp. 104–111. Springer (2022)
5. Han, K., Tzanetakis, M., Oginini, E., Matthew, A., Joseph, B.: Carbon Crest: A Personal Carbon Credit Model for Sustainable Transport Behaviour in Greater Manchester. Accepted for AI-2025: Forty-fifth SGAI International Conference on Artificial Intelligence, Cambridge, UK, 16–18 December 2025.
6. Manchester Evening News: Plan to expand Metrolink to Salford's Rugby area. Manchester Evening News, 28 March 2023. <https://www.manchestereveningnews.co.uk/news/greater-manchester-news/plan-expand-metrolink-salfords-rugby-30607273>

Bringing People and Luggage to the Airport

Bin Hu^{1[0000-0002-3184-8907]}, Patrick Stern², Alexander Schaffenberger², Theresa Hirsch³, Johanna Berndorfer³, Arno Fallast^{4[0000-0002-0712-8184]} and Samuel Lesak^{4[0009-0006-3072-0000]}

¹ AIT Austrian Institute of Technology, Giefinggasse 4, Vienna 1210, Austria

² Planum Fallast & Partner, Wastiengasse 14, Graz 8010, Austria

³ Brimatech Services, Lothringerstraße 14/3, Vienna 1030, Austria

⁴ FH Joanneum, Alte Poststraße 149, Graz 8020, Austria

Abstract

More than one billion passengers travel by airplanes annually in the EU. Since most airports are located in the outskirts, getting there may not be so straightforward. Nowadays, most frequently used transport modes are private car, taxi, and public transport (train, metro, bus, etc.). They all have their advantages but also critical drawbacks: Private car comes with high parking expenses at the airport, taxi is expensive for long journeys, and public transport has unflexible operation hours and usually requires transfers, which can be inconvenient for passengers with luggage.

In this paper, we examine three promising concepts designed to make airport trips more convenient, while remaining ecological and economical. Some are feasible today, and others will be possible in the near future. To estimate their practicability, we assess the concepts and their potential coverage in a case study for Austria. According to a survey we conducted, the most valued factors for air travelers are reliability, punctuality, number of transfers, availability (at any time), safety, travel time, comfort, transport costs, short walking distance, and environmental friendliness. Other studies show similar preferences [1, 2]. While most existing studies assess traveler preferences of current transport modes, we consider three future concepts in a case study for Austria: automated taxi, automated shuttle, and e-bike. We assess the operability and potential coverage based on GIS reachability modeling, cost estimation from operational data, and a logistics optimization model.

Automated taxi vs. conventional taxi vs. private car

The automated taxi (e.g., Tesla Robotaxi) offers 24/7 direct private transport and will be the most convenient choice. Its biggest advantage over conventional taxis is the affordability and thus the substantially increased potential coverage. Nowadays, a taxi trip from Vienna to the airport costs 30-40 Euro (fixed price tariff). Uber costs are comparable. For future automated taxies, estimated costs are around 0.16 Euro per km [3]. Even with a pessimistic estimate of 0.32 Euro per km, 30 Euro will enable trips up to 90 km. For comparison, driving a private car incurs parking costs of at least 35 Euro per day at the airport, making it the most expensive option.

Automated shuttle + automated taxi vs. train

If the distance to the airport is too far for an automated taxi, a combination with automated shuttle offers a reasonable alternative. The automated shuttle could act as a replacement for the train during off-peak times (e.g., at night, when trains are not profitable). An automated shuttle shared by 4-5 passengers can operate between any major railway station and a major airport (Vienna or Munich) for 30 Euro. Adding a 60-min trip via automated taxi from the railway station, this combination can cover most parts of Austria (see Figure 1 left) with only one transfer and is available 24/7.

E-bike or public transport + luggage service

For short distances, travelling to the airport by e-bike is a valid option if the following requirements are fulfilled: safe bike lanes to the airport, reasonable parking facilities for the e-bike at the airport, and a service to transport the luggage separately. If we assume that the e-bike can travel 20 km to the airport, around 40% of Vienna is reachable, see Figure 1 right. For the luggage, we assume that the service collects it on the day before and delivers it to the airport. Using a logistics optimization model, we obtain an OPEX of around 4-5 Euro per suitcase, which is reasonable. The luggage service is also an interesting option for people travelling via public transport.




Figure 1 Left: Reachability of automated shuttle from major airports (red diamonds) to major railway stations (black squares) + 60min automated taxi (blue area). *Right:* Reachability of Airport Vienna (red diamond) via 20 km ride on an e-bike (blue area).

To conclude, the case study shows three promising concepts for the airport transport in the future. The assumptions and results are not only valid in Austria but can be transferred to other countries in Europe.

Acknowledgments. This work received partially funding from the Austrian Research Promotion Agency (FFG) under grant number 5134303 (Sustainable Mobility for Airport Connectivity).

References

1. Gupta, S., Vovsha, P., & Donnelly, R. Air Passenger Preferences for Choice of Airport and Ground Access Mode in the New York City Metropolitan Region. *Transportation Research Record: Journal of the Transportation Research Board*, 2042(1), 3-11. (2008)
2. Makino, K., Compostella, J., Lee, Y., Circella, G. Dataset of air travelers' ground transportation choice at four airports in Northern California. Dryad. (2024)
3. ARK Invest. Big Ideas (2025)

Evaluating and Improving Data Management in Urban Data Platform Development Research Projects

Ignacio Marco-Pérez¹[0009-0003-7381-294X], Felix Larrinaga²[0000-0003-1971-0048], Alain Pérez²[0000-0002-6200-589X], Miren Illarramendi²[0000-0003-3770-1495], Beatriz Pérez¹[0000-0001-9235-7311], and Ángel Luis Rubio¹[0000-0002-3052-6404]

¹ Departamento de Matemáticas y Computación, Universidad de La Rioja, Spain

² EPS Electrónica e Informática, Mondragon Unibersitatea, Spain

Background and Objectives. Urban Data Platforms (UDPs) play a central role in the digital transformation of cities, enabling integrated management of energy, mobility, and infrastructure data to advance sustainability and climate goals [4]. Given their importance, it is essential to consider the perspective of validation and evaluation of these tools, an aspect that is often overlooked, particularly in research projects focused on developing this type of technology. Two examples of such projects are CITYFiED [1] and SmartEnCity [2], which, in alignment with the European objective of achieving climate neutrality by 2050, aimed to develop and validate UDPs supporting data-driven urban governance and energy efficiency.

Despite technological progress and lessons transferred from one project to the next, recurring challenges and opportunities in data management, quality, and governance became evident. In this context, this study aims to analyze the data management performance within both projects through the MAMD v4.0 framework, assessing the maturity and consistency of their data-related processes. The *Alarcos' Model for Data Improvement* [3] (MAMD by its acronym in Spanish) has a comprehensive process-based structure and a solid grounding in internationally recognized standards such as ISO 8000, ISO/IEC 33000, and the DAMA³ Data Management Book of Knowledge (DAMA-DMBoK) [5]. Although originally developed within the Spanish context, its methodological compatibility with UNE 178104:2017 and its alignment with global data management practices make it particularly suitable for evaluating data maturity in research-oriented smart city projects.

The research ultimately seeks to draw insights for improving data management practices in research and innovation projects, and to critically reflect on the suitability of data maturity models (originally designed for organizational settings) when applied to experimental and time-bounded research environments.

Research Methodology. A methodological triangulation was employed to assess data management maturity in the CITYFiED and SmartEnCity projects. With this design, project deliverables and internal documentation were first reviewed

³ Data Administration Management Association

to identify potential gaps in data management processes, following the structure of the DAMA-DMBoK. Second, these findings were refined through interviews with key members of the data platform development teams and, finally, complemented by the analysis of database and data flow schemas when required.

The study adopted an iterative analytical cycle, in which successive refinements allowed for the quantification of MAMD processes. Each task within the framework was assigned a degree of commitment (ranging from 0 to 1) reflecting the level of involvement. These values were aggregated by maturity level (from 1 to 5) and by component (*Data Management*, *Data Quality Management*, and *Data Governance*), also enabling comparison across projects. The resulting dataset supported both numerical and interpretive analysis, revealing patterns across both maturity levels and components.

Results and Discussion. The quantitative assessment revealed a consistent decline in the degree of commitment as maturity levels increased. While *Data Management* achieved the highest levels of engagement (particularly at intermediate maturity) both *Data Quality* and *Data Governance* showed weaker results, reflecting the limited long-term perspective typical of research projects.

Differences between the two initiatives indicate gradual improvement through experiential learning, especially in processes related to data infrastructure and historical data. Nevertheless, variations in data integration and data lifecycle suggest that project-specific objectives shaped the outcomes.

Although the UDP developments effectively covered operational data management, weaknesses remained in metadata and master data, data observability and traceability, as well as the lack of systematic strategies for quality and governance. These results emphasize the need to complement technical progress with organizational measures to ensure sustainability and data reuse.

Overall, the study demonstrates the feasibility of applying MAMD v4.0, originally designed for organizational settings, to the context of research-based urban data initiatives. This initiative has not only provided diagnostic information, but also supports decision-making for future R&D projects of a similar nature.

References

1. Cityfied: Replicable and Innovative Future Efficient Districts and Cities. <http://www.cityfied.eu/>, last visited on November, 2025
2. Smartencity – towards smart zero co2 cities across europe. <https://smartencity.eu/>, last visited on November, 2025
3. Alarcos Research Group, AQC Lab, D.: Mamd: Modelo alarcos de madurez de datos, diciembre 2023 (Dec 2023)
4. de Reuver, M., Ofe, H., Agahari, W., Abbas, A.E., Zuiderwijk, A.: The openness of data platforms: A research agenda. In: Proceedings of the 1st International Workshop on Data Economy. pp. 34–41. DE '22, Association for Computing Machinery, New York, NY, USA (Dec 2022). <https://doi.org/10.1145/3565011.3569056>
5. International, D.: DAMA-DMBOK: Data Management Body of Knowledge (2nd Edition). Technics Publications, LLC, Denville, NJ, USA (Jun 2017)

CLG-Smooth: Temporal Regularization for Smoother and Safer VLM-RL Driving Policies

Gabriela Cordova¹[0009-0001-1421-7898], Hafsa Iqbal¹[0000-0002-1773-8066],
 Abdulla Al-Kaff¹[0000-0003-0212-6400], Beatriz López
 Boada¹[0000-0002-8061-068X], and Fernando García
 Fernández¹[0000-0002-6291-5009]

Universidad Carlos III de Madrid, Spain
 {gcordova,hiqbal,akaff,bboada,fegarcia}@ing.uc3m.es

Abstract. Autonomous driving requires smooth, human-like control, and Vision-Language Reinforcement Learning (VLM-RL), provides dense, interpretable rewards by aligning visual observations with textual goals under physical constraints. We propose a *directional temporal smoothness* term to the VLM-RL synthesized reward to penalize abrupt steering changes between consecutive steps. The shaping is applied at rollout as stored reward and at replay as recomputed reward without double counting, improving low-level stability, safety, and training consistency.

Keywords: Reinforcement Learning · Vision-Language Models · Temporal Smoothness · Autonomous Driving

1 Introduction

Hand-crafted RL rewards are brittle for encoding smooth, human-like driving [1]. VLM-RL leverages Contrastive Language–Image Pretraining (CLIP) [2] to align visual observations with text and combines Contrasting Language Goals (CLG) [3, 4] semantics with bounded physical factors to yield dense, interpretable rewards. However, the temporal structure of control often remains under-specified, which induces noise in steering commands [5]. We present *CLG-Smooth*: a lightweight penalty on consecutive steering changes, which is applied during rollout and replay, that integrates VLM with RL seamlessly.

2 Method

The proposed synthesized driving reward $R_{\text{synthesis}}(s_t)$ is defined as follows:

$$R_{\text{synthesis}}(s_t) = r_{\text{speed}}(s_t) \cdot f_{\text{center}}(s_t) \cdot f_{\text{angle}}(s_t) \cdot f_{\text{stability}}(s_t), \quad (1)$$

where s_t is the vehicle state; $r_{\text{speed}}(s_t)$ measures speed alignment to a safe target; $f_{\text{center}}(s_t)$ is lateral position wrt lane center; $f_{\text{angle}}(s_t)$ is the heading alignment; and $f_{\text{stability}}(s_t)$ is the lateral-position consistency; all terms lie in $[0, 1]$.




Fig. 1. Performance metrics: (A) CPM, (B) Collision rate, (C) Routes completed, and (D) Entropy coefficient.

$R_{\text{synthesis}}(s_t)$ is employed to compute the reward R_t by including a quadratic temporal smoothness term, as follows:

$$R_t = R_{\text{synthesis}}(s_t) - \lambda_{\text{steer}} (\Delta a_t)^2, \quad \lambda_{\text{steer}} > 0, a_t \in [-1, 1], \quad (2)$$

where a_t denote the scalar steering at time t and define $\Delta a_t = a_t - a_{t-1}$. R_t is computed step-by-step during rollout and, for off-policy updates, re-applied from consecutive steers during replay to form critic targets. To avoid double counting, we employ a single shaped signal per transition: if stored rollout reward is negative, we keep it; otherwise, we use Eq. (2).

3 Results & Outlook

Figure 1 provides key tracked metrics: **Safety** (A-B) *Collisions per million* CPM and **collision rate** drop early and remain below the baseline, with steadier trajectories after the initial transient, indicating reduced inter-step steering oscillations. **Operational & training dynamics** (C-D): **routes completed** reaches higher final levels and **entropy coefficient** converges lower, indicating a less stochastic policy. Overall, penalizing rapid steering changes leads to smoother control, safer behavior, and better task throughput. Future work will enable high-level waypoint/curvature planning and dynamic, contextual multi-channel language goals aligned to physical factors.

Acknowledgments. Work is supported by PID2021-128327OA-I00 (MCIN/AEI/10.13039/501100011033), ERDF, and EU NextGenerationEU/PRTR.

Disclosure of Interests. No competing interests.

References

1. Doroudian, H., Taghavifar, P.: CLIP–RLDrive: Reward Shaping in Autonomous Driving with CLIP. arXiv:2411.09876 (2024)
2. Radford, A., et al.: Learning Transferable Visual Models from Natural Language Supervision. ICML (2021)
3. Huang, Z., Sheng, Z., Qu, Y., You, J., Chen, S.: VLM–RL: A Unified Vision–Language Models and Reinforcement Learning Framework for Safe Autonomous Driving. arXiv:2412.15544 (2024)
4. Sheng, Z.: VLM-RL Repository. GitHub. (2024)
5. Hsu, B.-J., Cao, H.-G., Lee, I., Kao, C.-Y., Huang, J.-B., Wu, I.-C.: Vision-Based Regularizing Action Policies for Smoothing Control in Autonomous Miniature Car Racing. Workshop on AI for Autonomous Driving @ IJCAI (2022)

Data-Driven Human-AI Co-Simulation and Digital Twin Integration for Smart Urban Mobility

Yuzhou LIU^[0000-0003-2785-5976], Amirhesam Aghanouri^[0000-0003-2785-5976], Cristina Olaverri-Monreal^[0000-0002-5211-3598]

Department Intelligent Transport Systems, Johannes Kepler University Linz, Austria
 {yuzhou.liu, amirhesam.aghanouri, cristina.olaverri-monreal}@jku.at

Abstract. This work extends 3DCoAutoSim platform into a data-driven framework for intelligent urban mobility. Integrating Unreal Engine 5 and SUMO, it provides a realistic 3D environment where humans and AI agents interact in real time. A new module captures and analyzes behavioral data from human drivers and pedestrians to enhance agent realism and study human influence on traffic flow and safety. By aligning simulations with real-world data from Linz, the platform acts as a digital twin for continuous calibration. Combining human participation, behavioral modeling, and data-driven analysis, it supports research on adaptive last-mile logistics, pedestrian safety, and modular vehicle deployment.

Keywords: Human-in-the-loop simulation, Digital twin, Last-mile transport, 3DCoAutoSim, SUMO

1 Introduction and Method

The simulator 3DCoAutoSim [1] was enhanced by adopting Unreal Engine 5 [2] for immersive visualization and SUMO [3] for large-scale traffic modeling. Both are synchronized through a custom TraCI-based [4] middleware enabling real-time interaction of vehicles and pedestrians. Fig.1 shows the overall architecture of the system.




Fig. 1: Architecture diagram illustrating the interaction between Unreal Engine, SUMO and TraCI.

The extended framework introduces human-in-the-loop capability: multiple participants act as drivers or pedestrians in mixed traffic, while their reactions, paths, and gestures are captured as behavioral datasets. Lightweight pattern recognition and clustering derive representative behavior models that are transferred to AI agents to increase realism. Fig. 2 shows a comparison of the simulation scene with real world.




Fig. 2: Comparison between highly realistic simulation scenes and real-world environments.

2 Results

Initial studies indicate that human-in-the-loop participation improves route adaptability, traffic smoothness, and realism in mixed mobility scenarios. Future work will refine behavioral modeling and strengthen digital-twin calibration to support data-driven policy analysis and human-centered mobility design.

Acknowledgment. This work was supported by the Ergodic project, co-financed by the European Union and Austrian Research Promotion Agency (FFG).

References

1. Hussein, A., Diaz-Álvarez, A., Armingol, J. M., & Olaverri-Monreal, C. (2018, November). 3DCoAutoSim: Simulator for cooperative ADAS and automated vehicles. In 2018 21st International Conference on Intelligent Transportation Systems (ITSC) (pp. 3014-3019). IEEE.
2. Epic Games. (2019). Unreal Engine. Retrieved from <https://www.unrealengine.com>
3. Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y. P., Hilbrich, R., ... & Wießner, E. (2018, November). Microscopic traffic simulation using sumo. In 2018 21st international conference on intelligent transportation systems (ITSC) (pp. 2575-2582). Ieee.
4. Wegener, A., Piórkowski, M., Raya, M., Hellbrück, H., Fischer, S., & Hubaux, J. P. (2008, April). TraCI: an interface for coupling road traffic and network simulators. In Proceedings of the 11th communications and networking simulation symposium (pp. 155-163).

Improving Highway Construction Site Safety with the Traffic Monitoring and Warning System

Alexander Maletzky^[0000-0003-4378-7854], Felix Oberhauser^[0000-0002-9278-2497],
and Roxana-Maria Holom^[0000-0002-1321-0475]

RISC Software GmbH, Softwarepark 32a, 4232 Hagenberg, Austria
alexander.maletzky@risc-software.at

Background and Motivation. Construction sites are one of the most dangerous places on highways, for workers and road users alike. For instance, drivers may be inattentive and miss the beginning of a construction site or may be stressed by the unfamiliar traffic routing, narrow lanes and construction vehicles that enter and exit the highway in unusual places. The proposed research demonstrator TMWS—*Traffic Monitoring and Warning System*—aims to automatically detect dangerous situations in highway construction sites and issue warnings to the workers in real time.

Methods. TMWS consists of one or more monocular RGB-cameras and edge computing devices. The cameras are placed near critical locations, e.g., the beginning of construction sites, highway exits within construction sites, or lane shifts (where individual lanes are temporarily re-routed to the opposing carriageway). The cameras record the oncoming traffic, detect and track vehicles, determine the lane they are using and estimate their distance and speed in real time. If a vehicle is driving too fast, keeps changing lanes, or does not use any of the designated lanes (e.g., the emergency lane), this indicates a potentially dangerous situation. Based on this information, TMWS may either trigger warnings or help construction vehicles exit the construction site in a safe moment.

In more concrete terms, TMWS operates in two phases: an initial automatic calibration phase, and the subsequent operational phase. During calibration, vehicles and license plates are detected using a state-of-the-art detection model, like YOLO [4], and tracked over consecutive frames using the model-agnostic ByteTrack algorithm [5]. At the end of the calibration, vehicle trajectories are clustered using the DBSCAN algorithm wrt. Hausdorff distance to identify individual lanes. A more direct approach to lane detection, employing pre-trained models such as CLRerNet [3], proved to be infeasible due to the different camera perspectives (dash-cam vs. camera mounted on a pole at a height of 3-4 m) and the complex lane routing in construction sites. Then, a centroid trajectory is calculated for each cluster, representing the ‘average’ trajectory of the corresponding lane. Simultaneously, the pixel-height of each detected license plate is computed and then compared to the known standardised height (12 cm in Austria) to estimate the distance of the license plate and the corresponding vehicle, in a similar fashion as in [2]. Thus, all estimated distances and inferred

vehicle speeds are finally averaged to obtain robust distance and (average) speed estimates for all points on all centroid trajectories.

Vehicles are also detected and tracked during the operational phase. Each vehicle is then assigned to the closest centroid trajectory to determine the lane it is using. Its current distance and speed are inferred from the corresponding distance on the trajectory assigned during calibration. Hence, no license plate detection is necessary in this phase, significantly reducing the computational complexity. Although this approach to distance and speed estimation lacks accuracy in absolute terms (especially compared to other vision-based vehicle speed estimation methods [1]), it allows to compute the speed of vehicles relative to the average speed on a lane during calibration and is thus sufficient for identifying safety-critical situations.

Conclusion and Outlook. TMWS is developed as part of an ongoing research project. Although most concepts and ideas have already been implemented in actual hardware and software, there are still some challenges and open problems. For instance, the precise rule-based definition of safety-critical situations based on the high-level features extracted by our system must be continuously refined with domain experts, taking both retrospective and prospective validation results into account. Likewise, the optimal warning system for construction site workers (acoustic and/or visual signals, wearable devices, etc.) is currently being researched.

Acknowledgments. This work was funded by the FFG (Austrian Research Promotion Agency) under the grant 910803 (SafeRoadWorks).

References

1. Fernández Llorca, D., Hernández Martínez, A., García Daza, I.: Vision-based vehicle speed estimation: A survey. *IET Intelligent Transport Systems* **15**(8), 987–1005 (2021). <https://doi.org/10.1049/itr2.12079>
2. Filipiak, P., Golenko, B., Dolega, C.: NSGA-II Based Auto-Calibration of Automatic Number Plate Recognition Camera for Vehicle Speed Measurement. In: Squillero, G., Burelli, P. (eds.) *Applications of Evolutionary Computation*. pp. 803–818. Springer International Publishing, Cham (2016). https://doi.org/10.1007/978-3-319-31204-0_51
3. Honda, H., Uchida, Y.: CLRerNet: Improving Confidence of Lane Detection with LaneIoU. In: 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1165–1174 (2024). <https://doi.org/10.1109/WACV57701.2024.00121>
4. Tian, Y., Ye, Q., Doermann, D.: YOLOv12: Attention-Centric Real-Time Object Detectors (2025). <https://doi.org/10.48550/arXiv.2502.12524>
5. Zhang, Y., Sun, P., Jiang, Y., Yu, D., Weng, F., Yuan, Z., Luo, P., Liu, W., Wang, X.: ByteTrack: Multi-object Tracking by Associating Every Detection Box. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 1–21. Springer Nature Switzerland, Cham (2022). https://doi.org/10.1007/978-3-031-20047-2_1

Bridging the Planning-Execution Gap with a Construction Site Revision System

Felix Oberhauser^[0000-0002-9278-2497], Alexander Maletzky^[0000-0003-4378-7854],
 Ahmed Alshenoudy^[0009-0004-1431-9255], Bertram
 Sabrowsky-Hirsch^[0000-0001-7081-5182], Markus Steindl^[0009-0008-8088-2264],
 Roxana-Maria Holom^[0000-0002-1321-0475], and Evans Doe
 Ocansey^[0009-0007-4281-2567]

RISC Software GmbH, Softwarepark 32a, 4232 Hagenberg, Austria
 felix.oberhauser@risc-software.com

Background and Motivation. Road construction sites are essential for maintaining critical infrastructure, yet they often require adjustments to layouts, such as lane width modifications or temporary object placements (e.g., barriers, signage). Minor placement deviations from the original plan can lead to hazardous scenarios for workers and road users, underscoring the need for frequent, accurate monitoring. Traditional methods, however, are labour-intensive, time-consuming, and limited in capturing real-time site conditions, creating a gap between planning and reality. This challenge necessitates an efficient, scalable solution to close this gap and improve safety, accuracy, and documentation for construction sites.

Methods. To address these challenges, we present a Construction Site Revision System (CSRS), a cost-effective solution designed to streamline the digitalisation of construction sites. CSRS integrates a consumer-grade calibrated stereo camera system with a synchronized GPS antenna mounted on a vehicle, enabling precise georeferencing of key road entities (e.g. delineators, traffic cones). Critical road objects are detected via a fine-tuned object detector, e.g. RF-DETR [3], and then combined with depth maps from the stereo camera to estimate their position relative to the vehicle. Road boundaries are estimated via a fine-tuned CLRerNet model [1]. All objects are then projected into real-world coordinates using the GPS signal and the vehicle's heading direction. Thus, the data can be utilised and visualised in geospatial applications and used for condition monitoring (e.g. orientation or absence of objects) and assertion of minimal road widths (important for special transports passing the site).

CSRS also addresses the scarcity of annotated datasets for Austrian construction site objects through a semi-automated annotation pipeline that accelerates the preparation of detection datasets by reducing manual annotation effort. The pipeline leverages open-vocabulary detection models (e.g., GroundingDINO [2], OmDet [7]) and foundation models (e.g., DINOv3 [4]) to generate pseudo-labels from text prompts and visual queries. To further enhance efficiency, we evaluate the informativeness of these pseudo-labels with respect to detector model training, drawing on active learning-inspired strategies for object detection [6]

and principles from object detection under sparse annotations [5]. The pipeline limits the user's role to verification and minor refinements by providing a curated subset of pre-annotated images, substantially reducing the human effort required to produce a high-quality dataset.

Conclusion and Outlook. While CSRS achieves reliable results with limited hardware budget and vehicle requirements, its research still faces open challenges. Each drive-through can vary due to object occlusions and changes in lighting, and the system's performance in these conditions remains to be explored. Additionally, reliable re-identification of road entities despite inaccurate positioning and efficient visualization of georeferenced data with user feedback remain areas for improvement.

Via the semi-automated annotation pipeline, we hope to contribute to overcoming dataset scarcity, enabling rapid development of high-quality training data for construction site applications. By integrating georeferenced data with real-time monitoring, CSRS addresses the planning-reality gap, providing practitioners with actionable insights into site conditions. This system not only enhances safety and reduces errors but also streamlines documentation processes, making it a valuable tool for critical infrastructure projects.

Acknowledgments. This work was funded by the FFG (Austrian Research Promotion Agency) under the grant 910803 (SafeRoadWorks).

References

1. Honda, H., Uchida, Y.: CLRerNet: Improving Confidence of Lane Detection with LaneIoU. In: 2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). pp. 1165–1174 (2024). <https://doi.org/10.1109/WACV57701.2024.00121>
2. Ren, T., Jiang, Q., Liu, S., Zeng, Z., Liu, W., Gao, H., Huang, H., Ma, Z., Jiang, X., Chen, Y., Xiong, Y., Zhang, H., Li, F., Tang, P., Yu, K., Zhang, L.: Grounding DINO 1.5: Advance the "Edge" of Open-Set Object Detection (May 2024). <https://doi.org/10.48550/arXiv.2405.10300>
3. Robinson, I., Robicheaux, P., Popov, M., Ramanan, D., Peri, N.: RF-DETR: Neural Architecture Search for Real-Time Detection Transformers (Nov 2025)
4. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: DINOv3 (Aug 2025). <https://doi.org/10.48550/arXiv.2508.10104>
5. Suri, S., Rambhatla, S.S., Chellappa, R., Shrivastava, A.: SparseDet: Improving Sparsely Annotated Object Detection with Pseudo-positive Mining (Aug 2023). <https://doi.org/10.48550/arXiv.2201.04620>
6. Yang, C., Huang, L., Crowley, E.J.: Plug and Play Active Learning for Object Detection (Mar 2024). <https://doi.org/10.48550/arXiv.2211.11612>
7. Zhao, T., Liu, P., Lee, K.: OmDet: Large-scale vision-language multi-dataset pre-training with multimodal detection network. IET Computer Vision **18**(5), 626–639 (Aug 2024). <https://doi.org/10.1049/cvi2.12268>

Intelligent Parking Management System: An Infrastructure-Based and CARLA Simulation Approach

Borja Pérez^[0009–0008–1372–4274], Mario Resino^[0009–0001–0900–7052], Elena Serrano^[0009–0001–4404–1301], Abdulla Al-Kaff^[0000–0003–0212–6400], and Fernando García^[0000–0002–6291–5009]

University Carlos III of Madrid, Leganés, Madrid, Spain

Abstract. We propose a closed-loop, infrastructure-based framework for intelligent parking management, capable of simple and misuse occupancy detection and active vehicle assignment and verification. The system is designed and validated in the CARLA simulator. Our framework comprises an Assignment Module at the entrance, which classifies vehicle types and assigns a specific parking lot and a Verification Module using infrastructure cameras which classify the status of the lot in real-time using deep learning. A central logic handler verifies if the vehicles are correctly parked, notifying misuse or invasions events. This paper presents the complete system architecture and validation methodology, demonstrating how simulation enables the development and test of robust management loop that is impractical to train and test using static public datasets.

Keywords: Smart Parking Management · Computer Vision · Deep Learning · Domain Gap.

1 Introduction

Current Smart Parking Systems (SPS), as in [1], they predominantly focus on binary occupancy detection: free or occupied. This detection is often performed using IoT sensor networks on each slot [2] or via machine learning models that predict this binary availability [3]. While Computer Vision provides a scalable alternative to IoT sensors, both approaches fail to manage critical misuse scenarios: i.e. bay invasions or wrong assigned slot. To address this gap, we propose a holistic CV-based framework developed in the CARLA simulator. Our system uses synthetic data to train a perception model that explicitly identifies misuse, integrating it with a logic module for vehicle assignment and verification.

2 Methodology

Our methodology leverages the CARLA simulator as an integrated development and validation environment. First, we generate a custom synthetic dataset from the simulator, including explicit labels for multiple types of misuse events. This dataset is used to train our infrastructure-based Perception Module, a deep learning model that identifies slot state between: empty, occupied and invade.

This module is then integrated into our overarching Management Framework, which handles high-level logic, such as vehicle-type slot assignment and parking verification, forming a complete, closed-loop system.

3 Results & Conclusion

Our framework's performance was validated on its ability to manage complex misuse scenarios, which are a key limitation of existing systems. Figure 1 demonstrates the simple pipeline. The perception model, trained on our custom synthetic dataset, successfully determines the final state of a parking slot $S(B_i)$ by analyzing detected vehicle V_i in relation to its assigned parking slot B_i and the set of adjacent slots B_{adj} . A misuse state is identified when the overlap with an adjacent slot exceeds a threshold τ . Therefore we propose a new more descriptive Invasion Ratio I_R metric that calculates the proportion of vehicle's own area outside its assigned slot B_i in adjacent slots B_j .

$$I_R(V_i, B_i) = \frac{\sum_{j \in \{B_{adj}\}} |V_i \cap B_j|}{|V_i|}. \quad (1)$$

The final state $S(B_i)$ is then defined as misuse if $I_R(V_i, B_i) > \tau_{invade}$. The full paper will analyze the sensitivity of our system to this new metric and determine an optimal τ_{invade} through experimentation in CARLA. This multi-class detection is a direct result of our methodology. We conclude that this approach provides a solution for building and validating scalable infrastructure-based smart parking systems even in misuse events that traditional benchmarks and IoT-based systems fail to address.




Fig. 1: Our framework pipeline with the final detection results, successfully classifying empty (green), occupied (red) and misuse (yellow) states.

Acknowledgments. The work carried out in this paper is part of the Ecomobility project <https://www.ecomobility-project.eu/>. This project is supported by the CHIPS Joint Undertaking and its members, including top-up funding from the national authorities of Türkiye, Spain, the Netherlands, Latvia, Italy, Greece, Germany, Belgium, and Austria under grant agreement number 101112306. This work in Ecomobility is cofounded by the MCIN/AEI/10.13039/501100011033 under the grant number PCI2023-143425

References

1. Channamallu, S.S., et al.: Smart Parking Systems: A comprehensive review of Digitalization of Parking Services. *Green Energy and Intelligent Transportation* (2025)
2. Ditta, A., et al.: Number plate recognition smart parking management system using IoT. *Measurement: Sensors* **37**, 101409 (2025)
3. Dahiya, A., et al.: Machine Learning-Based Prediction of Parking Space Availability in IoT-Enabled Smart Parking Management Systems. *Journal of Advanced Transportation* **2024**, 8474973 (2024)

It's about time: measuring time error propagation between V2X agents

Jaime Godoy Calvo, Alberto Gil-Cuadrado, Abdulla Al-Kaff, and Fernando García

University Carlos III of Madrid, Leganés, Madrid, Spain

Abstract. Reliable vehicular V2X communication requires precise temporal synchronization among vehicles, infrastructure, and servers. This work introduces a low-cost, high-precision time server chain that provides a unified time reference across agents. Time at each node is modeled probabilistically, with uncertainty quantified as a function of the chain's quality. Experimental results show that conventional, unsynchronized measurements can produce physically inconsistent latencies, including negative values. In contrast, the proposed approach enables accurate estimation of communication delays and allows uncertainty to be propagated into multi-agent data fusion, improving the reliability of time-dependent measurements in high-mobility scenarios

Keywords: V2X · time measurement · error propagation.

1 Introduction

In the context of vehicular V2X communication, temporal synchronization among the different agents is essential to ensure effective information exchange [2] [1]. For detection data to be fully exploitable, vehicles, infrastructure, and servers must share a common time reference. While synchronization may be overlooked in everyday scenarios, it becomes critical in high-mobility environments, where timing errors can lead to significant inaccuracies in obstacle position and orientation estimation [3]. Therefore, establishing a unified time reference is necessary to accurately measure communication delays between system agents.

2 Methodology

This work proposes the design of a high-precision, low-cost time server chain to be deployed across the various agents in a vehicular environment. Communication latency between machines is analyzed within the proposed chain. Time is modeled as a random variable T with an associated probability distribution $P(T)$, where the confidence in the temporal information decreases according to the quality of the chain. Formally, the measured time at agent i can be expressed as:

$$T_i = T_{true} + \delta_i, \quad (1)$$

where T_{true} is the true reference time and $\delta_i \sim \mathcal{N}(0, \sigma^2)$ represents the timing uncertainty introduced by the chain. The effective communication latency between agents i and j is then defined as:

$$L_{i,j} = \mathbb{E} [T_j - T_i], \quad (2)$$

with its confidence bounded by the variance $\sigma_i^2 + \sigma_j^2$. This probabilistic framework allows for a quantitative assessment of temporal accuracy as a function of the chain's quality.

3 Experimental Results




Fig. 1: (a) Histogram of communication latency without applying probabilistic time; (b) Time evolution of one agent with its associated chained confidence.

Figure 1 illustrates the measured latency in standard communication without applying the probabilistic time approach. Although the resulting histogram may appear reasonable at first glance, it is of limited practical use, as the temporal measurements at each agent are taken without a synchronized time server chain. This can lead to physically inconsistent results, such as negative latencies, where the clock of agent 2 records the message before it is actually sent by agent 1.

In contrast, the right-hand plot demonstrates that by modeling time with uncertainty propagation, we can derive a distribution that allows for more accurate estimation of time-dependent quantities, such as communication latency. This approach enables the propagation of latency uncertainty into data fusion processes across agents, improving the reliability of time-sensitive measurements in high-mobility environments.

Acknowledgments. Work is supported by PID2021-128327OA-I00 (MCIN/AEI/10.13039/501100011033), ERDF, and EU NextGenerationEU/PRTR. Disclosure of Interests. No competing interests.

References

1. Faizullin, M., Kornilova, A., Ferrer, G.: Open-source lidar time synchronization system by mimicking gnss-clock (2022). <https://arxiv.org/abs/2107.02625>
2. Hasan, K.F., Wang, C., Feng, Y., Tian, Y.C.: Time synchronization in vehicular ad-hoc networks: A survey on theory and practice. *Vehicular Communications* **14**, 39–51 (Oct 2018). <https://doi.org/10.1016/j.vehcom.2018.09.001>, <http://dx.doi.org/10.1016/j.vehcom.2018.09.001>
3. Zhao, Y., Zhou, L., Feng, W., Xu, S.: Fast time synchronization on tens of picoseconds level using uncombined gnss carrier phase of zero/short baseline. *Sensors* **20**(17), 4882 (2020). <https://doi.org/10.3390/s20174882>

Comparative Evaluation of Real-Time Communication Protocols in a Real Vehicle Teleoperation Platform

Javier Gómez^[0009-0001-7434-7459], Abdulla Al-Kaff^[0000-0003-0212-6400], and Fernando García^[0000-0002-6291-5009]

Universidad Carlos III de Madrid, Madrid, Spain

Abstract. This work presents a comparative evaluation of real-time communication protocols for the teleoperation of autonomous vehicles. Using a real vehicle and a fully operational teleoperation station, performance is tested under identical conditions. Metrics such as latency, jitter, and packet loss are measured to support the selection of safe and efficient teleoperation systems for real-world use.

Keywords: Teleoperation · Real-time communication · Video streaming

1 Introduction

Autonomous vehicles rely on sensors and AI-based perception to navigate, but their performance can encounter difficulties in extreme or unpredictable situations. Teleoperation can serve as an essential fallback, enabling human operators to remotely control vehicles through real-time video and command transmission.

Previous studies have examined communication technologies for teleoperation [1, 2], but most study only a small number of protocols or rely on simulation-based environments. This work advances the state of the art by comparing a large range of transport and streaming technologies under identical conditions using a real autonomous vehicle testbed, helping to better understand which protocols are most suitable for safe teleoperation.

2 Methodology

2.1 System Overview

As shown in Fig. 1, the experimental setup consists of a real vehicle equipped with drive-by-wire, connected over a 5G network to a teleoperation station with driving controls and a monitoring interface. Video is sent from the vehicle to the station, while control commands are transmitted back to close the loop.

2.2 Evaluated Protocols

Transport Baselines: UDP provides low-latency transmission and may experience packet loss, while TCP delivers data reliably at the cost of higher latency.

Streaming Protocols: RTP provides real-time streaming with jitter buffering, while SRT adds packet recovery to maintain low latency.




Fig. 1: System architecture with the vehicle, network, and teleoperation station.

Web-Based Protocols: WebSockets offer wide compatibility at the expense of latency, while WebRTC and WebTransport provide low-latency bidirectional communication.

3 Results and Conclusion

The tests were performed over a 5G network following a predefined route, covering three scenarios: video only, commands only, and simultaneous transmission. Latency, jitter, and packet loss were measured to compare protocol performance under real conditions.




Fig. 2: Heatmap for WebRTC (left) showing areas of higher latency, alongside the teleoperation interface (right) displaying real-time video and commands.

The results highlight general differences in latency, reliability, and responsiveness across protocols. As an example, Fig. 2 shows a heatmap for WebRTC illustrating areas with higher latency along the test route. Full quantitative results for all protocols will be presented in the final paper.

Acknowledgments. Work is supported by PID2021-128327OA-I00 (MCIN/AEI/10.13039/501100011033), ERDF, and EU NextGenerationEU/PRTR. Disclosure of Interests. No competing interests.

References

1. Karla, T., e.a.: Soft real-time communication with websocket and webrtc for web-based control loops. In: Proc. 24th Int. Conf. Methods Models Autom. Robot. (MMAR). pp. 634–639 (2019), <https://doi.org/10.1109/MMAR.2019.8864680>
2. Sahraoui, Y., e.a.: Performance evaluation of tcp and udp video streaming in vanets. In: Proc. 2018 Int. Conf. Smart Commun. Network Technol. (SaCoNeT). pp. 67–72 (2018), <https://doi.org/10.1109/SaCoNeT.2018.8585447>

Field Test Methodology for V2X Cooperative Perception between Heterogeneous Vehicles

Mario Resino¹, Jaime Godoy¹, Abdulla Al-Kaff¹, Fernando García¹, Jesús Espelosín², Lidia Sanchez², and Teresa Seco²

¹ University Carlos III of Madrid, Leganés, Madrid, Spain

² Aragon Institute of Technology (ITA), Zaragoza, Spain

Abstract. Validating the interoperability of Cooperative Perception (CP) V2X systems, key for overcoming autonomous vehicle sensor occlusions, remains a key challenge for heterogeneous platforms. This paper presents the field test methodology of a V2X CP system, validated over a 5G network between two different vehicle platforms in an urban environment. We detail how an ego vehicle transmits processed data to a second vehicle, providing it with occluded object information. The core of this work is the definition and execution of critical urban scenarios, demonstrating the methodology's viability in validating multi-vendor interoperability.

Keywords: Cooperative Perception · V2X · Autonomous Vehicles · 5G

1 Introduction

Intelligent Transportation Systems (ITS) are foundational to the Smart City vision, aiming to improve urban mobility and safety [1]. Autonomous Vehicles (AVs) require high-fidelity situational awareness, a challenge due to urban occlusions. V2X-enabled CP addresses this by sharing sensor data, extending perception range. However, while CP algorithm development has matured in simulation [2], validation relies heavily on simulation. As highlighted in [3], real-world deployment challenges remain largely unaddressed. This work presents the field test methodology from a V2X trial in Zaragoza. Key contributions are: (1) end-to-end validation of a heterogeneous V2X CP system over a 5G network; (2) definition of a field-test methodology of critical urban scenarios and a demonstration of the viability of sharing complex perception data.

2 Methodology




Fig. 1: Experimental framework for validating multi-vendor interoperability.

Our methodology has two key components (Fig 1): the end-to-end system architecture and the field test validation framework. The architecture's pipeline

involves an ego vehicle (LiDAR/GNSS) detecting and tracking objects in global coordinates. This JSON data is sent via a WireGuard 5G tunnel to a heterogeneous partner for integration. The validation framework, executed in ITA - Instituto Tecnológico de Aragón, Zaragoza, focused on critical scenarios (e.g. pedestrian occlusions, blind overtaking) where CP provides unique safety benefits. Systematic logging and post-analysis verified data integration, confirming functional interoperability.

3 Test Field

To validate the proposed methodology and demonstrate the benefits of V2X-enabled CP between heterogeneous platforms, a series of critical urban scenarios were executed at a dedicated test facility, as you can see in Fig 2. These scenarios were specifically designed to challenge individual vehicle perception, highlighting situations where shared data significantly enhances situational awareness, such as pedestrians crossing from behind occlusions or overtaking maneuvers at a crosswalk that occlude a crossing pedestrian. The successful execution and data exchange across these diverse scenarios confirmed the functional interoperability of the system, allowing the receiver vehicle to maintain full situational awareness.



(a) Pedestrian occlusion by vegetation. (b) Overtaking at occluded crosswalk.

Fig. 2: Field test scenarios validating interoperability in Zaragoza.

Acknowledgments. The work carried out in this paper is part of the Ecomobility project <https://www.ecomobility-project.eu/>. This project is supported by the CHIPS Joint Undertaking and its members, including top-up funding from the national authorities of Türkiye, Spain, the Netherlands, Latvia, Italy, Greece, Germany, Belgium, and Austria under grant agreement number 101112306. This work in Ecomobility is cofounded by the MCIN/AEI/10.13039/501100011033 under the grant number PCI2023-143425.

References

1. Elassy, M., Al-Hattab, M., Takruri, M., Badawi, S.: Intelligent Transportation Systems for Sustainable Smart Cities. *Transportation Engineering* **8**, 100252 (2024).
2. Zhang, Y., Fan, Y., Li, X., Yang, X., Li, H., Liu, W., Guo, J.: CooPercept: Cooperative Perception for 3D Object Detection of Autonomous Vehicles. *Drones* **8**(6), 228 (2024). <https://doi.org/10.3390/drones8060228>
3. Guo, A., Yuan, J., Wang, C., Sun, W., Wang, L.: When Autonomous Vehicle Meets V2X Cooperative Perception: How Far Are We?. In: ASE 2024. <https://arxiv.org/abs/2509.24927>

A Learning-Based Metaheuristic for the Vehicle Routing of Multi-Compartment Silo Trucks

Pit Schulenburg, Ulrich Göhner,, and Sebnem Gülficici

University of Applied Sciences Kempten, Bahnhofstraße 61, 87435 Kempten Germany
 pit.schulenburg@stud.hs-Kempten.de, ulrich.goehner@hs-kempten.de,
 sebnem.guel-ficici@hs-kempten.de

Abstract. This work presents a learning-based extension of the classical Ant System (AS) for a realistic Vehicle Routing Problem (VRP) in the context of bulk feed logistics. The considered problem combines multiple real-world constraints such as heterogeneous fleets, multi-compartment trucks, trailers, and daily time windows. The aim is to enhance route quality and computation efficiency through intelligent guidance based on a Transformer model. The proposed hybrid, called Neural Ant System (NAS), integrates learned decision-making into the pheromone-driven construction process and achieves consistent performance gains over the standard AS.

Keywords: Vehicle Routing Problem · Transformer · Ant Colony Optimization · Neural Combinatorial Optimization · Reinforcement-learning · Metaheuristics

1 Introduction

The Vehicle Routing Problem (VRP) represents one of the most fundamental combinatorial optimization problems in logistics and has been studied since Dantzig and Ramser's seminal work [2]. Modern logistics systems, however, incorporate numerous additional constraints such as heterogeneous fleets, time windows, and multi-compartment vehicles, making classical deterministic methods insufficient for daily dispatch planning. In this study, a rich VRP was examined, based on real silo truck transport scenarios in feed logistics. The objective was to optimize the assignment and sequencing of deliveries across multiple compartments and trips while maintaining operational feasibility.

Metaheuristics such as Ant Colony Optimization (ACO) [3] have proven effective for large-scale routing problems due to their robustness and adaptability. Yet, their performance still depends on static heuristics and manually tuned exploration-exploitation balances. In contrast, current developments in Neural Combinatorial Optimization (NCO) [1] have shown that machine learning can enhance such heuristics through adaptive decision-making. Research by Nazari et al. [8], Kool et al. [6], and Joshi et al. [5] demonstrated that Transformer-based architectures can learn routing policies directly from instance data. Building on

these advances, the present work proposes a hybrid metaheuristic that combines the interpretability and convergence stability of the Ant System with the learning capability of a Transformer model.

2 Learning-Based Metaheuristic System

We extend the classical Ant System by adding intelligent ants that are implemented with a Transformer encoder–decoder. These intelligent ants act as full colony members: they construct tours and update pheromones analogously to normal ants. The global pheromone structure and evaporation remain unchanged. All ants, classical and intelligent, write to the same matrix, which preserves the convergence behavior and interpretability of the algorithm.

The idea of an intelligent ant follows the work of Liu et al. from 2023 [7]. The Transformer architecture itself is adapted from Hu & Wang [4], an encoder–decoder that produces probabilities over Vehicles and Nodes. Masking enforces feasibility. To balance exploration and exploitation, we deploy multiple intelligent ants per iteration with complementary policies: one greedy ant selects the best action, while a second stochastic ant samples from the top- k candidates. Their completed tours deposit pheromone using the standard AS rule.

Problem-wise, we target a rich VRP from feed logistics: heterogeneous trucks with multi-compartment loading, trailers that remain coupled for the day, up to three tours per vehicle, and delivery-day time windows. Test scenarios comprise 70–80 customers per instance. On one hundred such instances, the learning-augmented colony achieves consistent distance reductions of 1.1 – 2.4% over the classical AS at comparable runtimes. A purely neural variant yields shorter inference times but lower route quality. The results indicate that integrating Transformer-based intelligent ants into an ACO framework combines the stability of swarm search with adaptive, data-driven guidance.

References

1. Bengio, Y., Lodi, A., Prouvost, A.: Machine learning for combinatorial optimization: A methodological tour d’horizon. *European Journal of Operational Research* 290(2), 405–421 (2021)
2. Dantzig, G.B., Ramser, J.H.: The truck dispatching problem. *Management Science*, 6(1), 80–91 (1959)
3. Dorigo, M., Maniezzo, V., Colomi, A.: Ant System: Optimization by a colony of cooperating agents. *IEEE Trans. Syst. Man Cybern. B* 26(1), 29–41 (1996)
4. Hu, J., Wang, C.: A Deep Reinforcement-Learning-Based Route Optimization Model for Multi-Compartment Cold Chain Distribution. *Mathematics*, 13(13), 2039 (2025)
5. Joshi, C.K., Cappart, Q., Rousseau, L.-M., Laurent, T., Bresson, X.: Learning TSP requires rethinking generalization. *arXiv:2006.07054* (2020)
6. Kool, W., van Hoof, H., Welling, M.: Attention, learn to solve routing problems! *ICLR* (2019)
7. Liu, Y., Qiu, J., Hart, E., Yu, Y., Gan, Z., Li, W.: Learning-Based Neural Ant Colony Optimization. *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 47–55. ACM, Lisbon (2023)
8. Nazari, M., Oroojlooy, A., Snyder, L.V., Takáć, M.: Reinforcement learning for solving the vehicle routing problem. *NeurIPS Workshop on Deep Reinforcement Learning* (2018)

Complex Phasor Representations on the Gaussian Plane to Annotate Arrhythmia Episodes in the ECG Signal

Jan Nikodem¹[0000-0002-3748-9255], Ryszard Klempous¹[0000-0003-4889-0930],
 Jerzy Rozenblit²[0000-0002-7348-4128], Konrad Kluwak¹[0000-0003-0983-2228],
 Marek Kulbacki^{3,4}[0000-0003-4609-106X], Dorota Zyśko¹[0000-0001-9190-0052]
 , and Dariusz Jagielski¹[0000-0002-3066-4030]

¹ Wroclaw University of Science and Technology, Wroclaw, Poland

² Depart. of Electrical and Computer Engineering University of Arizona Tucson, USA

³ Faculty of Computer Science, Polish-Japanese Academy of Information Technology,
 Warsaw, Poland

⁴ DIVE IN AI, Wroclaw, Poland

Extended Abstract

The ECG signal is an electrical recording of heart activity, i.e., the progressive depolarization of the myocardium that moves in a specific direction. Electrical stimulation of the heart muscle fibers generates a time-varying electric field represented by both the electrical potential and its direction. Therefore, to fully describe the electrical activity of the heart, it is necessary to know not only the electrical potential (scalar) but also the potential and direction (heart electrical vector \vec{E}) [5].

The heart's electrical axis is a “dominant” vector \vec{E} that describes the main direction of the electrical impulse spreading throughout the heart. It is important to detect early warning signs of changes in this vector \vec{E} , as they reflect the patient's cardiovascular condition [2].

In order to determine the vector \vec{E} , electrodes are placed on the patient's body to define the projection of \vec{E} onto a selected plane (e.g., frontal, horizontal, or transverse plane) determined by the electrode positions [7]. With a 12-lead electrocardiogram, selected projections of \vec{E} are presented as an illustrative logical sequence (hexaxial reference system) or through the limb leads anatomically (Cabrera system), which enables two-dimensional tracing of the heart's electrical activity and improves ECG interpretation [3].

This methodology has been refined over the years by increasing the dimensionality of the recorded ECG signal and thus enhancing the diagnostic information, forms the basis of our proposal. In standard ECG signal analysis, rhythm disturbances are identified as diagnostic information manifested through episodes of frequency change. Basic ECG signal analysis methods include time and frequency analysis, typically performed separately [4]. The separate analysis of time and frequency components in certain situations prevents comprehensive diagnosis due to the variability of their frequency structure over time. Therefore,

we propose extending ECG signal analysis into the complex number domain, where time-frequency characteristics are analyzed simultaneously.

We employ the Discrete Complex Fourier Transform (DCFT), which extends the ECG signal to the complex domain, where the standard ECG signal represents a projection onto the real axis $\overrightarrow{\text{Re}}$ (i.e., $\text{ECG}(t) = \text{Re}(\text{ECG}_{\text{complex}}(t))$) [6].

In this paper, we present the method for obtaining a complex phasor representation [8] of ECG signal segments between consecutive QRS complexes using the DCFT transformation. The dataset used in this study was obtained from the MIT/BIH arrhythmia database. For different heartbeat types, we calculate unique phasors and perform, for each arrhythmia type, unique geometric interpretations of complex exponentials on the Gaussian plane.

The principles for recognizing individual heart rhythm disturbances in ECG signals are already described in the literature [1]; however, these principles focus on human-based annotation. Our work focuses on transforming these principles to align with the capabilities and specifics of computer-based operation. This applies both to the classifiers and decision-making schemes employed, as well as to the disease process patterns used. For the classifiers and decision-making schemes developed based on them, we seek dichotomous solutions (YES/NO) and acyclic decision trees. In this work, we aim to explore aspects that others may overlook and to discuss understanding sufficiently to identify new genuine opportunities in computer-aided ECG signal annotation.

References

1. Gacek A, Pedrycz W,; ECG Signal Processing, Classification and Interpretation, A Comprehensive Framework of Computational Intelligence, Springer-Verlag London Limited 2012, ISBN978-0-85729-867-6, DOI: <https://doi.org/10.1007/978-0-85729-868-3>
2. Hoseini SS, Moeeny A, Shoar S, Shoar S, Naderan M, Gharibzadeh S, Dehpour AR.; Designing Nomogram for Determining the Heart's QRS Axis. Journal of Clinical and Basic Cardiology. 2012;14:12-15
3. Jaros, R., Martinek, R., Danys, L. (2019). Comparison of Different Electrocardiography with Vectorcardiography Transformations. Sensors, 19(14), 3072.
4. Kohler B-U, Hennig C, Orglmeister R,; The principles of software QRS detection, in IEEE Engineering in Medicine and Biology Magazine, vol. 21, no. 1, pp. 42-57, Jan.-Feb. 2002, doi: 10.1109/51.993193.
5. Kuhn L, Rose L,; ECG interpretation part 1: Understanding mean electrical axis. J Emerg Nurs. 2008 Dec;34(6):530-4. doi: 10.1016/j.jen.2008.01.007.
6. Kumar M A, Chakrapani A,; Classification of ECG signal using FFT based improved Alexnet classifier. PLoS One. 2022 Sep 27;17(9):e0274225. doi: 10.1371/journal.pone.0274225. PMID: 36166430; PMCID: PMC9514660.
7. Stafford M, Stafford D, (1993). Determination of the electrical axis of the heart. The Journal of the American Osteopathic Association. 93. 103-5. 10.7556/jaoa.1993.93.1.103.
8. Steinmetz C.P, Complex Quantities and Their Use in Electrical Engineering. Proceedings of the International Electrical Congress Held in the City of Chicago, August 21st to 25th, 1893. American Institute of Electrical Engineers. pp. 33–74

Early Detection of Alzheimer's Disease Using Deep Learning Models on MRI Data: A Study Based on ADNI

Brian Suárez Santiago^{*}, Ayoze González Hernández⁺, Gabriele Salvatore de Blasio^{*},
C.R. García-Rodríguez^{*}, A. Quesada-Arencibia^{*}

^{*}University of Las Palmas de Gran Canaria, Spain

⁺Hospital Universitario de Gran Canaria Dr. Negrín, Spain

briansuarezsantiago@gmail.com, ayozegonzalez@hotmail.com,
gabriel.deblasio@ulpgc.es, ruben.garcia@ulpgc.es,
alexis.quesada@ulpgc.es

Abstract. Alzheimer's disease (AD) is the most prevalent neurodegenerative condition worldwide, and its early detection remains one of the greatest challenges in modern medicine. This study explores the application of advanced artificial intelligence techniques—particularly convolutional neural networks (CNNs)—to the classification and staging of Alzheimer's disease from magnetic resonance imaging (MRI) data. The project leverages data from the Alzheimer's Disease Neuroimaging Initiative (ADNI) to train and validate multiple deep learning architectures, including VGG16, DenseNet201, and ResNet50, with the aim of identifying structural brain patterns associated with different stages of the disease, from mild cognitive impairment (MCI) to advanced AD.

A series of preprocessing and model tuning strategies were employed to enhance performance, including data augmentation, normalization, and evaluation through confusion matrices and accuracy metrics. The best performing model achieved an accuracy of 99.68% in binary classification (AD vs. healthy controls), and promising results were also obtained in multi-class classification tasks aimed at predicting disease stages.

This work contributes to the growing field of AI-based medical diagnosis, highlighting the potential of deep learning to support clinicians in early intervention strategies. Moreover, it addresses the ethical and legal considerations of handling medical data, aligning the implementation with GDPR and fair AI principles. By enhancing diagnostic accuracy at early stages, this approach may help improve treatment planning and patient outcomes in neurodegenerative diseases.

Keywords. Aging, Alzheimer's, Artificial Intelligence, Machine Learning, Predictive Modeling, Medical Imaging, Convolutional Neural Network, CNN

Framework development for remote health monitoring during manned space missions

Miklos Kozlovszky^{1,3}, Orsolya Szabó Nagyné², András Koleszár², Peter Kristof¹, Edit Csanák², Levente Kovács³, Virág Némethy²

¹ Space Lab, ² Product Design Institute, Sándor Rejtő Faculty of Light Industry and Environmental Protection Engineering, ³ John von Neumann Faculty of Informatics Obuda University Budapest, Hungary

{kozlovszky.miklos,szabo.orsolya,koleszar.andras,kristof.peter,csanak.edit,kovacs.le-
vente,nemethy.virag}@uni-obuda.hu

Space is a hostile environment. Manned space missions introduce substantial physiological and psychological risks. Current astronaut health-monitoring strategies rely primarily on episodic assessments and bulky measurement devices, which are inadequate for continuous monitoring in microgravity. This paper proposes a systems-level framework for remote, continuous, multimodal health monitoring tailored to the constraints and mission requirements of crewed spaceflight. The framework integrates wearable smart garments, embedded signal processing, secure communication pathways, and ground-based analytics. Preliminary results include an initial architecture design, hardware selection, and definition of modular subsystems that together support scalable, mission-ready physiological monitoring.

Keywords: astronaut health monitoring, wearable sensing, smart garments, microgravity, biomedical signal processing, embedded sensing

1 Introduction

Human space exploration is transitioning toward longer-duration missions, including prolonged space station occupation, sustained Lunar and Mars operations. These missions impose cumulative physiological stressors that require continuous health monitoring to support early anomaly detection and informed medical decision-making. Traditional biomedical monitoring methods—periodic measurements, rigid sensors, and experiment-specific instrumentation—do not provide sufficient temporal resolution or operational practicality for missions beyond low Earth orbit. To address this gap, we propose a framework for remote, continuous physiological monitoring through wearable smart garments integrated into the daily routines of astronauts. The framework emphasizes ergonomics, multimodality, embedded intelligence. Existing astronaut health-monitoring systems include biomedical belts, adhesive patches, and experiment-specific devices, providing valuable but limited physiological insights. Current limitations include: (i) episodic rather than continuous data collection; (ii) restricted sensing modalities; (iii) mechanical discomfort or operational interference; and (iv) limited onboard preprocessing, which leads to data inefficiencies.

In contrast, terrestrial research has demonstrated rapid advancements in smart textiles, textile-based electrodes, flexible circuits, and low-power biosignal acquisition systems. [1-3] However, adaptations for spaceflight remain largely unexplored, especially concerning microgravity ergonomics, reliability, and safety requirements.

A monitoring system suitable for space missions must support:

- Multimodal physiological sensing (ECG, respiration, temperature, EDA).
- Long-duration wearability, including comfort, durability, and seamless integration with mission garments.
- Embedded preprocessing, including motion-artifact reduction, feature extraction, and adaptive compression.
- System modularity, enabling seamless incremental system upgrades.
- Integration with onboard computing and communication infrastructure.
- Medical relevance, enabling trend analysis, early diagnosis, decision support.

2 Results

The initial phase of the project focused on developing a modular architecture integrating wearable sensing, embedded computation, communication interfaces, and remote analytics. The proposed architecture comprises four layers:

- Sensing Layer: Textile-based ECG electrodes, respiratory stretch sensors, IMUs, temperature sensors, and electrodermal activity modules. Sensors are embedded into a soft, ergonomic garment.
- Embedded Processing Layer: A low-power microcontroller and analog front-end perform real-time filtering, motion-artifact removal, and feature extraction. Local buffering ensures data integrity during communication disruptions.
- Communication and Integration Layer: Short-range wireless transmission links the garment with the spacecraft's onboard computing node. The architecture supports prioritization of medically significant events and delayed ground transmission.
- Remote Analytics Layer: Time-series modeling, trend analysis, and predictive algorithms support medical interpretation, anomaly detection, and decision support for mission operators.

The layered structure ensures maintainability, scalability, and subsystem independence.

Acknowledgments. The paper received financial support from the project "Supporting the development of the OE SpaceLab research centre" in the frame of the HUNOR Hungarian Space Programme by the Hungarian Ministry of Foreign Affairs and Trade.

References

1. Norsk, P., et al.: Human Adaptation to Spaceflight: The Role of Physiological Monitoring. *Acta Astronautica* (2020).
2. Piqueras, M., et al.: Wearable Biomedical Systems for Remote Monitoring in Extreme Environments. *IEEE Trans. Biomed. Eng.* (2021).
3. NASA Human Research Program: Human Health and Performance Risks for Space Exploration Missions. NASA HRP Publication (2022).

Rapid Prototyping Workflow for Robotic Gripper Development through CAD-Python Integration

Guillem Diaz Peralta^{1,2[0009-0003-6601-0201]}, Bernhard Höfig^{2[0009-0007-6646-9089]}, and Markus Glück^{2[0000-0002-7631-5389]}

¹ Universitat Politècnica de Catalunya, Colom str. 1, 08221 Terrassa, Spain
guillem.diaz.peralta@estudiantat.upc.edu

² Aalen University of Applied Sciences, Beethovenstr. 1, 73430 Aalen, Germany

Abstract. This paper presents an agile development process for robotic end-effectors that combines rapid prototyping with a lightweight CAD & Python loop that accelerates iteration from a Minimum Viable Product to industrial readiness, turning the final robot integration into a small step. A case study of an internal-expansion gripper illustrates short design-build-test cycles with additive manufacturing and incremental validation.

Keywords: Rapid Prototyping · Mechatronic Design · Product Design Methods · Robotic Handling · Agile Product Development.

1 Motivation

For a successful robotic application, very often the end-effector is the crucial component. When an appropriate end-effector is missing, training and industrial activities are constrained and delayed. Therefore, a design process is required to deliver functional results quickly while maintaining quality and controlling risks. In the Industry 4.0 context [1] Digital Twins can unify virtual design [2], physical experiments and robot adaptation into a repeatable loop with clear decision points. Agile product development has shown many advantages for innovative mechatronic design [3].

2 Agile development approach

The proposed approach uses short define-simulate-build-test-adjust cycles that allow to use a Minimum Viable Product to validate essential key points before applying further optimization. The core mechanism (CAD + Python) manages parameter changes, motion and clearance checks and adapts the Arduino code needed for the mechanism actuation. In this process, the final adaptation to the target robot becomes a re-parameterization and not a full redesign. Agile innovation is mainly driven by short iteration cycles which allow for a rapid prototyping, testing and validation.

3 Use case

To enable demonstrations and testing for a 50kg payload robot, a gripper was developed through the agile product design strategy. That moved the first design of an MVP to industrial readiness with minimal rework. Firstly, a hand-crank MVP quickly validated the mechanism principle, the auto-centering features and the motion transmission parts, allowing to create an actuation baseline for the next step. Following, a motorized iteration was 3d printed using off the shelf components and allowing for testing on smaller robots. This first autonomous test revealed tilting issues in the workpiece, which was afterwards solved. This validation on a smaller scale robot allowed material savings, reduced manufacturing times and overall enabled the quick iteration process required by the technique. Finally, a full-scale gripper was manufactured, which supposed 57h of 3d printing, and tested in combination with a KUKA KR 50 R2100 (see Fig.1).




Fig. 1. Iterations made following the dual approach workflow for the design of a robotic gripper.

4 Conclusions

The presented technique demonstrated that the integration of CAD-driven parameterization with Python scripts enabled to create a lightweight Digital Twin framework suitable for end-effector development. By adopting such iterative cycles, functional validations can be achieved quickly during the process. Using small-scale prototypes proved effective for this before scaling up to industrial applications.

References

1. Hermann, M., Pentek, T., & Otto, B. (2016). Design Principles for Industrie 4.0 Scenarios. 2016 49th Hawaii International Conference on System Sciences (HICSS), 3928-3937. <https://doi.org/10.1109/HICSS.2016.488>
2. Tao, F., Zhang, H., Liu, A., & Nee, A. Y. C. (2019). Digital Twin in Industry: State-of-the-Art. IEEE Transactions on Industrial Informatics, 15(4), 2405-2415. <https://doi.org/10.1109/TII.2018.2873186>
3. Glück, M (2025). Agile Innovation: Mit neuem Schwung zum Erfolg, 2. Aufl. Wiesbaden; Springer Vieweg. <https://doi.org/10.1007/978-3-658-46584-1>

Automated Tool Wear Detection for Industry 4.0: Image Segmentation of Photometric Stereo Data

Thomas Jäkel^[0009-0008-6523-0022] and Frank Schirmeier^[0009-0002-2842-0376]

Institute for Data-Optimized Manufacturing (IDF), University of Applied Sciences Kempten,
Bahnhofstraße 61, 87435 Kempten, Germany
thomas.jaekel@hs-kempten.de

Keywords: Tool Wear, Image Segmentation, Photometric Stereo

1 Introduction

Tool wear detection represents a critical challenge in Industry 4.0 subtractive manufacturing systems. To avoid defects, industry uses conservative tool life values with large safety margins, causing premature tool replacement and reduced efficiency, particularly in flexible production environments with high product variety. Automated detection methods address this issue through indirect approaches which evaluate process parameters affected by tool wear and direct methods that analyse the tool itself.

Recent advances in photometric stereo (PS), a technique bridging 2D imaging and full 3D acquisition, particularly deep learning-based models like SDM-UniPS [1], enable the capture of detailed 2.5D surface information from cutting tools using photographs under varying illumination. This work investigates whether the photometric stereo data provides advantages over conventional two-dimensional imaging for automated wear detection using the nnU-Net framework [2] for semantic segmentation, contributing to intelligent production systems through automated quality control.

2 Methodology

The experimental setup [3] utilized an industrial lens from Navitar mounted on a measuring microscope, a Raspberry Pi HQ camera (12.3 MP) and a custom ring light with individually addressable SK6812 LEDs. This system enabled varying lighting conditions required for photometric stereo analysis. The tool is subsequently photographed with nine different lighting directions employed. For each of the lighting directions, three exposure levels were captured, which were merged into HDR images. Twenty end mills were photographed, showing various degrees of wear ranging from minimal to severe.

The captured images were processed using SDM-UniPS. This model reconstructs surface normals without assumptions about surface, geometry, or lighting properties and outputs Normal, Base Colour, Roughness and Metallic Maps.

For automated wear segmentation the nnU-Net framework was selected due to its ability to automatically configure training pipelines based on dataset characteristics.

The cutting edges were manually labeled with wear classified into cutting-edge wear and flank wear. Additionally, the cutting edge itself was highlighted with a three-pixel-wide line.

The training loss decreased rapidly in early epochs and stabilized, while the pseudo Dice coefficient on validation data reached 0.90, indicating good generalization without overfitting. For comparison, an identical nnU-Net architecture was trained using the raw input images. The same augmentation and training parameters were used to ensure direct comparability. This training achieved a validation Dice coefficient of 0.95, slightly outperforming the 2.5D approach.

3 Results

Both segmentation models were evaluated on unseen end mills. The raw data model demonstrated higher sensitivity, successfully identifying even small defects (0.05×0.02 mm). However, this came with a tendency toward over-segmentation. In contrast, the PS-based model produced more conservative segmentations, yielding fewer false positives but occasionally missing worn regions. On uniformly worn areas, the raw data model provided more accurate wear area estimates and more complete edge recognition. Overall, while PS data contributed useful geometric context, conventional imaging yielded slightly superior segmentation accuracy with the chosen nnU-Net setup.

4 Conclusion

The finding that raw 2D data outperformed processed 2.5D data seems counterintuitive at first and requires careful interpretation. The raw data approach retained all input channels, whereas photometric stereo processing reduced the number of channels. While this reduction creates a semantically more meaningful representation, information loss is inevitable. The superior performance of the raw data model suggests that the additional information across multiple lighting directions provided valuable discriminative features that were not fully preserved in the compact PS outputs.

However, this comparison has important limitations. The nnU-Net architecture treated both input types as conventional images without explicitly exploiting the geometric meaning encoded in Normal Maps. Future approaches that explicitly model this geometric relationship might unlock the potential of PS data more effectively.

References

- [1] S. Ikehata, „Scalable, Detailed and Mask-Free Universal Photometric Stereo,“ 2023.
- [2] F. Isensee, J. Petersen, A. Klein, D. Zimmerer, P. F. Jaeger, S. Kohl, J. Wasserthal, G. Köhler, T. Norajitra, S. J. Wirkert und K. H. Maier-Hein, „nnU-Net: Self-adapting Framework for U-Net-Based Medical Image Segmentation,“ *CoRR*, Bd. abs/1809.10486, 2018.
- [3] T. Jäkel und F. Schirmeier, „Deep Photometric Stereo for Tool Wear Inspection,“ in *Proceedings of AI4IP 2025: Database and Expert Systems Applications Workshops (DEXA 2025)*, Cham, 2025.

SOMP - Stiffness Optimized Motion Planner

David Seyser^{[0009-0004-3793-3469]1}, Micheal Schebek^{[0009-0009-7996-3128]2}, and
Lucas Pammer^{[0009-0005-9881-2908]3}

University of Applied Sciences Technikum Wien, Höchstädtplatz 6, 1200 Vienna,
Austria

{david.seyser, michael.schebek, lucas.pammer}@technikum-wien.at
<https://www.technikum-wien.at/>

Abstract. Keywords: Stiffness Optimization · Robot-based Manufacturing · Motion Planner.

1 Introduction

Within the *Processing* category of robotic applications, milling has emerged as part of a rapidly growing subgroup, showing a steady increase in the number of newly installed robotic systems worldwide each year [1]. Compared to CNC machines, industrial robots offer greater flexibility, lower acquisition cost, and broader range of motion [2]. However, their low rigidity and susceptibility to vibration result in suboptimal part quality, impeding adoption in precision manufacturing [3]. Addressing these limitations through motion planning remains challenging, as no generalized solution exists and methods are often application-specific [4, 5]. Sampling-based planners struggle with external forces and dynamics, while compliance is typically handled as post-processing compensation rather than integrated into planning [6]. Therefore, we propose a motion planning framework for robotic milling that maintains maximum structural stiffness throughout machining, aiming to enhance precision and surface quality.

2 Methods

We sample inverse kinematics solutions at each waypoint and select optimal robot configurations, via multi-objective optimization considering manipulability, joint limits and stiffness. Given a target pose $\mathbf{T}_{\text{target}}$, we evaluate multiple valid joint configurations \mathbf{q} that achieve the desired end-effector position. The stiffness objective is based on the Cartesian stiffness matrix:

$$f_{\text{stiffness}}(\mathbf{q}) = g(\mathbf{K}_c(\mathbf{q})), \quad \mathbf{K}_c = (\mathbf{J}\mathbf{K}_j^{-1}\mathbf{J}^T)^{-1} \quad (1)$$

where \mathbf{J} is the geometric Jacobian, \mathbf{K}_j is the joint stiffness matrix, and $g(\cdot)$ is a scalar metric that quantifies the robot's resistance to deflection under external loads. Different formulations of $g(\cdot)$ capture various aspects of stiffness behavior, such as worst-case directional compliance or overall structural rigidity. The motion planner is applied to CAM-generated waypoints.

3 Experiments

We evaluate our stiffness-optimized motion planner against industry-standard motion planner in 2D milling operations. Testing includes pocket milling and hole drilling on aluminum to assess continuous and point-to-point trajectories. We compare shape tolerance, surface quality, vibration, and acoustic emissions between both approaches. Shape tolerance is measured using a micrometer, surface quality with a surface roughness gauge, vibration via accelerometers mounted on the milling table, and acoustic emissions with a condenser microphone positioned 50 cm from the vise. We hypothesize that maximizing robot stiffness during cutting reduces vibrations and noise, improving surface finish and extending tool life.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article

References

1. Müller, Christopher: World Robotics 2022 – Industrial Robots, IFR Statistical Department, VDMA Services GmbH, Frankfurt am Main, Germany, 2022.
2. Wenbo Wang, Qiang Guo, Zhibo Yang, Yan Jiang, Jinting Xu: A state-of-the-art review on robotic milling of complex parts with high efficiency and precision. In: *Robotics and Computer-Integrated Manufacturing*, Volume 79, 2023, 102436, ISSN 0736-5845, <https://doi.org/10.1016/j.rcim.2022.102436>.
3. Zhao, H., Li, X., Ge, K., Ding, H.: A contour error definition, estimation approach and control structure for six-dimensional robotic machining tasks. In: *Robotics and Computer-Integrated Manufacturing*, vol. 73, pp. 102235. Elsevier, Amsterdam (2022). <https://doi.org/10.1016/j.rcim.2021.102235>
4. Orthey, A., Chamzas, C., Kavraki, L.: Sampling-Based Motion Planning: A Comparative Review. In: Editor, F., Editor, S. (eds.) *Annual Review of Control, Robotics, and Autonomous Systems 2023*, LNCS, vol. 7, pp. 1–13. Springer, Heidelberg (2023). <https://doi.org/10.1146/annurev-control-061623-094742>
5. Jagadeesan, S., Paul, E., Klimchik, A., Arunachalam, N.: A Comprehensive Framework for Stiffness Modeling of Industrial Robot for Machining Applications. In: Editor, F., Editor, S. (eds.) **Journal of Manufacturing Processes 2025**, LNCS, vol. 152, pp. 1111–1122. Springer, Heidelberg (2025). <https://doi.org/10.1016/j.jmapro.2025.07.085>
6. Kim, S.H., Nam, E., Ha, T.I., Hwang, S.H., Lee, J.-H., Park, S.-H., Min, B.-K.: Robotic Machining: A Review of Recent Progress. In: Editor, F., Editor, S. (eds.) **International Journal of Precision Engineering and Manufacturing 2019**, LNCS, vol. 20, pp. 1629–1642. Springer, Heidelberg (2019).

AI-Driven Defect Detection and Adaptive Control for Winding Process Optimization

Raphael Seliger^{[0009-0002-4158-6696]*} and Ulrich Göhner

Institute for Data-Optimized Manufacturing, Kempten University of Applied Sciences, Bahnhofstraße 61, 87435 Kempten, Germany

Abstract. Winding processes are used in a wide range of applications, such as shipping, cranes, and yarn production. However, compensating for typical defects, such as *valleys*, *hills*, and *crossovers*, remains a major challenge and can cause severe material damage. This work aims to enhance winding quality and filling factor while minimizing manual operator adjustments and downtime by integrating real-time computer vision with classical control logic. While previous studies have addressed defect classification or adaptive control separately, their combination using modern AI algorithms has not yet been explored. Especially for complex winding processes in large-scale wire spools with high material diameters, there is little prior research. The current state of practice still relies on manual monitoring of winding quality, requiring operators to adjust traverse speed and direction continuously. This improves the defects locally but cannot prevent the root causes of irregularities.

Keywords: Deep Learning · Defect Detection · Adaptive Control · Industrial Automation · Machine Vision

1 Introduction and Motivation

Automated quality monitoring and adaptive control have become key factors in the digital transformation of manufacturing processes. In winding machines for wires, cables, or coils, production quality is often highly dependent on manual adjustments by the operator and, therefore, on the ability to visually detect and correct defects during production. These methods rely on manual control inputs through physical interfaces (e.g., joystick or rotary switches), where operators modify traverse velocity based on visual feedback. This approach is practical but limited by human reaction time and is not scalable for continuous or high-speed production. Even skilled operators can only improve the winding patterns after the defects appear, but cannot prevent the root causes of the irregularities. This is especially true for complex winding processes involving large-scale wire spools with high material diameters; it is difficult, if not impossible, to optimize the winding algorithms to the extent that they can compensate for the unpredictable behavior of the winding material without direct observation. Defects such as

* Corresponding author: raphael.seliger@hs-kempten.de

valleys, hills, and crossovers occur frequently, especially at flange transitions or in the middle of the spool, and can accumulate over multiple layers if not corrected in their origin.

Recent advances in computer vision and deep learning now enable robust real-time analysis of visual patterns in dynamic industrial environments. Similar approaches have been successfully applied in transformer, motor, and rope winding [1–3]. Nevertheless, these were often not combined with adaptive control to prevent the defects in their origin. Additionally, for large-scale wire spools with high material diameters, there is a notable lack of prior research regarding their unique material properties. Therefore, this work proposes a hybrid perception-control architecture that integrates real-time defect detection using a YOLO-based model combined with classical PD control strategies. The system continuously monitors the winding process and derives adaptive traverse speed corrections through a hysteresis-based decision layer.

2 Project Objectives

1. **Defect Detection:** How can deep-learning-based algorithms be applied to detect and classify winding defects in real-time industrial processes automatically, and which approaches are most effective?
2. **Tracking and Decision Logic:** What strategies can be used to enhance the temporal consistency and robustness of defect detection in deep-learning-based industrial vision systems?
3. **Adaptive Traverse Control:** How can adaptive traverse control be designed to automatically correct detected deviations while maintaining process stability in winding operations?

3 Expected Project Outcomes

The expected outcome is a prototype system capable of identifying and mitigating winding defects autonomously in real time. By combining modern AI-based perception with classical control, the approach tries to close the research gap in this area for winding processes. The anticipated benefits primarily include improvements in winding quality and filling factor through early defect detection and continuous correction on large-scale wire spools. Additionally, we lay the foundation for future extensions, such as imitation learning or reinforcement learning, for even more optimized autonomous winding control.

References

1. Xu, J., et al.: *Research and Application of Contactless Measurement of Transformer Winding Tilt Angle Based on Machine Vision*. Sensors, 2023.
2. Farnsworth, M., et al.: *Augmented classification for electrical coil winding defects*. J. Manuf. Syst., 2022.
3. Wang, C., et al.: *Surface defect detection of wire rope using laser point cloud data and deep learning*. Optics & Laser Technology, 2025.

End-to-End Bill of Process Generation Using Multi-Modal Transformers

David Ulrich[0009-0005-8144-4328], Raphael Seliger[0009-0002-4158-6696]* and Frank Schirmeier[0009-0002-2842-0376]

Institute for Data-Optimised Manufacturing, UAS Kempten, Bahnhofstr. 61, Kempten, Germany

Abstract. In modern production systems, the high flexibility in production processes and the reduction of lot sizes make efficient automation strategies in the production planning necessary. This work presents ongoing research on automatic generation of manufacturing process plans, also known as Bill of Process (BoP) generation. Prediction strategies based on either structured tabular enterprise data or geometric CAD data are discussed and implemented. The study compares classical decision-tree-based approaches, such as random forests, with modern, transformer-based deep learning architectures. The data used in this work originates from proprietary industrial sources and represents real-world manufacturing scenarios.

Keywords: Bill of Process Generation, Computer-Aided Process Planning, UV-Net, Transformer, CAD, Tabular Manufacturing Data

1 Introduction

The generation of a manufacturing plan, or Bill of Process (BoP), is a key step linking product design to production execution. It defines the sequence of operations that transform raw materials into finished products. Traditionally, this process relies on manual or semi-automatic expert input. At manufacturing companies, this expertise is encoded in historical datasets consisting of manufacturing specifications and BoPs that have been aggregated over time. Based on this aggregated knowledge, Automatic BoP generation systems could be built and offer the promise of substantial gains in efficiency, scalability, and consistency.

Given the sequential and structured nature of manufacturing operations, transformer-based models are a natural choice. Transformers [4] have achieved state-of-the-art results across sequence modeling tasks, effectively capturing dependencies in heterogeneous manufacturing data. However, tree-based models such as random forests remain strong baselines for tabular data due to their robustness and interpretability, often outperforming deep models [2]. This work explores the trade-offs between these approaches for BoP prediction and investigates how incorporating CAD-derived structural features can enhance predictive performance.

* Corresponding author: raphael.seliger@hs-kempten.de

2 Methodology

We use proprietary, real-world manufacturing data from industry partners to predict BoPs, considering two input modalities: tabular manufacturing data and CAD models in STEP format. Structured, attribute-based data is modeled using an FT-Transformer-inspired architecture [1], which applies transformer layers directly to tabular features. This enables the network to capture complex feature interactions while preserving inductive biases for heterogeneous column types. Since the data contains many sparsely populated features, missing values are not imputed but omitted from token generation, reducing token count. Preprocessing includes feature normalization, categorical embedding, and optional positional encodings derived from manufacturing sequence indices.

For geometric input, STEP files are converted into graph-based representations, where nodes represent faces and edges encode adjacency information. A UV-Net-like architecture [3] processes these graphs, capturing both geometric and topological context. The resulting latent embedding provides a compact, process-relevant representation of the part geometry. We compare graph- and node-level embeddings as well as architectural variants. A transformer decoder then generates the manufacturing plan. Decision-tree-based models serve as baselines to assess whether transformer and graph-based models achieve superior generalization on manufacturing datasets.

3 Results

Our experiments show that the baseline — tree-based models — already shows strong performance. The novel approach to apply the transformer architecture on sparse tabular data is more delicate; it needs several steps to achieve high quality. Especially in situations in which the tabular data is insufficient or incomplete, the integration of STEP files can achieve a major performance boost. Here, the strength of the transformer approach comes to play, since it is unclear how graph embeddings can be incorporated into a tree-based model in a natural way.

References

1. Gorishniy, Y., Rubachev, I., Khrulkov, V., Babenko, A.: Revisiting Deep Learning Models for Tabular Data. In: Advances in Neural Information Processing Systems. vol. 34, pp. 18932–18943. Curran Associates, Inc. (2021)
2. Grinsztajn, L., Oyallon, E., Varoquaux, G.: Why do tree-based models still outperform deep learning on typical tabular data? *Advances in Neural Information Processing Systems* **35**, 507–520 (Dec 2022)
3. Jayaraman, P.K., Sanghi, A., Lambourne, J.G., Willis, K.D.D., Davies, T., Shayani, H., Morris, N.: UV-Net: Learning from Boundary Representations. In: 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 11698–11707 (Jun 2021). <https://doi.org/10.1109/CVPR46437.2021.01153>
4. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., ukasz Kaiser, L., Polosukhin, I.: Attention is All you Need. In: Advances in Neural Information Processing Systems. vol. 30. Curran Associates, Inc. (2017)

Safe Reinforcement Learning using Ideas from Model Predictive Control

Georg Schäfer^{1,2}, Jakob Rehrl¹, Stefan Huber¹, and Simon Hirlaender²

¹ Josef Ressel Centre for Intelligent and Secure Industrial Automation,
Salzburg University of Applied Sciences, Salzburg, Austria

² Department of Artificial Intelligence and Human Interfaces,
Paris Lodron University of Salzburg, Salzburg, Austria
georg.schaefer@fh-salzburg.ac.at

1 Introduction

Reinforcement learning (RL) enables learning control policies directly from data, making it appealing for robotics and other cyber-physical systems (CPSs) [6, 4]. A persistent challenge is ensuring safety constraints during learning [3]. In real-world systems, violating mechanical limits can cause irreversible damage; thus, exploration must remain within safe regions.

We propose a framework that combines RL with offline model predictive control (MPC) to ensure safety constraints. Using a model of the system dynamics, MPC determines whether an action may lead to unsafe outcomes. Offline MPC computations define a *feasible state-action space*, representing all safe combinations of system states and control inputs. During training and deployment, the agent's actions are projected onto this feasible set, providing formal safety guarantees without impeding learning.

Our approach relates to shielding [1], where a reactive mechanism corrects unsafe actions. In contrast, we precompute feasibility using MPC, enabling a projections-based safety filter that adjusts unsafe actions to the nearest feasible point. By penalizing deviations between original and corrected actions, the RL agent learns inherently safe policies.

2 Methodology

We consider continuous state and action spaces, $\mathcal{S} \subseteq \mathbb{R}^n$ and $\mathcal{A} \subseteq \mathbb{R}^m$. The core of our method is the precomputation of a *feasible state-action space* \mathcal{F} , which is the union of all safe state-action pairs. A state s is feasible if there exists at least one safe action, defining the feasible state space $\mathcal{F}_{\mathcal{S}}$. For each state $s \in \mathcal{F}_{\mathcal{S}}$, the set of safe actions is denoted as $\mathcal{F}_{\mathcal{A}}(s)$. Thus,

$$s \in \mathcal{F}_{\mathcal{S}} \Leftrightarrow \mathcal{F}_{\mathcal{A}}(s) \neq \emptyset, \quad \mathcal{F} = \bigcup_{s \in \mathcal{F}_{\mathcal{S}}} \{s\} \times \mathcal{F}_{\mathcal{A}}(s).$$

We call an RL policy π safe if $\pi(s) \in \mathcal{F}_{\mathcal{A}}(s)$ for all feasible states.

To construct the feasible set, we employ MPC as a *feasibility oracle*. For a given state s and action a , MPC checks whether a constraint-satisfying trajectory exists. This defines a binary oracle function:

$$\text{Oracle}(s, a) = \begin{cases} 1, & (s, a) \in \mathcal{F}, \\ 0, & \text{otherwise.} \end{cases}$$

By systematically sampling this oracle across the state-action space, we approximate \mathcal{F} . For high-dimensional systems, this boundary can be learned more efficiently using function approximators such as Gaussian Processes [5] or neural networks.

This learned feasibility boundary enables the construction of safe policies. To ensure safety for any learning policy π , we employ a projection filter P_s that maps proposed actions to the nearest feasible alternative:

$$\tilde{\pi}(s) = P_s(\pi(s)) = \begin{cases} \pi(s), & \pi(s) \in \mathcal{F}_{\mathcal{A}}(s), \\ \arg \min_{a' \in \mathcal{F}_{\mathcal{A}}(s)} \|\pi(s) - a'\|, & \text{otherwise.} \end{cases}$$

The resulting filtered policy $\tilde{\pi}$ ensures $\tilde{\pi}(s) \in \mathcal{F}_{\mathcal{A}}(s)$ for all $s \in \mathcal{F}_{\mathcal{S}}$.

3 Discussion and Outlook

Offline MPC enables precomputation of feasible regions that guarantee safety during learning. This approach is general and complementary to online safety mechanisms such as barrier functions [2]. Future work focuses on:

- **Scalability:** Adaptive sampling (e.g., Bayesian optimization) to query the oracle near \mathcal{F} ’s boundary.
- **Model Uncertainty:** Incorporating conservative margins or online model adaptations to maintain safety.
- **Experimental Validation:** Implementation on a laboratory helicopter testbed to evaluate performance and safety during real training.

References

1. Alshiekh, M., Bloem, R., Ehlers, R., Könighofer, B., Niekum, S., Topcu, U.: Safe reinforcement learning via shielding. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
2. Ames, A.D., Coogan, S., Egerstedt, M., Notomista, G., Sreenath, K., Tabuada, P.: Control barrier functions: Theory and applications. In: 2019 18th European control conference (ECC). pp. 3420–3431. IEEE (2019)
3. García, J., Fernández, F.: A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research* **16**(1), 1437–1480 (2015)
4. Kober, J., Bagnell, J.A., Peters, J.: Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research* **32**(11), 1238–1274 (2013)
5. Rasmussen, C.E.: Gaussian processes in machine learning. In: Summer school on machine learning, pp. 63–71. Springer (2003)
6. Sutton, R.S.: Reinforcement learning: An introduction. A Bradford Book (2018)

Reducing Downtimes in Process Control Systems through the use of Digital Twins

Manuel Taferner¹, Matthäus Horn¹, Simon Kranzer¹, Hanna Trenkler¹, and Reinhard Portenkirchner¹

¹Salzburg University of Applied Sciences, Austria,
Department Information Technology
<https://www.fh-salzburg.ac.at>

Abstract. Minimizing unplanned downtimes in process control systems is a central requirement in highly automated production environments. As part of commissioned research conducted in collaboration with an industrial partner, this work addresses the challenge that arises when system updates, configuration errors, or delayed fault diagnosis interrupt production and result in substantial costs. Process control systems such as *zenon* from *COPA-DATA* are central to the operation of many industrial plants, yet they still often require temporary shutdowns when new control logic or visualization changes are implemented and tested. At the same time, virtual commissioning [4] and digital twin technologies together with concepts of digital shadows [1] offer promising means to validate changes before deployment. Furthermore, 3D operator guidance systems such as *twin* from *digifai* have the potential to accelerate troubleshooting by providing visual assistance during fault recovery [2].

This work investigates how a digital shadow, that is tightly integrated with a *zenon* based process control system, can contribute to the reduction of downtime in industrial environments. The guiding research question is how changes to *zenon* projects can be virtually implemented and tested without interrupting operation and whether an integrated 3D simulation system can further reduce downtime during fault situations. The underlying hypothesis is that by connecting a virtual replica of a production line, realized as a digital shadow in *twin*, with a *zenon* Soft PLC and a real production demonstrator, the duration and frequency of downtimes can be significantly reduced.

Building on the concepts of digital shadows and virtual commissioning, a proof of concept was realized for a cereal bar production line. The setup combines three elements. First, a three dimensional simulation in *twin* that represents the key process stages of material transport, cutting, and robotic handling. Second, a *zenon* Soft PLC and visualization project that controls both the virtual line and the physical demonstrator. Third, a physical demonstrator implementing the final segment of the line. All components communicate bidirectionally via the OPC UA protocol, which ensures real time data synchronization between digital and physical system parts. This hybrid architecture enables virtual testing of control logic and visualization changes in the digital shadow before

deployment to the live environment and allows the same 3D model to be used for operator guidance during fault handling.

Initial results from scenario based experiments show that typical control modifications and visualization updates can be developed and validated completely with the digital shadow and then transferred to the real system without additional downtime. Running the simulation and the physical demonstrator in parallel enables safe testing, training, and optimization while maintaining continuous operation of the real plant segment. In fault scenarios, the 3D simulation increased transparency of machine states and supported faster localization of the root cause in the observed cases compared to relying on a classical two dimensional HMI only.

The study also reveals challenges that need to be addressed in future work. These include the effort required to create and maintain sufficiently accurate three dimensional models, the management of configuration variants between virtual and real systems, and the integration of three dimensional operator guidance into existing engineering and maintenance workflows. The contributions of this work are threefold. It introduces a concrete hybrid architecture that connects *zenon*, *twin*, and a real demonstrator for virtual commissioning and three dimensional operator guidance. It reports qualitative findings on the potential of digital shadows and three dimensional guidance to reduce downtimes in *zenon* based environments. Given that related studies have demonstrated downtime reductions of over 80% [3] when virtual models are used for testing and diagnostics, similar effects are expected in the presented architecture. Quantifying these improvements will be part of future work.

Keywords: Process Control System · Digital Shadow · Downtime Reduction · Digital Twin · OPC UA

Acknowledgments. The results and findings presented here were developed in the *CoopLab* project. Special thanks to the funding bodies *COPA-DATA* and *digifai*.

References

1. Bergs, T., Gierlings, S., Auerbach, T., Klink, A., Schraknepper, D., Augspurger, T.: The concept of digital twin and digital shadow in manufacturing. *Procedia CIRP* **101**, 81–84 (2021)
2. Gurumurthy, R.: Digital twins: The key to manufacturing sustainability. *ISE ; Industrial and Systems Engineering at Work* **56**(9), 30–35 (2024)
3. Mendi, A.F.: A digital twin case study on automotive production line. *Sensors (Basel, Switzerland)* **22**(18), 6963 (2022)
4. Striffler, N., Voigt, T.: Concepts and trends of virtual commissioning – a comprehensive review. *Journal of manufacturing systems* **71**, 664–680 (2023)

Generating Approximately Periodic Multivariate Timeseries with Gaussian Processes

Elias Reich^[0009-0000-4561-9067] and Stefan Huber^[0000-0002-8871-5814]

Josef Ressel Centre for Intelligent and Secure Industrial Automation
Salzburg University of Applied Sciences, Salzburg, Austria
eliassteffen.reich@fh-salzburg.ac.at

1 Motivation

Discrete automated processes, e.g. in industrial production machines, display a strongly repetitive behavior. That is, a multivariate time series $x: I \rightarrow \mathbb{R}^D$ over a finite time interval $I = [0, T]$, coming from observations of those machines, roughly repeat themselves. As no two repetitions are identical, this motivates an underlying stochastic process as a model from a machine learning (ML) perspective. Such a generative model enables us to populate cybersecurity honeypots, generate mockups for unit testing and digital shadows, and the like. Moreover it can be used for anomaly detection through the likelihood of a given timeseries.

Let $x: I \rightarrow \mathbb{R}^D$ such that I can be split into consecutive segments I_1, \dots, I_k , each representing one repetition of the process. By introducing reparameterizations $\pi_i: I_i \rightarrow [0, 1]$ that normalize time within each segment, we define the *time-normalized cycle* as $y_i(t) = x(\pi_i^{-1}(t))$ for all $t \in [0, 1]$. If all pairs (y_i, y_j) are ε -approximately equivalent, i.e., $\|y_i - y_j\|_\infty \leq \varepsilon$, we call x *repetitive* and y_1, \dots, y_k repetitions of an *approximately periodic* time series. How x is split in its repetitions is shown in [2].

In this paper we introduce a generative model for *approximately periodic* time series based on the Gaussian process (GP), that can generate time series of arbitrarily many repetitions, preserving continuity over the whole time span while still allowing variation between repetitions. Existing approximately or quasi-periodic GPs focus on uncertainty estimation or short term forecasting [1] instead of generating completely new time series. Our model provides long term stability, interpretability and uncertainty quantification, unlike popular generative models such as NeuralODEs.

2 Methodology

We refer to the standard textbook [3] on GPs. In Bayesian linear regression, the prior for infinitely many, normally distributed parameters is a GP. It is parameterized by a covariance $k(\cdot, \cdot)$ and mean function $\mu(\cdot)$. Given input–output pairs $\{(t_i, y_i)\}_{i=1}^N$ with $y_i \in \mathbb{R}^D$ and $t_i \in \mathbb{R}$, the GP-regression model is formalized as $y_i = f(t_i) + \epsilon$, where $f \sim \mathcal{GP}(\mu(t), k(t, t'))$ and $\epsilon \sim \mathcal{N}(0, \sigma^2)$.


The core modelling challenge is to obtain a GP posterior that is approximately periodic in the sense that (i) the posterior mean is equivalent for all cycles, while (ii) the posterior covariance permits independent, smoothly decaying variation between cycles. A prior for periodic processes with period length p is an exponential kernel with inputs embedded on the circle [3]. It simplifies to

$$k(t, t') = \exp \left(\frac{2}{l^2} \sin \left(\frac{\pi}{p} (t - t') \right)^2 \right). \quad (1)$$

To model approximate periodicity, we introduce a smoothly decaying correlation envelope w that decreases the covariance between cycles

$$w(t, t') = \exp \left(-\frac{\|\phi(t) - \phi(t')\|^2}{2l^2} \right), \quad \phi(t) = \int \sin \left(\frac{\pi}{p} t \right)^2 dt. \quad (2)$$

The kernel w is a Gaussian kernel in an embedding space, which has vanishing derivatives at integer multiples of p to avoid unwanted phase shifts in the weighted kernel $g = w \cdot k$. Naively using g as a prior will have the effect of a decaying posterior mean $\hat{\mu}$ and increasing posterior variances $\text{diag}(\hat{\Sigma})$. We solve this problem by using the periodic prior k for the single repetition model $y = f + \epsilon$, $f \sim \mathcal{GP}(0, k(t, t'))$, $\epsilon \sim \mathcal{N}(0, \sigma^2)$ and only weight the covariance matrix of the posterior $y_*|y \sim \mathcal{N}(\hat{\mu}, \hat{\Sigma})$ at test inputs T_* with the kernel matrix W of the weight function w . From here, we can then sample approximately periodic trajectories $\gamma \sim \mathcal{N}(\hat{\mu}, W \circ \hat{\Sigma})$. Below is an example $\gamma: [0, 4p] \rightarrow \mathbb{R}^2$, $p = 1$.



References

1. Li, Y., Zhang, Y., Xiao, Q., Wu, J.: Quasi-periodic gaussian process modeling of pseudo-periodic signals. *IEEE Transactions on Signal Processing* **71**, 3548–3561 (2023). <https://doi.org/10.1109/TSP.2023.3316589>
2. Schindler, S., Reich, E.S., Messineo, S., Huber, S.: Topology-driven identification of repetitions in multi-variate time series. In: Proceedings of the 6th Interdisciplinary Data Science Conference – iDSC2025 (2025), <https://arxiv.org/abs/2505.10004>
3. Williams, C.K., Rasmussen, C.E.: Gaussian processes for machine learning. MIT Press: Cambridge, MA, USA (2006)

Empirical Optimization of Computer Vision Model for Real-World Part Verification in Manual Order Picking: A Technology Transfer Perspective

Md Khalid Siddiqui¹[0009-0000-2235-3585], Supratim Chakraborty¹[0009-0006-3559-0188],
Jonathan Kressel¹[0009-0002-4727-4116], and Jürgen Grinninger¹

¹ Hochschule Neu-Ulm, Wileystraße 1, 89231 Neu-Ulm
md.khalid.siddiqui@hnu.de

Abstract.

Context and Motivation. Manual order picking remains a critical yet error-prone process in manufacturing and logistics [1-3]. Despite the growing maturity of digital assistance systems, small and medium-sized enterprises (SMEs) continue to rely on human verification due to the high cost and technical complexity of existing solutions [1, 2, 4]. With the recent progress of compact object detection models such as You Only Look Once (YOLO) [5, 6], there is a huge potential to develop low-cost visual verification systems for SMEs in need of low cost alternatives. However, most academic studies evaluate such models with validation and test datasets, and little is known about how their training behavior translates to actual performance in physical production environments [1, 4, 7].

Approach and Setup. This work-in-progress investigates how training characteristics influence real-world detection quality at an operational picking station. The study builds upon two earlier prototypes for AI-assisted order picking[1, 4], which demonstrated the technical feasibility of integrating visual part verification into manual workflows. A dedicated picking station was implemented at the TTZ Smart Production & Logistics Lab [1, 4], replicating a small-scale industrial setup with labeled KANBAN bins, an overhead camera, and a display guiding the operator. This research extends these efforts by focusing on how variations in training epochs and dataset composition influence observable detection reliability during live operation.

For the object detection task the YOLOv11n model [5-7], a lightweight variant of the YOLOv11 architecture, was trained on a labelled dataset of 25 components, each with approximately 300 images, captured under varying lighting conditions. Several training runs were conducted with varying epoch counts while keeping other hyper-parameters constant. The resulting models were validated on held-out data and then deployed on the physical station to observe their real-time detection behavior during picking tasks.

Observations and Outlook. Preliminary results suggest that increasing the number of training epochs (from 128 to 256) improves inference accuracy and reduces the cases of missed detection significantly which does not change much on further increasing the training epochs, particularly for smaller or visually similar

parts. A positive effect of higher image resolution on detection confidence was also noted. The system currently achieves real-time responsiveness (~23 FPS) on the lab's picking station, providing continuous operator feedback.

The ongoing work focuses on aligning validation metrics from training with practical detection consistency observed in the physical setup. We wanted to identify cases where validation accuracy does not correspond to live performance, and to examine whether additional training images for difficult parts improve real-world reliability.

Our experiments aim to compare model validation results with their actual performance in a real picking environment. Rather than pursuing algorithmic novelty, this contribution seeks to document empirical relationships that can inform practitioners particularly SMEs in developing affordable computer vision picking assistance systems instead of costly alternatives [8]. With this, we provide the first empirical correlation study on the reliability of real-world manual-picking detection using computer vision.

Keywords: Object Detection, YOLO, Picking Assistance Systems, Smart Manufacturing

References

1. Siddiqui, M.K., Kressel, J., Grinninger, J.: Technologies for Assisting Manual Order Picking. *Industry 4.0 Science* 41, 6–19 (2025)
2. Heinin, J., Sohlberg, R.: Challenges and opportunities with manual order picking. (2023)
3. Baechler, A., Baechler, L., Autenrieth, S., Kurtz, P., Hoerz, T., Heidenreich, T., Kruell, G.: A comparative study of an assistance system for manual order picking--called pick-by-projection--with the guiding systems pick-by-paper, pick-by-light and pick-by-display. In: 2016 49th Hawaii International Conference on System Sciences (HICSS), pp. 523–531. IEEE, (2016)
4. Siddiqui, M.K., Hofmann, B., Engel, T.: Development and Implementation of a Real Time AI-Based Part Verification System in Picking Process Using Low-Cost Equipment. (2025)
5. Hidayatullah, P., Syakrani, N., Sholahuddin, M.R., Gelar, T., Tubagus, R.: YOLOv8 to YOLOv11: A comprehensive architecture in-depth comparative review. arXiv preprint arXiv:2501.13400 (2025)
6. Ali, M.L., Zhang, Z.: The YOLO framework: A comprehensive review of evolution, applications, and benchmarks in object detection. *Computers* 13, 336 (2024)
7. Jegham, N., Koh, C.Y., Abdelatti, M., Hendawi, A.: Yolo evolution: A comprehensive benchmark and architectural review of yolov12, yolo11, and their previous versions. arXiv preprint arXiv:2411.00201 (2024)
8. Mandar, E.M., Dachry, W., Bensassi, B.: Toward a real-time picking errors prevention system based on RFID technology. *Advances on Smart and Soft Computing: Proceedings of ICACIn 2020*, pp. 303–318. Springer (2020)

An Embedded On-Device Training Framework for Neural Networks

Michael Kargl^{1[0009-0009-7753-3335]}, David Muttenthaler¹, Florian Eibensteiner^{1[0000-0002-5311-4082]}, Christoph Dalpiaz¹, Phillip Kastner^{1[0000-0002-7943-7226]}, and Josef Langer¹

University of Applied Sciences Upper Austria, Hagenberg, Austria
 {michael.kargl, david.muttenthaler, florian.eibensteiner, christoph.dalpiaz, phillip.kastner, josef.langer}@fh-hagenberg.at

Abstract. We present a portable framework that enables inference and training of neural networks directly on embedded devices. It consists of a PyTorch model parser that exports networks to a binary format and a C++ runtime capable of executing inference and training on this binary format. A run of a minimal CNN on MNIST achieved numerical parity with PyTorch in inference and demonstrated the feasibility of our framework for on-device fine-tuning in non-time-constrained scenarios on an STM32F469 microcontroller.

Keywords: embedded ML · on-device training · tinyML

1 Introduction

Embedded AI has in recent years been a topic of increasing interest as it allows to process complex data effectively and locally on low-power devices. Most advancements in this area have been focused only on enabling the inference of models trained on more capable machines on lightweight embedded hardware [1]. Publications on enabling on-device training arise rarely and many of them are only suitable for specific use-cases [2] or are not geared towards bare-metal programming on small embedded systems [3]. We therefore implemented the to our knowledge first general bare-metal C++ neural network training framework that can take unmodified, non-proprietary Python models as an input.

2 Concept

The implemented framework consists of two stages as seen in Figure 1.

Stage 1 implements a parser in Python that takes any PyTorch model, that only consists of the implemented layers (convolution, fully connected, batch normalization, dropout, max & average pooling) and activation functions (ReLU and Softmax). The parser also supports frozen weights for memory efficient training. All this information about network architecture and weights is converted into a binary format that the C++ runtime uses.




Fig. 1. Overview of the model conversion pipeline.

Stage 2 implements the C++ runtime that interprets the data in the binary format to rebuild the structure of the neural network on the target platform. For training a neural network the weight data is loaded into SRAM memory, for inference the runtime can use Flash memory and only reads the currently used weights to keep SRAM usage minimal.

3 Results

The framework was evaluated on a simple CNN training on the MNIST dataset. It was confirmed that the runtime execution of a converted model has numerical parity to the reference model in inference and that the model converges properly during training. On an STM32F469 clocked at 168 MHz the average inference time per MNIST sample was \approx 53 ms and a training run with batch size 10 took \approx 2.9 s, demonstrating that the framework while not high in performance, is suitable for long-term on-device tuning in non-time-critical tasks.

In a comparison to the reference PyTorch implementation on a Desktop system it was observed that our framework only consumed 2.41 MB of additional memory opposed to PyTorch with 52.78 MB. Our first simple test programs for microcontrollers showed an over-all memory footprint of roughly 400 KB showing the reduced memory footprint that enables our framework for embedded use.

Acknowledgments. The financial support by the Austrian Federal Ministry of Labour and Economy, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged.

References

1. S. Han, H. Mao, and W. J. Dally, *Deep Compression: Compressing Deep Neural Networks with Pruning, Trained Quantization and Huffman Coding*, arXiv preprint arXiv:1510.00149, 2016.
2. H. Ren, D. Anicic, and T. Runkler, *TinyOL: TinyML with Online-Learning on Microcontrollers*, arXiv preprint arXiv:2103.08295, 2021.
3. TensorFlow Lite, *TensorFlow Lite: An Open-Source Deep Learning Framework for On-Device Inference*, Google, 2017. Available at: <https://www.tensorflow.org/lite>.

Human-Centered AI for Industrial Maintenance: Prototyping an LLM-Based Knowledge Platform

Gerald A. Zwettler¹, Martin Hanreich¹, Karin Kloiber³, Johannes Schönböck², and Josef Altmann²

¹ University of Applied Sciences Upper Austria, Research Group Advanced Information Systems and Technology, Hagenberg, Austria
 {gerald.zwettler,martin.hanreich}@fh-hagenberg.at

² University of Applied Sciences Upper Austria, Research Group Personalized Environments and Collaborative Systems, Hagenberg, Austria

³ Competence Center CHASE GmbH, Digitalization Chemical Systems, Linz, Austria

Introduction and Problem Definition. Industrial maintenance in small and medium-sized enterprises (SMEs) is being reshaped by demographic shifts, increasing system complexity, and the demand for sustainable operations. A key challenge is the retention and transfer of expert knowledge, which is often undocumented and tacit. As experienced technicians retire or change roles, organizations risk losing critical insights, leading to inefficiencies and downtime. This issue is especially pronounced in domains like injection molding, where maintenance decisions require deep contextual understanding and quick response. To address these challenges, we introduce a human-centered AI platform that applies recent advances in large language models (LLMs) and Retrieval-Augmented Generation (RAG). The platform is designed to digitize expert knowledge, offer contextual assistance, and support knowledge sharing across teams, thereby improving operational resilience and aligning with Industry 5.0 principles.

Related Work. The integration of AI into industrial maintenance builds on foundational work in human-centered design and intelligent systems. [1] emphasize the importance of expert-driven contextual design in industrial environments, advocating for deep involvement of domain experts to ensure usability and relevance. [2] argue for a shift toward Industry 5.0, highlighting the need for resilient and ethically grounded systems that prioritize human-centricity. Recent advances in LLMs have shown promise in digital manufacturing, with [3] identifying key benefits in process optimization, data structuring, and human-machine interaction. Additionally, [4] demonstrate the limitations of rule-based systems in dynamic maintenance scenarios, underscoring the value of adaptive, data-driven approaches for real-time reasoning and continuous learning.

Methodology. The platform was developed using a User-Centered Design (UCD) approach, ensuring that domain experts were actively involved throughout the prototyping process. The system architecture integrates LLMs with a RAG framework, enabling the retrieval of domain-specific information from technical documents, manuals, and historical maintenance logs. This retrieved content is then used to generate context-aware responses tailored to the technician's current task. Multimodal interaction capabilities—including text, voice, and image inputs—allow users to engage with the system in flexible and intuitive ways.

Document reasoning features enable the platform to interpret and explain complex technical content, while feedback loops support continuous learning and adaptation. The prototype was deployed in a real-world injection molding environment, evaluated through task-based scenarios and expert feedback.

Results. The evaluation demonstrated that the platform effectively supports technicians by providing relevant, situation-specific guidance during maintenance tasks. Users reported improved confidence in decision-making and appreciated the system's ability to interpret technical documentation and respond to queries in natural language. The multimodal interface proved valuable in noisy environments and when visual inspection was required. Moreover, the platform facilitated the digitization and preservation of expert knowledge, enabling scalable knowledge sharing across teams and reducing reliance on individual expertise. These outcomes suggest that the system can play a significant role in mitigating the effects of workforce aging and skill shortages in industrial settings.

Conclusion and Outlook. This work introduces a novel approach to industrial maintenance that combines human-centered design with state-of-the-art AI technologies. By embedding LLMs within a RAG architecture and supporting multimodal interaction, the proposed platform addresses key challenges in knowledge retention, contextual decision support, and technician empowerment. The system aligns with Industry 5.0 goals by emphasizing human-centricity, sustainability, and resilience. Future work will focus on expanding the platform's domain coverage, enhancing its reasoning capabilities, and integrating predictive maintenance features. The initial results of the expert user study underscore the potential of human-centered AI to transform industrial maintenance and support SMEs in their digital transformation journey.

Acknowledgments. The authors acknowledge financial support from the COMET Centre CHASE, funded within the COMET - Competence Centers for Excellent Technologies program by the BMIMI, the BMWET and the Federal Provinces of Upper Austria and Vienna, managed by the Austrian Research Promotion Agency (FFG).

Further thanks to our research partner RICO Elastomere Projecting GmbH for providing expertise, test persons, and valuable discussion.

References

1. Augstein, M., Neumayr, T., Pimminger, S., Ebner, C., Altmann, J., Kurschl, Contextual design in industrial settings: Experiences and recommendations. In: Proceedings of the 20th International Conference on Enterprise Information Systems (ICEIS 2018), pp. 429-440, Funchal, Madeira, Portugal (2018)
2. Grosse, E. H., Sgarbossa, F., Berlin, C., Neumann, W. P. Human-centric production and logistics system design and management: transitioning from Industry 4.0 to Industry 5.0. *Int. Journal of Production Research*, 61(22), pp. 7749–7759, (2023)
3. Ouerghemmi, C., Ertz, M. Integrating Large Language Models into Digital Manufacturing: A Systematic Review and Research Agenda, *Computers*, 14:8 (2025)
4. De Gasperis, G., and Facchini, S.D., A Comparative Study of Rule-Based and Data-Driven Approaches in Industrial Monitoring, *arXiv* (2025)

Finetuning LLMs on Synthetic Tasks for Improved STEP Understanding

Fabio Lischka^{1[0009-0003-7617-2407]}, Frank Schirmeier^{1[0009-0002-2842-0376]}, and Ulrich Göhner^{1[0000-0002-4795-4798]}

¹ Institute for Data-Optimised Manufacturing, UAS Kempten, Bahnhofstr. 61, Kempten, Germany
 {fabio.lischka, frank.schirmeier, ulrich.goehner}@hs-kempten.de

1 Introduction

Large Language Models (LLMs) have shown remarkable flexibility in reasoning tasks across domains and promise scalable benefits for Industry 4.0 applications. However, domain adaptation is necessary to access their full potential. In this paper, we evaluate how Reinforcement Finetuning (RFT) can be used to adapt an LLM to the input format of STEP files, a text-based format that stores 3D objects in Boundary Representation format (BRep). STEP is commonly used as an exchange format in manufacturing. Making these files accessible to LLM-based agents holds immense potential for industrial automation.

2 Objectives

The goal of this line of research is to obtain finetuned LLMs that exhibit improved performance on tasks involving STEP files. Such LLMs would be of great value for downstream applications like manufacturing cost estimation or even agentic pipelines for CAD generation. While synthetic CAD tasks such as the one in this paper could be solved more efficiently by rule-based approaches or graph-based Machine Learning, LLMs promise superior generalisation to unseen tasks and interpretability via reasoning traces.

3 Literature Review

There are already frameworks that seek to enable LLMs to answer questions about CAD models, e.g. [2]. Another line of work focuses on finding condensed textual representation of CAD models, e.g. in the form of code snippets or command sequences that reproduce the CAD model. [3] Findings from both areas of research can be combined with our finetuning approach.

4 Methods

We use a synthetic CAD dataset consisting of CAD files with annotated blind holes. The dataset is based on MFCAD++ [1], but simplified to include only blind holes as features for simplicity. Our synthetic task is to count the number of blind holes in a file from this dataset. The prompt for the LLM includes the entire STEP file as text.

We use Qwen3-0.6B as our base LLM for finetuning. Without finetuning, this model performs no better than random guessing at our synthetic task.

We employ GRPO as our finetuning algorithm. We compare two reward functions: A binary reward based on correctness of the LLM's answer alone, and a reward function including the difference between the LLM's answer and the true number of blind holes (which we refer to as "L1 penalty").

5 Results

The binary reward function does not improve accuracy on the synthetic task significantly. The reward function with L1 penalty, however, increases accuracy from 10% for the model without finetuning to about 50% on the training dataset. Hold-out evaluation confirms substantial improvements.

6 Discussion

The experiment highlights the importance of choosing reward functions that assign different rewards to different LLM-generated answers.

Our experiment further shows that Reinforcement Finetuning is a viable path to teach LLMs basic understanding of structured data, like STEP files. The natural language understanding of LLMs appears to be a useful asset for solving problems with structured language as input in a sample-efficient manner. It should be noted, though, that our synthetic task can be solved by counting keywords alone and does not strictly require geometric understanding.

Further research is required to evaluate the performance of the finetuned LLM on downstream tasks. Future research directions include using a vision encoder to capture geometric information from a rendering of the STEP file.

References

1. Colligan, A.R., Robinson, T.T., Nolan, D.C., Hua, Y., Cao, W.: Hierarchical cadnet: Learning from b-reps for machining feature recognition. *Computer-Aided Design* **147**, 103226 (Jun 2022). <https://doi.org/10.1016/j.cad.2022.103226>
2. Kienle, C., Alt, B., Katic, D., Jäkel, R.: Quercad: Grounded question answering for CAD models. *CoRR* **abs/2409.08704** (2024). <https://doi.org/10.48550/ARXIV.2409.08704>, <https://claudius-kienle.github.io/papers/p/quercad.html>
3. Xie, H., Ju, F.: Text-to-cadquery: A new paradigm for CAD generation with scalable large model capabilities. *CoRR* **abs/2505.06507**(technical) (2025). <https://doi.org/10.48550/ARXIV.2505.06507>

Comparing Generalisation in Neural ODEs under Different Regularisation Schemes

Olaf Sassnick^{1,2,3}, Thomas Rosenstatter^{1,2},
Andreas Unterweger^{1,2,3}, and Stefan Huber^{1,2}

¹ Josef Ressel Centre for Intelligent and Secure Industrial Automation

² Salzburg University of Applied Sciences, Salzburg, Austria

³ Paris Lodron University of Salzburg, Salzburg, Austria

olaf.sassnick@fh-salzburg.ac.at

1 Introduction

Neural Ordinary Differential Equations (Neural ODEs) [1] represent a family of machine learning models that describe the dynamics of a system by learning its state derivatives. As their mathematical formulation resembles the laws of physics, Neural ODEs are particularly suited for modeling dynamic systems. This makes Neural ODEs interesting for industrial applications, such as predictive maintenance, anomaly detection, and data-driven machine simulations.

In their simplest form, Neural ODEs are defined as $\dot{\mathbf{h}}(t) = f_\theta(\mathbf{h}(t))$, where f is a neural network parametrised by θ and \mathbf{h} represents the system state. The function f_θ can be interpreted as a multidimensional continuous vector field describing how the system evolves from any possible state \mathbf{h} . During training, the parameters θ are fitted, while during inference, the system is numerically forward-integrated in multiple steps to compute the next system state $\mathbf{h}(t + \Delta t)$.

The *regularisation* plays an important role when training Neural ODEs, as it can stabilize the learned dynamics and control model complexity. Most existing research has focused on how regularisation influences training stability or the numerical behaviour of the ODE solver [2,4,3], whereas less attention has been paid to how different regularisation strategies affect the *generalisation capability* of Neural ODEs.

In the context of machine learning, generalisation refers to capability of a model to perform well on unseen data. Improving generalisation helps building robust models, particularly when training data are scarce or noisy. Regarding Neural ODEs, the parameters θ are optimised using only the derivatives observed in the training. Nevertheless, the resulting model f_θ defines a continuous vector field that, in principle, could generate correct derivatives for previously unseen states. The extent to which these extrapolated dynamics remain correct determines the *generalisation capability*.

While recent theoretical works, like [5], derive generalisation bounds under Lipschitz constraints, a wider range of empirical comparisons across different regularisation schemes is not yet available. Therefore, in this work, we systematically investigate multiple regularisation approaches in Neural ODEs, analysing their effect on the *generalisation capability* across different scenarios.

2 Approach

For the comparison we consider a baseline Neural ODE and several established regularised variants, including Regularized NODEs (RNODE) with Jacobian and kinetic regularisation [2], temporal-based regularisation [3], and Neural ODEs with stochastic noise introduced to the hidden layer [4]. Wherever applicable, we also evaluate hybrid approaches by combining different regularisation schemes.

The models are trained and tested on a set of physical systems, such as synthetic mass-spring systems, which are commonly used as benchmark problems in the existing literature. For well-comparable results across different regularisation schemes, the network structure, like layer count and size, is remaining unchanged within each experiment. The evaluation is performed by measuring trajectory errors and the number of solver steps required during numerical forward-integration. Preliminary results show that fewer steps indicate less stiff vector fields, which correlates with improved generalisation and numerical stability. The *generalisation capability* of each trained model is evaluated by (i) temporal extrapolation beyond the training time, (ii) unseen initial conditions, and (iii) a combination of both. For (ii), the initial conditions used for training and testing are varied to different extents, ensuring non-overlapping state-space regions. This allows a direct assessment of how well the models extrapolate beyond the training trajectories.

The full implementation and all experiment scripts will be published on GitHub to support reproducibility.

Acknowledgment

The financial support by the Austrian Federal Ministry of Economy, Energy and Tourism, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged.

References

1. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. In: Advances in Neural Information Processing Systems. vol. 31 (2018)
2. Finlay, C., Jacobsen, J.H., Nurbekyan, L., Oberman, A.: How to train your neural ode: the world of jacobian and kinetic regularization. In: International conference on machine learning. pp. 3154–3164. PMLR (2020)
3. Ghosh, A., Behl, H., Dupont, E., Torr, P., Namboodiri, V.: Steer : Simple temporal regularization for neural ode. In: Advances in Neural Information Processing Systems. vol. 33 (2020)
4. Liu, X., Xiao, T., Si, S., Cao, Q., Kumar, S., Hsieh, C.J.: Neural SDE: Stabilizing Neural ODE Networks with Stochastic Noise (Jun 2019). <https://doi.org/10.48550/arXiv.1906.02355>, arXiv:1906.02355
5. Marion, P.: Generalization bounds for neural ordinary differential equations and deep residual networks. In: Advances in Neural Information Processing Systems. vol. 36 (2023)

BLE 5.4 Periodic Advertising with Responses: Eliminating Limitations of Advertisement-Based BLE Communication

Marek A. Bawiec* and Antoni Duda

Department of Computer Engineering, Faculty of Electronics,
Wrocław University of Science and Technology, Poland
marek.bawiec@pwr.edu.pl, antoni.duda@programmer.net

Abstract. The advertisement-based Bluetooth Low Energy (BLE) communication model offers low overhead and purely connectionless operation, but in its plain broadcast form provides no structural protection against flooding, spoofing, replay, or unsolicited multi-receiver triggering. This extended abstract revisits the design evaluated in *Experimental Evaluation of Advertisement-Based Bluetooth Low Energy Communication* (Sensors, 2020) and shows how the BLE 5.4 feature *Periodic Advertising with Responses* (PAwR) introduces deterministic timing and slotting that constrain responder behaviour at the MAC level. Although PAwR itself does not provide authentication or access control, its time- and slot-scoped operation enables a lightweight security layer when combined with a whitelist and per-packet HMAC. A compliant implementation on Nordic nRF SoCs, using the Zephyr-based nRF Connect SDK, demonstrates that this combination eliminates the main limitations of advertisement-based BLE communication while preserving its connectionless model.

Keywords: BLE · PAwR · DoS resilience · Connectionless IoT

1 Introduction and Context

Advertisement-mode BLE was systematically evaluated in *Experimental Evaluation of Advertisement-Based Bluetooth Low Energy Communication* [1], which introduced a reversed-role model where advertisers transmit application data and scanners act as passive receivers. This enabled low-latency, low-energy, fully connectionless communication for dense IoT deployments. However, systems built on legacy BLE 4.x/5.0 advertisements suffer from structural limitations: no acknowledgments, no timing discipline, unrestricted receiver reactivity, and no built-in access control. As a result, attackers may induce *unsolicited multi-receiver triggering* (one crafted frame prompting reactions from many scanners), saturate scanning resources with bogus traffic, or inject spoofed and replayed frames without temporal or identity validation.

* Corresponding author.

BLE 5.4 introduces *Periodic Advertising with Responses* (PAwR) [2], which retains the connectionless model but provides deterministic periodic events and subevents with dedicated response windows. These timing constraints provide the discipline that the 2020 advertisement-based concept implicitly relied on but could not enforce with earlier BLE standards. When combined with lightweight authentication—a whitelist and per-packet HMAC—PAwR enables secure, time-scoped two-way exchanges that eliminate the sources of broadcast-induced DoS vulnerabilities.

2 PAwR Timing and Lightweight Authentication

PAwR structures time into periodic events B_k containing subevents S_i of duration Δ and period $T_B = \text{period}(B_k)$. Each responder is deterministically mapped $\text{map}(\text{TagID}) \rightarrow S_i$ and may transmit only after a valid reception within its assigned window $W_i = [t_i, t_i + \Delta]$. Channel Selection Algorithm #2 yields predictable hopping $c_k = f(k)$, while the PAwR controller maintains timing alignment under bounded drift [2]. Thus, receivers cannot transmit outside W_i , and unsolicited responses are structurally prevented.

We add a minimal security layer: a tag whitelist and an HMAC computed over `header || payload || seq`. A scanner responds only if: (i) the frame is received within W_i , (ii) the TagID is present in the whitelist, (iii) the HMAC verifies, and (iv) the sequence number is fresh (preventing replay). Frames failing any clause generate *no response*. Hence, unsolicited multi-receiver triggering, spoofing, and replay are blocked by timing + whitelist + HMAC, and flooding attempts outside response windows become ineffective.

3 Implementation and Summary

A reference implementation on nRF54-class hardware [3], using the Zephyr-based nRF Connect SDK [4], operates exclusively in PAwR mode (legacy/extended advertisements and GATT disabled). In a two-board setup with adversarial traffic injected on advertising channels, we observe stable periodic synchronization, deterministic slotting, and zero unauthorized responses. Traffic outside assigned windows produces no transmissions, confirming that PAwR supplies strong temporal guardrails while whitelist+HMAC provide identity and integrity checks. Together, they eliminate the limitations inherent to broadcast-only BLE and enable secure, scalable, connectionless IoT operation.

References

1. Nikodem, M., Bawiec, M.: Experimental Evaluation of Advertisement-Based Bluetooth Low Energy Communication. *Sensors* **20**(1), 107 (2020).
2. Bluetooth SIG: Bluetooth Core Specification Version 5.4. Kirkland, WA (2023).
3. Nordic Semiconductor: nRF54L15 Product Specification, Rev. 1.2 (2024).
4. Nordic Semiconductor: nRF Connect SDK Documentation, Version 2.9.x (2024).

Disclosure of Interests. The authors declare no competing interests.

What is Systems Theory: Core Concepts for Understanding and Teaching Systems Theory and Cybernetics in Further Education

Stefan N. Groesser

Bern University of Applied Sciences
School of Engineering and Computer Sciences
stefan.groesser@bfh.ch

Extended Abstract

Keywords: Systems Thinking · Cybernetics · Systems Theory · Systems Education · Transdisciplinary Learning · Reflexive Systems

Multiple recent studies support the view that traditional disciplinary approaches are inadequate for addressing the complexity, uncertainty, and interconnectedness of contemporary challenges, and that systems thinking and cybernetics are essential for educational innovation and transdisciplinary learning [1]. Research highlights how cybernetic principles, such as feedback loops and adaptability, can transform education by making curricula more responsive to rapid changes in technology and labor markets, and by fostering collaboration between academia, industry, and policymakers to equip students with future-proof skills [2],[3]. Systems thinking and cybernetics, components of the field known as systems theory, are also shown to enhance the development of transversal competencies like creativity, teamwork, and communication, which are crucial for navigating complex environments and are best cultivated through holistic, action-oriented learning approaches [1], [4], [5]. Systems theory provides a universal framework for analyzing, modeling, and understanding entities composed of interacting components. It is inherently multidisciplinary, offering abstract principles and mathematical tools applicable to physical, biological, social, and engineered systems [6], [7] [8]. Similarly, cybernetics contributes insights into regulation, feedback, learning, and communication processes across domains, bridging the different fields of research and sciences [9]. Historically, systems theory and cybernetics emerged as attempts to overcome the fragmentation of science and to create a unified framework for studying complex phenomena [4]. Recent academic discussions on systems approaches reveal significant divergence in how systems thinking is perceived: some participants view it as outdated or overly mechanistic, while others primarily associate it with psychological or therapeutic contexts. This diversity of perspectives highlights a persistent lack of shared understanding and consensus about what constitutes a “systemic” approach. Research confirms that the term “systemic” is interpreted in multiple, sometimes conflicting ways—ranging from technological integration and policy development to collaborative processes and personal transformation—making it challenging to define or communicate a unified systems perspective [10]. Despite the expanding use of systems thinking in various domains—engineering, management, ecology,

education, and social work—there remains a lack of coherence in how and what is taught under the label of “systemic”. No shared framework currently maps the core concepts, traditions, and educational implications of systems theory and cybernetics across disciplines. Consequently, a simple but fundamental question –“What does it mean to think systemically?” – cannot be answered consistently. This gap points to the need for a discipline-agnostic, integrative perspective: a conceptual structure that can serve as a common reference for teaching, learning, and applying systems ideas in diverse educational and professional contexts. A shared “systemic language” would enable communication across fields and support both specialization and synthesis. The objective of this study is to develop a conceptual framework, referred to as the Systems Navigator, that synthesizes key concepts from different traditions of systems theory and cybernetics. The aim is to reconstruct the original intention of systems education: to serve as a meta-discipline capable of integrating knowledge across scientific, professional, and societal domains. The research is explorative and indicative rather than confirmatory. It seeks to identify, analyze, and systematize the conceptual elements. By mapping the conceptual landscape of systems theory and cybernetics, this paper contributes to the creation of a “Systems Navigator” – a framework that identifies both the commonalities and the distinct contributions of different systemic traditions. The outcome is intended to (1) provide educators and learners with an overview of key conceptual elements, (2) support integrative curriculum design across disciplines, and (3) foster transdisciplinary dialogue in line with the foundational vision of the systems movement. Ultimately, the purpose is to restore systems theory’s original ambition: to overcome the fragmentation of knowledge by providing a holistic approach to inquiry, design, and action.

References

1. Choukran, H., Ghneim-Herrera, T. Advancing sustainable futures through education: A Montpellier case study on student-led advocacy for transdisciplinary approaches. *Journal of Adult and Continuing Education*, 31(1) (2025)
2. Oberer, B., Erkollar, A. Leadership 4.0: Digital Leaders in the Age of Industry 4.0. *International Journal of Organizational Leadership*, 7, 404–412 (2018)
3. Reséndiz-Castro, M. et al. Transdisciplinary Cyber-systemic Design of Instruments to Measure Academic Performance in Middle and Higher Education Systems. *Systemic Practice and Action Research*, Springer, 35(3), 395–440 (2022)
4. Schwaninger, M. Intelligent organizations: an integrative framework. *Syst. Res.*, 18, 137–158 (2001)
5. Štrukelj, T. et al. A Strategic Model for Sustainable Business Policy Development. *Sustainability*, 12(2):526 (2020)
6. Adams, K.M. et al. Systems Theory as the Foundation for Understanding Systems. *Syst. Engin.*, 17: 112–123 (2014)
7. Caws, P. General Systems Theory: Its Past and Potential. *Syst. Res.*, 32 (2015)
8. Meehan, E. J. Explanation in Social Science: A System Paradigm. Homewood, Ill.: The Dorsey Press, pp. 128. (1968)
9. Lissack, M. Requisite Variety, Cognition, and Scientific Change. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 42 (2020)
10. Midgley, G., Lindhult, E. A systems perspective on systemic innovation. *Systems Research and Behavioral Science*, Wiley Blackwell, 38(5), 635–670 (2021)

First Empirical Evidence of Human Decision-Making in an Interactive Water–Food Simulator

Martin FG. Schaffernicht^{1[0000-0002-3656-1059]} and Stefan N. Groesser^{2[0000-0003-3412-2711]}

¹ Facultad de Economía y Negocios, Universidad de Talca, Av. Lircay, 3461334 Talca, Chile
martin@utalca.cl

² Bern University of Applied Sciences, Quellgasse 21, 2501 Biel, Switzerland

In a previous study [1], we addressed water scarcity from the perspective of dynamic decision-making. UNESCO reported that agriculture accounts for $\frac{3}{4}$ of the total and 50% of the water for agriculture is drawn from groundwater. Human water use is expected to continue to increase [2]. Because agriculture's supply function, we argue that consumer choice and demand are important drivers for agricultural agents to decide what to produce. Water is a comparatively cheap resource for agriculture. Therefore, water may be overused to achieve profit. We conjectured several possible origins of misrepresentations leading to production decisions with unsustainable consequences:

PC 1: Groundwater is not considered because it is invisible.

PC 2: The slowness of water infiltration towards groundwater was not considered.

PC 3: Producers prioritize the short-term.

We developed an interactive simulator containing an agricultural decision problem: players must achieve the greatest sustainable profit from producing almonds over ten rounds of decisions: *planned production* and *planned pumping of groundwater* for irrigation. They can monitor the effects on the important system variables. Independently from player decisions, *precipitation* decreases from year to year, *diminishing surface water*. Decreasing water supply reduces what is sustainable.

In the current contribution, we present and discuss the first dataset generated with 304 interpretable runs carried out by 32 participants in a master engineering course on Modelling and Simulation in Switzerland. Participants carried out between 1 and 25 instances of the game (“runs”). Runs per participant varied between 1 run (3 individuals) and 15 or more (6 individuals). The average time per run was between 10 seconds and 20 minutes. We found no interpretable relationship between the number of runs and time per run. 10 individuals' time per run was below 1 minute, leaving less than 6 seconds per decision round.

We interpret each respective last run as the submitted solution, and the previous runs as part of the search phase. In the last runs, 14 individuals maintained constant values in both decision variables, and six individuals maintained one decision variable constant, choosing the highest allowed values in 8 of the 14 cases. We found no interpretable correlation between the number of runs and whether the values of the decision variables were constant in the final run. Most of the 272 “search” runs were with constant values for both decision variables (216 for planned production and 201 for planned pumping).

Turning to the outcomes, we focus on *accumulated profits* and *water reserves* as the fundamentals for the double goal of achieving the highest sustainable profit. Over the complete runs, the results spread over a wide range of values: *Accumulated profit* went

from \$47,500 to \$551,047, and *water reserves* ranged between 353,553 and 1,427 m³. It stands to reason that if *production* generates *profit* and consumes water, higher *profit* comes with decreased *water reserves*. The negative relation between *profit* and *water reserves* is apparent in our dataset ($y = -0.9441x + 498171$, $R^2 = 0.8618$).

Since participants do not receive instructions on how to choose between *profits* and *sustainability*, each participant makes their personal choice. With respect to PC 3, these students did not systematically prioritize *profit* over *sustainability*. Only half of the players accounted for the ongoing decrease in precipitations. To use a financial metaphor, if *water reserves* are like a cash balance, and minding *sustainability* is analogous to caring for the cash balance, *production* should react to decreasing *precipitation* by a reduction. Indeed, a fraction of the students diminished their *production* over time. If this was in response to decreasing *precipitation* is not established yet. However, many students did not reduce their consumption (or not enough): in 154 of the 304 complete runs, the accumulated water consumption surpassed the accumulated precipitation.

We also looked at pumping. Players who value sustainable water use will avoid consuming water needlessly. Water consumption for almond production is needed to make profit, but the pumped water is lost from groundwater, which is part of the water reserves; it may be replaced by infiltration from the surface to groundwater, but decreasing precipitation implies diminished infiltration, and during the long infiltration process, infiltrated water is not available in the groundwater. Accordingly, there are reasons for disciplined pumping, especially under decreasing precipitation. During the 304 complete runs, pumping increased in 62 cases, decreased in 61 cases, and was kept constant in the remaining 181 cases. Most runs show no trace of accounting for loss of groundwater (PC 1) or the slowness of infiltration (PC 2). However, if misperception of groundwater or of the infiltration delay are the causes remains to be answered. Players' decisions reveal the consequences of their understanding and reasoning. Patterned relationships between decisions and system behavior may suggest certain courses of reasoning but are not in themselves evidence of it. In the current phase of this research line, we use eye-tracking to elicit data on what a player is attending to before recording their next decision.

In this presentation, we will show important details of the findings outlined in this abstract and focus on how the participants' pumping decisions respond to decreasing precipitation. There are two extreme stances: one can sustain a high production level and increase pumping, or one can maintain a low pumping flow and decrease planned production. The 32 participants were on the line between these stances and will be discussed at the conference.

References

1. Schaffernicht, M. and S. Groesser. *Interacting with the Water Cycle - Towards an Experimental Paradigm*. in *19th International Conference on Computer Aided Systems Theory*. 2024. Las Palmas, Spain: Springer Nature.
2. Khan, Z., I. Thompson, and C.R. Vernon, *Global monthly sectoral water use for 2010–2100 at 0.5° resolution across alternative futures*. *Scientific Data*, 2023. **10**.

Global Warming: A Systemic Strategy for Weathering Mammoth Complexity

Markus Schwaninger¹ and Stefan Ott²

¹ University of St. Gallen

² School of Management, Eastern Switzerland University of Applied Sciences

1 Extended Abstract

The purpose of this study is to propose a systemic framework for dealing with global warming effectively. Our research question is: "How can society bring about a holistic approach to mastering the complexity of global warming it is confronted with?"

The evidence of global warming is overwhelming. The World Meteorological Organization (2025) states that 2024 was the hottest year in its 175-year observational record. The International Panel on Climate Change keeps updating the world about continually rising temperatures in the atmosphere.

The dominant approaches to coping with climate change are one-dimensional. The focus is on combating greenhouse gases. In contrast, we advance a more holistic, multi-layered perspective in this paper. Our study relies on the dimensions of scope and scale with their sub-dimensions. The dimension of scope embraces the two types of strategies – mitigation M and adaptation A – for coping with climate change.

Strategies M and A are considered to manifest themselves in different expressions and intensities:

$$\text{Scope: } (M, A), M = (m_1, \dots, m_p), A = (a_1, \dots, a_r), \quad (1)$$

with m_i, a_j : Intensity of the expressions of the sub-dimensions.

Mitigation aims to diminish and eventually stop the growth in temperature of the atmosphere that way, by diminishing the emission of greenhouse gases. Adaptation is the strategy of humanity taking measures to live with the effects of global warming, and ultimately to survive. Evolutionary history tells us that adaptation is related to the fundamental capability of humans to cope with catastrophes and changes in general: developing new skills by means of learning, intuition, fantasy, intelligence and the will to survive.

A set of specific measures for both M and A strategies is presented. The two kinds of strategy are complementary. They produce synergies, covering a territory which promises extending the effectiveness of measures dedicated to coping with climate change.

The dimension of scale leverages the cybernetic principle of recursion to organize efforts and actions at multiple organizational levels to addressing the complexities of global warming.

Scale has a set of sub-dimensions, which range from the smallest unit to the largest unit of reference, – individual, organization, municipality, region, etc., to world. The reference units of a given level are embedded in units of the next higher level, and they are at the same time composed by the units of the lower level:

$$\text{Scale: } R_1 \subset R_2 \subset \dots \subset R_n, \quad (2)$$

with R_i denoting the reference systems on levels i ($i=1, \dots, n$). Normally, the reference system consists of $k > 1$ units.

$$R_i: R_{i,1}, \dots, R_{i,k}$$

The organization-in-focus copes with complexity along the fronts where it arises. Recursive organization enables a system to cope with huge complexities. The duet of M and A strategies is applicable to each one of the levels of a recursive system. The two strategies M and A open a spectrum of responses that is wider than the recipes for “fighting” and “defeating” climate change, as they are oftentimes given. They are the components of an ambidexterity that embraces both combat and protection.

Next, we introduce the dimension of time (t) as a third axis that captures the dynamic evolution of interactions among the levels of recursion. Thus, sustainability processes can be represented as recursive functions evolving over time intervals.

In continuous time, the dynamics of each level of recursion are represented by a delay differential equation:

$$dR_i(t)/dt = F(R_i(t), R_{i-1}(t-\tau_1), R_{i+1}(t-\tau_2)), \quad (3)$$

where $F(\cdot)$ denotes the feedback function that governs the rate of change of R_i over time, while τ_1 and τ_2 capture the temporal delays between actions at one level and their effects at adjacent levels. The term captures the non-instantaneous propagation of interventions across different levels of recursion. The formulation of the process in continuous time aligns with the System Dynamics paradigm, in which change rates accumulate into stocks representing the evolving state of the system.

Examples include changes at the supranational level R_2 (such as CO₂ pricing), which can only be implemented in national law (R_3) after a time delay, or incentives at the middle recursion levels for stopping deforestation and monocultures, replanting forests, and restoring peatlands, which are implemented at lower levels after a time delay. This represents top-down influences. Feedback is also possible in the opposite direction, for example when incentives prove to be insufficiently effective at lower levels and are not sufficiently implemented, which in turn – with a time lag – can prompt higher levels to make corrections.

The framework presented here introduces a new perspective and strategies that chart sustainable paths into the future. If implemented, it will contribute to curbing global warming and strengthening humanity's resilience in the face of climate change.

Systems Thinking and Symbols: Reflections on Symbolic Qualities of System Archetypes

Meike Tilebein and Li Zheng

University of Stuttgart, Institute for Diversity Studies in Engineering, Germany

Keywords: Systems Thinking, System Archetypes, Symbols

Systems thinking is considered as essential for solving complex problems. In a very generic way it can be defined as “a cognitive paradigm with which people come to perceive themselves and the world to be dynamic entities that display continually emerging patterns arising from the interactions among many inter-dependent connecting components” (Randle and Stroink 2018, p. 646–647).

Within the broad range of different systems thinking and systems approaches, System Dynamics, with its different tools and methods, has the potentials to foster understanding of complex systems in numerous different application areas and to be applied by a broad variety of people, from schoolchildren to decision makers in science and industry (Schwaninger 2009).

From the System Dynamics perspective, and going more into detail, Richmond (1993) identifies seven distinct systems thinking skills, of which the first three, (1) dynamic thinking as thinking in behavioral patterns rather than in events, (2) closed-loop thinking as seeing the feedback loop structure of a system rather than thinking in one-way relations, and (3) generic thinking as approaching systems from a generic view to first see their commonalities rather than their specifics, are addressed by the qualitative modeling approach within System Dynamics, causal loop diagramming.

Causal loop diagrams (CLDs) are graphical-verbal models that form a dynamic hypothesis about the origin of the dynamic phenomenon under consideration. They describe the inherent structures or mechanisms that drive the behavior of the respective system. A focus here is on feedback loops. The two basic loop types that CLDs consist of are balancing and reinforcing loops. Balancing loops (or negative feedback loops), which stabilize change, are often indicated by using a B, a minus, or a scale icon. For reinforcing loops (or positive feedback loops), which amplify change, an R, a plus, or a snowball icon are used. With

the icons symbolizing the generic nature of the respective loop, their behavior can be easily understood. However, usually more complex CLDs are needed to model a specific system, in which various and intertwined forms of the two basic loop types can be found, and the resulting overall behavior cannot be deduced. Apart from CLDs that describe specific problems, there are also more generic CLDs consisting of two or more loops, the archetypes. Each of the archetypes describes the common mechanisms behind a dynamic phenomenon that, as a behavioral pattern over time, can be found similarly in many different real-world contexts. Moreover, each archetype comes with generic hints on how to influence its problematic dynamics. (Kim 1992, Senge 1990)

With the hypothesis that the symbol concept can further enrich systems thinking, our guiding research question in this paper is whether we can extend the symbols concept from basic loops to the system archetypes. This question is embedded in the overarching aim to foster understanding of more complex CLDs. To answer this question, we build on ideas derived from the concept of philosopher Ernst Cassirer and his influential work “The Philosophy of Symbolic Forms” (1923, 1925, 1929). In this work, Cassirer conceptualizes human beings as most fundamentally “symbolic animals,” interposing systems of signs or systems of expression between themselves and the world (Friedman 2023).

We explore the System Dynamics literature for initial evidence for our hypothesis, report our preliminary findings and derive avenues for further research.

References

1. Friedman, M. (2023). Ernst Cassirer, In: Zalta, E.N., Nodelman, U. (eds.), The Stanford Encyclopedia of Philosophy (Fall 2023 Edition), <https://plato.stanford.edu/archives/fall2023/entries/cassirer>
2. Cassirer, E. (2020). The Philosophy of Symbolic Forms. Three Volume Set. Routledge
3. Kim, D. H. (1992). Systems Archetypes I - Diagnosing Systemic Issues and Designing High-Leverage Interventions. Pegasus Communications
4. Randle, J. M., Stroink, M. L. (2018). The Development and Initial Validation of the Paradigm of Systems Thinking. *Systems Research and Behavioral Science*, 35(6), 645–657. <https://doi.org/10.1002/sres.2508>
5. Richmond, B. (1993). Systems thinking: Critical thinking skills for the 1990s and beyond. *System Dynamics Review*, 9(2), 113–133. <https://doi.org/10.1002/sdr.4260090203>
6. Schwaninger, M. (2009). System Dynamics in the Evolution of the Systems Approach. In: Meyers, R. A. (ed.), Encyclopedia of Complexity and Systems Science, 753–766. Springer
7. Senge, P. (1990). The Fifth Discipline: The Art and Practice of the Learning Organization. Doubleday Currence

Exploring Predictors of Employee Inclusion in Hybrid Teams: A Machine Learning Analysis of Complex Work Dynamics

Susanne Kerschbaumer^{1[0009-0006-8224-3154]}, Stephan A. Boehm^{2[0000-0001-9220-7241]}, and Martina Hartner-Tiefenthaler^{1[0000-0003-1597-3975]}

¹ Institute of Management Science, TU Wien, Austria

² Center for Disability and Integration, University of St. Gallen, Switzerland
suzanne.kerschbaumer@tuwien.ac.at

Keywords: Hybrid Work, Workplace Inclusion, Sociotechnical Systems.

Hybrid work has reshaped how teams collaborate, introducing new complexities as employees work across organizational, geographical, cultural, and temporal boundaries—trends already noted two decades ago but now intensified through changes in the work environment during and after the COVID-19 pandemic [1]. Defined as a combination of remote and in-office work [2], hybrid work requires organizations to balance flexibility and collaboration to ensure that all employees feel like valued members of their teams. Deloitte's 2022 global report found that 80% of companies now offer some form of remote or hybrid work [3], illustrating that hybrid work is here to stay. While remote work offers benefits, including new opportunities for historically marginalized employees [4], it can also pose risks such as isolation, which affects employee outcomes such as job satisfaction [5]. As a result, fostering employee inclusion is more critical – yet also more challenging – than ever.

Hybrid teams can be conceptualized as sociotechnical subsystems in which social dynamics (e.g., trust, relatedness, leadership) and technology-enabled ways of working (e.g., digital communication practices) jointly shape emergent team outcomes [6–7]. Employee inclusion is influenced not by isolated factors but by the interaction of multiple elements, including nonlinear and sometimes unpredictable relationships. This framing aligns with work-system perspectives highlighting the complexity inherent in distributed and boundary-crossing work arrangements [1]. Against this backdrop, understanding inclusion in hybrid teams requires attending to the interplay of multiple interdependent factors rather than relying solely on linear, additive predictors.

Workplace inclusion has long been recognized as a key mechanism linking workforce diversity to positive organizational outcomes [8]. Shore et al.'s [8] model conceptualizes inclusion through belongingness and uniqueness and outlines several antecedents of inclusion. However, given the profound transformation of work arrangements in recent years, new predictors may have emerged, and existing models may require adaptation. Although research on hybrid-team inclusion has grown since the COVID-19 pandemic, findings remain scattered, often based on qualitative studies or small samples. What is missing is an overarching, quantitative investigation that synthesizes prior insights while accounting for complex, nonlinear patterns characteristic of hybrid work.

To address this gap, we examined predictors of employee inclusion in hybrid teams by integrating previous findings with a machine-learning-based exploration using a large organizational dataset ($N \approx 1,400$) collected in collaboration with a German manufacturing company. Machine learning approaches such as random forests support inductive exploration and enable the detection of non-linear relationships and higher-order interactions among predictors [9]. The dataset included 58 potential predictors spanning leadership, teamwork, well-being, sociodemographics, and hybrid-work-

specific factors. We used random forests with cross-validation, calculated feature importance values, derived Accumulated Local Effects (ALE) plots to visualize predictor influence, and computed H-statistics to quantify predictor interactions.

Our findings show that a wide variety of factors contribute to employees' feelings of belongingness and authenticity in hybrid teams. Among the most important predictors were intrateam trust, diversity in perspectives, justice, relatedness, autonomy, virtual leadership, and shared mental models regarding hybrid work. Many of these predictors exhibited non-linear effects, such that increases in belongingness or authenticity were small at low levels of the predictor but became considerably larger once a certain threshold was exceeded. Additionally, several predictors displayed substantial interaction effects, underscoring the complex, interdependent nature of inclusion in hybrid teams. Trust, relatedness, and justice were particularly interaction-dominant for belongingness, whereas diversity in perspectives, justice, and trust showed strong interactions for authenticity.

These findings align with sociotechnical principles emphasizing the interplay between the social context of teamwork and the technology-mediated conditions under which hybrid collaboration occurs. They also highlight how digitalized, dispersed work arrangements reshape the antecedents of inclusion.

Overall, this study advances current understanding of workplace inclusion by integrating prior findings with a large-scale, machine-learning-based analysis that reflects the complexity, nonlinearity, and interdependence characteristic of hybrid teams. By uncovering nuanced patterns that traditional linear models might overlook, this work offers new hypotheses for future confirmatory research and provides actionable insights for organizations aiming to foster inclusion in modern hybrid workplaces.

References

1. Carayon, P.: Human factors of complex sociotechnical systems. *Applied Ergonomics* **37**(4), 525–535 (2006)
2. Vartiainen, M.: Flexible hybrid work. Edward Elgar, Cheltenham (2024)
3. Deloitte: Global remote work survey: Remote, controlled—This changes everything. <https://www.deloitte.com/content/dam/assets-shared/docs/services/tax/2022/dttl-tax-remote-controlled.pdf>, last accessed 2025/11/10
4. Patole, S., Carpenter, R.E., McWhorter, R.R.: Digital inclusion and systemic barriers: The role of technology in shaping the participation of women of color in hybrid workspaces. *International Journal on Integrating Technology in Education* **14**(1), 39–56 (2025)
5. Gajendran, R.S., Ponnappalli, A.R., Wang, C., Javalagi, A.A.: A dual pathway model of remote work intensity: A meta-analysis of its simultaneous positive and negative effects. *Personnel Psychology* **77**(4), 1351–1386 (2024)
6. Emery, F.: Designing socio-technical systems for “greenfield” sites. *Journal of Occupational Behaviour* **1**(1), 19–27 (1980)
7. Pasmore, W.A., Francis, C., Haldeman, J., Shani, A.B.: Sociotechnical systems: A North American reflection on empirical studies of the seventies. *Human Relations* **35**(12), 1179–1204 (1982)
8. Shore, L.M., Randel, A.E., Chung, B.G., Dean, M.A., Ehrhart, K.H., Singh, G.: Inclusion and diversity in work groups: A review and model for future research. *Journal of Management* **37**(4), 1262–1289 (2011)
9. Woo, S.E., Tay, L., Oswald, F.: Artificial intelligence, machine learning, and big data: Improvements to the science of people at work and applications to practice. *Personnel Psychology* **77**(4), 1387–1402 (2024)

Data-driven decision support in complex situations-

IRIS Integrated Reachback Information System: Case Study: Data-driven Decision Support and Sector-based Optimization with IDEA4C

IDEA4C - Model Diagnosis, Socio-Economic Impacts and Future Operations using special Integrated Assessment Scenarios to optimize a Systemic Risk Analysis

Stefan Wolfgang Pickl

DKKV, COMTESSA, Munich, Germany
stefan.pickl@unibw.de

Extended Abstract

The current crisis situations faced both the decision-makers in politics in particular as well as the players in the logistics industry in generally a multitude of complex challenges. The supply chains considered currently need special attention to be optimized on complex disruptions that affect the general security of supply [1].

As part of the project »Living Lab - Future Operations« an interdisciplinary consortium develops a trend-setting interactive visualization tool to support decision makers politics, to adapt and to explore different strategies [2]. This talk gives an overview on that project. Furthermore it presents the special sector-based approach IDEA4C as part of the IRIS project [2]:

I - Identification of critical regions, sectors and coupling principles

Cost-Benefit Analysis for critical sectors.

D - Data-driven constraint optimization

Coupled Sector-based Models

E - Exploratory Strategies/ Recognition of critical pathways and sectors

Characterization of Pathways

A - Adaptation and Quantifying Analysis

Coordination of Strategies

This 4C-approach describes a new integrated modelling suite for developing and assessing relief distribution strategies to support complex disaster management problems [2, 3].

This hybrid optimization framework as example for IRIS has been conceptualized and is currently being developed with and for experts in the field of disaster and emergency management, in order to tackle the real issues arising during such this crisis [3].

References

1. Wang, Z., Nistor, M.S., Pickl, S.W.: Analysis of the Definitions of Resilience. *IFAC-PapersOnLine* 50(1), 10649–10657 (2017). doi:10.1016/j.ifacol.2017.08.1756
2. Pickl, S.: Integration of Human Factors' Aspects into Multilayered Reachback Operations via IRIS Integrated Reachback Information System. In: *Advances in Human Factors, Software, and Systems Engineering* (Conference Proceedings), AHFE Open Access (2018). doi:<http://doi.org/10.54941/ahfe100128>
3. Raap, M., Meyer-Nieberg, S., Pickl, S., Zsifkovits, M.: Aerial Vehicle Search-Path Optimization: A Novel Method for Emergency Operations. *Journal of Optimization Theory and Applications* 172(3), 965–983 (2017). doi:10.1007/s10957-016-1014-y

Federated Machine Learning for Risk Prediction in Nursing Care

Elisabeth Mayrhuber¹[0009-0006-0471-9997], Philip Stampfer²[0000-0002-0019-2140], Sai Pavan Kumar Veeranki³[0000-0003-2009-6428], and Stephan Winkler¹[0000-0002-5196-4294]

¹ University of Applied Sciences Upper Austria, Hagenberg, Austria

² Joanneum Research Forschungsgesellschaft mbH, Graz, Austria

³ Steiermärkische Krankenanstaltengesellschaft mbH, Graz, Austria

Nursing care presents unique challenges for development of artificial intelligence (AI) due to fragmented, heterogeneous data and the critical nature of decisions. The N!CA project addresses these challenges by harmonizing nursing data from different health care providers, including hospitals and nursing homes, into the Observational Medical Outcomes Partnership (OMOP) Common Data Model (CDM), enabling standardized feature extraction and model development. Prediction models to identify patients at high risk of dehydration [6] and pressure injuries [1] are developed to support nurses in delivering efficient and personalised care. By focusing on transparency, the project ensures that the model outputs are interpretable for clinical staff, fostering trust and usability in real-world care settings.

To enable robust and privacy-preserving AI development in nursing care, data from hospital and nursing home settings were anonymized or pseudonymized, depending on the healthcare provider, and harmonized using OHDSI tools [3], including WhiteRabbit for profiling, Rabbit-in-a-Hat for structural mapping, and dbt for the implementation of the ETL pipeline. We established semantic mappings to apply OMOP standard vocabularies by using existing mappings (e.g. for medication), and creating mappings from a proprietary coding system to SNOMED-CT for nursing diagnoses and interventions. The harmonized data was transformed into the OMOP CDM, facilitating standardized feature engineering across institutions. Previous research has demonstrated the effectiveness of machine learning algorithms for predicting pressure injuries in clinical settings, achieving high performance metrics (e.g., pooled AUROC ≈ 0.94) across multiple studies [5]. These models typically leverage structured clinical data but lack interoperability and standardization across institutions [2]. Feature engineering included cohort definition, temporal predictors, and clinically meaningful aggregations. Predictive models are developed using the PatientLevelPrediction (PLP) framework [8]. As part of this work, we are extending PLP by implementing symbolic regression, a transparent modeling technique that generates human-readable mathematical expressions [4]. This addition supports interpretable model training, which is essential in the critical domain of nursing care [7]. Model training is conducted in a federated learning setup, allowing institutions to train locally on while aggregating model parameters centrally. This approach avoids the transfer of raw data, to ensure data privacy.

As of now, model development is ongoing. The harmonized dataset has been prepared using the OMOP CDM, and feature engineering pipelines are in place. We are currently implementing symbolic regression within the PLP framework to enable interpretable model training. The models will target risk prediction for dehydration and pressure injuries in approximately 20,000 patients. Future work includes training and validating models in a federated setting, involving several nursing homes and hospitals from two different organizations. This work establishes the foundation for interpretable AI in federated machine learning in nursing care. By harmonizing data across institutions and extending the PLP framework with symbolic regression, our aim is to develop transparent and clinically relevant models. The use of the OMOP CDM ensures reproducibility and supports external validation, while federated learning preserves privacy. Future evaluations will focus on model performance, explainability, and integration into nursing workflows to support real-world decision-making.

References

1. Chung, M.L., Widdel, M., Kirchhoff, J., Sellin, J., Jelali, M., Geiser, F., Mücke, M., Conrad, R.: Risk Factors for Pressure Injuries in Adult Patients: A Narrative Synthesis. *International Journal of Environmental Research and Public Health* **19**(2), 761 (Jan 2022). <https://doi.org/10.3390/ijerph19020761>
2. Gutheil, J., Stampfer, P., Kramer, D., Schrempf, M., Jauk, S., Veeranki, S.P.K., Mrak, P., Pinter, G., Feichtner, F.: It's All About Specific Features: Development of Prediction Models for Key Geriatric Risks Across Three Hospital Providers. In: *Studies in Health Technology and Informatics*. IOS Press (Apr 2025). <https://doi.org/10.3233/SHTI250186>
3. Henke, E., Zoch, M., Peng, Y., Reinecke, I., Sedlmayr, M., Bathelt, F.: Conceptual design of a generic data harmonization process for OMOP common data model. *BMC Medical Informatics and Decision Making* **24**(1), 58 (Feb 2024). <https://doi.org/10.1186/s12911-024-02458-7>
4. Kronberger, G., Burlacu, B., Kommenda, M., Winkler, S.M., Affenzeller, M.: Symbolic regression. CRC Press, Taylor & Francis Group, Boca Raton, first edition edn. (2025)
5. Ma, Y., He, X., Yang, T., Yang, Y., Yang, Z., Gao, T., Yan, F., Yan, B., Wang, J., Han, L.: Evaluation of the risk prediction model of pressure injuries in hospitalized patient: A systematic review and meta-analysis. *Journal of Clinical Nursing* **34**(6), 2117–2137 (Jun 2025). <https://doi.org/10.1111/jocn.17367>
6. Ogbolu, M.O., Eniade, O.D., Kozlovszky, M.: Systematic Review and Meta-Analysis of Risk Factors for Dehydration and the Development of a Predictive Scoring System. *Healthcare* **13**(16), 1974 (Aug 2025). <https://doi.org/10.3390/healthcare13161974>
7. Polat Erdeniz, S., Veeranki, S., Schrempf, M., Jauk, S., Ngoc Trang Tran, T., Felfernig, A., Kramer, D., Leadolter, W.: Explaining Machine Learning Predictions of Decision Support Systems in Healthcare. *Current Directions in Biomedical Engineering* **8**(2), 117–120 (Sep 2022). <https://doi.org/10.1515/cdbme-2022-1031>
8. Schuemie, M., et. al.: Health-Analytics Data to Evidence Suite (HADES): Open-Source Software for Observational Research. In: *Studies in Health Technology and Informatics*. IOS Press (Jan 2024). <https://doi.org/10.3233/SHTI231108>

Assessing ECG signal quality via patient-specific embedding similarity

Martina Pletz, Alexander Maletzky^[0000-0003-4378-7854], and Michael Giretzlehner^[0000-0003-1620-1002]

RISC Software GmbH, Softwarepark 32a, 4232 Hagenberg, Austria
martina.pletz@risc-software.at

Assessing the usability of biomedical signals relies on accurately estimating their quality using signal quality indices (SQIs). Especially in the medical field, poor signal quality may lead to errors or inappropriate clinical decisions [1]. In non-ICU settings (e.g. the general ward), ECG recordings are often affected by various artifacts, such as motion artifacts from walking or muscle artifacts due to sudden body movement [1]. These artifacts degrade signal quality, making it essential to assess signal quality in these environments. To that end, we propose a novel approach that employs a subject recognition model for comparing a given query snippet to high-quality reference snippets from the same patient (e.g., from a prior ICU stay).




Fig. 1. When comparing ECG snippets from the same patient, a subject recognition model is applied to assess signal quality. ECGs acquired in ICU and in the general ward may differ in lead configuration and quantity. The similarity score is hypothesized to correspond to the signal quality.

Methods. A subject recognition model was trained on the public Code-15% dataset to distinguish patients based on their ECGs. Building upon the framework proposed by Seják et al. [2], the neural network architecture consists of an embedding module and a discriminator head, processing 10-second ECG signals sampled at 500 Hz and preprocessed as described in the original paper. The embedding module projects ECG signals into a latent space in which signals from the same patient are positioned proximally, while the discriminator assesses whether a given pair of ECG embeddings originates from the same subject. We used their model as a foundation for signal quality assessment. Specifically, when applied to ECGs known to originate from the same patient, we hypothesize that the similarity scores produced by the model serve as indications for signal quality.

Furthermore, two key technical challenges had to be solved to accommodate our use case: handling asymmetrical lead configurations and comparing multiple ECG snippets per patient. A key feature is its adaptability to different numbers of ECG leads (which correspond to input channels), enabling its use in both ICU and general wards, where the (number of) leads often differ. Lead configuration can be specified independently during training and evaluation, and we investigated strategies for handling asymmetrical lead inputs, for example, filling non-existing leads with zero or using only intersections of leads. Additionally, the model supports training with multiple reference snippets, which, in our use case, allows to account for varying conditions and ECG morphologies during the ICU stay. We researched different strategies to compare the embedding vectors of several ECGs with each other.

We evaluated the subject recognition performance of our models on the public datasets MIMIC-IV-ECG, MIMIC-III-WDB and PTB-XL.

Results. The subject recognition model demonstrates robust performance across various ECG lead configurations and snippet strategies. Identification accuracy is high when using the full 12-lead setup, and the model retains strong performance under asymmetrical lead conditions and multiple reference snippets as well (AUROC for single-lead multi-snippet per dataset between 0.87 and 0.96). To evaluate the model's behavior with signal degradation, synthetic noise (e.g. powerline interference, linear drift) was introduced to the query ECGs to simulate real-world conditions. Analysis of the model's similarity scores revealed a positive correlation with the signal-to-noise ratio (Pearson product-moment correlation coefficient per dataset between 0.19 and 0.40), supporting our hypothesis that similarity scores serve as indicators of signal quality.

Limitations and Future Work. In this work, we assume that the reference snippets from the ICU meet certain quality standards, which may not be guaranteed in real-life settings. In the future, we plan to extend the validation of our approach from synthetic noise to datasets with ground-truth signal quality labels, as well as prospectively acquired records from general wards.

Acknowledgments. This work was funded by the FFG (Austrian Research Promotion Agency) under the grants 902677 (AIMS) and 918015 (HEART). This project is also financed by research subsidies granted by the government of Upper Austria. RISC Software GmbH is Member of UAR (Upper Austrian Research) Innovation Network.

References

1. Satija, U., Ramkumar, B., Manikandan, M.S.: A review of signal processing techniques for electrocardiogram signal quality assessment. *IEEE Reviews in Biomedical Engineering* **11**, 36–52 (2018). <https://doi.org/10.1109/RBME.2018.2810957>, <https://ieeexplore.ieee.org/document/8304761/>
2. Seják, M., Sido, J., Žahour, D.: ElectroCardioGuard: Preventing patient misidentification in electrocardiogram databases through neural networks. *Knowledge-Based Systems* **280**, 111014 (2023). <https://doi.org/10.1016/j.knosys.2023.111014>, <https://www.sciencedirect.com/science/article/pii/S0950705123007645>

Retina-Focused Preprocessing for Enhanced Deep Learning Classification of Multiple Sclerosis from OCT Scans

Lorena Álvarez-Rodríguez^{1[0009-0004-5791-5663]}, Beatriz Cordón^{2[0000-0001-5470-8977]}, Elena García-Martin^{2[0000-0001-6258-2489]}, Jorge Novo^{1[0000-0002-0125-3064]}, Joaquim de Moura^{1[0000-0002-2050-3786]}, and Marcos Ortega^{1[0000-0002-2798-0788]}

¹ VARPA group, Instituto de Investigación Biomédica de A Coruña (INIBIC), Universidade da Coruña, A Coruña (Spain)

{lorena.alvarezr, jnovo, joaquim.demoura, mortega}@udc.es

² Aragon Institute for Health Research (IIS Aragon), Miguel Servet Ophthalmology Innovation and Research Group (GIMSO), University of Zaragoza (Spain)

beatrizcordon@gmail.com, egmvivax@unizar.es

Abstract. This work researches the impact of retina-focused preprocessing on OCT-based deep learning classification of Multiple Sclerosis (MS). Raw OCT B-scans were compared with anatomically preprocessed images, where the retina was segmented, flattened, and isolated from background structures. Experiments with several convolutional architectures showed that focusing on the retinal region enhances classification accuracy and consistency across repetitions, supporting the relevance of anatomically guided preprocessing in OCT-based neurological screening.

Keywords: Multiple Sclerosis · OCT · Computer Vision · Medical image processing · Deep Learning

1 Introduction

Optical Coherence Tomography (OCT) has become a valuable non-invasive imaging modality to assess retinal alterations associated with neurological disorders such as Multiple Sclerosis (MS). The retina, as an accessible extension of the central nervous system, provides structural biomarkers that can reflect neurodegenerative processes. Deep learning has shown strong potential in OCT-based disease screening, but its performance can be influenced by image artifacts, anatomical variability, and non-retinal background content.

This study researches how a retina-focused preprocessing strategy based on segmentation and curvature-correction affects the ability of deep convolutional networks to distinguish MS patients from controls. By isolating the retinal region and removing curvature and peripheral noise, the approach aims to promote more consistent representations that do not contain artifacts that could hinder training. The performance of several CNNs is compared between raw and preprocessed OCT B-scans to assess the impact of retinal isolation on screening accuracy and robustness.

2 Methodology

Two independent configurations were evaluated: (i) raw images, directly taken from the OCT cubes without any anatomical correction, and (ii) retina-focused images, where each B-scan was automatically segmented to isolate the retinal layers and subsequently flattened, thus removing the patient's curvature and masking all non-retinal regions in black. Then, several convolutional backbones were trained under both conditions. Figure 1 summarizes the proposal.




Fig. 1. Diagram of the proposed methodology.

3 Results and Conclusions

Both configurations, using raw and retina-focused OCT images, yielded promising results in the classification of Multiple Sclerosis. The best retina-isolated model (ResNet-18) archived a mean 0.774 ± 0.009 F1-score, while the same architecture with raw B-scans obtained 0.754 ± 0.031 . The analysis indicates that each preprocessing strategy provides complementary insights: while raw B-scans preserve the global retinal context, anatomically processed images emphasize layer-specific information and reduce peripheral variability which produces in some cases a slight increase in performance. These findings highlight the relevance of evaluating different retinal representations to better understand how anatomical normalization influences deep learning-based screening of neurological disorders.

Acknowledgments. This work was supported by the Instituto de Salud Carlos III (ISCIII), Government of Spain [grant numbers PI17/01726, PI20/00437, PI23/00935, RD21/0007/0022 (Inflammatory Disease Network - RICORS), FORT23/00010 (Programa FORTALECE)]; MICIU/AEI/10.13039/501100011033 and by “ERDF A way of making Europe” [grant numbers PID2023-148913OB-I00, PID2024-161024OB-I00]; the Consellería de Educación, Universidade, e Formación Profesional, Xunta de Galicia, Grupos de Referencia Competitiva [grant number ED431C 2024/33]; and by the Government of Aragon [group B23_23R and PROY_B50_24].

Tracing Clinical Pathways from Unstructured Medical Documents: A Case Study on a Standard Operating Procedure for Chest Pain*

Natasha Trajkovska^{1[0000-0001-8265-9041]}, Stefan Kastenhofer^{1[0009-0003-2295-3742]}, Simone Sandler^{1[0000-0001-8919-7804]}, Sophie Bauernfeind^{1[0009-0005-5290-3900]}, Christoph Praschl^{1[0000-0002-9711-4818]}, Oliver Krauss^{1[0000-0002-8136-2606]}, Ralf Hecktor², Daniel Herzmanek², Michael Roiss², Michael Haider², Maximilian Rechenmacher^{3[0009-0002-0444-6483]}, and Dominik Roth^{3[0000-0002-3384-3675]}

¹ Research and Development Department, University of Applied Sciences Upper Austria, Hagenberg i. M., Austria, {firstname.lastname}@fh-hagenberg.at

² Treetop Medical GmbH, Vienna, Austria,
{firstname.lastname}@treetop-medical.com

³ Department of Emergency Medicine, Medical University of Vienna, Vienna, Austria, max.rechenmacher@meduniwien.ac.at, dominik.roth@meduniwien.ac.at

Introduction: Healthcare systems aim to deliver consistent, high-quality care through standardized clinical guidelines. However, real-world adherence remains difficult due to dynamic clinical environments and inconsistent documentation. A major challenge is that essential clinical information is predominantly stored in unstructured text, making it difficult to systematically analyze care pathways and evaluate compliance with standard operating procedures (SOPs) [5]. To close this gap, we build on our prior work in semantic extraction [6] and introduce a pipeline that transforms unstructured medical documents into HL7 FHIR AuditEvents. This standardized representation allows the application of process mining (PM) and conformance checking to reconstruct real-world care pathways and quantitatively assess adherence to SOPs, demonstrated on a treatment workflow for chest-pain in emergency medicine.

Related Work: PM has increasingly been applied to healthcare settings to analyze care pathways and guideline adherence. Muñoz-Gama et al. [3] demonstrated that conformance checking can reveal both compliance with clinical recommendations and justified deviations due to emergency situations. Heidemeyer et al. [1] proposed a pipeline that transforms structured FHIR hospital data into event logs for PM. Their approach presupposes well-structured routine data within hospital systems. By comparison, we generate event logs directly from unstructured documents, allowing PM even with limited systematic data.

Methodology: We transform unstructured clinical documents into a semi-structured JSON representation capturing metadata, clinical events, observations, and key diagnostic values [6]. Based on this representation, we apply a rule-based mapping strategy to convert all extracted entities into standardized

* Supported by Austrian Research Promotion Agency FFG under project number 927650.

HL7 FHIR R5 AuditEvent resources following a defined order, because of unreliable timestamps in clinical contexts [2]. The mapping uses predefined rules that match event types, terminology cues, and laboratory values to FHIR attributes.

Procedural information such as ECG, laboratory tests, imaging, and catheterization is mapped to AuditEvents using structured fields. Clinical observations and critical decision variables (e.g., ST-segment elevation, high-sensitivity troponin values) trigger additional AuditEvents based on SOP-aligned clinical thresholds. This enables us to represent both explicitly documented procedures and clinically inferred steps relevant to chest-pain triage.

The resulting FHIR event stream serves as a canonical process log, automatically converted into XES for PM [4]. Process discovery is performed at three levels of abstraction: (i) metadata-based workflow reconstruction, (ii) content-informed process reconstruction, and (iii) value-based decision tracking.

Results: We analyzed 90 medical documents from 10 patients presenting with chest pain to reconstruct their clinical pathways at three abstraction levels. The methodology allows extracting structured clinical events from this unstructured text and transform them into FHIR AuditEvents for subsequent PM at different abstraction levels. At level (i), we recover the overall workflow and key care phases (10 patients/78 events); at level (ii), the enriched event extraction allows us to fill documentation gaps (10 patients/82 events); and at level (iii), detailed value-based information enables us to closely follow the diagnostic decision-making pathway (4 patients/17 events).

Conclusion: Using the proposed pipeline, we successfully reconstruct the full clinical trajectories of the analyzed patients, giving details into the medical treatment path with adherence to the underlying SOP. The insights gained provide a foundation for future research using larger and more diverse datasets and a deeper analysis of time-related factors.

References

1. Heidemeyer, H., Auhagen, L., Majeed, R.W., Pegoraro, M., Bienzeisler, J., Peeva, V., et al.: A pipeline for the usage of the core data set of the medical informatics initiative for process mining—a technical case report. In: German Medical Data Sciences 2024. IOS Press (2024)
2. Mans, R.S., Van der Aalst, W.M., Vanwersch, R.J., Moleman, A.J.: Process mining in healthcare: Data challenges when answering frequently posed questions. In: Int. Workshop on Process-oriented Information Systems in Healthcare. Springer (2012)
3. Munoz-Gama, J., Martin, N., Fernandez-Llatas, C., Johnson, O.A., Sepúlveda, M., Helm, E., et al.: Process mining for healthcare: Characteristics and challenges. Journal of Biomedical Informatics (2022)
4. Pointner, A., Krauss, O., Erhard, A., Schuler, A., Helm, E.: Multi-perspective process mining interfaces for hl7 auditevent repositories: Xes and ocel. In: dHealth 2023. IOS Press (2023)
5. Pradhan, S.K., Jans, M., Martin, N.: Getting the data in shape for your process mining analysis: An in-depth analysis of the pre-analysis stage. ACM Computing Surveys **57**(6) (2025)
6. Trajkovska, N., Roiss, M., Bauernfeind, S., Alnajdawi, M., Sandler, S., Herzmanek, D., et al.: Generating actionable insights from patient medical records and structured clinical knowledge. In: dHealth 2024. IOS Press (2024)

Feature-Based Clustering of Shockwave Signals for Understanding Variability in Electrohydraulic Extracorporeal Shockwave Therapy

Hannah Janout^{1,2,3}[0000-0002-0294-3585], Jonas Flatscher²[0009-0002-3147-9539], Stephan Winkler³[0000-0002-5196-4294], Paul Slezak²[0000-0002-1000-4393], and Cyril Slezak^{2,3}[0000-0003-4718-9966]

¹ University of Applied Sciences Upper Austria, 4232 Hagenberg, Austria

² Ludwig Boltzmann Institute for Traumatology, 1200 Vienna, Austria

³ Johannes Kepler University, 4040 Linz, Austria

Hannah.Janout@fh-hagenberg.at

Electrohydraulic devices (EH) used in Extracorporeal Shockwave Therapy (ESWT) display stochastic variability in the shockwave signals they generate [2]. The source of this variability is not fully understood and still underresearched, with hypotheses suggesting fluctuations in initial amplitude, variations in source position, or changes in the geometry of the plasma bubble. Understanding these mechanisms is essential for ensuring consistent therapeutic outcomes, establishing a baseline for comparing shockwave pulses across devices, and for quantitatively assessing differences between manufacturer norms (based on single-shot measurements) and the actual signals generated during standard treatments.

Direct observation of the source position or plasma bubble geometry is not possible during either laboratory or in situ measurements. Although several thousand shockwave signals were recorded and their origin positions estimated using a specialized setup, variations in plasma bubble geometry could not be captured. Early attempts to identify patterns in measured signals for source position analysis relied on shape-based algorithms such as S-Shape, dynamic time warping, and hierarchical clustering. However, excessive noise, signal length, and high variability made these approaches computationally impractical and insufficiently accurate.

We implemented high-fidelity 3D simulations, generating a noise-free dataset with known ground truth. This approach enables systematic variation of source parameters and provides clean signals for analysis and workflow optimization before returning to the measured shockwave signals. Feature-based methods were employed to reduce dimensionality and ensure interpretability, followed by clustering to detect underlying patterns in the signals. This combination offers a robust and explainable workflow for investigating the mechanisms driving variability in EH ESWT signals.

3D acoustic simulations were conducted in a homogeneous water bath to eliminate tissue-related heterogeneities and isolate source-driven effects. Controlled variations in the source generating the shockwave pulse included position shifts of up to 3 mm along all axes and modifications to the source shape, from spherical to ellipsoidal and egg-shaped geometries [1].

Feature extraction followed two strategies: automated time-series feature computation, which yielded over 600 statistical descriptors, and a custom feature set consisting of 21 domain-specific descriptors designed to capture clinically relevant characteristics such as rise time, peak count, maximum/minimum pressure, and signal variability.




Fig. 1. Feature distribution, representative signals per cluster, and spatial coordinates of shockwave origins grouped into four clusters.

Clustering was performed using k-means on the computed feature sets. The performance of the clustering was evaluated through the silhouette score and the Davies-Bouldin index, with expert-labeled subsets serving as qualitative benchmarks to compare the validity of the clustering metrics with expert knowledge.

Preliminary results indicate that variation in source position is the dominant factor influencing signal morphology, while variations in source shape had a negligible effect. Statistical tests (Kruskal-Wallis and post-hoc Dunn's tests) also revealed a significant relationship between subsets of the custom features and the individual axis directions.

This study presents a computational framework for analyzing ESWT signal variability, combining high-fidelity simulations with custom feature-based techniques tailored to adapt to characteristics prominent in shockwave signals. By identifying source position as the primary driver of stochastic behavior, these findings inform future work integrating real-life measurements with insights gained from simulations to validate and enhance these outcomes.

Acknowledgments. This research was funded by the Austrian Research Promotion Agency (project number 40079333) and Trauma Care Consult.

References

1. Narushin, V.G., Romanov, M.N., Griffin, D.K.: Egg and math: introducing a universal formula for egg shape. *Annals of the New York Academy of Sciences* **1505**(1), 169–177 (2021). <https://doi.org/10.1111/nyas.14680>
2. Ogden, J.A., Tóth-Kischkat, A., Schultheiss, R.: Principles of shock wave therapy. *Clinical Orthopaedics and Related Research* (387), 8–17 (2001). <https://doi.org/10.1097/00003086-200106000-00003>, pMID: 11400898

Towards Improving Medical Image Semantic Segmentation using Model Soups*

Joaquín Ortiz de Murua, Jónathan Heras^[0000-0003-4775-1306], and Adrián Inés^[0000-0002-6989-1636]

Dpto. Matemáticas y Computación, Universidad de La Rioja, Logroño, La Rioja, Spain {joaquin.ortiz-demurua, jonathan.heras, adrian.ines}@unirioja.es

Semantic segmentation is a computer vision task that involves assigning each pixel in an image to a predefined category or class, and nowadays is mainly approached by developing deep learning models. These models play a crucial role in medical imaging [2]; however, given the critical nature of medical applications, it is essential that each pixel of an image is correctly classified, as even a small segmentation error can lead to significant diagnostic consequences in a real-world scenario. In order to tackle this problem, several techniques are commonly used, such as K-Fold cross validation, hyperparameter optimization and testing a variety of different architectures. Nevertheless, these approaches often entail high computational costs and require the generation of intermediate models that are usually discarded if they do not obtain the best results.

Model soups [4] take advantage of intermediate models to improve their performance. This approach aims to maximize evaluation metrics by averaging the weights of multiple models trained with different hyperparameters. By doing so, no model is discarded during the process of hyperparameter optimization or K-Fold cross validation. Although model soups do not usually improve the performance of traditional model ensembles, they do not require additional memory or inference costs. In the original work [4], several fine-tuned transformer models for both text and image classification were employed, including CLIP, ALIGN and BERT. Other studies have extended this approach to semantic segmentation of Light Detection and Ranging (LiDAR) data [3]. However, to our knowledge, no prior research has shown the benefits of applying this strategy to semantic segmentation models such as traditional CNN-based architecture models or to vision transformer models applied in the medical context.

In this work, we have developed a library to facilitate the creation of model soups for semantic segmentation independently of the underlying architecture of the segmentation model. Currently, we allow two distinct strategies of intermediate models as ingredients: checkpoints saved every N epochs during training with the hold-out strategy, and models obtained from a K-Fold cross validation process. Furthermore, three “recipes” are used for model souping with the hold-out strategy *uniform soup* (the average of all fine-tuned models), *weighted soup* (giving higher weights to later, more trained models) and *greedy soup* (sequentially

* This work is partially supported by Agencia de Desarrollo Económico de La Rioja [ADER 2022-I-IDI-00015], by the Government of La Rioja [Inicia 2023/01, Afianza 2024/01], and by Grant PID2024-155834NB-I00 funded by MI-CIU/AEI/10.13039/501100011033 and by ERDF/EU.

adding models to the soup and keeping them only if the validation performance improves). In the case of ingredients from K-Fold cross validation, only the uniform recipe is available, since no independent validation split is shared across folds.

Our library for model soups has been validated in the Breast Ultrasound Image Dataset (BUSI) [1] using four different segmentation architectures (CNN-based models, UNet and FPN; and transformer-based models, MaskFormer and its evolution Mask2Former). For the hold-out strategy, we used 60% of the data for training, 20% for validation and 20% for testing; and checkpoints were saved every 10 epochs. For the K-Fold scenario, we used a 5-fold cross-validation setup, with a separate 20% test split.

The results for our experiments are shown in Table 1. After evaluating the model soups on the test split, the greedy soup consistently demonstrates a good performance with transformer-based models in the hold-out scenario, while the weighted soup also achieves strong performance with CNN-based models. Future work will explore additional datasets as well as alternative strategies and recipes.

Architecture	Hold-out			K-Fold		
	Best ingredient	Uniform	Weighted	Greedy	Best ingredient	Uniform
U-Net	0.7655	0.7556	0.7634	0.7624	0.7589	0.7230
FPN	0.7577	0.7784	0.7710	0.7727	0.7697	0.7277
MaskFormer	0.7604	0.7645	0.7605	0.7604	0.7747	0.6527
Mask2Former	0.7735	0.7590	0.7557	0.7735	0.7712	0.7725

Table 1. Comparison of the Dice coefficient obtained from the best individual checkpoints on the validation split and from model soups evaluated on the test split. Best results for each model in bold.

References

1. Al-Dhabayani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. Data in Brief **28**, 104863 (2020). <https://doi.org/https://doi.org/10.1016/j.dib.2019.104863>, <https://www.sciencedirect.com/science/article/pii/S2352340919312181>
2. Liu, X., Song, L., Liu, S., Zhang, Y.: A review of deep-learning-based medical image segmentation methods. Sustainability **13**(3), 1224 (2021)
3. Wasfy, O., Basiony, S., Torki, M.: Enhancing lidar semantic segmentation using model soups: Aiccasa23 (11 2023)
4. Wortsman, M., Ilharco, G., Gadre, S.Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A.S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., Schmidt, L.: Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time (2022), <https://arxiv.org/abs/2203.05482>

A Learning-Embedded Hyper-heuristic Framework for the Integrated Healthcare Timetabling Competition 2024

Kate Han^{1,2}  and Theodore Papamarkou² 

¹ Salford Business School, University of Salford, Greater Manchester, UK

² PolyShape Ltd, UK / PolyShape Greece

k.han3@salford.ac.uk

Workshop: Data Science in Medical and Bio-Informatics

Keywords: Healthcare Timetabling · Hyper-heuristics · Reinforcement Learning · Scheduling Optimization

Extended Abstract

The Integrated Healthcare Timetabling Competition 2024 (IHTC 2024) poses a multiresource scheduling problem involving patient admission planning, operating theater scheduling, and nurse-to-room assignments under shared temporal and resource constraints [3]. Compared to the classic academic timetabling and scheduling problem, this problem is characterized by higher complexity due to interdependent decisions across patient flows, hospital capacity, and staff rosters. Hard restrictions include gender-specific room allocation, admission deadlines, surgeon availability, and maximum working hours, while soft constraints address continuity of care, workload balance, and minimization of idle operating theater time [4].

This paper proposes a learning-embedded hyper-heuristic framework for solving IHTC 2024 benchmark instances. The proposed framework extends the reinforcement learning hyper-heuristics developed by Han et al. for examination timetabling[1,2] to a broader healthcare context, where multiple interdependent decisions and feasibility constraints interact at scale. Existing approaches for the IHTC 2024 benchmark include multi-neighbourhood simulated annealing [5] and parallelised hyper-heuristics [6], yet these methodologies rely predominantly on manually designed decision logic and do not embed adaptive learning mechanisms to guide heuristic selection autonomously.

In the proposed approach, a high-level hyper-heuristic learns over time which low-level heuristics are most effective for improving solution quality or restoring feasibility. Candidate solutions encode patient-to-room assignments, admission dates, nurse-to-patient rosters, and surgical theater schedules. Each candidate is evaluated via deterministic simulation to ensure hard constraint satisfaction, followed by a weighted penalty function to capture soft constraint violations. A reinforcement learning mechanism maintains the utility values for each heuristic and updates them according to the solution improvements, the stability of the

satisfaction of the constraint, and the search stage. Move acceptance is regulated using simulated annealing with reheating phases, thereby balancing intensification and diversification while retaining feasibility. Evolution-inspired perturbations and sequence-based neighbourhood transitions may be integrated to maintain search diversity and scalability. The proposed approach will be tested using a time limit calculated by the IHTC 2024 benchmark tool.

The expected contribution of this research lies in establishing a generalisable, learning-driven methodology for healthcare scheduling that does not require bespoke heuristic design for each problem instance. This framework is designed to autonomously adapt heuristic selection strategies to varying hospital configurations and constraint structures. Future research will include an empirical assessment of both public and hidden IHTC 2024 datasets, a sensitivity analysis of learning rate parameters, and integration with hybrid optimization strategies, such as simulated annealing and adaptive threshold search, to improve the robustness and scalability of the solution.

Acknowledgment

The authors appreciate the support of Salford Business School and PolyShape Ltd. for providing computational resources and domain expertise.

References

1. Han, K., Oginni, E., McCall, J.A.W., McMullan, P., Chen, Y.: Exam Timetabling Problem Using Cooperative Hyper-heuristics: A Case Study of the ITC2007 Dataset. MISTA (2017).
2. Han, K., McMullan, P.: Hyper-heuristics using Reinforcement Learning for the Examination Timetabling Problem. MISTA (2017).
3. Ceschia, S., Rosati, R.M., Schaerf, A., Smet, P., Vanden Berghe, G., Zanazzo, E.: The Integrated Healthcare Timetabling Competition 2024: Problem Description and Rules. In: PATAT 2024.
4. Brandt, T., Klein, T.L., Reuter-Oppermann, M., Schäfer, F., Thielen, C., van de Vrugt, M., Viana, J.: Integrated patient-to-room and nurse-to-patient assignment in hospital wards. arXiv preprint arXiv:2309.10739 (2023).
5. Zanazzo, E., Smet, P., Ceschia, S., Rosati, R.M., Schaerf, A., Vanden Berghe, G.: A Multi-Neighbourhood Simulated Annealing Approach to the Integrated Healthcare Timetabling Problem. SSRN 5376065.
6. Davison, M., Burgess, G., Hamm, R.M., Lowery, B., Page, A.: A Parallelised Hyper-heuristic Framework for the Integrated Healthcare Timetabling Competition 2024. SSRN 5601273.

Clinical Data Preparation for Process Mining in Hospitals with Routine Data

Klaus Arthofer¹ [0000-0001-6311-0088], Charles Walcher¹ [0009-0007-0856-7254],
Selina Ardlberger¹ and Tilman Königswieser²

¹ University of Applied Sciences Upper Austria

² Oberösterreichische Gesundheitsholding GmbH, Linz, Austria

Introduction: Current research highlights the potential of process mining in hospitals, yet practical implementation remains challenging [1, 2]. This work investigates how real-world clinical data can be prepared and used for process mining. Within the PICA¹ project, we developed a data management guideline, involving the Fast Healthcare Interoperability Resources (FHIR) exchange standard. We build upon prior work showing FHIR's suitability for process mining as it enables semantic interoperability of clinical data through a consistent data structure [3].

Data and Setting: We used data from LeiVMed (“Leistungsvergleich Medizin”), a medical benchmarking program developed by a research group at the authors faculty. Since 2015 LeiVMed has monitored and analyzed patient treatment data for benchmarking across nine hospitals in Upper Austria. To date, it comprises approximately 300,000 benchmarking cases, across roughly 40 medical case categories, combining structured, mostly administrative data, including patient, encounter, procedure and laboratory data into unified medical case records. For this work we focused on cholecystectomy cases as a manageable subset and excluded complicated cases to reduce heterogeneity. The resulting event log contains 4555 cases recorded between 2015 and 2023.

Methodology: This methodology is based on the above-mentioned data management guideline and the medical case scope of LeiVMed. The prepared event log contains patient ID, timestamp, activity, performing as well as requesting organizational unit and the performing hospital location. Organizational units and the hospital location are important filter criteria for constructing treatment cases in process mining. A descriptive analysis was performed to assess data scope and complexity, including the number of unique cases, total activities, and average number of activities per case, which served as an early indicator of overcomplexity, also referred to as spaghetti models [4, 5].

To ensure consistency, LeiVMed uses SNOMED CT for coding medical concepts. We coded activities during the creation of FHIR AuditEvents, using SNOMED CT

¹Process Intelligence and Conformance Auditing (PICA) is a project of the University of Applied Sciences Upper Austria. This project is financed by research subsidies granted by the government of Upper Austria.

concepts of the semantic type “procedure”. Standardized codes ensure uniform labeling and enable cross-site comparisons in process mining [7].

Our analysis reveals that the real-world data quality varies. Typical issues include missing values and incomplete or identical timestamps. Differing levels of activity label specificity (e.g. “CT” vs “CT upper abdomen”) stem from medical procedure data focusing treatment costs, which are to be considered during the data preparation step. To reduce noise, incomplete or incorrect entries were excluded, infrequent activities below a predefined threshold were pruned. We consulted domain experts to evaluate the clinical relevance of activities and to identify implausible or misleading entries for removal. Existing research likewise emphasizes their involvement [1, 2, 7].

Result: The data in the medical case examined proved promising but did not yet allow for detailed process mining, as structured routine data does not provide a full picture of clinical pathways.

Conclusion and Outlook: Available routine data offers a starting point for process mining, but many clinically relevant activities remain uncaptured. Future work should expand data extraction to include unstructured data sources or documents such as discharge letters, which contain relevant information not available in structured data. Large Language Models (LLMs) could assist in extracting information from documents and also in preparing data for process mining [8].

References

1. Munoz-Gama, J., et al.: Process mining for healthcare: characteristics and challenges. *J. Biomed. Inform.* 127, 103994 (2022). doi:10.1016/j.jbi.2022.103994
2. De Roock, E., Martin, N.: Process mining in healthcare – an updated perspective on the state of the art. *J. Biomed. Inform.* 127, 103995 (2022). doi:10.1016/j.jbi.2022.103995
3. Pointner, A., Krauss, O., Erhard, A., Schuler, A., Helm, E.: Multi-perspective process mining interfaces for HL7 AuditEvent repositories: XES and OCEL. *Stud. Health Technol. Inform.* 301, 168–173 (2023). doi:10.3233/SHTI230034
4. Van der Aalst, W.: Process mining. *Commun. ACM* 55(8), 76–83 (2012). doi:10.1145/2240236.2240257
5. Fernandez-Llatas, C., Martinez-Millana, A., Martinez-Romero, A., Benedi, J.M., Traver, V.: Diabetes care related process modelling using process mining techniques: lessons learned in the application of interactive pattern recognition coping with the spaghetti effect. *Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.* 2015, 2127–2130 (2015). doi:10.1109/EMBC.2015.7318809
6. Helm, E., Lin, A.M., Baumgartner, D., Lin, A.C., Küng, J.: Towards the use of standardized terms in clinical case studies for process mining in healthcare. *Int. J. Environ. Res. Public Health* 17(4), 1348 (2020). doi:10.3390/ijerph17041348
7. Vathy-Fogarassy, Á., Vassányi, I., Kósa, I.: Multi-level process mining methodology for exploring disease-specific care processes. *J. Biomed. Inform.* 125, 103979 (2022). doi:10.1016/j.jbi.2021.103979
8. Krauss et al., Generating Actionable Insights from Patient Medical Records and Structured Clinical Knowledge, *Studies in Health Technology and Informatics*, 74–80 (2024), doi:10.3233/SHTI240015

ASA-Specific Detection and Prediction of Complications in Post-Surgery Records of Patients

Julia Vetter¹, Marina Strobl¹, Gerhard Halmerbauer²,
Tilman Königswieser³, and Stephan M. Winkler^{1,4*}

¹ School of Informatics, Communications and Media
University of Applied Sciences Upper Austria, Hagenberg, Austria

² School of Business and Management
University of Applied Sciences Upper Austria, Steyr, Austria

³ Salzkammergut-Klinikum, Oberösterreichische Gesundheitsholding GmbH

⁴ Department of Computer Science
Johannes Kepler University; Linz, Austria

For more than ten years, the University of Applied Sciences Upper Austria and Oberösterreichische Gesundheitsholding (OÖG) have developed the platform *LeiVMed* for the monitoring and analysis of data of treatments of patients. In the current research context, our main goal is to identify models that are able to detect complications (*a posteriori*) in the data of treatments as well as models that are able to predict the risk of complications after surgeries.

Our data base consists of more than 16 000 samples storing information about the surgeries and treatments of seven different medical case classes: herniotomy, cholecystectomy, rectal surgery, prostatectomy, hip replacement surgery, and thyroidectomy. Further, for each case risk factors such as smoking, diabetes, etc. and information about eventual complications are available. We here distinguish between minor complications (such as hematoma and seroma) and major complications (such as heart attacks and pneumonia).

As described previously ([1], [2]), *LeiVMed* research is dedicated to medical data science such as the detection of complications in patient records as well as the data based prediction of the length of stay of surgery patients. In our current research, the main objective is to identify (1) retrospective and (2) predictive detection models for each medical class separately as well as all data aggregated. We see a clear correlation between the American Society of Anesthesiologists (ASA) score classification and the complication rate; as this indicates the ASA score as a strong predictor for complication probabilities, we train models for each ASA class separately.

We use machine learning (ML) for identifying models that are able to detect complications in patient records. In general, machine learning is understood as that branch of computer science that is dedicated to the development of meth-

* corresponding author

The work described in this paper was supported by *LeiVMed*, a joint research project of FH OÖ and OÖG.

ods for learning knowledge and models from given data. In *supervised machine learning* the modeling algorithm is given a set of samples with input features as well as target / output values - i.e., the data are “labeled”. The goal then is to find a model (or a set of models) that is (are) able to map the input values to the designated target variable(s). In our particular case there are two target variables, namely on the one hand whether there was a complication at all and on the other hand whether the complication was minor or major. As both these variables are binary, we are facing *classification problems* that are to be solved.

In this research we especially analyze the performance of several modern machine learning methods, namely black box methods as well as explainable AI methods. While black box methods produce models which are functions of the inputs and produce outputs, where the internal functioning of the model is either hidden or too complicated to be analyzed, explainable AI methods produce models whose structure is not hidden, but can be analyzed and understood in detail.

In detail, use a multi-model ensemble (MME) approach for model performance evaluation and apply the following machine learning methods: Symbolic regression / classification, artificial neural networks, random forests, gradient boosting algorithms, and quadratic discriminant analysis.

First preliminary results show that

- retrospective models are able to correctly classify 85.11 % of ASA-1 the samples, 81.29 % of ASA-2, 75.79 % of ASA-3, and 67.24 % of ASA-4;
- predictive models produce results that are 53.57 % correct for ASA-1 patients, 56.64 % for ASA-2 patients, 44.05 % for ASA-3, and 68.97 % for ASA-4 patient samples.

In the final paper we will give particular results, namely retrospective as well as prospective classification accuracies for the following surgery types:Herniotomy, cholecystectomy, rectal surgery, prostatectomy, hip replacement surgery, and thyroidectomy.

In future research we will focus on the analysis of variable impact factors and variable relevances; we are especially interested in how far variables that are available at the beginning or in the course of the treatment can be used for detecting complications. Furthermore, we also aim at creating models that are able to detect treatments for which anomalies (complications) can be ruled out.

References

1. Strobl, M., Vetter, J., Königswieser, T., Halmerbauer, G., Winkler, S.M.: Using explainable artificial intelligence for data based detection of complications in records of patient treatments. (Computer Aided Systems Theory – EUROCAST 2022) 173–180
2. Vetter, J., Strobl, M., Buur, L., Königswieser, T., Halmerbauer, G., Winkler, S.M.: Data-based prediction of the duration of the postoperative stay of patients. (Computer Aided Systems Theory – EUROCAST 2024) 335–344

Estimation of Glycated Hemoglobin from Continuous Glucose Monitoring data by Symbolic Regression

J. Ignacio Hidalgo^{1,3}, Esther Maqueda², Oscar Garnica^{1,3},
Wenbo Sun³, Daniel Parra¹, and J. Manuel Velasco^{1,3}

¹ Universidad Complutense de Madrid

² Hospital Universitario de Toledo

³ Bioinspired Intelligence S.L.

absys@ucm.es

Abstract. This paper presents a machine learning method for estimating glycated hemoglobin (HbA1c) levels from data collected by Continuous Glucose Monitoring (CGM) systems. Conventional approaches to determining HbA1c levels require invasive blood tests, which pose challenges such as high costs, discomfort, and limited accessibility in certain geographic areas. The search for noninvasive, economical, and efficient computational models for the estimation of HbA1c is therefore of interest and importance. Our methodology is based on Grammatical Evolution, a type of genetic programming that evolves symbolic expressions to discover models capable of estimating HbA1c levels from CGM data. GE is well known for its ability to explore various mathematical models and incorporate domain-specific knowledge via customizable grammars, making it highly suitable for this application. The study utilizes a specialized grammar designed to reflect the variables involved in diabetes monitoring, thereby enabling the generation of symbolic models that accurately represent the relationships between CGM data parameters—average glucose, median glucose values, Glucose Management Indicator (GMI), Coefficient of Variation (CV), and Time in Range (TIR)—and HbA1c levels. The research dataset includes CGM records from a reduced cohort as a preliminary proof-of-concept. Our research indicates that GE-generated models achieve remarkable precision in estimating HbA1c, outperforming GMI models.

Keywords: First keyword · Second keyword · Another keyword.

1 Introduction

Glycated Hemoglobin (HbA1c) is a form of hemoglobin chemically linked to glucose. Hemoglobin is a protein in red blood cells that carries oxygen throughout the body. When blood glucose levels are elevated over time, glucose molecules bind to hemoglobin in red blood cells through a process called glycation. The HbA1c level measures the amount of glucose that has become attached to hemoglobin

over the life of the red blood cell, typically about 120 days. Therefore, HbA1c reflects average blood glucose levels over the past two to three months. HbA1c is expressed as a percentage. For example, an HbA1c level of 7% indicates that 7% of the hemoglobin in the blood is glycated. The American Diabetes Association (ADA) provides the following guidelines for HbA1c levels. If the value is lower than 5.7%, it is considered normal; if it is between 5.7% and 6.4%, it indicates prediabetes; and if it is higher than 6.5%, the person has diabetes.

Monitoring HbA1c levels is crucial in managing diabetes. It helps assess the effectiveness of diabetes treatment plans and adjust medications, diet, and exercise to better control blood glucose levels. Unlike daily blood glucose testing, which provides a snapshot of glucose levels at a specific time, HbA1c levels offer a long-term view of glucose control. This makes HbA1c an essential tool for managing both type 1 and type 2 diabetes. Several studies have demonstrated the correlation between high Hb1Ac values and microvascular complications (retinopathy, nephropathy, neuropathy, etc.).

This paper presents a novel approach to estimating Glycated Hemoglobin (HbA1c) levels using Grammatical Evolution (GE) for symbolic regression [1]. Traditional methods for estimating HbA1c involve direct measurement through blood tests or capillary measures, which are invasive, painful, and inaccessible in resource-limited settings. As such, there is a growing interest in developing non-invasive, cost-effective, and efficient computational models for HbA1c estimation. We use structured GE, a form of genetic programming, to evolve symbolic expressions that accurately estimate HbA1c levels from readily available clinical parameters. GE's flexibility in exploring a diverse space of mathematical models and its ability to incorporate domain-specific knowledge through a customizable grammar make it ideally suited for this task. Our approach involves developing a grammatical structure tailored to the variables obtained from glucose monitoring systems, facilitating the generation of symbolic models that capture possible relationships between various clinical indicators and HbA1c levels.

The data set used in this study comprises clinical records from a diverse patient population, including fasting and postprandial blood sugar levels and other relevant clinical markers. We detail the preprocessing, GE algorithm configuration, and criteria used for model selection and validation. Our findings demonstrate that the GE models generated can provide accurate estimates of HbA1c levels, with performance metrics surpassing those of traditional regression techniques and existing computational models. The best-performing models exhibit a strong correlation with actual HbA1c measurements, highlighting the potential of GE for symbolic regression in medical applications.

References

1. Lourenço, N., Colmenar, J.M., Hidalgo, J.I., Garnica, Ó.: Structured grammatical evolution for glucose prediction in diabetic patients. In: Proceedings of the Genetic and Evolutionary Computation Conference. pp. 1250–1257. ACM (2019)

Multi-Disease Explainable Deep Learning System for Automated Chest X-ray Screening

Elena Goyanes^{1,2[0009-0001-5328-2262]}, Breogán Fernández-Moreira^{1,2}, Joaquim de Moura^{1,2[0000-0002-2050-3786]}, Jorge Novo^{1,2[0000-0002-0125-3064]}, and Marcos Ortega^{1,2[0000-0002-2798-0788]}

¹ VARPA Research Group, Instituto de Investigación Biomédica de A Coruña (INIBIC), University of A Coruña, A Coruña, Spain

² Department of Computer Science and Information Technologies, University of A Coruña, A Coruña, Spain

{elena.goyanes.gonzalez, breogan.fernandez, joaquim.demoura, jnovo, mortega}@udc.es

Abstract. This paper presents an explainable deep learning system for the simultaneous screening of multiple pulmonary diseases in chest X-rays. Ten state-of-the-art convolutional and transformer-based architectures were evaluated to identify the most effective model for multi-label classification, with the best-performing model achieving an average AUC of 0.841. The interpretability of the system was enhanced using Grad-CAM, providing visual explanations consistent with clinical findings.

Keywords: Chest X-ray, Deep Learning, Explainable AI, Medical Image Processing, Pulmonary Pathologies

1 Introduction

Chest radiography remains a fundamental and widely accessible imaging modality for detecting and evaluating pulmonary diseases such as pneumonia, pleural effusion, pneumothorax, and cardiomegaly. Despite its diagnostic importance, the increasing global incidence of respiratory diseases combined with growing patient volumes and limited availability of expert radiologists, highlight the need for automated, accurate, and interpretable diagnostic support systems. Recent advances in deep learning have demonstrated promising performance in chest X-ray analysis; however, their lack of transparency complicates clinical trust and deployment. Explainable AI methods, such as Gradient-weighted Class Activation Mapping (Grad-CAM), allow visualization of model reasoning and help clinicians to validate the reliability of automated predictions.

In this work, we propose an explainable and fully automated system for multi-disease screening in chest X-rays. A comparative study of ten state-of-the-art deep learning architectures was conducted to identify models that balance accuracy and interpretability, aiming to support real-world clinical decision.




Fig. 1. Examples of Grad-CAM visual explanations for representative pathologies.

2 Methodology

Experiments were conducted using the NIH ChestX-ray14 dataset, comprising 112,120 radiographs from 30,805 patients annotated with 14 thoracic pathologies. A patient-wise data split was applied to ensure independent training, validation, and testing. Images were resized and center cropped to 224×224 pixels, normalized, and augmented through horizontal flips and random resize crop to enhance generalization. Ten representative architectures were trained under a unified multi-label classification protocol: ConvNeXt, DenseNet, EfficientNetV2, MaxViT, MNASNet, ResNet, ShuffleNet, VGG, ViT, and WideResNet. Model interpretability was evaluated through Grad-CAM heatmaps (Fig. 1), highlighting the regions contributing to each pathology prediction. Classification performance was assessed by per-class and average area under the ROC curve (AUC).

3 Results and Conclusions

Across the evaluated architectures, MaxViT achieved the highest average AUC (0.841). Grad-CAM visualizations showed that the best-performing networks focused on clinically relevant pulmonary regions, with MaxViT providing the most coherent and interpretable activation maps. Overall, the system demonstrated that attention-based architectures offer an effective balance between accuracy and explainability, supporting their potential for reliable and interpretable computer-aided diagnosis. Future work will focus on external validation with expert-annotated datasets to enhance clinical applicability and reliability.

Acknowledgments. This work was supported by Government of Spain through the research projects with [grant number PID2023-148913OB-I00] and [grant number PID2024-161024OB-I00], funded by MICIU/AEI/10.13039/501100011033 and by ERDF, EU; Consellería de Educación, Universidade e Formación Profesional, Xunta de Galicia, Grupos de Referencia Competitiva, [grant number ED431C 2024/33], predoctoral grant [grant number ED481A-2023-152]. Also supported by the ISCIII under the grant [FORT23/00010] as part of the Programa FORTALECE of MICIU.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Data Science Approaches to Health Workforce Planning: A Scoping Review and Taxonomy

Hakim Guenther¹, Katja Zeume¹, and Laura Anderle¹[0000-0002-4225-4741]

Westphalian University of Applied Sciences, Gelsenkirchen, Germany
 {hakim.guenther, katja.zeume, laura.anderle}@w-hs.de
<https://www.w-hs.de/informatik-ge/>

1 Background

The global health landscape is evolving rapidly, marked by aging populations and a rising burden of chronic diseases. These changes place increasing demands on health systems and underscore the critical importance of having adequate physician workforce. However, persistent challenges — including geographic maldistribution, increasing specialization, and rising physician burnout — continue to hinder health system performance and equal access to care. [1] Social determinants of health, but also different usage behavior with regard to health care services, further increase the complexity of health care planning, especially in geographical areas (countries, regions or municipalities) with sub-areas that are subject to different social factors. Anticipating future physician workforce needs is therefore essential for informed policy making and sustainable health workforce planning [2]. Forecasting the supply and demand of physicians is inherently complex, influenced by numerous dynamic and interrelated factors such as demographic trends, policy changes, healthcare utilization patterns, and educational system capacities. Despite increasing interest in data-driven approaches—such as machine learning, simulation modeling, and predictive analytics—there has been no systematic synthesis of their application in physician workforce forecasting. Previous reviews have focused primarily on conceptual frameworks or traditional planning models, leaving a gap in our understanding of how advanced computational methods are being used in this space, and whether they are used for improving health workforce planning with regard to more complex data sources than classical approaches, for instance social determinants. Therefore, the aim of this scoping review is to systematically identify, map, and analyze data-driven approaches used to forecast the demand for physicians. In doing so, we seek to characterize the methodologies employed, examine their underlying assumptions and data sources, assess their strengths, limitations, and capability for the inclusion of social health determinants, and identify gaps to inform future research and planning efforts.

2 Material and Methods

A scoping review following the PRISMA-ScR guidelines is performed by two independent reviewers, using PubMed via MedLine and Scopus as databases,

and all discrepancies resolved. The identified studies are classified manually with regard to the depth of data they use (e.g. global population statistics, disease incidences, health service usage data), and the sophistication and variety of data science methods employed. It is evaluated whether any validation routine is provided in the publications, and how these criteria interfere.

3 Results

After 1132 papers were identified using the search strings, 45 articles matching the study selection criteria are determined, only 4 of which provided evaluation metrics of their results. Method categories employed were fixed ratio application (5 studies), rate projection (26 studies), regression-based approaches (5 studies), and time-series forecasting (10 studies). 43 papers employ at least two of the data categories (expert opinion, basic population data, demographic stratification, composite indicators, service utilization data).

4 Conclusion

Despite the complexity and heterogeneity of available data sources, only few approaches to needs-based health workforce planning rely on advanced data science methods. Most published results employ simple data and methods and provide no evaluation of their prognoses. Employing more sophisticated forecasting methods and more complex data is promising in achieving better prognoses of the need for medical personnel, which is crucial in times of demographic challenges. Additionally, a standardization of validation routines for doctoral needs prognoses would facilitate research as well as comparability of the methods employed.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Jacobzone, S.: Ageing and the Challenges of New Technologies: Can OECD Social and Healthcare Systems Provide for the Future?. *Geneva Pap Risk Insur Issues Pract* 28, 254–274 (2003), <https://doi.org/10.1111/1468-0440.00222>.
2. Roberfroid, D., Leonard, C., Stordeur, S.: Physician supply forecast: better than peering in a crystal ball? *Hum. Resour. Health* 7(1), 10 (2009), <https://doi.org/10.1186/1478-4491-7-10>.

Extending the Piloted-Beam Search with Marginal Survivors

Stefan Habringer¹, Florian Holzinger¹, and Stefan Wagner¹

Heuristic and Evolutionary Algorithms Laboratory, University of Applied Sciences
Upper Austria, stefan.habringer@fh-hagenberg.at

Effective problem-solving in large or complex search spaces relies on algorithms that can balance exploration of new possibilities with exploitation of promising candidates. For many combinatorial and structured decision-making problems, beam search and its variants provide an effective compromise between breadth and depth by maintaining a fixed-width frontier of promising partial states in the search tree. Piloted Beam Search (PBS), as described in [1], augments beam search with a domain-independent lookahead (the pilot method) to improve guidance. In certain scenarios, the search concludes prematurely when no further candidate solutions can be generated, even though additional computational budget remains, leading to an incomplete utilization of available resources. This paper extends PBS with the concept of marginal survivors to improve robustness and generalizability.

One issue with beam search is its dependence on an appropriate heuristic for estimating the quality of solution candidates [4]. This can be addressed either by introducing policies learned during runtime [2] or by using problem-independent look-ahead methods [1,5]. Conventional beam search discards nodes outside the beam and evaluates only the nodes inside the beam via its lookahead. Such rigid pruning may limit exploration and reduce the algorithm's ability to find promising solutions, especially when multiple candidate solutions have comparable quality. Adding stochasticity to the ordering of candidates, together with small perturbations in the lookahead, has been shown to improve performance [3].

To address these limitations, we propose two extensions to PBS. First, we introduce a fitness-sensitive selection strategy inspired by the "lucky losers" concept in sports, where competitors advance despite narrowly missing the initial cutoff. These marginal survivors promote broader exploration while still prioritizing high-quality solutions. The number of marginal survivors is bounded by the chosen beam width, ensuring controlled growth of the candidate pool. Compared to simply increasing the beam width, this approach offers finer control over the number of expanded nodes and is applied only when the algorithm still has computational budget after the initial PBS run completes.

Figure 1 illustrates the selection of marginal and random survivors for an example with a beam width of two. At the first decision, two candidate nodes (5 and 10) have promising heuristic estimates; at the next level, three nodes share the same heuristic value (15). By selecting all candidates with comparable heuristic estimates, a larger portion of the search tree is explored.




Fig. 1: Search tree showing heuristic look-ahead values $h(x)$. Candidates selected within a beam width of two (blue) at levels two and three. Marginal (left) and random (right) survivors at level three.

The primary objective of this paper is to evaluate the extended PBS with two extensions: (i) marginal survivors and (ii) random survivors. We evaluate these approaches on multiple instances of the knapsack problem, using both single and multi-dimensional instances. Because the pilot method is sensitive to node ordering, we also evaluate several established heuristic ordering strategies. Preliminary results indicate that marginal survivors are especially beneficial when the branching factor is large and the beam width is comparatively small. The approach enables investigation of how tree width (the number of choices at each level) and depth (the number of items that can be included in the knapsack) affect solution quality.

References

- Bindreiter, O., Werth, B., Wagner, S.: Using the pilot method as a problem-independent metaheuristic for multi-objective beam search. In: Computer Aided Systems Theory – EUROCAST 2024. pp. 61–68. Springer Nature Switzerland (2025). https://doi.org/10.1007/978-3-031-83885-9_6, http://dx.doi.org/10.1007/978-3-031-83885-9_6
- Ettrich, R., Huber, M., Raidl, G.R.: A Policy-Based Learning Beam Search for Combinatorial Optimization, p. 130–145. Springer Nature Switzerland (2023). https://doi.org/10.1007/978-3-031-30035-6_9, http://dx.doi.org/10.1007/978-3-031-30035-6_9
- Frohner, N., Neumann, B., Pace, G., Raidl, G.R.: Approaching the traveling tournament problem with randomized beam search. Evolutionary Computation **31**(3), 233–257 (2023). https://doi.org/10.1162/evco_a_00319, http://dx.doi.org/10.1162/evco_a_00319
- Sabuncuoglu, I., Bayiz, M.: Job shop scheduling with beam search. European Journal of Operational Research **118**(2), 390–412 (Oct 1999). [https://doi.org/10.1016/s0377-2217\(98\)00319-1](https://doi.org/10.1016/s0377-2217(98)00319-1), [http://dx.doi.org/10.1016/S0377-2217\(98\)00319-1](http://dx.doi.org/10.1016/S0377-2217(98)00319-1)
- Voß, S., Fink, A., Duin, C.: Looking ahead with the pilot method. Annals of Operations Research **136**(1), 285–302 (Apr 2005). <https://doi.org/10.1007/s10479-005-2060-2>, <http://dx.doi.org/10.1007/s10479-005-2060-2>

Robust Reaction Strategies for Worker-Flexible Job Shop Scheduling Problems under Uncertainty

David Hutter[0009-0001-9132-1349], Thomas Steinberger[0000-0002-5717-7873],
and Michael Hellwig[0000-0002-6731-8166]

Vorarlberg University of Applied Sciences, 6850 Dornbirn, Austria
david.hutter@fhv.at, thomas.steinberger@fhv.at, michael.hellwig@fhv.at

The Flexible Job Shop Scheduling Problem (FJSSP) is an NP-complete combinatorial optimization problem in which a set of jobs needs to be scheduled for production. Each job consists of a sequence of operations that must be scheduled on one of the eligible machines without preemption. The processing time of each operation depends on the assigned machine [2]. An extension is the FJSSP with Worker Flexibility (FJSSP-W). Here, every operation is additionally required to have one of multiple eligible workers assigned. The processing time of the operation then depends on the assigned machine and the assigned worker. The FJSSP-W are commonly approached with techniques from Artificial Intelligence (AI), Operations Research (OR), or Simulation-based Optimization (SBO) [3].

In addition, this paper considers the FJSSP-W in the presence of stochastic processing times, machine breakdowns, and worker unavailability. If a resource (i.e., machine or worker) becomes unavailable during the execution of a schedule, all operations assigned to this resource need to be either delayed until the resource is available again or reassigned to different resources, if available. For the FJSSP-W, this increases the complexity significantly. As the machine-worker combinations influence the processing time of an operation, the originally assigned worker may no longer be a good choice on a different machine.

Initial schedules can be made robust through different optimization methods, like predictive or proactive scheduling [1, 2]. Robust optimization methods typically lead to less efficient but more reliable schedules [1]. In real life use cases, a compromise between efficiency and robustness is often preferable. This means that even robust schedules still need to resolve conflicts arising during schedule execution as uncertain events cause delays. While a simple right-shift strategy, i.e. delaying all operations until their resources are available again, guarantees a feasible solution and is easy to determine, the resulting schedules are inefficient. More complex reaction strategies can lead to more efficient resolutions of conflicts, but are computationally expensive. Since these changes often need to be made during schedule execution, methods need to balance result quality and computational efficiency. Including robust optimization methods in the initial planning reduces the amount of conflicts occurring during execution, at the cost of less efficient schedules. To achieve a good balance, different predictive-reactive or proactive-reactive strategies can be used. Common reaction strategies include rule-based decision making, machine learning models, or metaheuristics [2, 1]. These methods can be used for either full or partial rescheduling. In a full resched-

ule, all remaining operations are re-optimized to retrieve a new schedule. While this is more computationally expensive, it can lead to higher efficiency for the remaining operations. Using partial rescheduling, more of the original schedule can be preserved. The aim is to only move the operations affected by the uncertain events to different resources, resulting in a more reliable execution of the overall schedule, yet potentially less efficient schedules for the remaining operations.

Using a simulation-based approach to generate uncertain events in FJSSP-W production environments, this paper compares different reaction strategies in a proactive-reactive optimization system with the initial scheduling handled by a robust genetic algorithm variant based on [4]. In particular, the aim is to investigate which approaches do preserve the original schedule as much as possible during partial rescheduling. To this end, we compare rule-based, metaheuristic, and reinforcement learning methods to achieve a reliable schedule execution, while trying to keep the overall execution time as low as possible. Preliminary tests show that the proactive GA can reduce the number of conflicts. For the remaining conflicts, considering the performance requirements, a metaheuristic approach for the reaction strategy is expected to produce good results in a short amount of time, especially compared to the simplistic right-shift strategy. While rule-based approaches are very fast, formulating a well working set of rules applicable to many problem instances typically is very difficult. Similarly, a reinforcement learning approach is expected to struggle to adapt to many different production environments. Additionally, the work presents insights into robustness indicators for created schedules, given that the used GA tends to find multiple solutions with similar execution times.

Acknowledgements

The financial support by the Austrian Federal Ministry of Labour and Economy, the National Foundation for Research, Technology and Development, the Christian Doppler Research Association is gratefully acknowledged.

References

1. Bakon, K., Holcinger, T., Süle, Z., Jaskó, S., Abonyi, J.: Scheduling Under Uncertainty for Industry 4.0 and 5.0. *IEEE Access* **10**, 74977–75017 (2022). <https://doi.org/10.1109/ACCESS.2022.3191426>
2. Dauzère-Pérès, S., Ding, J., Shen, L., Tamssouet, K.: The flexible job shop scheduling problem: A review. *European Journal of Operational Research* **314**(2), 409–432 (Apr 2024). <https://doi.org/10.1016/j.ejor.2023.05.017>
3. Ghasemi, A., Farajzadeh, F., Heavey, C., Fowler, J., Papadopoulos, C.T.: Simulation optimization applied to production scheduling in the era of industry 4.0: A review and future roadmap. *Journal of Industrial Information Integration* **39**, 100599 (2024). <https://doi.org/https://doi.org/10.1016/j.jii.2024.100599>
4. Hutter, D., Steinberger, T., Hellwig, M.: A Parameter Adaptive Genetic Algorithm With Restarts for Flexible Job Shop Scheduling Problems with Worker Flexibility. In: Proceedings of the 2025 ACM Conference on Genetic and Evolutionary Computation (GECCO'25). pp. 1–4 (2025)

Model-Based Optimisation of the Energy Efficiency of Machining Processes

Doris Bernroider¹[0009-0007-7089-8708], Fabian Spitzer¹[0009-0007-7827-6866],
Jochen Giedenbacher¹, and Holger Groening^{1,2}[0000-0001-8241-591X]

¹ University of Applied Sciences Upper Austria, School of Engineering,
Stelzhamerstraße 23, 4600 Wels, Austria

² Center of Excellence for Smart Production, Roseggerstraße 15, 4600 Wels, Austria

The manufacturing sector is the backbone of modern industrialised societies and plays a decisive role in driving economic growth and prosperity. At the same time, the manufacturing industry is also responsible for a significant share of global energy consumption. Estimates from international energy agencies and scientific studies indicate that the industry accounts for between 25% and 40% of final energy consumption, with a significant portion attributable to manufacturing processes. Since 2000, global industrial energy demand has increased by around 70%, growing at an average rate of approximately 1.3% per year between 2010 and 2022 [1]. Consequently, the manufacturing sector bears a major responsibility in achieving the transition to an energy-neutral economy.

Machine tools are key components of industrial manufacturing and, due to their high power consumption and diverse operating conditions, represent a major source of energy use. Studies, such as those analysed in Denkena et al. [2] demonstrates that machine tools account for a significant share of energy demand in production and therefore constitute a crucial lever for improving efficiency.

Based on the geometric characteristics of the workpieces and the associated machining processes, this study focuses on analysing the energy consumption of machine tools. The target is to identify optimisation potential based on a differentiated acquisition of data from all relevant energy consumers. Three key optimisation approaches are considered: the machine's design, its operation, and the machining process. The aim is to evaluate energy consumption and identify opportunities to reduce it, thereby improving sustainability in manufacturing.

To systematically assess the identified optimisation potentials and derive practical strategies, a structured representation of energy consumption is required. For this purpose, two perspectives are combined:

Component level:

ISO 14955-1 provides a standardised basis for classifying energy consumers to ensure the comparability of measurements. In line with this recommendation, the individual consumers of a machine are grouped into energy-relevant functional categories.

Process level:

Machining processes are represented based on cutting forces. For this purpose, the model of Kienzle is applied, which determines cutting and chip formation forces as a function of the material, tool, machining parameters, and

the specific operation. Using the physical relationship $P = F \cdot v$, forces can be converted into mechanical power and, through efficiency factors, into electrical power. In this way, the energy demand of the process can also be estimated using a model-based approach.

To empirically support and validate the model development, energy measurements are carried out while a defined reference part is machined from different materials. This series of experiments not only reveals the machine's energy consumption during machining but also enables the evaluation of the influence of tool wear on energy demand. According to Denkena et al. [3], [4], wear phenomena on the tool lead to an increase in cutting forces, which is also reflected in the machine's power consumption.




Fig. 1. Schematic visualisation of the process from data acquisition to the identification of optimisation potentials.

The analysis makes it possible to quantify the share of energy consumed by the machining process itself and the portion attributable to auxiliary and support units. Furthermore, it highlights the relationship between theoretically calculated cutting forces and the machine's energy demand. Evaluating these results contributes to energy-efficient manufacturing by enabling the well-founded identification of saving potentials and thus the optimisation of machining processes.

References

1. International Energy Agency (IEA): Energy Efficiency 2023. Paris (2023). <https://www.iea.org/reports/energy-efficiency-2023>, last accessed 2025/06/26
2. Denkena, B., Abele, E., Brecher, C., Dittrich, M.-A., Kara, S., Mori, M.: Energy efficient machine tools. *CIRP Annals*, 69(2), 646–667 (2020). <https://doi.org/10.1016/j.cirp.2020.05.008>
3. Denkena, B., Krüger, M., Schmidt, J.: Condition-based tool management for small batch production. *International Journal of Advanced Manufacturing Technology* 74, 471–480 (2014). <https://doi.org/10.1007/s00170-014-6013-2>
4. Pittalà, G.M., Rizzuti, S.: An investigation of the effect of tool wear on cutting force coefficients for solid end mills. *Procedia CIRP*, vol. 117, pp. 444–449 (2023). <https://doi.org/10.1016/j.procir.2023.03.075>

Closing the Feedback Loop: Adaptive Constraint Generation for Tree-Search Algorithms through Large Language Models*

Stefan Habringer, Sebastian Leitner^[0000-0003-4261-2868], and Johannes Karder^[0000-0002-7006-8663]

Heuristic and Evolutionary Algorithms Laboratory
University of Applied Sciences Upper Austria, 4232 Hagenberg, Austria

Human-in-the-loop machine learning (HITL-ML) [3,4] encompasses a broad spectrum of research focused on integrating human expertise into algorithmic decision-making. The interaction of domain experts with software components represents challenges when applying these approaches in non-academic settings. Interactions with human experts are typically confined to the training and evaluation phase, with no intermediate process to make necessary adjustment when real-world conditions change. Adoption of these systems further depends on interpretability and trust in the underlying system. Domain experts must understand and observe how their input influences the decision of the underlying model [1]. To address these shortcomings, we propose a framework that establishes a continuous interaction workflow between domain experts and the optimization system. Within this framework, domain experts can express new or revised constraints in natural language, which the system subsequently integrates into the optimization process. We demonstrate the approach in a real-world logistics scenario, where steel plates must be loaded onto transportation vehicles while maximizing the shipped weight and throughput of the facility and simultaneously minimizing storage areas and costs. Domain experts impose the structure and constraints on the optimization process. Since constraints can dynamically change, the system must flexibly and reliably integrate new information. Constraints encompass product, transportation and customer specific restrictions and can either be generally applicable or highly specific. Subsequently, the constraint set for a given problem is highly changeable, necessitating a system that is capable of incorporating new constraints effectively and reliably.

The workflow of the proposed system is depicted in Figure 1. Initially, a domain expert starts the optimization system (1), which solves the optimization problem at hand via tree-search [2]. The system then returns a set of solutions (2), which are all feasible given a currently active set of constraints. The domain expert can then inspect these solutions and formulate new constraints (3), based on the observations that are made, in natural language. A large language model (LLM) translates these natural language constraints into executable code

* This work was done within the project LogAI^d (#FO999921379), funded by the state of Upper Austria as part of the #upperVISION2030 initiative AI Region Upper Austria.




Fig. 1. Interactions of the human operator with the optimizer and LLM components.

constraints (4). Within this setup, the LLM serves as a semantic translator. It is provided with a predefined class structure and interface specification, which define a contract that any generated C# code must respect. Based on this contract, the model produces source code that encapsulates additional constraints. The resulting code is then integrated during runtime using the .NET Compiler Platform (Roslyn), thus incorporating the new constraints. These newly generated constraints modify the optimization behavior directly, thereby creating a dynamic interaction between the machine and human expert. To evaluate this approach, several different LLMs are compared. Evaluation metrics include a) the ability to produce executable code that adheres to the defined contract, b) the ability to represent the intended constraint logic, which is evaluated by a suite of predefined validation tests, and c) the computational efficiency.

Preliminary results using four different language models and four representative constraints, indicate differences in contextual understanding. Smaller models fail to capture the full domain context and are inconsistent with regard to code generation. In contrast, larger models reliably produce correct code, both syntactically and semantically. Furthermore, more complex and multi-layered constraints comparable to real-world domain restrictions are evaluated and a comprehensive analysis of runtime and performance metrics is conducted.

References

1. Adadi, A., Berrada, M.: Peeking inside the black-box: A survey on explainable artificial intelligence (xai). *IEEE Access* **6**, 52138–52160 (2018). <https://doi.org/10.1109/ACCESS.2018.2870052>, <http://dx.doi.org/10.1109/ACCESS.2018.2870052>
2. Leitner, S., Fleck, P., Wagner, S., Affenzeller, M.: Dynamic transport-lot assignment for the hot-storage area. In: Proceedings of the 36th European Modeling & Simulation Symposium. vol. 36. CAL-TEK srl (2024). <https://doi.org/10.46354/i3m.2024.emss.031>, <http://dx.doi.org/10.46354/i3m.2024.emss.031>
3. Mosqueira-Rey, E., Hernández-Pereira, E., Alonso-Ríos, D., Bobes-Bascarán, J., Fernández-Leal, A.: Human-in-the-loop machine learning: a state of the art. *Artificial Intelligence Review* **56**(4), 3005–3054 (Aug 2022). <https://doi.org/10.1007/s10462-022-10246-w>, <http://dx.doi.org/10.1007/s10462-022-10246-w>
4. Wu, X., Xiao, L., Sun, Y., Zhang, J., Ma, T., He, L.: A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems* **135**, 364–381 (Oct 2022). <https://doi.org/10.1016/j.future.2022.05.014>, <http://dx.doi.org/10.1016/j.future.2022.05.014>

Towards Less Energy Reliance - A Computational Study on Benchmark Instances for Selected Locations in Austria

Alexander Kinast¹[0000-0002-0361-1136], Dominik Falkner¹[0009-0002-7741-2225], Michael Bög¹[0000-0001-8921-8548], and Roland Braune²[0000-0003-4086-229X]

¹ RISC Software GmbH, Softwarepark 32a, 4232 Hagenberg im Mühlkreis, Austria
office@risc-software.at

² Department of Business Decisions and Analytics, Oskar-Morgenstern-Platz 1 1090 Vienna, Austria

Background and Motivation Due to reliance on Russian gas, Europe has had high and volatile energy prices since the war in Ukraine began. In particular, Austria and Germany rely heavily on Russian Gas [1].

In [4], it is stated that 80% of the Austrian gas was imported from Russia in 2021. If this supply is suddenly stopped, this is called a gas supply shock. Based on various studies, this would have significantly impacted the Austrian companies' profits and, consequently, the Austrian GDP. The article describes that this scenario did not happen, even though the politicians had to be prepared. In 2024, Russia was still the largest source of gas imports in Austria.

An extensive study has analyzed the impact of these increased energy prices on the profitability of companies in 16 different European countries. One of the key findings of this study is that the increased energy prices have resulted in decreased profits for companies in the European Union. Another key finding is that these increased energy prices have a higher impact on the manufacturing sector, as this sector typically has energy-intensive production processes. [2]

In [3], it is explained that power contracts for factories are unlike those of ordinary homes. One way to define these contracts is by consumption and peak power. Peak power is the maximum power demand at any given moment, while consumption is the total electricity used. Existing scheduling algorithms have to be adjusted to consider this peak power as well as other typical optimization goals.

One of the most important optimization problems in the manufacturing sector is the hybrid flow shop scheduling problem (HFSSP). In [5], it is described that this problem is important in many industrial production areas, such as electronics, paper, textiles, and concrete.

In this paper, we extend an existing HFS benchmark data set to include energy-related aspects. Additionally, we provide photovoltaic data for a set of current industry hotspots in Austria. The benchmark data sets are merged with the photovoltaic data to form the foundation for the computational study.

Method The first part of the paper focuses on the creation of an Industrie 4.0 benchmark dataset. Therefore, an existing flow shop scheduling benchmark dataset with 480 instances is extended with PV-cell power production, variable

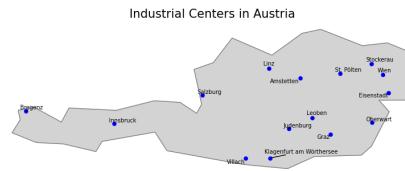


Fig. 1: Selected industrial cities for which PV module data are provided.

production speed, and variable power consumption. The key component to couple production speed and power consumption is an adjustable formula that can adapt to different real-world scenarios. To create this dataset, data about key industrial areas in Austria and their solar radiation levels are included in the benchmark dataset. These key industrial areas are shown in Figure 1.

The created benchmark instances consider conflicting objectives, the minimization of makespan, and the minimization of the peak power consumption. The final part of the paper focuses on the development of an efficient metaheuristic and encoding for this problem. Therefore, a machine assignment-based encoding is developed.

The paper concludes with a large computational study utilizing the developed machine encoding to create reference solutions for the new benchmark problem.

Acknowledgements The work was done within the project “Secure Prescriptive Analytics”, funded by the state of Upper Austria as part of the program “#upperVISION2030”.

References

1. Aliu, F., Mulaj, I., Hašková, S.: Consequences of the russia-ukraine war: evidence from dax, atx, and ftsemib. *Studies in Economics and Finance* **40**(3), 549–568 (2023)
2. Herman, R., Nistor, C., Jula, N.M.: The influence of the increase in energy prices on the profitability of companies in the european union. *Sustainability* **15**(21) (2023). <https://doi.org/10.3390/su152115404>, <https://www.mdpi.com/2071-1050/15/21/15404>
3. Nagasawa, K., Ikeda, Y., Irohara, T.: Robust flow shop scheduling with random processing times for reduction of peak power consumption. *Simulation Modelling Practice and Theory* **59**, 102–113 (2015). <https://doi.org/https://doi.org/10.1016/j.smpat.2015.08.001>, <https://www.sciencedirect.com/science/article/pii/S1569190X1500115X>
4. Pichler, A., Hurt, J., Reisch, T., Stangl, J., Thurner, S.: Economic impacts of a drastic gas supply shock and short-term mitigation strategies. *Journal of Economic Behavior & Organization* **227**, 106750 (2024). <https://doi.org/https://doi.org/10.1016/j.jebo.2024.106750>, <https://www.sciencedirect.com/science/article/pii/S0167268124003640>
5. Ruiz, R., Vázquez-Rodríguez, J.A.: The hybrid flow shop scheduling problem. *European Journal of Operational Research* **205**(1), 1–18 (2010). <https://doi.org/https://doi.org/10.1016/j.ejor.2009.09.024>, <https://www.sciencedirect.com/science/article/pii/S0377221709006390>

An Incremental Algorithm for Optimized Container Transport and Stacking under Dynamic and Uncertain Conditions

Nico Ralfs¹ and Matthias Becker¹

Leibniz University Hannover, Welfengarten 1, 30167 Hannover, Germany
 xmb@hci.uni-hannover.de

Abstract. This paper addresses the dynamic optimization of container transport and stacking under stochastic arrival, release, and movement times. Four algorithms are evaluated: a random baseline, a Greedy heuristic, an A* search, and a reactive Incremental algorithm. Using a simulation with uncertain arrivals and crane operations, performance is assessed via blocked arrival time (BAT), timely delivery (TBT), and crane manipulations (CM). Results show that the Incremental algorithm achieves the most robust performance across varying scenarios.

Keywords: dynamic optimization · container stacking · uncertainty

1 Introduction

Container stacking in real-world logistics is shaped by uncertainty in arrival, processing, and transport times. Deterministic models thus often fail to remain feasible or efficient once stochastic deviations occur [2, 1]. This work, developed for the GECCO DynStack Competition 2025, studies algorithmic strategies for moving containers from an arrival stack to a target stack while reacting to unpredictable system changes.

The simulation environment hides underlying probability distributions, allowing algorithms to observe only the current state. Performance is measured using three KPIs: blocked arrival time (BAT), containers delivered on time (TBT), and crane manipulations (CM). The goal is to identify which algorithm handles uncertainty most effectively.

2 Approach

We implemented four strategies: a random decision baseline, a Greedy algorithm, an A* search, and an Incremental algorithm.

2.1 Incremental Algorithm

The Incremental method continuously monitors the system state and updates its plan only when relevant events occur, such as:

- arrival of a prioritized container,
- a ready container becoming blocked,
- congestion of the arrival stack,
- exhaustion of valid moves.

The algorithm adopts the Greedy logic for prioritization and relocation, but enhances it with a state manager that reacts to changes during execution. It does not attempt to predict arrival or movement durations; uncertainty is addressed implicitly through continuous plan adaptation.

This reactive design enables robust behavior under fluctuating arrival patterns without requiring explicit stochastic modeling.

3 Conclusion

Across all scenarios, the Incremental algorithm demonstrates the best overall performance. It maintains a BAT of 0.0 and achieves strong TBT values while minimizing unnecessary crane movements. Greedy shows competitive but inconsistent results, A* suffers from computational overhead in complex cases, and the random strategy serves as a weak baseline.

Reactive planning therefore proves superior in dynamic environments, as continuous monitoring enables timely adjustments without forecasting. Future work may explore estimating container ready times, enhancing A* heuristics, or combining reactive and predictive strategies.

The Incremental algorithm achieved second place in the DynStack 2025 competition with BAT=0.0, TBT=128, and CM=336 for 132 processed containers.

References

1. Leitner, S., Wagner, S., Affenzeller, M.: Incrementally solving the dynamic stacking problem. In: EUROCAST 2024. Springer (2025)
2. Raggl, S., Beham, A., Wagner, S.: Effects of arrival uncertainty on solver performance in dynamic stacking. EMSS (2020)
3. Hart, P., Nilsson, N., Raphael, B.: A formal basis for heuristic path search. IEEE SSC (1968)
4. Jin, X., Duan, Z., Song, W.: Container stacking optimization via deep RL. Eng. Appl. AI (2023)

Real-Time Data Stream Processing for Interpretable Machine Learning

Philipp Neuhauser^{1[0009-0000-8756-8181]}, Florian Holzinger^{1[0009-0005-0858-3233]}, and Stefan Wagner^{1[0000-0002-0323-9274]}

University of Applied Sciences Upper Austria, 4232 Hagenberg, Austria

Industrial production facilities and the associated data landscape have become increasingly complex since digitalisation in the age of Industry 4.0. Machines are constantly being upgraded with heterogeneous sensor suites that deliver an increasing amount of data in the form of event streams at ever shorter intervals. This high-frequency data recording should ultimately provide deeper insights into the respective production processes, enabling optimisation, automation, and cost reduction through the application of machine learning models.

Real-time capability is becoming increasingly important, especially in mission-critical areas where decisions must be made immediately. Therefore, machine learning and data processing pipelines must be low-latency to respond quickly to events. However, the ability to use such information in near real-time remains a challenge. Most supervised ML algorithms expect static batch data sets, where all observations are synchronised to a single time axis. In contrast, industrial data streams have varying frequencies, delays, and missing values.

Without appropriate preprocessing, ML models might deliver unreliable predictions. Therefore, a certain degree of interpretability of the trained models is also of enormous importance to validate their reliability. White-box approaches address this need by providing transparent, human-readable structures that allow experts to trace predictions back to interpretable mathematical expressions. A promising representative of this category is symbolic regression [2], which uses genetic programming (GP) to evolve a set of mathematical primitives and input variables to approximate a target function.

In recent years, the technology stack built around Apache Kafka [3] and Apache Flink [1] has become the de facto standard for high-throughput, distributed event stream processing in industrial contexts. These frameworks provide a robust foundation for real-time analytics, enabling continuous data flows that can handle heterogeneous, asynchronous, and high-frequency sensor signals. Heterogeneous sensor streams are aligned through adaptive windowing mechanisms that synchronise data in both time and frequency domains, as shown in Figure 1. Within these windows, statistical and temporal aggregation functions (e.g., mean, median, minimum, maximum) can be applied to summarise high-frequency sensor signals while preserving key process dynamics and minimising information loss. Missing or delayed sensor events are compensated through online imputation methods to ensure data continuity under low-latency constraints.

Despite this technological maturity, there is still limited systematic understanding of how different stream-processing configurations affect the quality and usability of data for machine learning in production environments. This paper ad-




Fig. 1. Real-time Data Stream Processing Pipeline for Interpretable Machine Learning

addresses this gap by investigating how real-time preprocessing strategies influence the predictive performance and interpretability of models applied to industrial data streams. In our work, we introduce an integrated architecture that combines event ingestion, synchronisation, and real-time model inference within the Apache Kafka and Apache Flink ecosystem. The preprocessed and synchronised streaming data are then used to train interpretable, white-box machine learning models based on symbolic regression. A series of experiments systematically test stream-processing configurations, such as window types, aggregation strategies, and imputation methods, to quantify their impact on data quality, model accuracy, and interpretability. The results are expected to provide guidelines for designing data pipelines that balance real-time performance, predictive accuracy, and transparency in production environments.

Acknowledgments. This work was supported by the research project ENPro. The project ENPro is financed by research subsidies granted by the government of Upper Austria.

References

1. Hueske, F., Kalavri, V.: Stream Processing with Apache Flink. O'Reilly Media, Sebastopol, CA (apr 2019)
2. Kronberger, G., Burlacu, B., Kommenda, M., Winkler, S.M., Affenzeller, M.: Symbolic Regression. Chapman and Hall (2024)
3. Shapira, G., Palino, T., Sivaram, R.: Kafka - The Definitive Guide. O'Reilly Media, Sebastopol, CA, 2 edn. (nov 2021)

Towards Improving Substation Load Forecasts in Renewable-Driven Grid Environments

Dominik Falkner^{1,2,*[0009-0002-7741-2225]}, Felix Oberhauser¹, Franz Ebner¹,
and Michael Affenzeller^{2,3}

¹ RISC Software GmbH, Softwarepark 32a, 4232 Hagenberg im Mühlkreis, Austria
dominik.falkner@risc-software.at

² Institute for Formal Models and Verification, Johannes Kepler University Linz,
Austria

³ University of Applied Sciences Upper Austria, Softwarepark 11, 4232 Hagenberg,
Austria

Background and Motivation: Modelling substation load is essential for ensuring grid security and stability. Reliable forecasts help prevent overloads and blackouts, tasks that have become increasingly challenging due to the growing volatility introduced by renewable energy sources. As the electrification of transport and heating accelerates, and as flexible demand and distributed generation expand, substation-level load modelling becomes even more critical for network planning, power flow optimisation, and operational decision-making.

Substations play a central role in the power grid. They provide voltage transformation, act as aggregation points for consumption and generation, and cover extensive geographical areas. Historically, load forecasting relied on relatively homogeneous household profiles and a limited number of large consumers, supported by predictable and controllable thermal power generation. Today, however, load profiles are shaped by a wide variety of factors. Industrial demand often follows regular schedules, while renewable generation introduces strong non-linear and weather-dependent fluctuations. Additional complexity arises from demographic shifts, urban growth, and increased adoption of electric vehicles, heat pumps, and battery storage systems [4]. Renewable generation affects nearly all grid levels, producing characteristic midday peaks and pronounced seasonal patterns. Dynamic pricing and flexible loads further shift demand into specific periods [6, 7]. To investigate these challenges empirically, we analyse a real-world substation dataset.

Data: Our dataset originates from a real-world 110 kV substation, an infrastructure level for which publicly available information is extremely scarce. The substation represents a typical Austrian supply area with strong and increasing photovoltaic penetration, complemented by a smaller contribution from hydropower. The dataset includes regional weather variables and spans a continuous three-year period at a 15-minute resolution without missing values.

Methods: This study focuses on day-ahead and up to three-day-ahead forecasts, which are essential for informed operational decision-making. To model substation consumption, we consider a diverse set of forecasting techniques that reflect both established industry practices and recent methodological advances. Statistical baselines such as ARIMA are included because they are widely used,

easy to interpret, and provide a meaningful performance reference. Explainable approaches such as Symbolic Regression [1] allow the derivation of explicit analytical expressions that can offer insights into underlying system dynamics. We further evaluate modern pre-trained foundation models, including Chronos 2 [2] and TiRex [3]. To ensure a fair comparison, we apply a consistent training, validation, and testing procedure and use standard accuracy metrics relevant to operational planning. In addition to point forecasts, we examine the probabilistic intervals produced by the models to evaluate their suitability for decision-making under uncertainty.

Contribution: Our contributions include a systematic comparison of several state-of-the-art forecasting models, evaluated under the contextual shifts characteristic of today's evolving energy system. To our knowledge, this is the first systematic evaluation of time series foundation models for substation-level data. Additionally, we incorporate regional photovoltaic generation through data fusion and photovoltaic output simulation [5]. Overall, we assess the practical suitability of these approaches for supporting the decision-making of a distribution system operator, particularly in the context of increasing renewable energy integration.

References

1. Affenzeller, M., Winkler, S.M., Kronberger, G., Kommenda, M., Burlacu, B., Wagner, S.: Gaining Deeper Insights in Symbolic Regression, p. 175–190. Springer New York (2014). https://doi.org/10.1007/978-1-4939-0375-7_10
2. Ansari, A.F., Shchur, O., Küken, J., Auer, A., Han, B., Mercado, P., Rangapuram, S.S., Shen, H., Stella, L., Zhang, X., Goswami, M., Kapoor, S., Maddix, D.C., Guerrou, P., Hu, T., Yin, J., Erickson, N., Desai, P.M., Wang, H., Rangwala, H., Karypis, G., Wang, Y., Bohlke-Schneider, M.: Chronos-2: From univariate to universal forecasting (2025), <https://arxiv.org/abs/2510.15821>
3. Auer, A., Podest, P., Klotz, D., Böck, S., Klambauer, G., Hochreiter, S.: Tirex: Zero-shot forecasting across long and short horizons with enhanced in-context learning (2025), <https://arxiv.org/abs/2505.23719>
4. Hong, T., Xie, J., Black, J.: Global energy forecasting competition 2017: Hierarchical probabilistic load forecasting. International Journal of Forecasting **35**(4), 1389–1399 (Oct 2019). <https://doi.org/10.1016/j.ijforecast.2019.02.006>
5. Joint Research Centre: Photovoltaic geographical information system (pvgis). <http://re.jrc.ec.europa.eu/pvgis> (2023), accessed: 2025-11-13
6. NERC: Developing load model composition data for dynamic load modeling. Tech. rep., North American Electric Reliability Corporation (NERC) (February 2017)
7. Peng, H., Lou, Y., Li, F., Sun, H., Liu, R., Jin, B., Li, Y.: Decomposition framework for long term load forecasting on temperature insensitive area (2024). <https://doi.org/10.2139/ssrn.4938668>, <http://dx.doi.org/10.2139/ssrn.4938668>

Modeling of a DELTA-Robot for AI-based path planning

Andreas Weyermann, Tobias Weiser, and Dirk Jacob

University of Applied Sciences Kempten, Bahnhofstraße 61, 87435 Kempten,
Germany andreas.weyermann@hs-kempten.de, tobias.weiser@hs-kempten.de,
dirk.jacob@hs-kempten.de

Abstract. The present document was created as part of the European funding programme AI4Green and details the development of a numerical kinematic model for a DELTA-robot. The resulting model is subsequently employed to train a reinforcement-learning (RL) algorithm. By evaluating power consumption, the RL agent can be guided to exploit the workspace in a more energy-efficient manner.

Keywords: Delta Robot · Modeling · Digital twin · Kinematics · Trajectory planning · Reinforcement-learning · Energy Efficiency

1 Introduction

Linear motion of the tool center point (TCP) of a DELTA-robot in task space result in energetically inefficient non-linear motion in the joint space. They introduce idle movements that unnecessarily increase energy consumption. Direct point-to-point movements (PTP) are unsafe due to potential collision risks and prevent path planning in this regard. The optimisation performed with a reinforcement-learning algorithm seeks to minimise idle motion, while training the RL agent to make optimal use of locally available drive power and to efficiently utilise the kinetic energy stored in the robot's kinematics. In the preceding Plot-Bot project [1], a kinematic system for a Delta robot had already been designed, implemented, and validated. It is intended to serve as the technical platform for subsequent optimisation studies. The digital twin generated from this geometry enables simultaneous kinematic simulation for model training and guarantees that any agent misbehaviour immediately effect a negative reward signal. The following chapters describe the kinematic modeling of the robot used for modeling purposes from [1] based on [2, 3].

2 Modeling

The study in [2] outlines various computational methods for modeling the dynamics of a robot. The chosen approach employs an internal force balance - specifically, a Newton–Euler formulation- for the mechanical components. The increased computational effort incurred in this approach also yields the internal forces of the kinematics, which can be utilized later on.

2.1 Geometric correlations, acceleration and internal forces

To determine the robot arm configuration, the upper-arm and fore-arm vectors are computed using inverse kinematics. If a workspace violation occurs, an initial reward can be generated at that point. With the arm vectors, the Jacobian matrix required for velocity correlation can be determined. Subsequently, based on the position, the now known velocities, and the desired workspace acceleration, the joint accelerations as well as the drive torques for each arm are computed (see Fig.1). This allows additional rewards to be generated for exceeding the maximum acceleration and maximum torque limits.

2.2 Reinforcement learning and time integration

With the returned values (red arrow, Fig.1) from the numerical intermediate steps, the RL agent can be trained from this point onward. Inverse kinematics, the joint velocities derived from the Jacobian matrix, the joint accelerations and drive torques provide the necessary information to maintain compliance with the workspace, local velocities, local accelerations and the maximum drive torque. The temporal integration of the actuator power yields the energy demand for the generated trajectory and enabling its energetic assessment. The reinforcement-learning agent thus learns to optimize its trajectories with respect to energy consumption.




Fig. 1. Training loop of reinforcement learning

Acknowledgments. The research reported in this paper has been funded by European Interreg Austria-Bavaria project "AI4Green".

References

1. Jacob, D., Haberstroh, P., Neidhardt, D. & Timmermann, B. Modular 3D-Printed Robots for Education and Training for Industrie 4.0. *Computer Aided Systems Theory – EUROCAST 2017*. pp. 214-219 (2018)
2. Brinker, J., Wahle, M. & Corves, B. A Comparative Study of Inverse Dynamics based on Clavel's Delta robot. (2015),
3. Wang, J., Liu, X. Analysis of a novel cylindrical 3-DoF parallel robot. (2003)

Optimizing the energy efficiency of a delta robot through AI-based trajectory planning

Nico Weber and Dirk Jacob

University of Applied Sciences Kempten, Bahnhofstraße 61, 87435 Kempten,
Germany nico.weber@hs-kempten.de, dirk.jacob@hs-kempten.de

Abstract. This work aims to increase the energy efficiency of a delta kinematic robot using reinforcement learning (RL) by examining and optimizing the robot's movement sequences between variable positions. An RL environment was developed to train a model that is able to automatically generate an energy efficient trajectory between two random points in a 3D-workspace, while avoiding a static cubic obstacle with random position and size. Using a digital twin, the robot's limits and energy consumption were simulated during training. By balancing the reward function between the task of reaching the goal and minimizing energy consumption, a final AI model was created that generates trajectories using around 28 % less energy than manually created ones.

Keywords: Delta Robot · Digital twin · Kinematics · Trajectory Planning · Reinforcement Learning · Energy Efficiency

1 Introduction

With the increasing use of robots in industrial production, the energy consumption - and with it, the cost of operation and impact on the environment - are becoming an ever-growing concern. To counteract this trend as part of the AI4GREEN project, an AI-based system is envisioned to optimize the energy consumption of robots through more efficient control. A common industrial task where robot movement is essential is the pick-and-place application [1], in which an object is moved from one position to another. Optimizing the trajectory between these positions offers a promising area for improvement, as manually generating energy efficient trajectories can be tedious for humans. One robot type used for such tasks are delta kinematic robots, one of which was previously developed at Kempten University of Applied Sciences [2]. This 3-DoF delta robot called "PlotBot" was used as the robot the model is trained for. To optimize the movement between two positions during a pick-and-place task, a model was developed using reinforcement learning (RL) to automatically generate energy-efficient trajectories between two random positions for a 3-DoF delta robot. In order to increase the complexity and realism of the trajectory planning, a cubic obstacle with randomized position and size was added to the workspace, which the RL agent had to avoid.

2 AI-Based System

To train the AI model, a RL environment was created that simulates the robot's surrounding workspace and generates the random obstacles as well as the start and goal positions to help the agent with generalization. A digital twin of the robot was used to simulate the inner constraints and calculate the energy consumption during training. The agent controls the robot by applying acceleration vectors to the tool center point (TCP) at 0.02 s intervals, generating a continuous trajectory while having a fine level of control leading to optimization potential in every step. The resulting trajectory consists of sequential points, each containing the position, speed and acceleration data of the TCP every 0.02 s. For the reward function, a dense reward with reward shaping is used to guide the agent during training. The reward function is split into two objectives: (1) reaching the goal while avoiding the obstacle and keeping actions inside the robot's limits; and (2) minimizing energy consumption by incorporating terms that historically reduce energy use and adding rewards based on the energy spent. Balancing both is essential to achieving an agent that reaches the goal consistently while minimizing energy use. Curriculum learning was used to further improve performance by gradually increasing task difficulty, resulting in a model capable of generating energy efficient trajectories even in complex scenarios.

To determine the optimal model, three different algorithms consisting of PPO [3], TRPO [4], and TQC [5] were used on the environment. After comparing the resulting models based on multiple metrics such as success rate and energy consumption, TQC produced the best model. It was able to reach the goal position 97.36 % of the time while avoiding obstacles placed directly between the start and goal position. In the same scenario, a reduction in energy consumption of around 29.60 % compared to manually created trajectories is achieved. This is reduced to 27.36 % when no obstacle needed to be avoided. Overall, the trajectories created by the model reduce the energy consumption compared to manually created trajectories by around 28 %.

Acknowledgments. The research reported in this paper has been funded by European Interreg Austria-Bavaria project "AI4Green".

References

1. Müller, C, World Robotics 2024 - Industrial Robots, IFR Statistical Department, VDMA Services GmbH, Frankfurt am Main, Germany, 2024
2. Jacob, D., Haberstroh, P., Neidhardt, D. & Timmermann, B. Modular 3D-Printed Robots for Education and Training for Industrie 4.0. *Computer Aided Systems Theory - EUROCAST 2017*. pp. 214-219 (2018)
3. Schulman, J., Wolski, F., Dhariwal, P., Radford, A. & Klimov, O. Proximal Policy Optimization Algorithms. (2017)
4. Schulman, J., Levine, S., Abbeel, P., Jordan, M. & Moritz, P. Trust Region Policy Optimization. (2015)
5. Kuznetsov, A., Shvechikov, P., Grishin, A. & Vetrov, D. Controlling Overestimation Bias with Truncated Mixture of Continuous Distributional Quantile Critics. (2020)

Generic HVAC model evaluation for predictive maintenance application

David Sengl¹, Edit Paráda¹, Aleksey Bratukhin² and Markus Kobelrausch³

¹ University of Applied Sciences Technikum Wien, Vienna, Austria

² University for Continuing Education Krems, Wiener Neustadt, Austria

³ Institute of Computer Technology TU Wien Vienna, Austria

david.sengl@technikum-wien.at

Abstract. In this work, a generic ventilation simulation model is developed using the dynamic simulation program TRNSYS. The model is validated using real-life monitoring data and analysed for efficiency improvements as well as maintenance behaviour. By incorporating comfort boundaries in a reinforcement learning model, the system achieves an efficiency increase of approx. 30 % while improving the efficiency with minimal effort in the running system drastically. While the validation confirms the model's accuracy, deviations are observed due to missing sensory data. The findings demonstrate the potential of simulation-based optimization for ventilation systems, improving energy efficiency while maintaining indoor comfort and keeping disruptions for users to a minimum.

Keywords: HVAC, Minimal effort, Predictive maintenance, Energy Efficiency

1 Introduction and Motivation

Private households contribute approx. 27% of Europe's final energy demand, largely due to Heating, Ventilation, and Air Conditioning (HVAC) systems [1]. Improving HVAC efficiency is critical for reducing energy consumption and environmental impact. This study develops a generic HVAC simulation model to optimize energy efficiency and enable predictive maintenance, using advanced simulation and machine learning to optimize the impact to effort ratio of building operations.

2 Methodology

The generic HVAC model was developed using the widely utilized TRNSYS simulation tool for dynamic energy systems. The model incorporates six essential components (fan, filter, heating register, cooling register, humidifier, and recirculation air valve) to simulate various ventilation units. The model was simplified using a decision tree to determine the most sustainable maintenance action (replacement vs. inefficient operation) to reduce complexity while preserving accuracy.

The simulation model was integrated with reinforcement learning (RL) to optimize HVAC settings and maintenance prediction. The RL model, trained using Gymnasium,

processes TRNSYS input data and additional engineered features (e.g., time of day, seasonality) to maximize a custom reward function balancing energy efficiency and user comfort [2]. The system operates in a multi-zone environment, representing different building segment.

3 Validation and Results

The model was validated against real-world data from an office building, focusing on indoor temperature, humidity, pressure differences and electrical energy demand. Five control logics (time-based, temperature-based, CO₂-based, mixed and AI-based) were evaluated. The AI-based control logic [3] demonstrated better performance, achieving a 30% increase in energy efficiency while maintaining comfort boundaries compared to the implemented temperature-based logic, resulting in energy flexibility potential.

Life Cycle Assessment and Life Cycle Cost analysis revealed that the AI-based control logic reduced CO₂ emissions by approx. 30% compared to failure-based maintenance. However, frequent component replacements in this approach led to higher costs, highlighting a trade-off between energy efficiency, flexibility and economic feasibility.

4 Conclusion

The paper demonstrated the potential of simulation-based optimization for HVAC systems, achieving significant energy efficiency improvements while maintaining indoor comfort. Though confirmed as accurate and reliable, model refinement in data acquisition and boundary implementation is needed. Future work will therefore focus on these areas to showcase real efficiency improvements and reduce user effort.

5 Acknowledgments

This project was funded by the Federal Ministry for Economy, Energy and Tourism Republic of Austria under the research Program FH Forschung für die Wirtschaft 2024. The authors acknowledge the support of the Austrian Research Promotion Agency (FFG) for the EnerFlexBuildings project.

References

1. Eurostat, Energy statistics - an overview, https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Energy_statistics_-_an_overview (2024), last accessed 2025/11/07
2. P. O. Fanger, Thermal Comfort: Analysis and Applications in Environmental Engineering, McGraw-Hill Book Company, New York (1972)
3. C. Stippel, R. Sterzinger, D. Sengl, A. Bratukhin, M. Kobelrausch, S. Wilker und T. Sauter, Online HVAC Optimization under Comfort Constraints via Reinforcement Learning, in 7th International Conference on Industrial Cyber-Physical Systems, St. Louis, USA (2024)

Optimization Strategies for Deploying Symbolic Regression Models on Embedded Hardware for Energy Management

Alexander Prielinger¹, Christoph Dalpiaz², Florian Eibensteiner², Kathrin Preiner², and Josef Langer¹

¹ Fronius International GmbH, Thalheim, Austria

{prielinger.alexander, preiner.kathrin}@fronius.com

² Josef Ressel Centre for AI on Resource Limited Devices

University of Applied Sciences Upper Austria, Hagenberg, Austria

{christoph.dalpiaz, florian.eibensteiner, josef.langer}@fh-hagenberg.at

Abstract. In this paper, we propose an optimization for deploying symbolic regression-based energy management algorithms on embedded hardware. We investigate two strategies: parameter reduction using Spearman correlation analysis, and the use of fixed-point arithmetic to replace floating-point operations. Both approaches are evaluated on an ARM Cortex-A55 platform using real household power measurements. The results show significant reductions in computational effort with minimal loss in control performance, enabling efficient embedded deployment of symbolic regression energy management systems.

Keywords: Symbolic Regression, Energy Management System, Optimization, Embedded Systems

1 Introduction

The increasing adoption of photovoltaic (PV) systems is driven by declining costs and growing environmental awareness. While PV technology has reached high efficiency levels, further improvements now focus on efficient energy utilization through optimized control systems. Symbolic regression-based energy management systems [1] offer sophisticated control capabilities but adaptation for resource-constrained embedded hardware is desired to reduce computational complexity. This work investigates porting and optimizing such algorithms for embedded deployment while maintaining control performance.

2 Concept

This work presents a two-fold optimization strategy for deploying symbolic regression-based energy management on embedded hardware. First, Spearman

correlation analysis identifies the most influential parameters of the original 15-parameter model, enabling parameter reduction while maintaining control quality. Second, fixed-point arithmetic (Q16.16) replaces traditional floating-point operations, significantly reducing computational overhead and memory requirements. Both optimization approaches were systematically evaluated on an ARM Cortex-A55 platform using comprehensive household power data in the Master's thesis by Prielinger [2].




Fig. 1. Conceptual flow of the optimization process for embedded energy management systems.

3 Results

The optimization approach was evaluated using three distinct energy management controllers with varying complexity and performance characteristics. Fixed-point arithmetic implementation (Q16.16) achieves substantial performance gains, reducing execution cycles by up to 23.89 % and instruction count by 17.54 % while maintaining prediction accuracy within acceptable bounds. These improvements are achieved with minimal impact on the controllers' energy management performance.

Parameter reduction through correlation analysis yields comparable computational benefits with acceptable accuracy decrease. This approach successfully identifies and eliminates low-impact parameters while preserving essential model characteristics. Each optimization strategy was evaluated independently, demonstrating that both fixed-point arithmetic and parameter reduction offer viable paths for deploying energy management algorithms on resource-constrained embedded platforms.

Acknowledgments. The financial support by the Austrian Federal Ministry of Labour and Economy, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association is gratefully acknowledged.

References

1. K. Kefer, R. Hanghofer, P. Kefer, M. Stöger, M. Affenzeller, S. Winkler, S. Wagner, and B. Hofer, *A Model-Based Learning Approach for Controlling the Energy Flows of a Residential Household Using Genetic Programming to Perform Symbolic Regression*,
2. A. Prielinger, *Optimierung eines KI-basierten Energiemanagementsystems für Embedded Hardware Portierung symbolisch-regressiver Modelle, Performance-Steigerung und Retraining*, Master's Thesis, 2025.

Smart Forecasts for Smart Factories

Simone Arnold¹[0009-0008-0905-1599], Gianluca Podann¹[0009-0000-0525-4083],
 Andreas Kneißler²[0000-0002-4072-5044], Katja Zeume²[0009-0000-3946-8514], and
 Kaan Celik²[0009-0004-6882-0643]

¹ University of Applied Sciences Dortmund simone.arnold@fh-dortmund.de

² Westphalian University of Applied Sciences Gelsenkirchen

Abstract. The increasing integration of renewable energy into industrial production systems creates new challenges for production planning and energy management. This paper presents a machine learning-based approach for forecasting locally generated renewable energy, with a particular focus on photovoltaic power. Using long short-term memory networks trained on long-term local weather and energy production data, the study investigates how both the length and the composition of the training dataset influence prediction accuracy. In addition, the impact of data imputation using measurements from nearby weather stations of the German Weather Service is analyzed to address data gaps in local datasets. The results contribute to the development of energy-aware production planning strategies, aiming to improve synchronization between renewable energy availability and industrial demand. The proposed method supports more efficient, flexible, and sustainable manufacturing processes through intelligent, data-driven energy forecasting.

Keywords: Machine Learning · AI · Prediction · Renewable Energy · Production Planning.

1 Introduction

Industrial production systems face increasing pressure to operate efficiently while reducing their environmental footprint. The transition to sustainable manufacturing requires not only renewable energy integration but also intelligent mechanisms to align energy supply and demand. The DIANE^{PRO} project [2] addressed this challenge by developing data-driven and AI-assisted planning tools to optimize production under fluctuating renewable energy conditions. It linked locally generated renewable energy with industrial consumption through advanced machine learning models and digital energy management strategies. Building upon these results, the present study focuses on addressing the key challenges identified in DIANE^{PRO}, particularly those related to data quality, forecasting accuracy, and the robustness of AI models under incomplete or variable datasets.

2 Research Questions

This work focuses on forecasting renewable energy availability, with an emphasis on solar power. It builds on previous research using machine learning for

wind and solar prediction. As shown in [1], long short-term memory (LSTM) networks can effectively model local energy generation. Building on these results, this study examines how the size and composition of the training dataset affect predictive accuracy. A major challenge is the availability and quality of training data, as LSTM models are sensitive to missing values that can degrade performance. Consequently, the study addresses two research questions:

How can data gaps be filled intelligently to stabilize AI training while minimizing errors from imputation? What is the minimal training period required to achieve reliable one-day-ahead forecasts?

3 Methodology

The dataset consists of local weather measurements recorded near a PV installation at an industrial site. These data form the primary input for the renewable energy forecasting model. Since the local measurements contain occasional gaps, missing values are replaced using data from the nearest German Weather Service station. This approach maintains data continuity and local relevance. Model performance is then compared with results from complete datasets to assess the effect of imputation on forecast accuracy. Building on previous work in energy-aware production planning, this study extends the concept to dynamic scheduling supported by localized energy predictions [1, 3], highlighting the advantages of site-specific over generalized models.

Acknowledgments. The authors would like to express their sincere gratitude to the Applied Excellence Department for supporting and funding the development of a new master's degree program with a strong emphasis on problem-based learning.

Disclosure of Interests. The authors declare that they have no competing interests. The research was conducted in the context of publicly funded academic curriculum development with no commercial influence.

References

1. Arnold, S., Podann, G., Faller, C., Kneißler, A.: LSTM-Based Renewable and Local Energy Forecasting Compared to Commercial Forecasting Tools. 24th Wind & Solar Integration Workshop, Berlin, Germany, 07–10 October 2025.
2. Arnold, S., Podann, G., Faller, C., Kneißler, A.: Data-driven integration and optimized use of different energy flows for sustainable production AALE 2025, Dresden, Germany, 12–14 March 2025.
3. Dobschinski, J.: Forecasting Beyond the Storm - Machine Learning Models for Offshore Windpower under Operational Constraints. 24th Wind & Solar Integration Workshop, Berlin, Germany, 07–10 October 2025.

Sample-Efficient Pareto Front Modeling for Energy-Aware Reinforcement Learning Using Bayesian Optimization

Georg Schäfer^{1,2}, Jakob Rehrl¹, Stefan Huber¹, and Simon Hirlaender²

¹ Josef Ressel Centre for Intelligent and Secure Industrial Automation,
Salzburg University of Applied Sciences, Salzburg, Austria

² Department of Artificial Intelligence and Human Interfaces,
Paris Lodron University of Salzburg, Salzburg, Austria
georg.schaefer@fh-salzburg.ac.at

1 Introduction

Industrial automation increasingly demands control strategies that balance performance with energy efficiency. A common approach to this multi-objective problem, particularly in reinforcement learning (RL), is to use a single, scalar reward function R_t that combines the competing objectives. In our previous work on the Quanser Aero 2 testbed [1], we used a linear scalarization of the form: $R_t = -(1 - \alpha) \cdot |\Delta_t| - \alpha \cdot P_t$ where Δ_t is the tracking error and P_t is the power consumption at time step t . The weight $\alpha \in [0, 1]$ balances the trade-off: $\alpha = 0$ prioritizes perfect tracking, while $\alpha = 1$ prioritizes minimal power use.

The central challenge lies in selecting the value of α . In our prior work [1], this was done through *manual weighting* – a time-consuming process where a practitioner intuitively selects and tests a few α values, which is prone to user bias and may miss optimal trade-offs.

This work automates the weight selection process to systematically model the Pareto front – the set of optimal trade-off solutions. We compare two methods for exploring this front: uniform grid search and multi-objective Bayesian optimization (MOBO). Our results demonstrate that the MOBO approach achieves superior performance across multiple quality metrics while requiring fewer evaluations to match or exceed the performance of uniform grid search.

2 Methodology and Results

We employ the same experimental setup as in our previous work [1], using the Quanser Aero 2 in 1-DoF configuration. To automate the exploration of the Pareto front implied by the reward function R_t , we compare two distinct sampling strategies for the weight parameter α : (1) uniform grid search with 11 equally spaced α values, and (2) Bayesian optimization (BO) using qEHVI acquisition function with 5 initial Sobol points and 6 BO-guided selections, each trained and assessed under identical conditions.

We evaluate the quality of the resulting Pareto front using three standard metrics: hypervolume (measuring the dominated area), spacing (evaluating the uniformity of the solution distribution), and maximum spread (assessing the range coverage). Table 1 summarizes the comparative results.

Table 1. Comparison of Pareto front quality metrics

Metric	Grid Search	BO	Difference
Hypervolume	762.7772 ↓	769.4843 ↑	+6.7071
Spacing	12.2595 ↓	11.2163 ↑	-1.0432
Maximum Spread	51.3134 ↓	51.3259 ↑	+0.0124

The results demonstrate that the automated BO approach consistently outperforms uniform grid sampling across all metrics. Notably, BO achieves superior hypervolume and a better maximum spread, indicating a higher-quality and broader Pareto front approximation. Fig. 1 illustrates the progression of hypervolume and maximum spread relative to the final grid search performance. The rapid improvement in early trials highlights BO's sample efficiency in intelligently identifying the most promising regions of the weight space.




Fig. 1. Progression of hypervolume and maximum spread vs. grid search baseline.

These findings validate BO as an effective method for automating multi-objective weight selection in industrial control applications, significantly reducing the number of evaluations needed to discover high-quality trade-off policies.

References

1. Schäfer, G., Seliger, R., Rehrl, J., Huber, S., Hirlaender, S.: Multi-Objective Reinforcement Learning for Energy-Efficient Industrial Control. In: Database and Expert Systems Applications - DEXA 2025 Workshops. pp. 67–72. Communications in Computer and Information Science, Springer Charm, Bangkok, Thailand (Aug 2025). <https://doi.org/10.1007/978-3-032-02003-1>

Powder Bed Defect Detection in L-PBF Using YOLO RT-DETR

Pit Schulenburg and Cindy Buhl

University of Applied Sciences Kempten, Bahnhofstraße 61, 87435 Kempten Germany
 pit.schulenburg@stud.hs-Kempten.de, cindy.buhl@hs-kempten.de

Abstract. Laser Powder Bed Fusion (L-PBF) relies on consistent powder distribution and stable melt dynamics, yet small deviations in the powder bed can lead to significant defects in the final part. This work presents a defect-detection approach based on YOLO RT-DETR applied to layer-wise imaging of the powder bed. The model reached strong detection performance and outperformed a manually labeled reference set. The results show that real-time defect recognition in L-PBF is feasible with limited data, and non-specialized camera hardware.

Keywords: Laser Powder Bed Fusion · Additive Manufacturing · Powder Bed Monitoring · Defect Detection · YOLO RT-DETR · In-process Quality Assurance · Metal 3D Printing

1 Introduction

Laser Powder Bed Fusion (L-PBF) is commonly used for producing metal components with high geometric precision [2]. The process depends on stable powder spreading and controlled laser melting, yet small disturbances at the layer level can introduce defects that affect part quality [1]. Early identification of such irregularities is essential to prevent defects from propagating through subsequent layers.

Optical monitoring has become a practical tool for observing powder-bed conditions, as it allows non-invasive tracking of the build surface and visible powder-bed anomalies. However, the appearance of defects varies strongly with geometry and process parameters, which complicates robust detection.

This work examines whether YOLO RT-DETR [3] can reliably detect layer-wise powder-bed defects under realistic constraints. The focus is on real-time applicability, minimal data requirements, and robustness despite uneven class distributions. Even with a small training set, the model demonstrates that lightweight architectures can support in-process monitoring without specialized imaging hardware or extensive annotation efforts.

2 Learning-Based Detection Framework

Accurate defect detection in L-PBF requires capturing subtle variations in powder distribution and melt behavior [1]. Our approach combines standard optical

imaging of each layer with a transformer-based real-time detector. The aim is to assess whether such a compact system performs reliably when training data are limited and defect classes are unevenly represented.

Layer-wise images were recorded using a camera positioned inside the machine over the powder bed. Typical defect patterns included insufficient spreading, local elevations, pores, recoater streaks, and displaced debris, which have been reported as critical surface-visible anomalies in L-PBF [1, 2]. A dataset of 1,200 images was manually labeled, reflecting the irregular occurrence of defects in practical builds.

YOLO RT-DETR was trained from pretrained weights with fixed-resolution inputs and standard augmentation. Training converged quickly and required modest computational resources. Performance was evaluated on unseen layers from different builds to ensure generalization across geometries and powder-bed conditions.

The model consistently recognized the most relevant defect types and remained stable even in borderline cases. Despite the small dataset, it surpassed a manually curated reference set. These results indicate that compact transformer-based detectors can provide reliable insight into powder-bed anomalies and serve as a foundation for future in-process control strategies.

References

1. Grasso, M., Colosimo, B.M.: Process defects and in situ monitoring methods in metal powder bed fusion: a review. *Measurement Science and Technology* 28(4), 044005 (2017)
2. Zitelli, C., Folgarait, P., Di Schino, A.: Laser Powder Bed Fusion of Stainless Steel Grades: A Review. *Metals* 9(7), 731 (2019)
3. Zhao, Y., Lv, W., Xu, S., Wei, J., Wang, G., Dang, Q., Liu, Y., Chen, J.: DETRs Beat YOLOs on Real-time Object Detection. arXiv preprint arXiv:2304.08069 (2023)

Exploring Zero-Shot Foundation Models for Multivariate Time Series Anomaly Detection

Martin Uray^{1,2}, Saverio Messineo¹, Roland Kwitt², and Stefan Huber¹

¹ Josef Ressel Centre for Intelligent and Secure Industrial Automation,
Salzburg University of Applied Sciences, Austria

² Department of Artificial Intelligence and Human Interfaces,
Paris Lodron University of Salzburg, Austria
martin.uray@fh-salzburg.ac.at

1 Motivation and Methodology

Multivariate Time Series Anomaly Detection (MTSAD) plays a vital role in ensuring reliability and safety across numerous application domains, such as industrial process monitoring and financial risk management [2]. Forecasting-based approaches are commonly used for MTSAD, where models learn to predict future observations from historical data, e.g. $\hat{X}_{t+1} = f_\theta(X_{t-w:t})$, with the parametrized model f_θ learning to forecast future values based on a context window of length w . Here, the multivariate time series of length T is denoted by $X = \{x_t\}_{t=1}^T$, $x_t \in \mathbb{R}^N$ with N as the number of features. $X_{t,i} \in \mathbb{R}$ denotes the value of feature i observed at time step t . Anomalies are then identified by quantifying the normed deviation between the predicted and the true values, e.g. $e_t = \|\hat{X}_{t+1} - X_{t+1}\|$, and by applying a threshold to this score.

A Foundation Model (FM) is a large-scale, pre-trained model trained on diverse data sources, designed to be adaptable to a wide range of downstream tasks with minimal or even no fine-tuning. Moreover, such models have demonstrated remarkable zero-shot performance across various domains. Recently, similar models have also become available for univariate time series forecasting [1].

A naive zero-shot application of FMs to MTSAD offers potential advantages. First, these models are trained on diverse datasets, enabling strong generalization to previously unseen data. Second, their ability in a zero-shot setting can eliminate the need for costly and energy-intensive training from scratch. However, unlike their usage with univariate data, their zero-shot application to MTSAD remains largely underexplored.

2 Experimental Evaluation

MTSAD via Forecasting. To naively apply a univariate forecasting FM f_θ with pre-trained model weights θ to MTSAD in zero-shot, each feature n of the time series X can be treated independently, $\hat{X}_{t+1,n} = f_\theta(X_{t-w:t,n})$, while accepting to lose the spatial dependency between features. The quantitative results on benchmark datasets indicate that this approach does not outperform state-of-the-art

or even standard baseline methods. A qualitative analysis further reveals that the error score increases primarily during transitions into anomalous segments. Once the context window contains only anomalous samples, the error signal drops again, failing to consistently indicate the presence of anomalies. The same effect is observed during transitions from anomalous to benign segments.

MTSAD Using Pre-Trained Representations. To better capture spatial dependencies between features, we further evaluate the anomaly detection task using intermediate representations of the time series. We decompose the FM f_θ into its components $f_\theta = h_\theta \circ g_\theta$, where g_θ generates embeddings from the input time series and h_θ performs forecasting based on these embedded representations. In this setup, we use only g_θ to extract the embeddings for further analysis. For a joint representation, we concatenate all features embeddings $z_t = [g_\theta(X_{t-w:t,0}), g_\theta(X_{t-w:t,1}), \dots, g_\theta(X_{t-w:t,n})]$, where $z_t \in \mathbb{R}^d$ with d being the dimension of the embeddings, and each $g_\theta(X_{t-w:t,i})$ represents the embedding of the univariate feature i over the context window $[t-w, t]$.

Using generated toy data, we computed the embeddings and analyzed them visually. The high-dimensional embeddings were reduced to three dimensions, revealing that benign and anomalous data can be readily separated in this lower-dimensional space. Applying the same procedure to benchmark datasets yields a different outcome. A clear visual separation between benign and anomalous data is no longer observable. Furthermore, standard outlier detection methods, such as Isolation Forest, fail to effectively learn the benign distribution from the training data and consequently struggle to identify anomalies in the test data.

3 Discussion

The results indicate that a naive zero-shot application of FMs is not competitive for the task of MTSAD. One possible explanation is that the employed FM is *too effective* at capturing the temporal dynamics of the time series, leaving little forecasting error once the model has adapted to the data. This becomes evident in the qualitative analysis, where forecasts based on context windows consisting entirely of anomalous data yield minimal forecasting error. Nevertheless, the same experiment provides valuable insights: FMs are indeed capable of detecting changes in the underlying data distribution, making them promising candidates for change-point detection. To further assess their applicability to MTSAD, fine-tuning the FM on the target training data should be considered.

References

1. Das, A., Kong, W., Sen, R., Zhou, Y.: A decoder-only foundation model for time-series forecasting. In: Proceedings of the 41st International Conference on Machine Learning. ICML'24, JMLR (2024)
2. Sarfraz, M.S., Chen, M.Y., Layer, L., Peng, K., Koulakis, M.: Position: quo vadis, unsupervised time series anomaly detection? In: Proceedings of the 41st International Conference on Machine Learning. ICML'24, JMLR (2024)

Operator-in-Loop Validation and Annotation Framework for AI-Driven QA in Metal AM*

Faiza Waheed, Kai Höfig, and Fabian Riß

Technische Hochschule Rosenheim, Germany

{faiza.waheed, kai.hoefig, fabian.riss}@th-rosenheim.de

Abstract. Metal additive manufacturing (AM), particularly laser powder bed fusion (L-PBF), is widely used in high-reliability sectors such as aerospace, medical, and automotive due to its design freedom and ability to produce complex geometries. However, it also poses major quality assurance (QA) challenges, with defects such as porosity, cracks, or lack of fusion occurring inconsistently. Recent research has explored in-situ monitoring and deep learning-based defect classification using layer-wise imaging data [2], yet many systems lack human validation, limiting model performance and operator trust. In our previous work [1], we presented a prototype AI-driven defect detection framework for the TRUMPF TruPrint 1000, combining visual layer monitoring with real-time defect inference. However, the previous system did not incorporate human validation or retraining, which we address in this work.

Keywords: Defect Detection · Quality Assurance · Deep Learning · Operator in Loop · Annotation Framework · Additive Manufacturing

1 Methods and Results

We present a human-in-the-loop validation and annotation framework to bridge AI-based defect detection and production-level QA requirements. The system architecture includes a dashboard-based interface for operators to verify model predictions in real-time and on historical data. When the deep learning model detects a potential defect, the operator reviews the prediction via the dashboard and uses a checkbox mechanism to confirm or override it. Each validated inference is stored as a JSON document in MongoDB, along with metadata (Production_ID, Slice_number, Defect_class, Operator_feedback, and other machine parameters), forming a curated annotation pool for iterative retraining.

In our pilot deployment with L-PBF print jobs on the TRUMPF TruPrint 1000, the model's recall on the groove defect class improved from 50% to 65%. Operators reported that the checkbox validation took 40-50 seconds per slice, with minimal fatigue. The annotation schema enabled efficient indexing and querying in MongoDB, and annotations were consistent with defect class labels and operator agreement. The system is low-latency, reliable, and operator-friendly. Figure 1 shows Precision-Recall curves before and after retraining with

* Funded by INTERREG Bayern-Österreich (AI4Green, 01.05.2024–30.04.2027).

validated annotations, illustrating improvement in recall, specifically for the groove defect class.

This work presents a practical method for integrating human oversight into an AI-based quality assurance workflow for metal AM. By capturing validated annotations directly from the operator and feeding them into the model retraining pipeline, the system supports a continuously improving defect detection model. Our contributions include: (1) a validation UI embedded in a live QA dashboard; (2) a structured, scalable annotation database schema; and (3) a pilot study demonstrating measurable improvements in model performance and user satisfaction. Future work will add real-time alerts for ambiguous predictions and leverage explainable AI to improve operator trust.




Fig. 1: PR-curves before and after retraining with validated annotations

2 Conclusions

By integrating human-validated annotations into the retraining pipeline, the system enables continuous improvement of defect detection in metal AM. Our contributions include: (i) a live QA dashboard with validation UI; (ii) a structured, scalable annotation database; and (iii) a pilot study demonstrating measurable model and user improvements. Future work will add real-time alerts for ambiguous predictions and leverage explainable AI to improve the system.

Acknowledgments. This work is funded by INTERREG Bayern-Österreich (AI4Green, 01.05.2024 - 30.04.2027)

References

1. Waheed, F., Höfig, K., Riß, F.: AI-Driven Data Management Framework for Quality Assurance in Additive Manufacturing: A Case Study on the TRUMPF TruPrint 1000. Springer (2025). https://doi.org/10.1007/978-3-032-05073-1_1
2. Zhao, Y., Ren, H., Zhang, Y., Wang, C., Long, Y.: Layer-wise multi-defect detection for laser powder bed fusion using deep learning algorithm. Optics and Laser Technology (2024), <https://doi.org/10.1016/j.optlastec.2024.110648>

AI4Green: Harnessing the Ecological Potential of AI in Industry*

Kai Höfig¹, Mario Döller², Faiza Waheed¹, Fabian Riß¹, Michael Scholz³, Joerg Bauer³, Hannes Waclawek⁴, Georg Schäfer⁴, Stefan Huber⁴, Bernhard Heinzl⁵, Michael Hellwig⁶, and Steffen Finck⁶

¹ Technische Hochschule Rosenheim, Germany

² Kufstein University of Applied Sciences, Austria

³ Technische Hochschule Deggendorf, Germany

⁴ Salzburg University of Applied Sciences, Austria

⁵ Software Competence Center Hagenberg GmbH, Austria

⁶ Vorarlberg University of Applied Sciences, Austria

Abstract. This paper summarizes our ongoing research from the *AI4Green* project funded by INTERREG Bayern-Österreich. We outline challenges and opportunities for AI-based resource optimizations in challenging domains such as agriculture, robotics and production processes. We conclude that the optimization potential is high and provide an outlook of our future research activities to develop resource-optimal industrial AI-based solutions.

Keywords: Defect Detection · Quality Assurance · Deep Learning · Operator in Loop · Annotation Framework · Additive Manufacturing

Pilot Projects and Results

With global challenges such as climate change, increasing environmental pollution and rising resource consumption, it is becoming increasingly important to find sustainable and also reliable solutions to ensure a liveable and safe future for generations to come [1]. Significant opportunities for savings lie in the field of business and industry, particularly by improving production processes [2]. Through three pilot initiatives in the fields of agriculture, robotics, and manufacturing, we aim to demonstrate existing optimization potential using modern AI-based methods.

The first pilot action focuses on CO₂ reduction in **agriculture**. Mobile AI-based sensor platforms, such as drones, UAVs, etc. are being used for tasks like monitoring crop health, planning irrigation, and even automating afforestation or harvesting, allowing for more precise and efficient use of resources. The coordination of multiple mobile sensor platforms as a swarm is a key factor in this automation, playing an increasingly crucial role in optimizing their performance. Using automated coordination and information exchange between aerial

* Funded by INTERREG Bayern-Österreich (AI4Green, 01.05.2024–30.04.2027).

and ground vehicles, drones perform aerial scanning and image data collection, for ground vehicles to carry out targeted actions like detection of storm-induced crop damage, large-scale weed mapping, and identification of *Jacobaea vulgaris* (Jakobskreuzkraut), a toxic plant dangerous to livestock. This cooperation ensures precise intervention, reduces overlaps and unnecessary operations, and ultimately saves resources.

In industries such as automotive manufacturing, **industrial robots** are major energy consumers. Optimizing their motion trajectories holds significant potential for reducing overall energy consumption. Classical control theory remains the de facto standard for trajectory planning. However, modern data-driven machine learning approaches – particularly Reinforcement Learning (RL) – have emerged as a promising alternative. We introduce novel methods to enable RL in industrial environments for energy-efficient trajectory planning demonstrating – in both simulation and on a real-world robotic system – that careful problem formulation can dramatically increase sample efficiency. In parallel, we investigate Physics-Informed Reinforcement Learning (PiRL) as a complementary approach that embeds domain knowledge. In robotics control, PiRL has shown promise across a range of applications, e.g., locomotion or manipulation. Here, physics-based surrogate models have been used to accelerate policy learning in high-dimensional control spaces.

In **production**, we collect data to predict the demand and the quality of products. We integrate meal participation forecasts into main dish prediction through the refinement of simulated extrapolation approaches to decrease food waste. These forecasting advancements will form the basis for developing meta-heuristic optimization methods tailored to meal production planning. Complementary to the food production optimization project, we address AI-driven resource savings in Additive Manufacturing (AM), specifically in the field of Laser Powder Bed Fusion (L-PBF) for 3D metal components. We prevent the production of defective parts by early detection of failure points, thus reducing associated material waste by an estimated $\sim 30\%$ and energy consumption by $\sim 15\%$ compared to post-build non-destructive inspection methods.

We addressed resource optimizations using AI-based algorithms in the aforementioned domains and have shown that the potential is high to improve the ecological footprint.

Acknowledgments. This work is funded by INTERREG Bayern-Österreich (AI4Green, 01.05.2024 - 30.04.2027)

References

1. European Commission: The european green deal. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en (2020), [Accessed: April 23, 2025]
2. European Commission: Clean industrial deal. https://commission.europa.eu/topics/eu-competitiveness/clean-industrial-deal_en (2024), accessed: April 23, 2025

Shock models driven by a two-state Markov modulated Poisson process

Nicola Giordano¹, Barbara Martinucci^{1[0000-0001-8340-4200]}, and Paola Paraggio^{1[0000-0002-3308-7937]}

Department of Mathematics, University of Salerno, Fisciano (SA), 84084, Italy
 {ngiordano, bmartinucci, pparaggio}@unisa.it

In the context of shock models, it is customary to assume that shocks arrive according to a homogeneous or non-homogeneous Poisson process. While such processes benefit from independent increments and thus offer substantial mathematical tractability, this assumption often proves inadequate in applications where external conditions evolve over time. This has motivated several authors to explore alternative counting processes for modeling shock occurrences. For instance, [2] examines a geometric counting process whereas generalized Pólya-type processes have been investigated in [1] and [5].

The present contribution focuses on shock models governed by a two-state Markov modulated Poisson process. The Markov-modulated Poisson process (MMPP) is a flexible stochastic extension of the classical Poisson process in which the rate parameter is governed by an underlying continuous-time Markov chain. This modulation mechanism enables the MMPP to capture complex real-world phenomena—such as bursty arrival patterns and time-varying event rates—commonly encountered in telecommunications, finance, and queueing theory (see, e.g., [3]). A notable special case of the MMPP is the Switched Poisson process (SPP), in which the underlying Markov chain is restricted to two states. Unlike the standard Poisson process, the SPP features dependent interarrival times due to the temporal correlation induced by the underlying Markov chain. Consequently, SPP-driven shock models provide a natural and effective framework for applications in which systems accumulate damage or experience failures at rates influenced by an unobservable or fluctuating environmental regime.

Building on the analytical results recently obtained in [4], we investigate both extreme and cumulative shock models driven by a Switched Poisson process. In particular, we derive certain analytical results concerning the failure time, including its survival function and the main moments.

References

1. Cha J.H.: Characterization of the generalized Pólya process and its applications. *Advances in Applied Probability* 46(4), 1148-1171 (2014).
2. Di Crescenzo A., Pellerey F: Some Results and Applications of Geometric Counting Processes. *Methodology and Computing in Applied Probability* 21, 203-233 (2019).
3. Fischer, W., Meier-Hellstern, K.: The Markov-modulated Poisson process (MMPP) cookbook. *Performance Evaluation* 18, 49–171 (1993)

4. Giordano, N., Martinucci, B., Paraggio, P.: On a two state Markov modulated Poisson process and its time changed variants. Submitted (2025)
5. Konno H.: On the exact solution of a generalized Pòlya process. *Advances in Mathematical Physics*, John Wiley & Sons (2010)

Varentropy-Based Analysis of Precipitation Variability

Antonio Di Crescenzo^{1[0000-0003-4751-7341]}, Serena Spina^{1[0000-0001-6408-7596]},
and Paolo Villani^{2[0000-0002-0951-0336]}

¹ Department of Mathematics, University of Salerno, Via Giovanni Paolo II 132,
84084 Fisciano (SA), Italy – email: {adicrescenzo,sspina}@unisa.it

² Department of Civil Engineering, University of Salerno, Via Giovanni Paolo II 132,
84084 Fisciano (SA), Italy – email: p.villani@unisa.it

Abstract

In recent decades, the application of Information Theory to problems of system identification in the Earth and Environmental Sciences has been characterized by renewed interest. Although the concept of entropy has found increasingly wide application, less attention has been devoted so far to varentropy, i.e., the variance of entropy.

For a discrete random variable X , with probability distribution $p(x)$, the varentropy is defined as:

$$V(X) = \text{Var}[-\log p(X)] = E[(-\log p(X))^2] - H(X)^2,$$

where $H(X) = -\sum_x p(x) \log p(x)$ represents the Shannon entropy. Recently, the concepts of weighted varentropy, residual varentropy and related estimators have been also introduced and studied (see, for example, [5] and reference therein).

While the Shannon entropy $H(X)$ expresses the mean information content of a set of observations, varentropy measures how much the information fluctuates around its mean value.

Shannon entropy has already been widely used to analyze the variability and complexity of daily rainfall data.

For example, in [4], the authors employed entropy to evaluate the maximum achievable information content of an urban rainfall network, analyzing the sensitivity of information to both measurement density and temporal resolution.

More recently, the Multiscale Sample Entropy has been applied in [6] to analyze the complexity of daily rainfall patterns from a network of over 600 meteorological stations, demonstrating that this measure can reveal significant differences in rainfall regimes across temporal scales.

These studies confirm that entropic measures represent powerful tools for investigating precipitation processes, providing a quantitative basis for the characterization of uncertainty and complexity in hydrological dynamics. However, entropy alone is not capable of describing the variability of informational uncertainty in time or space. In this context, varentropy represents a natural extension, useful to quantify the dispersion of information associated with rainfall processes

and offering a more comprehensive description of the statistical structure and informational stability of precipitation regimes.

Our study focuses on the analysis of the varentropy associated with daily rainfall amounts recorded at Italian stations, where the historical series is characterized by a high frequency of zero values (days without rain). Specifically, considering the random variable X representing daily rainfall amount, the probability distribution $p(x)$ is highly skewed and modeled through a random vector that highlights the component corresponding to zeros and the remaining component associated with positive rainfall values. In the literature, this random variable has been modeled using a discrete component for zeros and an absolutely continuous component for positive rainfall (see, for example, [2], [3] and references therein).

We propose a different mathematical formulation, by taking into account that this feature influences the informational structure of the data and, consequently, the estimated variance of entropy. We consider also a Bernoulli random variable $Y = \mathbf{1}_{\{X>0\}}$ such that

$$P(X = 0) = p(0, 0) = \gamma, \quad P(X > 0) = P(Y = 1) = 1 - \gamma = \sum_{x>0} p(x, 1),$$

where $p(x, y) = P(X = x, Y = y)$ and $0 < \gamma < 1$.

We analyze entropy- and varentropy-based measures of (X, Y) , highlighting the contributions due to zero-valued and positive data, respectively. This allows us to characterize not only the average level of uncertainty of the local atmospheric system but also its informational stability and coherence over time.

Finally, we focus on the problem of analyzing the rare events, specifically those related to intense rainfall (see [1] and reference therein). As a first step, we employ a weighted varentropy to focus attention on catastrophic events. Then, starting from the study of a residual varentropy estimator, we define a criterion to select the most significant data (i.e., the most intense rainfall events) based on a hypothesis testing approach.

References

1. Avino, A., Cimorelli, L., Furcolo, P., Noto, L.V., Pelosi, A., Pianese, D., Villani, P., Manfreda, S.: Are rainfall extremes increasing in southern Italy? *Journal of Hydrology* **631** 130684 (2024)
2. Gong, W., Yang, D., Gupta, H. V., Nearing, G.: Estimating information entropy for hydrological data: One-dimensional case. *Water Resour. Res.* **50**, 5003–5018 (2014)
3. Popuri, S. K., Nagaraj, K.N., Mehta, A., Mousavi, A.: Density estimation using entropy maximization for semi-continuous data. *Digital Signal Processing* **116** 103107 (2021)
4. Ridolfi, E., Alfonso, L., Di Baldassarre, G., Napolitano, F.: An entropy approach for evaluating the maximum information content achievable by an urban rainfall network. *Natural Hazards and Earth System Sciences* **11** 2075–2083 (2011)
5. Saha, S., Kayal, S.: Weighted (residual) varentropy and its applications. *Journal of Computational and Applied Mathematics* **442** 115710 (2024)
6. Zhou, J., Zhang, X., Wang, S., Ma, Z.: Rainfall patterns from multiscale sample entropy analysis. *Frontiers in Water* **4** 885456 (2022)

On the rejection sampling method for multidimensional diffusions through a drift-based transformation

Antonio Di Crescenzo^{1[0000-0003-4751-7341]}, Verdiana Mustaro^{1[0000-0003-4583-2612]}, and Serena Spina^{1[0000-0001-6408-7596]}

Dipartimento di Matematica, Università degli Studi di Salerno,
Via Giovanni Paolo II, 132, 84084 Fisciano (SA), Italy
{adicrescenzo,vmustaro,sspina}@unisa.it

Abstract. We consider a class of drift-based transformations for multidimensional diffusion processes, first studied in [1]. For such processes, the transition probability density function (p.d.f.) of the transformed process, namely $f(\mathbf{x}, t | \mathbf{y}, \tau)$, is related to the p.d.f. of the original one, i.e. $\hat{f}(\mathbf{x}, t | \mathbf{y}, \tau)$, through a product-form representation. The construction is driven by a weight function w , obtained as the solution of a suitable partial differential equation. In this framework, a key feature is that it yields a particularly transparent relation between the transition densities of the original and transformed diffusions. In fact, their ratio reduces to the simple expression

$$\frac{f(\mathbf{x}, t | \mathbf{y}, \tau)}{\hat{f}(\mathbf{x}, t | \mathbf{y}, \tau)} = \frac{w(\mathbf{x})}{w(\mathbf{y})}, \quad (1)$$

which highlights how the entire effect of the transformation can be captured through the weight function. This identity plays a key role in several analytical and computational settings: for instance, it suggests a practical application in rejection sampling schemes. The idea behind this scheme is to sample a relatively complex and nonstandard p.d.f., namely f , by exploiting a proposal p.d.f., such as \hat{f} , that we assume can be easily evaluated. The transformed process can be efficiently sampled, provided that a constant k exists such that $k \hat{f} \geq f$ on the domain. The acceptance probability reduces to $1/k$, and the representation (1) allows the acceptance rule to be evaluated in a simple and computable form. Although the framework applies broadly, we focus here on the Wiener and Ornstein-Uhlenbeck processes as tractable examples to illustrate the transformation. For these cases, numerical simulations are performed to assess the efficiency and practical tractability of the proposed sampling method.

References

1. A. Di Crescenzo, V. Mustaro, S. Spina, A framework for drift transformations of multidimensional diffusion processes with applications to Wiener and Ornstein-Uhlenbeck dynamics. Under review.

2. A. Di Crescenzo, V. Giorno, A. G. Nobile, L.M. Ricciardi, On a symmetry-based constructive approach to probability densities for two-dimensional diffusion processes, *J. Appl. Probab.* 32 (1995) 316-336. <https://doi.org/10.2307/3215291>
3. E. Di Nardo, G. D'Onofrio, T. Martini, Tommaso, Orthogonal gamma-based expansion for the CIR's first passage time distribution, *Appl. Math. Comput.* 480 (2024). <https://doi.org/10.1016/j.amc.2024.128911>
4. V. Giorno, A. G. Nobile, On the construction of a special class of time-inhomogeneous diffusion processes, *J. Stat. Phys.* 177 (2019) 299–323. <https://doi.org/10.1007/s10955-019-02369-2>
5. V. Giorno, A. G. Nobile, E. Pirozzi, L. Ricciardi, On the construction of first-passage-time densities for diffusion processes, *Sci. Math. Jpn.* 64 (2006).
6. H. Maatouk, X. Bay. A new rejection sampling method for truncated multivariate Gaussian random variables restricted to convex sets. In: Cools, R., Nuyens, D. (eds) Monte Carlo and Quasi-Monte Carlo Methods. Springer Proc. Math. Stat. (2016), vol 163. Springer, Cham. https://doi.org/10.1007/978-3-319-33507-0_27

Testing heteroscedasticity in diffusion processes*

Giuseppina Albano¹[0000-0002-2317-0331]

Dip. di Studi Politici e Sociali, Università di Salerno, 84084 Fisciano (SA), Italy,
pialbano@unisa.it

Abstract. Diffusion processes derived from deterministic growth equations are widely employed in modeling population dynamics. While classical deterministic models capture the average trajectory of the system, stochastic extensions provide a more realistic description of the observed dynamics. In particular, when the variance of the stochastic component changes with the level of the process or with time, the system exhibits heteroskedasticity. Traditional heteroskedasticity tests are based on regression residuals and assume a fixed model structure. However, in the context of diffusion processes, such tests are not directly applicable. We propose a heteroskedasticity test in diffusion processes that arise from stochastic versions of growth equations. The test relies on a distance measure between probability distributions, specifically the Jensen–Shannon divergence, by comparing the empirical distribution of simulated data to the theoretical normal distribution implied by the homoskedastic model.

Keywords: Growth equation · Heteroscedasticity test · Monte Carlo simulation.

1 Randomness and heteroscedasticity in growths

Let $N(t)$ be the population size at time $t \geq 0$. The exponential growth equation is described by the following Ordinary Differential Equation:

$$dN(t) = [f(N(t))]N(t) dt \quad (1)$$

in which we include the birth rate possible depending on $N(t)$. In order to include randomness in (1), we can consider the following two stochastic generalizations:

1. $f(N(t)) \implies f(N(t)) + \sigma \frac{dW_t}{dt}$
2. $[f(N(t))]N(t) \implies [f(N(t))]N(t) + \sigma \frac{dW_t}{dt}$

where $W(t)$ is a standard Brownian Motion. The resulting stochastic differential equations (SDEs) are:

1. $dN(t) = [f(N(t)) + \sigma \frac{dW_t}{dt}]N(t)dt = f(N(t))N(t)dt + \sigma N(t)dW_t$
2. $dN(t) = f(N(t))N(t)dt + \sigma dW_t$

* (Workshop “Stochastic Models, Statistical Methods, and Applied Systems Simulations”)

Clearly, the function $f(\cdot)$ and the parameter σ has to be such that the drift and the infinitesimal variance verify Lipschitz and linear growth conditions, in order to have the existence and the uniqueness of the solution of the SDEs. The analytical solution is determined by Ito's lemma.

Further, assuming that $P[N(0) = x_0] = 1$, we have that

1. the finite-dimensional distributions of the resulting process are lognormal;
2. the finite-dimensional distributions of the resulting process are normal;

We point out that the mean functions of the two processes are equal, and the only difference lies in the infinitesimal variances. Such variability leads to different probability distributions and consequently different inferences. In particular, case 1 is characterized by a process with variability that depends on the population size, while in case 2, the variability is constant over time.

2 The test

In our view, formulating the heteroskedasticity test is equivalent to identify the source of randomness-whether it lies in the drift term or solely in the birth rate function $f(\cdot)$. The null hypothesis H_0 corresponds to homoskedasticity, while the alternative hypothesis represents the general case of heteroskedasticity.

The test statistic is based on the Kullback–Leibler (KL) divergence, which here measures the distance between the normal distribution and the empirical distribution. The KL divergence between two probability distributions P and Q is defined as

$$D_{\text{KL}}(P|Q) = \int p(x) \log \frac{p(x)}{q(x)} dx, \quad (2)$$

where $p(x)$ and $q(x)$ denote the probability density functions of P and Q , respectively. Note that the KL divergence is not symmetric and to address this, we adopt its symmetric counterpart, the Jensen–Shannon (JS) divergence.

Under the null hypothesis H_0 of homoskedasticity, the theoretical distribution P follows a normal law, while the empirical distribution Q is obtained from sample data. Since the empirical distribution does not have a natural continuous density, we estimate the density $q(x)$ using a kernel density estimator. Finally, since the sampling distribution of the test statistic under H_0 is not available in closed form, the critical values are obtained via Monte Carlo simulations.

References

1. Chen, Q., Hu, M., Song, X. A nonparametric specification test for the volatility functions of diffusion processes. *Econometric Reviews*, 38(5), 557–576 (2017).
2. Osterreicher, F., Vajda, I., A new class of metric divergences on probability spaces and its statistical applications. *Annals of the Institute of Statistical Mathematics* **55** 3, 639–653 (2003).
3. White, H., A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica*, **48**(4):817–838 (1980).

Some time-inhomogeneous pseudo-first-order kinetic models in random environment*

Virginia Giorno^[0000-0001-6474-3952] and Amelia G. Nobile^[0000-0003-2363-7496]

Dipartimento di Informatica, Università di Salerno, Via Giovanni Paolo II, n. 132, 84084 Fisciano (SA), Italy, {giorno,nobile}@unisa.it

Abstract. We consider two time-inhomogeneous stochastic models constructed starting from a linear first-order ordinary differential equation. The deterministic version, known also as Lagergren model, is used in population dynamics and in chemistry as an empirical model for the kinetics of absorption and chemical reactions. For both stochastic models with constant carrying capacity, we determine the transition probability density functions and the related conditional moments. Particular attention is paid to the first-passage time problem through a constant boundary for time-homogeneous processes and for time-inhomogeneous processes with proportional intensity functions. Various numerical computations and simulations are performed to compare the proposed models.

Keywords: Growth curves · Transition distributions · First-passage times

1 The models

Various phenomena in a variety of scientific fields, such as biology, medicine, epidemiology, ecology, engineering and chemistry are usually modeled via ordinary differential equations (see Allen [1] and Banks [2]). The pseudo-first-order (PFO) kinetic model, also known as Lagergren model, is widely used in population dynamics (cf. Mir and Dubeau [3], Giorno and Nobile [4]) and in environmental studies as an empirical model for describing adsorption kinetics and chemical reactions (cf. Liu and Shen [5], Revellame et al. [6]).

Let $x(t)$ be the size of the population at time t . The deterministic version of the PFO model is described by the ordinary differential equation (ODE):

$$\frac{dx(t)}{dt} = \lambda(t) [K(t) - x(t)], \quad x(t_0) = x_0, \quad (1)$$

with $x_0 < K(t_0)$, where $\lambda(t)$ is a positive continuous function representing the intrinsic growth intensity function and $K(t)$ represents the time-varying environmental carrying capacity, i.e. the maximum sustainable size of the population. In the Lagergren model, $x(t)$ is the adsorption capacity at time t , $K(t) = K$ is

* (Workshop “Stochastic Models, Statistical Methods, and Applied Systems Simulations”)

the adsorption capacity at equilibrium and $\lambda(t) = \lambda > 0$ is the speed constant. Denoting by $\Lambda(t|t_0) = \int_{t_0}^t \lambda(\theta) d\theta$ the cumulative intensity function, the solution of (1) is

$$x(t) = K(t) - [K(t_0) - x_0] e^{-\Lambda(t|t_0)} - \int_{t_0}^t K'(\theta) e^{-\Lambda(t|\theta)} d\theta. \quad (2)$$

Several approaches can be used to construct the stochastic generalizations of $x(t)$ (see Cortés et al. [7], Giorno and Nobile [8]). We use the stochastic differential equations (SDE) approach by introducing the random environmental fluctuations in two different ways. In the first approach, the stochastic process $X_1(t)$ is constructed starting from the ODE (1) by changing $\lambda(t)$ with $\lambda(t) + \sigma(t) \xi(t)$, where $\xi(t)$ denotes a white noise process and $\sigma^2(t)$ is the noise intensity function. In the second approach the stochastic process $X_2(t)$ is obtained from the solution (2) by interpreting the cumulative intensity function $\Lambda(t|t_0)$ as the mean of a time-inhomogeneous Wiener process $Z(t) = \Lambda(t|t_0) + \int_{t_0}^t \sigma(\tau) dW(\tau)$, where $W(t)$ is the standard Wiener process. The stochastic processes $X_1(t)$ and $X_2(t)$ are time-inhomogeneous diffusion processes with state-space $(-\infty, K(t))$, characterized by infinitesimal drift $\alpha(t)[K(t) - x]$ and infinitesimal variance $\sigma^2(t)[K(t) - x]^2$, where $\alpha(t) = \lambda(t)$ for $X_1(t)$ and $\alpha(t) = \lambda(t) - \sigma^2(t)/2$ for $X_2(t)$.

Making use of theoretical and simulation methods, we compare the transition probability density functions, the first-passage time densities through a constant boundary and some statistically significant quantities for various choices of the parameters.

References

1. Allen, L.J.S.: An Introduction to Stochastic Processes with Applications to Biology. Chapman and Hall/CRC, Lubbock, Texas, USA (2010)
2. Banks, R.B.: Growth and Diffusion Phenomena. Mathematical Frameworks and Applications. Springer-Verlag, Berlin (1994)
3. Mir, Y., Dubeau, F.: Linear and logistic models with time dependent coefficients. *Electronic Journal of Differential Equations* **2016**(18), 1–17 (2016)
4. Giorno, V., Nobile A.G.: Some time-homogeneous diffusion models for population growth in random environments. *Commun Nonlinear Sci Numer Simulat* **142**, Article number: 108502 (2025)
5. Liu, Y., Shen, L.: From langmuir kinetics to first and second-order rate equations for adsorption. *Langmuir* **24**, 11625–11630 (2008)
6. Revellame, E.D. , Fortela, D.L., Sharp, W., Hernandez, R., Zappi, M.E.: Adsorption kinetic modeling using pseudo-first order and pseudo-second order rate laws: A review. *Cleaner Engineering and Technology* **1**, Article number: 100032 (2020)
7. Cortés, J.-C., Navarro-Quiles, A., Santonja, F.-J., Sferle, S.-M.: Statistical analysis of randomized pseudo-first/second order kinetic models. Application to study the adsorption on cadmium ions onto tree fern. *Chemometrics and Intelligent Laboratory Systems* **240**, Article number: 104910 (2023)
8. Giorno, V., Nobile A.G.: Richards and Gompertz stochastic growth models with time-varying carrying capacity. *Computational and Applied Mathematics* **44**, Article number: 387 (2025).

On time-changed Gauss-Markov processes, first passage times, and simulations

Luigia Caputo¹[0000-0002-1618-7506]

Maria Francesca Carfora²[0000-0002-4570-1690]

Kei Kobayashi³[0000-0002-4913-7185], Enrica Pirozzi⁴[0000-0001-8692-2759]

¹ Dipartimento di Matematica e Applicazioni “Renato Caccioppoli”, Università degli Studi di Napoli Federico II, Napoli, Italy

luigia.caputo@unina.it

² Istituto per le Applicazioni del Calcolo “Mauro Picone”, Consiglio Nazionale delle Ricerche, Napoli, Italy

mariafrancesca.carfora@cnr.it

³ Department of Mathematics, Fordham University, New York, US

kkobayashi5@fordham.edu

⁴ Dipartimento di Matematica e Fisica, Università degli Studi della Campania “Luigi Vanvitelli”, Caserta, Italy

enrica.pirozzi@unicampagna.it

Due to their central role in the theory of stochastic processes, Gauss-Markov (GM) processes are extremely useful not only from the mathematical point of view but also in stochastic modeling in many fields, such as biology, population dynamics, neuroscience, financial markets, and so on (see, for instance, [1]). Specifically, a GM process $X(t)$ can be identified as the solution of the following Stochastic Differential Equation (SDE) ([2])

$$dX(t) = [a(t)X(t) + b(t)]dt + g(t)dB(t), \quad t > 0, \quad X(0) = x_0,$$

where $B(t)$ is a standard Brownian motion and all coefficient functions $a(t)$, $b(t)$, $g(t)$ satisfy the well-known assumptions to guarantee the existence of the solution ([3]). The GM process can also be represented as ([4])

$$X(t) = m_X(t) + \eta(t)B(\rho(t))$$

with

$$m_X(t) = \mathbb{E}[X(t)]; \quad \text{Cov}[X(s), X(t)] = \gamma(s)\eta(t); \quad \rho(t) = \frac{\gamma(t)}{\eta(t)}$$

where $\rho(t)$ is a positive increasing function. The aim of the present work is to consider a time-changed GM (tcGM) process, i.e. the stochastic process obtained by composing a GM process $X(t)$ and the continuous inverse of an α -stable subordinator process $E_\alpha(t)$ ([5, 6]). The new process $X(E_\alpha(t))$ is neither Gaussian nor Markovian. However, it is the solution of the following time-changed SDE ([7])

$$dX(E_\alpha(t)) = [a(E_\alpha(t))X(E_\alpha(t)) + b(E_\alpha(t))]dE_\alpha(t) + g(E_\alpha(t))dB(E_\alpha(t)). \quad (1)$$

Very few results can be found about the first passage time (FPT) of such a process, and several different approaches have been adopted ([6, 8, 9]). We aim to investigate the FPT probability density function and to clarify its relationship with the subordinated FPT density ([6, 10]). As a first step, we perform simulations of sample paths of $X(E_\alpha(t))$ by an Euler-Maruyama discretization procedure applied to SDE (1), also based on an ad hoc routine for generating the random time $E_\alpha(t)$ ([11]). Such simulated paths will provide samples of FPT through a constant level S . Furthermore, it is also possible to consider

$$X(E_\alpha(t)) = m_X(E_\alpha(t)) + \eta(E_\alpha(t))B(\rho(E_\alpha(t))). \quad (2)$$

A simulation algorithm will be also applied to (2) in order to construct further sample estimates of FPT. Theoretical and numerical comparisons will be provided.

References

1. Leonenko, N., Pirozzi, E.: The time-changed stochastic approach and fractionally integrated processes to model the actin-myosin interaction and dwell times. Mathematical Biosciences and Engineering, 22(4): 1019-1054, (2025)
2. Buonocore, A., Caputo, L., Pirozzi, E., Ricciardi, L. M.: The first passage time problem for Gauss-diffusion processes: Algorithmic approaches and applications to LIF neuronal model. Methodol. Comput. Appl. Probab. 13(1), 29-57 (2011)
3. Baldi, P.: Stochastic Calculus: An Introduction Through Theory and Exercises. Cham, Switzerland: Universitext, Springer (2017)
4. Di Nardo, E., Nobile, A. G., Pirozzi, E., Ricciardi, L. M.: A computational approach to first-passage-time problems for Gauss-Markov processes. Adv. Appl. Probab. 33(2), 453-482 (2001)
5. Meerschaert, M. M. and Straka, P.: Inverse stable subordinators. Math. Model. Nat. Phenom. 8(2), 1-16 (2013)
6. Leonenko, N. and Pirozzi, E.: First passage times for some classes of fractional time-changed diffusions. Stochastic Analysis and Applications 40(4), 735-763 (2022)
7. Kobayashi, K.: Stochastic Calculus for a Time-Changed Semimartingale and the Associated Stochastic Differential Equations. Journal of Theoretical Probability 24, 789-820 (2011)
8. Ascione, G., Pirozzi, E. and Toaldo, B.: On the exit time from open sets of some semi-Markov processes, The Annals of Applied Probability 30(3), 1130-1163 (2020)
9. Ascione, G., Patie, P. and Toaldo, B.: Non-local heat equation with moving boundary and curve-crossing of delayed Brownian motion, arXiv preprint arXiv:2203.09850 (2025)
10. Caputo, L., Carfora, M.F. and Pirozzi, E.: Comparing some simulation strategies for first passage times of time-changed Brownian motion, In: Quesada-Arencibia, A., Affenzeller, M., Moreno-Díaz, R. (eds) Computer Aided Systems Theory – EU-ROCAST 2024. Lecture Notes in Computer Science, (2025)
11. Jum, E. and Kobayashi, K.: A strong and weak approximation scheme for stochastic differential equations driven by a time-changed Brownian motion. Probab. Math. Statist. 36(2), 201-220 (2016)

Fractional stochastic processes for risk models and simulations

Enrica Pirozzi[0000-0001-8692-2759]

Dipartimento di Matematica e Fisica, Università degli Studi della Campania “Luigi Vanvitelli”, Caserta, Italy
 enrica.pirozzi@unicampania.it

This contribution is devoted to provide a collection of a recent results about the use of fractional time-changed stochastic processes in risk models. (Instances of fractional stochastic processes can be found in [1], [3], [4],). The fractional approach for such kind of processes adopts the inverse of a stable subordinator ([5]) for the stochastic time-change. The sub-martingale property is studied to guarantee the net-profit condition. The long-range dependence behaviour is considered in order to include effects of memory. The infinite-horizon ruin probability, a specialized version of the Gerber–Shiu function, is considered and investigated.

For a fractional Poisson process a special focus is on the inter-arrival times, that are not exponentially distributed, but they follow a Mittag-Leffler distribution. The Cramér–Lundberg formula and a relationship between the continuous premium rate and the fractional claim frequency for this risk model are investigated. Closed-form expressions for the finite-time ruin probability are derived when the risk model involves a compound fractional risk process with exponential claims. Furthermore, for a general claim distribution, ruin probability estimates are also provided.

Simulation algorithms are implemented as a verification tool and investigation strategy, and also to show and compare different dynamics of such fractional stochastic processes for varying values of parameters.

This contribution is essentially based on papers [2] and [6] obtained by jointly working with Nikolai Leonenko, Andrey Pepelyshev, Alois Pichler and Xiangyun Meng.

References

1. Kobayashi, K.: Stochastic Calculus for a Time-Changed Semimartingale and the Associated Stochastic Differential Equations. *Journal of Theoretical Probability* 24, 789–820 (2011)
2. Leonenko, N., Pepelyshev, A., Pichler, A., Pirozzi, E., Meng, X., Probability of ruin within finite time and Cramér–Lundberg inequality for fractional risk processes. *TEST* (2025). <https://doi.org/10.1007/s11749-025-00991-9>
3. Leonenko, N., Pirozzi, E.: The time-changed stochastic approach and fractionally integrated processes to model the actin-myosin interaction and dwell times. *Mathematical Biosciences and Engineering*, 22(4): 1019-1054, (2025)
4. Leonenko, N. and Pirozzi, E.: First passage times for some classes of fractional time-changed diffusions. *Stochastic Analysis and Applications* 40(4), 735–763 (2022)

5. Meerschaert, M. M. and Straka, P.: Inverse stable subordinators. *Math. Model. Nat. Phenom.* 8(2), 1–16 (2013)
6. Pirozzi, E. On a Fractional Stochastic Risk Model with a Random Initial Surplus and a Multi-Layer Strategy. *Mathematics* 2022, 10, 570. <https://doi.org/10.3390/math10040570>

PKI on IIoT environments: requirements, design and compliance

I. Pereira-Domínguez¹[0009-0008-6046-1059] and C. Hernández-Goya¹[0000-0002-9468-708X]

Department of Engineering and Systems. University of La Laguna. 38271 La Laguna.
Tenerife. Spain
ipereira@ull.edu.es, mchgoya@ull.edu.es

Keywords: OPC UA · PKI · Post-Quantum Cryptography · IIoT

1 Abstract

This project addresses the difficulties of secure authentication management in industrial environments by implementing a dedicated Public Key Infrastructure (PKI) and Certification Authority (CA) based on Open Platform Communications Unified Architecture (OPC UA) [2]. The challenges posed by including post-quantum cryptographic primitives in actual industrial communication environments is also explored.

2 Introduction

Through the use of trusted CAs, PKIs allow endpoint verification and trust establishment. However, many industrial OPC UA systems still depend on self-signed certificates, increasing the risk of spoofing, unauthorized access and man-in-the-middle attacks. Such trust models are fundamental to establishing secure relationships between entities and are widely analyzed in the context of PKI-based infrastructures [1].

3 Methodology and Implementation

The suggested design replaces self-signed certificates with CA-signed certificates to create a centralized trust model that improves authentication, confidentiality, and data integrity. Several security parameters—such as key lengths, signature algorithms, and authentication policies—were configured to meet the OPC UA security standards and industry best practices. Additionally, the implementation explores the incorporation of post-quantum cryptographic (PQC) algorithms to assess their viability in industrial settings.

The methodology was divided into two phases: the deployment of a PKI framework aligned with OPC UA specifications, and the experimental incorporation of post-quantum cryptographic primitives to evaluate their viability in actual industrial communication environments.

First, a CA hierarchy was implemented using open-source tools to issue and administer X.509 certificates for OPC UA clients and servers. The configuration included secure key generation policies, certificate validation periods and revocation procedures using Certificate Revocation Lists (CRLs) as part of the setup. These certificates were then installed into OPC UA nodes to replace self-signed certificates, allowing encrypted communication channels and mutual authentication in compliance with OPC UA security profiles.

In the second phase, post-quantum algorithms, ML-DSA-44 (CRYSTALS-Dilithium2) for digital signatures and ML-KEM-768 (CRYSTALS-Kyber) for key establishment, were integrated as part of a custom security policy. The objective was to evaluate interoperability, performance, and robustness compared to traditional RSA or ECC schemes. Prior work in this area remains limited; for example, [3] outlines several challenges regarding certificate formats, algorithm negotiation, and deployment on constrained devices.

4 Conclusions

The findings indicate that integrating a dedicated PKI enhances authentication, confidentiality, and overall reliability in OPC UA communications, providing a viable solution for industrial settings evolving from self-signed configurations to a more secure and resilient IIoT ecosystem.

However, a number of difficulties surfaced during the post-quantum implementation stage. Integrating those brand-new algorithms required adapting the existing OPC UA cryptographic interfaces to support larger key and signature sizes, as well as to modify encoding and verification procedures. When using these algorithms in environments with limited resources, memory limitations and computational overhead also needed to be considered.

Acknowledgments. This work was supported by the project SCITALA C064/23 and the C065/23 Cybersecurity Chair of the University of La Laguna, both funded by INCIBE, with the Recovery, Transformation, and Resilience Plan (Next Generation) financed by the European Union; PID2022-138933OB-I00 ATQUE, funded by MCIN/AEI/10.13039/501100011033; and 2023DIG28 IACTA, funded by Cajacanarias and La Caixa Foundations.

References

1. Chandramouli, R., Iorga, M., Chokhani, S.: Trusted public key infrastructure. In: 2013 ICIT. IEEE (2013). <https://doi.org/10.1109/ICIT.2013.6489675>
2. Mahnke, W., Leitner, S.H., Damm, M.: OPC Unified Architecture. Springer (2009). <https://doi.org/10.1007/978-3-540-68899-0>
3. Scheible, P.: Hybrid X.509 Certificates for OPC UA – Combining Classical and Post-Quantum Cryptography. Master's thesis, Hochschule Offenburg (2020)

Federated Quantum Machine Learning for Privacy-preserving Mobile Malware Detection: A Proposal

Francesco Mercaldo^{1,*}, Fabio Martinelli⁴, Marta Petyx², Patrizia Agnello², and Antonella Santone¹

¹ Department of Medicine and Health Sciences “Vincenzo Tiberio”, University of Molise, Campobasso, Italy

² Istituto Nazionale per l’Assicurazione contro gli Infortuni sul Lavoro (INAIL), Roma, Italy

³ Department of Engineering, University of Sannio, Benevento, Italy

⁴ Institute of High Performance Computing and Networking, National Research Council of Italy (CNR), Rende, Italy

`francesco.mercaldo, antonella.santone@unimol.it, m.petyx, p.agnello@inail.it, fabio.martinelli@icar.cnr.it`

Considering the inefficiency of current antimalware in the detection of zero-day threats, researchers are proposing the adoption of deep learning for malware detection, with particular regard to the mobile ecosystem. Nonetheless, there are several issues limiting the adoption of these methods in the real-world domain, the first one is the high presence of false positive, the second one is related to privacy and the third one is the lack of explainability. For these reasons, in this paper we propose a framework for mobile malware detection aimed to overcome these limits, in particular we propose the adoption of quantum machine learning [1], that demonstrated its supremacy with respect to classic deep learning models, as for instance convolutional neural networks, in combination with federated machine learning [2], an emerging paradigm aimed to enable a set of clients to build a common model without sharing the data (only the updated weights) and considering explainability, by integrating the adoption of a way to visualise the model prediction.

Figure 1 shows the workflow of the proposal.

The workflow starts with decentralized client devices, such as smartphones and tablets, that collect mobile applications (e.g., Android/iOS). These data are transformed into images samples, which are then processed by a Quantum Model.

The dataset is divided into training, validation, and testing sets. During training, a federated server coordinates model updates from all clients. Each client performs local training on its own data using defined hyperparameters. The locally trained weights are periodically sent to the server, which aggregates them to update the global model until convergence is achieved.

Once training is complete, the global model is tested to assess its accuracy and ability to distinguish between benign and malicious applications. An explainability module then produces heatmaps highlighting the most influential image




Fig. 1. The workflow.

regions, enhancing transparency and trust in the predictions of the quantum model.

Acknowledgment

This work has been partially supported by EU DUCA, EU CyberSecPro, SYNAPSE, PTR 22-24 P2.01 (Cybersecurity) and SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the EU - NextGenerationEU projects and ALOHA: a framework for monitoring the physical and psychological health status of the Worker through Object detection and federated machine learning, Call for Collaborative Research BRiC -2024, INAIL.

References

1. J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning,” *Nature*, vol. 549, no. 7671, pp. 195–202, 2017.
2. Q. Yang, Y. Liu, T. Chen, and Y. Tong, “Federated machine learning: Concept and applications,” *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, 2019.

Low-Latency Conversational Ranker for Multi-Agent Voice Interaction

Marcos Rodriguez-Vega¹ and Pino Caballero-Gil¹

Universidad de La Laguna, San Cristóbal de La Laguna, Spain

Voice user interfaces are evolving from single-assistant dialogues to multi-agent scenes, where several specialized avatars listen and may respond simultaneously. In such environments, conversational coordination requires two immediate decisions after each user utterance: *who* is being addressed and *which* avatar(s) should respond, possibly with varying degrees of urgency. This challenge has been studied in the context of multi-party dialogue and addressee recognition [1,2], as well as in turn-taking models for interactive systems and human-robot communication. Traditional approaches rely on heavy natural language understanding pipelines or dialogue-policy models that trade accuracy for latency and often require large, supervised datasets built around fixed agent sets. To address these limitations, we present a **training-free, multilingual conversational ranker** that dynamically resolves addressees and estimates *speak-intent* in real time for flexible agent pools.

The problem can be defined as follows: given (i) a set of avatars $A = \{a_1, a_2, \dots, a_n\}$ each described by a concise profile (name or aliases, role, description, and keywords), and (ii) the current user utterance u together with a lightweight dialogue state, the goal is to compute:

1. A probability distribution $p(a_i \mid u, \text{ctx})$ over the avatars, indicating how likely each one is to be the addressee, including a possible *broadcast* case when the message is directed to all.
2. For each avatar a_i , an *intention to participate* score $I(a_i) \in [0, 1]$, reflecting its readiness or obligation to take the next conversational turn.

Our main contribution is a practical, deployment-ready ranker that infers conversational addressees and turn-taking intent without any supervised training. The system relies solely on avatar profiles and live user utterances, combining symbolic cues with lightweight embeddings to achieve real-time inference under noisy and multilingual conditions.

Each avatar is represented as a compact semantic card with four fields, *names*, *roles*, *description*, and *keywords*. These are pre-encoded into embeddings $E_i = \{e_{\text{names}}, e_{\text{roles}}, e_{\text{desc}}, e_{\text{kw}}\}$ and paired with a lexical bag-of-terms for quick pruning. This “index-time augmentation” enables avatars to be added or updated dynamically without retraining. Given a user utterance u , the model first narrows the search space from n avatars to a small candidate set K using simple signals: explicit vocatives or role mentions (e.g., “Hey, Andrea”, “Alex”), continuity from the previous turn, or lexical similarity (TF-IDF) when no clear cues are found. For each candidate a_i , the utterance embedding e_u is compared with

the avatar's fields to compute a similarity value:

$$\text{sim}_i = \text{max_cos}(e_u, E_i),$$

where $\text{max_cos}(\cdot)$ denotes the highest weighted cosine similarity across the avatar's name, role, description, and keywords embeddings.

This similarity is then combined with symbolic cues indicating vocatives, directed questions, or continuity from the previous turn:

$$\text{score}_i = \alpha \text{sim}_i + \beta v_i + \gamma q_i + \delta c_i,$$

where $v_i, q_i, c_i \in \{0, 1\}$ denote, respectively, the presence of a vocative or alias match (v_i), a direct question addressed to the avatar (q_i), and conversational continuity when the same avatar was the focus in the previous turn (c_i). The coefficients $\alpha, \beta, \gamma, \delta$ control the relative importance of semantic similarity and symbolic cues.

These raw scores provide a weighted estimate of conversational relevance for each avatar. If all scores remain below a threshold or plural vocatives are detected, the utterance is labeled as a *broadcast*, meaning that it is addressed to the group rather than to a single participant.

To manage turn-taking, each avatar is assigned a continuous intent value reflecting its readiness to respond:

$$I(a_i) = \sigma(w_1 \text{score}_i + w_2 Q_i + w_3 T_i - w_4 R_i + w_5 P_i),$$

where Q_i indicates whether the utterance is a question directed to a_i , T_i captures topical relevance between the utterance and the avatar's domain, R_i penalizes recent speech to avoid repeated turns, and P_i activates when a previous question or task remains pending. The logistic function $\sigma(x) = 1/(1 + e^{-x})$ maps these combined features into a normalized range $[0, 1]$, producing interpretable readiness values. A small hysteresis margin ($\tau_{\text{on}} > \tau_{\text{off}}$) is applied to stabilize turn-taking and prevent rapid oscillations or overlapping responses.

Overall, this pipeline provides a unified and efficient mechanism for real-time multi-agent voice interaction. The ranker and its evaluation scripts are available in the accompanying repository, as part of a broader framework for multi-agent conversational research. This framework enables controlled experiments where several specialized agents interact within a shared conversational context, allowing us to study user behavior and latent biases that emerge when people engage with large language models through natural, fast-paced dialogue. **Acknowledgement.** Cybersecurity Chair C065/23 INCIBE.

References

1. Gu, J.C., Tao, C., Ling, Z.H.: Who says what to whom: A survey of multi-party conversations. In: Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence (IJCAI). pp. 5486–5493 (2022)
2. Ouchi, H., Tsuboi, Y.: Addressee and response selection for multi-party conversation. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing. pp. 2133–2143 (Nov 2016). <https://doi.org/10.18653/v1/D16-1231>

Detection and Reconstruction of Traffic Signs Under Adversarial Conditions

Alejandro Barreto-Ramallo¹, Marcos Rodriguez-Vega¹, and Pino Caballero-Gil¹

Universidad de La Laguna, San Cristóbal de La Laguna, Spain

The use of autonomous driving systems has been increasing in the last few years. In such systems, one of the most important problems that must be solved is the speed of detection, as can be seen in [1]. Also, some initial solutions to this problem involve Convolutional Neural Networks (CNN), as in [2]. These studies highlight how critical fast and accurate detection is for ensuring real-time decision making in autonomous driving systems.

Therefore, there is a need to develop new tools that let us overcome these obstacles and make autonomous driving safer. To solve this, we developed a deep learning pipeline capable of effectively detecting and classifying traffic signs. We also aim to improve the pipeline's robustness so it can function correctly even under adverse environmental conditions and when signs are damaged.

We trained a YOLO-based detector/classifier on more than 20,000 images and evaluated it on over 12,000 images from the German Traffic Sign Recognition Benchmark (GTSRB), which contains signs captured under challenging visual conditions. Table 1 summarizes performance against two representative baselines: a classical model (logistic regression) and a compact CNN with tuned hyperparameters. Note that the CNN operates on pre-extracted Regions of Interest (ROI), whereas YOLO processes the full image and performs joint detection and classification. This difference highlights YOLO's broader capability compared with models limited to classification. The last three columns of Table 1 report inference times (in milliseconds) for each model at the 50th, 90th, and 95th percentiles. As expected, YOLO requires more time per inference due to its additional detection stage, whereas the other models focus solely on classification.

Table 1. Unified performance and inference statistics for tested models

Model	Precision	Recall	F1 Score	F1 Macro	Accuracy	p50 (ms)	p90 (ms)	p95 (ms)
Logistic Regression	0.825	0.798	0.811	0.803	0.858	0.071	0.111	0.164
CNN	0.944	0.931	0.937	0.933	0.956	0.360	0.417	0.426
YOLO	0.984	0.964	0.974	0.953	0.949	11.648	15.882	16.709

Table 1 reveals a clear trade-off between accuracy and latency. YOLO achieves the highest precision (0.984), recall (0.964), and F1 score (0.974), as well as the best F1 macro value (0.953), indicating balanced performance across all classes. However, its overall accuracy (0.949) is slightly lower than that of the CNN (0.956), which may be due to class imbalance: while accuracy is dominated by

frequent classes, the macro-averaged F1 better reflects per-class consistency. The CNN exhibits low and stable latency from its limited ROI-based task. In contrast, YOLO provides a full detection–classification solution at the expense of higher computation.

An analysis of YOLO’s confusion matrix revealed systematic misclassifications between visually symmetric traffic signs (e.g., left vs. right curves). To address this issue, we proposed a hybrid model that combines YOLO with an optimized CNN trained specifically for a binary discrimination task (left or right). This post-processing stage refines YOLO’s predictions and leads to a 3% increase in overall accuracy. The improvement achieved by the hybrid approach is illustrated in Figure 1.




Fig. 1. YOLO improved with an optimized CNN for post-processing ambiguous signs

This pipeline offers a robust starting point. Next, we will further enhance robustness to adverse and rare conditions (e.g., severe weather, motion blur, occlusions, vandalism) and explore calibration and uncertainty-aware decision making to increase the safety of autonomous driving systems.

Acknowledgments. This work was made possible by C065/23 Cybersecurity Chair of the University of La Laguna, funded by INCIBE, as part of the Recovery, Transformation, and Resilience Plan (Next Generation) funded by EU.

References

1. Yang, Y., Luo, H., Xu, H., Wu, F.: Towards real-time traffic sign detection and classification. *IEEE Transactions on Intelligent Transportation Systems* **17**(7), 2022–2031 (2016). <https://doi.org/10.1109/TITS.2015.2482461>
2. Zhu, Z., Liang, D., Zhang, S., Huang, X., Li, B., Hu, S.: Traffic-sign detection and classification in the wild. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2016)

A PQC-Improved Secure Aggregation Scheme for Federated Learning

L. Syne^[0009-0005-3805-1430], C. Hernández-Goya^[0000-0002-9468-708X], and P. Caballero-Gil^[0000-0002-0859-5876]

Department of Computer Engineering and Systems. University of La Laguna. 38271
La Laguna. Tenerife. Spain

alu0101533254@ull.edu.es mchgoya@ull.edu.es pcaballe@ull.edu.es

Extended Abstract

Federated Learning (FL) enables clients to collaboratively train a global Machine Learning model through an aggregation server without sharing their raw data. However, FL introduces new security threats: it is vulnerable to so-called inference attacks, in which a malicious aggregator may infer information about clients' private data from their model updates [1]. Additionally, malicious clients can submit falsified updates to the aggregation server [2]. Secure Aggregation (SA) and secret sharing-based Multi-Party Computation (MPC) are mitigation techniques designed to restrict the aggregation server to learning only the sum of clients' updates (i.e., the server cannot access individual updates) while maintaining robustness against client dropouts. Furthermore, with the advent of quantum computing, existing SA protocols mainly using secret sharing-based MPC and classical cryptographic primitives remain vulnerable to potential quantum attacks.

This paper presents a Secure Aggregation (SA) method that integrates post-quantum secure channels for client key exchange and Feldman's Verifiable Secret Sharing (VSS) [3], providing resilience against malicious clients. The proposed protocol is inspired by the approach presented in the work [4]. It combines private and pairwise masking between clients, Shamir's Secret Sharing (SSS) for client dropout tolerance, and a VSS to ensure the integrity of the aggregated data. Experiments conducted under simulated Federated Environments demonstrate the viability of the proposed approach. Figure 1 describes the proposal.




Fig. 1. PQC-Improved Secure Aggregation Protocol

References

1. H. Fereidooni et al., “SAFELearn: Secure Aggregation for private FEderated Learning,” in *IEEE Security and Privacy Workshops (SPW)*, pp. 56-62, 2021.
2. D. Kolasa, K. Pilch, and W. Mazurczyk, “Federated learning secure model: A framework for malicious clients detection,” *SoftwareX*, vol. 27, p. 101765, 2024.
3. Yi-Hsiu Chen and Yehuda Lindell. Feldman’s Verifiable Secret Sharing for a Dishonest Majority. Cryptology ePrint Archive, Paper 2024/031, 2024. <https://eprint.iacr.org/2024/031>
4. Bonawitz, K., Ivanov, V., Kreuter, B., Marcedone, A., McMahan, H. B., Patel, S., Ramage, D., Segal, A., and Seth, K. Practical secure aggregation for federated learning on user-held data. *arXiv preprint arXiv:1611.00220*, 2017.

Acknowledgments: This work was supported by the Cybersecurity Chair of the University of La Laguna (C065/23) and the strategic project SCITALA (C064/23), both funded by INCIBE through the Recovery, Transformation, and Resilience Plan (Next Generation EU), financed by the European Union.

Anomaly Detection in Smart IoT Environments Using Classical and Quantum Machine Learning

A. García Carrasco, C. Hernández-Goya^[0000-0002-9468-708X], and J. Molina-Gil^[000-0001-7702-9264]

Department of Computer Engineering and Systems. University of La Laguna. 38271 La Laguna. Tenerife. Spain

alu0101764023@ull.edu.es mchgoya@ull.edu.es jmmolina@ull.edu.es

Abstract. In this paper, we evaluate the usefulness of Quantum Machine Learning (QML) for anomaly detection in real Internet of Things (IoT) traffic involving DDoS, DoS, and Mirai attacks. Classical and quantum models are implemented and compared to assess whether quantum approaches offer a viable alternative to established Machine Learning methods(ML). Logistic Regression (LR), Random Forest (RF), and a Variational Quantum Classifier (VQC) are applied to the CIC IoT 2023 dataset.

Keywords: IoT · anomaly detection · Variational Quantum Classifier (VQC).

1 Introduction

The objective of this study is to provide an initial comparison of classical ML methods and newly designed QML approaches using the public dataset.

2 Methodology and Results

The dataset selected [1] is part of the CIC IoT 2023 data collection. This set has been designed using a wide topology composed of a multitude of real IoT devices. It includes smart home devices, cameras, sensors, and microcontrollers that are connected and configured to enable the execution of various attacks and to capture the corresponding traffic. It contains 40 attributes related to the traffic of each observation. The data is labelled since the type of attack to which each observation belongs, or whether it belongs to benign traffic, is indicated. The observations have been filtered, considering only those that belong to the analysed attacks (DDoS, DoS, or Mirai, in each case) and to benign traffic. Consequently, three new datasets were created and used to train classical ML models. In the context of QML techniques, 50 observations were sampled from the preceding datasets. The sets were randomly divided using a 75/25 ratio to define the training set and test set respectively. To address class imbalance, which is present in

the target variable, the random oversampling technique has been applied. The analysis carried out using classical and quantum ML models comprised several steps. For the logistic regression (LR) model, non-collinearity among variables was verified and the data was standardised to avoid scale-related effects. In the random forest (RF) model, hyperparameters were tuned to minimise the out-of-bag error. For the variational quantum classifier (VQC), training was performed on a quantum simulator using a ZZFeatureMap and a constructed ansatz, with COBYLA selected as the optimiser due to its gradient-free nature, intended to accelerate convergence [2].

The LR model correctly classified almost all instances and, importantly, did not misclassify any benign traffic as attacks, although some malicious samples were labelled as normal. Similarly, the RF model achieved a high classification rate, confirming the effectiveness of the proposed approach.

For the VQC, benign traffic was classified with higher accuracy; however, its overall performance remained notably inferior to that of LR and RF. Classical techniques consistently exceed 99% across all metrics, while the quantum classifier exhibits a substantial performance drop.

In summary, classical ML methods demonstrated strong generalisation across all attack types, with LR offering better interpretability and faster computation. The Mirai attack is particularly noteworthy, as it was correctly classified using only one predictor: the number of packets in the window. Conversely, the VQC showed significantly lower performance, which aligns with the early stage of development of quantum ML approaches. Future work should explore additional attack types, multi-class classification, and, especially, the execution of quantum models on real quantum hardware to better assess their potential.

Acknowledgments. This work was supported by the strategic project SCITALA (C064/23), funded by INCIBE through the Recovery, Transformation, and Resilience Plan (Next Generation EU); by the project PID2022-138933OBI00 “ATQUE,” funded by MCIN/AEI/10.13039/501100011033 and co-financed by the European Regional Development Fund (ERDF, EU); and by the project 2023DIG28, funded by CajaCanarias and the “la Caixa” Foundation.

References

1. Neto, E. C. P., Dadkhah, S., Ferreira, R., Zohourian, A., Lu, R., & Ghorbani, A. A. (2023). CICIoT2023: A real-time dataset and benchmark for large-scale attacks in IoT environment. *Sensors*, 23(13), 5941. <https://doi.org/10.3390/s23135941>.
2. Qiskit Machine Learning Development Team. (2024). *Training a Quantum Model on a Real Dataset*. Qiskit Machine Learning Tutorials. Retrieved May 8, 2025, from https://qiskit-community.github.io/qiskit-machine-learning/tutorials/02a_training_a_quantum_model_on_a_real_dataset.html

QUBO-Based Modeling for Biomedical Diagnosis: Application to Breast Cancer

Sonia Diaz-Santos^{1[0000-0002-8482-3957]}, Daniel Escanez-Exposito^{1[0000-0003-4215-0501]}, Pino Caballero-Gil^{1[0000-0002-0859-5876]}, and Cándido Caballero-Gil^{1[0000-0002-6910-6538]}

Department of Computer Engineering and Systems, University of La Laguna,
Tenerife San Cristobal de La Laguna, Spain.
 {sdiazsan, jescanez, pcaballe, ccabgil}@ull.edu.es

Abstract

Optimization techniques inspired by annealing have recently gained attention as promising approaches for solving complex learning tasks in cybersecurity and anomaly detection, where combinatorial structures and high-dimensional data often challenge traditional models. These methods have also shown growing potential in biomedical applications, particularly in breast cancer detection [1]. In this context, a binary classification framework is introduced, formulated as a Quadratic Unconstrained Binary Optimization (QUBO) problem and solved through classical simulated annealing using the D-Wave software stack. This formulation serves as a classical analogue of quantum annealing approaches implemented on D-Wave quantum hardware. The performance of the proposed QUBO-based model is compared with that of a grid-optimized Support Vector Machine (SVM). The SVM margin maximization problem is encoded into a binary quadratic form solvable by annealing-based samplers [2,3,4]. Experiments conducted on the Breast Cancer Wisconsin (Diagnostic) dataset [5] employed seven normalized features. The QUBO was instantiated and solved using `dimod` and `dwave-system` with a simulated annealing sampler [6], and the full implementation is publicly available in a Google Colab notebook [7]. As shown in Fig. 1, the QUBO classifier achieved an accuracy of 0.9162, a precision of 0.9650, and a recall of 0.8990, outperforming the SVM, which obtained 0.8978, 0.9573, and 0.8762, respectively. These outcomes suggest that the QUBO-based formulation provides comparable or slightly superior performance, characterized by fewer false positives and overall consistency with the classical baseline. Overall, these findings highlight the potential of QUBO-based formulations as interpretable, high-precision classifiers for complex data domains such as cybersecurity, and motivate further exploration on real quantum annealers in future work.

Keywords: Quantum Machine Learning · QUBO · Binary Classification · Qubit Resolution · Breast Cancer




Fig. 1: Combined visualization of QUBO model results.

Acknowledgments. This work was supported by the C065/23 Cybersecurity Chair of the University of La Laguna, funded by INCIBE, and by the strategic project C064/23 SCITALA, financed through the Recovery, Transformation and Resilience Plan by the European Union (Next Generation EU). This research was also co-financed by the Agencia Canaria de Investigación, Innovación y Sociedad de la Información of the Consejería de Universidades, Ciencia e Innovación y Cultura, and by the European Social Fund Plus (FSE+) under the Programa Operativo Integrado de Canarias 2021–2027, Axis 3, Priority Theme 74 (85%).

References

1. Díaz-Santos, S., Escámez-Expósito, D.: Classical vs. Quantum ML for Breast Cancer Detection. In: *Proc. DRCN 2023*, pp. 1–5. IEEE. DOI:10.1109/DRCN57075.2023.10108230
2. Kochenberger, G., Hao, J.K., Glover, F., Lewis, M., Lü, Z., Wang, H., Wang, Y.: The Unconstrained Binary Quadratic Programming Problem: A Survey. *J. Combin. Optim.* **28**(1), 58–81 (2014). <https://doi.org/10.1007/s10878-014-9734-0>
3. Glover, F., Kochenberger, G., Du, Y.: A Tutorial on Formulating and Using QUBO Models. *arXiv preprint arXiv:1811.11538* (2019). <https://doi.org/10.48550/arXiv.1811.11538>
4. Venturelli, D., Mandrà, S., Knysh, S., O’Gorman, B., Biswas, R., Smelyanskiy, V.: Quantum Optimization of Fully Connected Spin Glasses. *Phys. Rev. X* **5**(3), 031040 (2015). <https://doi.org/10.1103/PhysRevX.5.031040>
5. Wolberg, W.H., Street, W.N., Mangasarian, O.L.: Breast Cancer Wisconsin (Diagnostic) Data Set. *UCI Machine Learning Repository* (1995). <https://archive.ics.uci.edu/dataset/17/breast+cancer+wisconsin+diagnostic>
6. D-Wave Systems Inc.: Ocean SDK Documentation. <https://docs.ocean.dwavesys.com> (accessed 31 Oct 2025)
7. Díaz-Santos, S., Escámez-Expósito, D.: High-Precision QUBO Classification for Breast Cancer via Low-Qubit Resolution: Google Colab Notebook. *Google Colaboratory* (2025). <https://colab.research.google.com/drive/1Tuw5JuettLr0DFP3yoWYu2Y8AhbalJhqn#scrollTo=akzgRNuyrctk>

Quantum Key Distribution Simulation in Rust

Daniel Escanez-Exposito¹[0000-0003-4215-0501]
 and Pino Caballero-Gil¹[0000-0002-0859-5876]

Department of Computer Engineering and Systems, University of La Laguna,
 Tenerife San Cristobal de La Laguna, Spain.
 {jescanez,pcaballe}@ull.edu.es

Extended Abstract

The `qkd` crate is an open-source Rust library designed to simulate Quantum Key Distribution (QKD) protocols in a configurable and reproducible manner [1]. QKD enables two communicating parties to establish information-theoretically secure secret keys by exploiting fundamental quantum-mechanical principles such as the no-cloning theorem and measurement-induced disturbance. Since physical implementations of QKD require specialized quantum hardware and controlled optical setups, software-based simulators provide an accessible platform for evaluating protocol behaviour, analysing security metrics, and performing rapid prototyping. The crate logo can be seen in Fig. 1.



Fig. 1. `qkd` logo

The library currently supports several foundational QKD protocols, including BB84 [2], Six-State [3], and B92 [4]. Its modular architecture isolates participant models, protocol execution logic, quantum-state representations, and auxiliary utilities. Core abstractions such as `QKD` and `Participant` structure the simulation workflow from quantum-state preparation and measurement to basis reconciliation and key sifting. Lightweight representations of qubits and complex matrices enable the modelling of state transformations and measurement statistics while preserving computational efficiency.

Users may configure parameters such as the number of transmitted qubits, the level of bit-flip noise introduced in the quantum channel, the eavesdropping/interception rate, and the number of repeated simulation trials. The library computes multiple performance and security indicators, including two distinct Quantum Bit Error Rates (QBERs): the QBER estimated by Alice and Bob during the public discussion phase, and the QBER present in the final sifted key. Additionally, the simulator allows the specification of confidence levels required by the legitimate participants to determine whether the protocol execution remains secure under the observed error conditions. Results are exported in CSV format for downstream scientific analysis.

The documentation of `qkd` is clear, extensive, and well-structured, which greatly facilitates its adoption in both research and development contexts. The online API reference provides a detailed description of the library's modular architecture (covering participants, protocol execution, quantum-state abstractions, and utility components), allowing users to readily understand and extend its internal design. The public README further outlines the crate's objectives, supported protocols, configurable parameters, and available command-line and library interfaces, offering an accessible entry point for newcomers.

The `qkd` crate is intended for research, education, and early-stage prototyping. Unlike idealized QKD simulators, it enables the study of more realistic scenarios by incorporating channel noise and configurable disturbance models. Its modular design permits straightforward extension: developers may add new protocol implementations, incorporate alternative noise models, or enhance existing components thanks to the library's open-source nature. This makes `qkd` a robust and flexible tool for simulating quantum communication systems beyond ideal theoretical assumptions. Overall, `qkd` provides a memory-safe, performant, and flexible Rust-based framework for simulating QKD protocols, supporting the study of quantum-secure communication strategies when laboratory experimentation is impractical or cost-prohibitive.

Acknowledgments. This work was supported by the C065/23 Cybersecurity Chair of the University of La Laguna, funded by INCIBE, and by the strategic project C064/23 SCITALA, financed through the Recovery, Transformation and Resilience Plan by the European Union (Next Generation EU).

References

1. Escanez-Exposito, D.: *qkd: A Rust Library for Quantum Key Distribution Simulation*. Crates.io (2025). <https://crates.io/crates/qkd>, last accessed 2025/11/14
2. Bennett, C.H., Brassard, G.: Quantum Cryptography: Public Key Distribution and Con Toss. *Theoretical Computer Science* **560**, 7–11 (1984)
3. Bruß, D.: Optimal Eavesdropping in Quantum Cryptography with Six States. *Phys. Rev. Lett.* **81**(14), 3018–3021 (1998). <https://doi.org/10.1103/PhysRevLett.81.3018>
4. Bennett, C.H.: Quantum Cryptography Using Any Two Nonorthogonal States. *Phys. Rev. Lett.* **68**(21), 3121–3124 (1992). <https://doi.org/10.1103/PhysRevLett.68.3121>

qcrypto: Library for Scientific Development and Testing of Quantum Cryptographic Primitives

Jorge Garcia-Diaz^{1[0009-0006-4400-6836]}, Daniel Escanez-Exposito^{1[0000-0003-4215-0501]}, Javier Padilla-Pio^[0009-0002-8439-2875], and Pino Caballero-Gil^{1[0000-0002-0859-5876]}

University of La Laguna, Tenerife, Spain
 {jgarcid, jescanez, jpallip, pcaballe}@ull.edu.es

Extended Abstract

The recent advancement of quantum computing technologies necessitates a proactive development of robust cryptographic solutions that can withstand attacks from future quantum adversaries [9, 7]. This paper presents **qcrypto**, a novel, open-source software library tailored for the scientific exploration, implementation, and rigorous testing of Quantum Cryptographic (QC) primitives. The library is developed using the Rust programming language, leveraging its guarantees of memory safety, thread concurrency, and high performance; critical for simulating computationally intensive quantum protocols accurately [8].

The primary objective of **qcrypto** is to facilitate an open-source, efficient and simple to use tool for advanced research into quantum-secure communication and computation. The scope of **qcrypto** is broad, enhancing the development and experimentation with a diverse set of foundational QC primitives. These include well-established protocols such as Quantum Key Distribution (BB84, Six-State, E91, B92, BBM92 and SARG04) [2], which forms the basis for unconditionally secure key exchange, as well as more complex interaction protocols. Specifically, the library provides implementations for a Quantum Zero-Knowledge Proof, enabling user authentication without revealing secret information [5]; Quantum Bit Commitment, a fundamental building block for secure multi-party computation [4]; and Quantum Oblivious Transfer, which allows a sender to transfer one of several pieces of information to a receiver, who can choose which one to receive, without the sender knowing which piece was chosen [3]. It also incorporates implementations of Quantum Coin Flipping, a protocol for achieving fairness in a dispute between two parties [1], and Quantum Digital Signature, which offer information-theoretic security properties beyond those achievable with classical protocols [6].

This library allows researchers to easily integrate new protocols, modify existing parameters, and simulate quantum channel conditions concerning noise and decoherence. This feature is essential for accurately modelling real-world performance and assessing the practical feasibility of these theoretical primitives. By leveraging Rust, **qcrypto** ensures that the implemented simulation experiments

are fast and memory-efficient, therefore providing the high performance that is essential for simulating quantum communications. The library is intended to serve as a beneficial research tool, simplifying the transition of theoretical quantum cryptographic concepts into verifiable results, thereby facilitating research and experimentation of new ideas for the scientific community.

Acknowledgments.

This work has been possible thanks to the Cybersecurity Chair sponsored by INCIBE through an initiative carried out within the framework of the funds of the Recovery, Transformation and Resilience Plan, financed by the European Union (Next Generation). It is also part of the PID2022-138933OB-I00 and 2023DIG28 IACTA research projects funded by MCIN/AEI/10.13039/501100011033/FEDER EU, and the CajaCanarias la Caixa Foundation.

References

1. Ambainis, A.: A new protocol and lower bounds for quantum coin flipping. In: Proceedings of the thirty-third annual ACM symposium on Theory of computing. pp. 134–142 (2001)
2. Bennett, C.H., Brassard, G.: Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science* **560**, 7–11 (2014). <https://doi.org/https://doi.org/10.1016/j.tcs.2014.05.025>, theoretical Aspects of Quantum Cryptography – celebrating 30 years of BB84
3. Bennett, C.H., Brassard, G., Crépeau, C., Skubiszewska, M.H.: Practical quantum oblivious transfer. In: Annual international cryptology conference. pp. 351–366. Springer (1991)
4. Brassard, G., Crépeau, C.: Quantum bit commitment and coin tossing protocols. In: Conference on the Theory and Application of Cryptography. pp. 49–61. Springer (1990)
5. Garcia-Díaz, J., Escámez Expósito, J.D., Caballero-Gil, P., Molina-Gil, J.: BB84-inspired quantum zero-knowledge proof for user authentication over quantum channel. In: Workshop on Quantum-Resistant Cryptography (QuRCry). pp. 30–32. Universidad de La Laguna (2025)
6. Gottesman, D., Chuang, I.: Quantum digital signatures. arXiv preprint quant-ph/0105032 (2001)
7. Grover, L.K.: A fast quantum mechanical algorithm for database search. In: Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing. p. 212–219. STOC '96, Association for Computing Machinery, New York, NY, USA (1996). <https://doi.org/10.1145/237814.237866>, <https://doi.org/10.1145/237814.237866>
8. Klabnik, S., Nichols, C.: The Rust programming language. No Starch Press (2023)
9. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing* **26**(5), 1484–1509 (1997). <https://doi.org/10.1137/s0097539795293172>

Stochastic Noise Adaptative Quantum Data Locking (SNA-QDL) Protocol

Yauset Cabrera-Aparicio¹[0009-0001-9164-0374],
 Jorge Garcia-Diaz¹[0009-0006-4400-6836], and
 Pino Caballero-Gil¹[0000-0002-0859-5876]

University of La Laguna, Tenerife, Spain
 {alu0101428994, jgarcidi, pcaballe}@ull.edu.es

Extended Abstract

Quantum Data Locking (QDL) represents a striking departure from classical information expectations in a theoretical manner. It illustrates how the transition from classical to quantum information drastically reshapes the relationship between data, secrecy, and key size. On the positive side, QDL demonstrates that, unlike in classical Shannon security, large amounts of information can be *locked* using a remarkably *small* secret key, partially thanks to the significant basis-dependence and measurement sensitivity of quantum states.

Nevertheless, QDL phenomena also entails an important negative implication of quantum settings: the protection it offers is intrinsically vulnerable to classical information leakage [1]. Traditionally, security of QDL protocols hinges on manipulating what an adversary can effectively extract from the quantum states, rather than on enforcing an inherently information-theoretic notion of secrecy. Thus, any unintended disclosure of side information, such as partial key bits or correlations neglected in the classical description, can abruptly unlock sensitive data [2]. This dual nature makes QDL a compelling yet delicate paradigm, capturing both the enhanced capabilities and the subtle security challenges that arise when moving from classical to quantum information frameworks.

Recent research has expanded both the theoretical foundations and the practical interpretation of the phenomenon, revealing its potential as well as its limitations in realistic communication settings [3]. Although early works revealed the paradigm-shifting gap between classical and quantum accessible information, subsequent studies highlighted how this advantage deteriorates under noise, memory restrictions, side-channel leakage, or adversaries equipped with sophisticated quantum storage and collective measurement capabilities. Modern analyses now focus on quantifying how much *locking power* is preserved under imperfect channels, environmental decoherence, and adversaries who may delay measurements, exploit correlated noise, or dynamically adapt their strategy based on partial classical side information. Taking into account such considerations, it is incumbent upon us to prompt an active shift from static QDL constructions toward adaptive, context-aware protocols capable of reacting to fluctuations in

channel conditions and anticipating the capabilities of realistic eavesdroppers. In practice, we require protocols that can adjust encoding strategies, key selection mechanisms, and modulation schemes in response to estimated adversarial power, particularly in scenarios where an eavesdropper might discretely exploit partial information leakage to compromise security.

Our first prototype of the SNA-QDL protocol builds upon existing constructions that employ ensembles of Clifford unitary matrices [4], whose sampling properties are known to approximate second-order unitary designs (and third-order in specific scenarios, [5]), thereby ensuring strong locking behavior under *theoretically ideal* conditions. Given a security parameter $\kappa \gg 1$, our protocol introduces active, artificially generated noise components designed to interfere with the noise-correlation analyses typically exploited in standard QDL scenarios. These artificial noise blocks consist of pseudo-random rotation structures engineered from two key ingredients: a public family of functions f_1, \dots, f_κ , and a pair of probabilities $p_1, p_2 \in [0, 1]$ that allow to practically disrupt any localized noise-monitoring strategies on the qubits within. When properly orchestrated, such stochastic modulation ensures robustness even in situations where an eavesdropper gains partial access to intermediate noise-injection phases, as the effective structure of the locked data becomes highly sensitive to the adversary's uncertainty about the induced rotations.

A key innovation of SNA-QDL lies in its use of Bob's quantum memory to enhance reception and verification, allowing the protocol to preserve locking strength and limit side-information leakage, with an early prototype indicating that it can be effectively realized to protect quantum information against progressively stronger eavesdroppers.

Acknowledgments. This work has been possible thanks to the C065/23 Cybersecurity Chair of the University of La Laguna, funded by INCIBE, with the Recovery, Transformation, and Resilience Plan (Next Generation) financed by the European Union.

References

1. DiVincenzo, D.P., Horodecki, M., Leung, D.W., Smolin, J.A., Terhal, B.M.: Locking of accessible information and implications for the security of quantum cryptography. (2005) [Accessed: 15-10-2025]
2. Dupuis, F., Hayden, P., Winter, A.: Locking classical information in quantum states. (2010) [Accessed: 09-10-2025]
3. Rasit Onur Topaloglu: Quantum logic locking for security. *Machines* **6**(3), 27 (2022). <https://doi.org/10.3390/machines6030027>. [Accessed: 10-10-2025]
4. Zhao, Y., Li, K., Zhang, X.: Fault tolerant quantum data locking (2020). [Accessed: 10-10-2025]
5. Webb, Z.: The Clifford group forms a unitary 3-design. (2015) [Accessed: 10-10-2025]
6. Coles, P.J., Berta, M., Tomamichel, M., Wehner, S.: Entropy in quantum information theory – communication and cryptography. (2018). [Accessed: 14-10-2025]

The impact of BLE channel on RSSI-based ranging

Jakub Pietrucha¹ and Maciej Nikodem²[0000-0002-9242-2029]

¹ NOKIA, Wrocław, Polska, jakub.pietrucha@nokia.com

² Wrocław University of Science and technology, Wrocław, Polska,
maciej.nikodem@pwr.edu.pl

Abstract. The Received Signal Strength Indicator (RSSI) has long been used for communication range estimation and indoor localization, as the strength of the radio signal attenuates with distance. However, in indoor environments, RSSI is highly variable due to multipath propagation, reflections, and other radio phenomena, making the RSSI-distance relationship complex and non-trivial. This challenge is also evident in Bluetooth Low Energy (BLE) systems, which are often used for indoor positioning and asset tracking. To address this, many authors have proposed various RSSI processing methods aimed at reducing ranging errors and improving localization accuracy. Some approaches suggest measuring RSSI across multiple BLE channels, yet there is limited understanding of how channel selection influences the measurements and the resulting distance estimation. In this work, we analyze the impact of BLE channel selection on distance estimation between communicating devices. Through experimental evaluation, we demonstrate that choosing an appropriate channel can significantly improve ranging accuracy. Nevertheless, determining the optimal channel is not straightforward.

Keywords: Bluetooth Low Energy · RSSI-based ranging · Indoor positioning · Multichannel.

1 Extended Abstract

Ranging based on the received signal strength indicator (RSSI) is one of the most popular approaches used for indoor localization, as all radio transceivers measure and report RSSI for received packets [2]. Unfortunately, due to complex propagation in indoor settings and non-uniform antenna propagation patterns, the RSSI-distance relationship is not trivial, it is challenging to account for these effects, and improve ranging accuracy. Authors have proposed various approaches, including averaging, filtering, fingerprinting, and machine learning [4, 7]. While these methods improve accuracy, the change is not significant. In recent years, authors have started to benefit from BLE's transmission in 3 [1, 3, 6] and 40 [5] distinct frequency channels, demonstrating that the use of multiple channels can enhance ranging accuracy. Unfortunately, recent works do not fully explain how the use of distinct frequency channels affects ranging error.

In this work, we analyze the RSSI measurements collected in the experimental setup within the office environment [5]. The RSSI samples were collected in over 100 test points across the 100m² area, and radio packets were transmitted on all 40 BLE frequency channels. We analyze RSSI variability as a function of distance and channel (Fig. 1) to determine whether RSSI variability between channels and distances follows any patterns. We also analyze how different selections of frequency channels affect the resulting ranging accuracy in a search for the best possible selection of channels that will improve indoor localization quality for the entire area.




Fig. 1. Violin plots presenting variability of RSSI measurements accross 40 BLE channels for a fixed distance between communicating devices. Colors denote BLE channels.

References

1. Aranda, F.J., Parralejo, F., Aguilera, T., Álvarez, F.J., Torres-Sospedra, J.: RSS Channel-Based Integration for BLE Fingerprinting Positioning. In: 2023 13th International Conference on Indoor Positioning and Indoor Navigation (IPIN). pp. 1–6 (2023)
2. Aziz, T., Koo, I.: A Comprehensive Review of Indoor Localization Techniques and Applications in Various Sectors. *Applied Sciences* **15**(3) (2025)
3. De Blasio, G., Quesada-Arencibia, A., García, C.R., Rodríguez-Rodríguez, J.C., Moreno-Díaz, R.: A Protocol-Channel-Based Indoor Positioning Performance Study for Bluetooth Low Energy. *IEEE Access* **6**, 33440–33450 (2018)
4. Maduranga, M.W.P., Tilwari, V., Abeysekera, R.: Improved RSSI Indoor Localization in IoT Systems with Machine Learning Algorithms. *Signals* **4**(4), 651–668 (2023)
5. Nikodem, M., Szeliński, P.: Channel Diversity for Indoor Localization Using Bluetooth Low Energy and Extended Advertisements. *IEEE Access* **9**, 169261–169269 (2021)
6. Pastor, J.A., Gil-Martínez, A., Mateos, A.H., Algaba-Brazález, A., Tornero, J.L.G.: Bluetooth Low Energy separate-channel fingerprinting with frequency-scanned antennas. *Internet of Things* **33**, 101732 (2025)
7. Wu, T.: Review of RSSI-based Positioning Algorithm and Accuracy. *Frontiers in Computing and Intelligent Systems* **2**(1), 10–12 (Nov 2022)

Direction-finding of IoT devices method combining Amplitude and Phase monopulse comparison functions

Jose A. Lopez-Pastor^{1[0000-0001-8266-8669]}, Miroslav Hutar^{2[0000-0002-4505-1464]}, Peter Brida^{2[0000-0002-5442-9246]}, and Jose Luis Gomez-Tornero^{3[0000-0003-0488-3990]}

¹Department of Electronics, Computer Technology, and Projects, Technical University of Cartagena member of European University of Technology EUT+, 30203, Cartagena, Spain

²Faculty of Electrical Engineering and Information Technology, University of Zilina, 01026 Zilina, Slovakia

³Department of Information and Communication Technologies, Technical University of Cartagena member of European University of Technology EUT+, 30203, Cartagena, Spain

joseantonio.lopez@upct.es
 miroslav.hutar@feit.uniza.sk
 peter.brida@uniza.sk
 josel.gomez@upct.es

1 Extended abstract

This work introduces a novel direction finding (DF) technique for implementing an Indoor Real-Time Location System (IRTLS) for mobile Internet of Things (IoT) devices using WiFi technology. One of the easiest already developed DF systems is based on the amplitude monopulse function employing a pair of tilted antennas and the RSSI metric [1]. This technique is referred to as amplitude-only direction finding (AODF). The main drawback of the AODF technique is its narrow field of view (FoV), the region over which the DoA can be effectively estimated [2]. Because of the limited FoV of the AODF technique, we propose combining AODF with the phase-difference monopulse technique (PDDF) to create a novel DF system. For that, the PDDF technique is implemented by acquiring the phase difference between the antennas. To this end, we use CSI information from the OFDM modulation employed in the Wi-Fi protocol, acquired with the novel open-source FeitCSI tool [3]. It enables us to obtain CSI using a COTS solution that features an embedded computer and a WiFi MiMo card.

With this new proposal, the limited FoV of the AODF systems (typically 60°, as demonstrated in previous work [1]) is enhanced without the use of complex antenna systems, such as mechanically rotating antennas. Moreover, the proposed PDDF technique uses only two antennas connected to a single WiFi MiMo NIC, making it a much simpler topology than conventional PDDF Wi-Fi systems, which typically use four or more antennas and several synchronized receivers. Therefore, the main contribution of this work is the proposal of a single AP WiFi DoA system employing a pair of tilted directive antennas connected to a 2X2 MIMO WiFi card, which allows extending the FoV to 180° by acquiring the RSSI and CSI to combine amplitude and phase comparison, mixing the AODF and PDDF techniques.

The system has been tested in a highly challenging indoor environment, specifically in a classroom at the Polytechnic University of Cartagena. The test area is 4.4 m x 17 m, where we created a grid of 95 test points with an adjacent distance of 90 cm placed in the student desks and the main corridor. 200 samples of RSSI per antenna and 200 samples of phase-difference CSI were acquired at each of the 95 test points. To estimate the real DoA, the amplitude-only monopulse function ($AOMF_{RSSI}(\theta)$) and the phase-difference monopulse function ($PDMF_{CSI}(\theta)$) calibrated at the anechoic chamber were employed. This is an advantage of the monopulse DoA systems compared to the fingerprinting systems, which require calibration at each location where it is deployed. With each one of the 200 samples at each of the points, the estimated DoA with both systems is computed via the pseudospectrum $AOPS_{RSSI}(\theta)$ and the $PDPS_{CSI}(\theta)$, and with the multiplication of both, a single θ value is estimated.

As shown in Table 1, the mean average error (MAE) of the AO&PD-DF system combined is 8.4° , the Root Mean Square Error is 12.9° , and the 90^{th} percentile is 19.9° . These values are a significant improvement over the results of each system employed individually. The AODF system has a 27.4° MAE, a 33.3° RMSE, and a 48.1° 90^{th} percentile. Regarding PDDF, the MAE is 12.9° , the RMSE is 18.2° , and the 90^{th} percentile is 34.9° .

Table 1. DoA estimation errors for the different DF methods

FoV	Method	MAE	RMSE	90^{th} percentile
90° $\theta = [-45^\circ, 45^\circ]$	AODF	27.4°	33.3°	48.1°
	PDDF	12.9°	18.2°	34.9°
	AO&PD-DF	8.4°	12.9°	19.9°

References

- [1] J. L. Gomez-Tornero, D. Canete-Rebenaque, J. A. Lopez-Pastor, and A. S. Martinez-Sala, ‘Hybrid Analog-Digital Processing System for Amplitude-Monopulse RSSI-Based MIMO WiFi Direction-of-Arrival Estimation’, *IEEE Journal on Selected Topics in Signal Processing*, vol. 12, no. 3, 2018, doi: 10.1109/JSTSP.2018.2827701.
- [2] S. M. S. and D. K. Barton, ‘Monopulse principles and techniques’, *Artech House Radar Library*. Artech House, 2011, doi: 10.1109/MAP.1985.27828.
- [3] Miroslav Hutar, Peter Brida, and Juraj Machaj, ‘FeitCSI, the 802.11 CSI tool.’ Accessed: May 26, 2025. [Online]. Available: <https://feitcsi.kuskosoft.com/>

Analysis of Tools for Measuring Wi-Fi Network Performance with an Emphasis on Secure Indoor Localization and Coverage Optimization

Ivana Bridova^{1[0000-0002-7862-8792]}, Juraj Machaj^{1[0000-0002-7544-8796]}, and Peter Brida^{1[0000-0002-5442-9246]}

University of Žilina, Univerzitná 8215/1, 01026 Žilina, Slovakia

Abstract. Wi-Fi networks require reliable monitoring and optimization, as their quality is influenced not only by transmission speed but also by signal strength (RSSI), interference, stability, and channel utilization. This work evaluates available tools for indoor Wi-Fi network monitoring and proposes a solution for selecting an appropriate monitoring tool based on the required parameters in a specific environment. The comparison focuses on signal strength (RSSI), interference, connection stability, coverage visualization, and channel detection. The results showed that Acrylic Suite is suitable for professional analysis, while WiFiMan is practical for quick mobile diagnostics. AI-based tools provide automated visualization and offer potential for future network optimization. The choice of tool always depends on the intended goal—whether detailed analysis or rapid coverage assessment is required.

Keywords: Coverage analysis · Indoor localization · Network analysis tools · Network optimization · WiFi.

1 Introduction

The work focuses on comparing desktop tools, mobile tools, as well as applications utilizing artificial intelligence. The analysis included:

- **desktop tools** - NetSpot, inSSIDer, Homedale, WiFiMan, DL Network Monitor, Xirrus WiFi Inspektor, Vistumbler, WLAN monitor, RealWifi, Acrylic Suite.
- **mobile tools** - WiFi Analyzer, WiFi Scount, Wi-Fi Monitor, NetX, Net Analyzer, Netmonitor, Fing, Open Signal.
- **AI-based tools** - Ruijie Cloud AI Heatmap, Ekahau AI Pro

They were evaluated according to criteria such as RSSI measurement support, heatmap visualization, channel detection, data export, and suitability for practical field measurements. Special attention was given to tools capable of generating coverage maps and graphical outputs, as these visualizations are essential for identifying dead zones and interference. [1], [2]

1.1 Methodology of the work

The methodology included the selection of tools, the establishment of evaluation criteria, and the execution of measurements in an indoor environment across two floors with designated measurement points. At each point, the signal strength and available networks in the 2.4 GHz and 5 GHz bands were recorded. For detailed comparison, Acrylic Suite and WiFiMan were chosen for their heatmap support, complemented by the AI tool Ruijie Heatmap for automatic coverage map generation and noise analysis. [3]

1.2 Conclusion

In conclusion, it can be stated that the appropriate choice of tool depends on the purpose of the measurement. For technicians and administrators of large Wi-Fi infrastructures, Acrylic Suite is ideal, while WiFiMan is sufficient for quick field diagnostics. AI tools such as Ruijie Heatmap offer potential for analysis automation and represent a promising direction for future deployment in monitoring systems. The study confirmed that the combination of multiple tools enables more comprehensive network evaluation and significantly contributes to optimizing coverage and reducing interference in environments with high device density.

Acknowledgements This work was supported in part by the Slovak Scientific Grant Agency (VEGA) Grant Agency through the Research of Effective Localization and Communication Solutions in Heterogenous Wireless Networks 1/0580/25 and NATO SPS G6140 “Advanced technologies for Physical Resilience Of Critical Infrastructures (APRIORI)”.

References

1. S. Komatnani Govindan, H. Vijayaraghavan, K. Kishore Anthuvan Sahayaraj, and A. Mary Joy Kinol, “Optimizing Internet-Wide Port Scanning for IoT Security and Network Resilience: A Reinforcement Learning-Based Approach in WLANs with IEEE 802.11ah,” *Fiber and Integrated Optics*, vol. 43, no. 1, pp. 14–42, Jan. 2024, doi: 10.1080/01468030.2024.2342275.
2. M. R. Putri, M. Susanto, S. Purwiyanti, F. Hamdani, and A. A. Aziz, “Network Performance Evaluation based on RSSI for Inpatient Room Monitoring Application,” in 2024 FORTEI-International Conference on Electrical Engineering (FORTEI-ICEE), IEEE, Oct. 2024, pp. 148–153. doi: 10.1109/FORTEI-ICEE64706.2024.10824375.
3. M. R. Jivthesh, M. R. Gaushik, P. Adarsh, G. H. Niranga, and N. S. Rao, “A Comprehensive survey of WiFi Analyzer Tools,” in 2022 IEEE 3rd Global Conference for Advancement in Technology (GCAT), IEEE, Oct. 2022, pp. 1–8. doi: 10.1109/GCAT55367.2022.9972040.

Classical AI Search Algorithms Analysis for Optimized SAR Trajectory

R. Aguasca-Colomo¹[0000-0003-2217-8005], J.M. Molina-Gil²[0000-0001-7702-9264], M.A. Jiménez-Suárez¹, F.J. Navarro-Amador³

¹ Instituto Universitario de Sistemas Inteligentes y Aplicaciones Numéricas a la Ingeniería, Universidad de Las Palmas de Gran Canaria (ULPGC), Spain

² Departamento de Ingeniería Informática y Sistemas, Universidad de La Laguna (ULL), Spain

³ Ejército del Aire y del Espacio, Ala 46, 802 Escuadrón, Base Aérea de Gando, Gran Canaria, Spain

ricardo.aguasca@ulpgc.es

Abstract. Search and Rescue (SAR) operations in maritime environments pose major logistical and technical challenges due to the inherent uncertainty in target localization. This paper analyzes and compares four representative algorithms of Classical AI, Breadth-First Search (BFS), Depth-First Search (DFS), Uniform Cost Search (UCS), and Wavefront Expansion, applied to aerial SAR trajectory planning. Two controlled simulation scenarios were developed to evaluate these algorithms. Results demonstrate that deterministic search algorithms provide a robust foundation for designing efficient exploration patterns in SAR missions.

Keywords: SAR, AI, Trajectory Planning Algorithms.

1 Introduction and Motivation

Aerial Search and Rescue (SAR) missions are essential for maritime safety, particularly in high-traffic regions such as the Atlantic near the Canary Islands. Optimized trajectory planning and efficient search-area management are vital to reducing response times and increasing mission success rates. Classical Artificial Intelligence (AI) provides a systematic framework for improving traditional search techniques through structured state-space exploration and deterministic decision-making [1]. This study evaluates the effectiveness of classical AI search algorithms in generating optimized SAR trajectories. The objective is to establish a comparative basis that supports future integration of predictive and learning-based models for real-world operational use.

2 Methodology and Development

Four representative search algorithms—BFS, DFS, UCS, and Wavefront Expansion—were implemented in Python within a Google Colab environment. Each algorithm was tested on two maritime-inspired operational templates: (a) an open search area expanding from the last known position (LKP) and (b) a T-shaped inter-island environment with navigational constraints. Both templates included High-Probability Zones (HPZs)

representing the most likely target locations. A total of 400 simulations (100 per algorithm and per environment) were conducted, recording performance indicators such as execution time, explored nodes, total distance, and successful detection rate. This experimental design ensured a consistent comparative framework under equivalent operational conditions [2].

3 Main Results

The results demonstrate that all algorithms achieved a 100% success rate in target localization, although with substantial differences in efficiency and coverage. BFS ensured systematic area exploration and stable performance, ideal for full-coverage missions. DFS produces the shortest routes but with greater variability in node exploration, being efficient in constrained, high-priority environments. UCS provided balanced performance between cost and coverage, suitable for missions where movement costs vary (fuel, turbulence, risk). Wavefront Expansion achieved maximal coverage but incurred the highest computational cost and total distance, making it appropriate only for scenarios with high spatial uncertainty. Comparative analysis confirmed that environment topology and HPZ positioning strongly affect search behavior. BFS and UCS showed higher stability between runs, whereas DFS achieved faster convergence [3,4].

4 Conclusions and Future Work

The study confirms the applicability of classical AI search algorithms in SAR trajectory planning, emphasizing the balance between coverage, speed, and available information. BFS and UCS are suitable for systematic missions, DFS for predictable patterns, and Wavefront Expansion for exhaustive searches requiring full coverage. Future research will integrate environmental and drift-prediction models using recurrent neural networks and reinforcement learning (RL) to enhance adaptability in dynamic conditions.

Acknowledgments. Authors would like to thank Spanish Air Force ARCC team. This work was funded by SCITALA C064/23, funded by INCIBE, with the Recovery, Transformation, and Resilience Plan (Next Generation) financed by EU.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. IAMSAR Manual, Vol. II – Mission Coordination (ICAO/IMO, 2022)
2. Datacamp, “Breadth-First Search in Python”, <https://www.datacamp.com/es/tutorial/breadth-first-search-in-python>, last accessed 2025/04/02
3. GeeksforGeeks, “Uniform Cost Search Algorithm”, <https://www.geeksforgeeks.org/dsa/uniform-cost-search-dijkstra-for-large-graphs/>, last accessed 2025/04/02
4. M. Soulignac, et alt. "Adapting the wavefront expansion in presence of strong currents", IEEE Int. Conf. on Robotics and Automation, Pasadena, CA, USA, 2008, pp. 1352-1358

Detecting Object Consistency Errors in Reference-Conditioned Image Generation

Anna Maschek^[0009-0006-5090-7569], Ricardo Serban^[0009-0005-5284-842X],
 Andreas Stöckl^[0000-0003-1646-0514], and David C. Schedl*^[0000-0002-7621-3526]

University of Applied Sciences Upper Austria, Hagenberg, Austria

*david.schedl@fh-hagenberg.at

Introduction: While recent advances in generative image editing models have enabled remarkable capabilities for controlling and editing generated scenes, a critical challenge remains in verifying the accurate display of objects from a given reference image. Although prior work has addressed facial consistency [2] and detected inconsistencies related to the text prompt [1], non-human reference objects remain underexplored. In this work, we introduce a dataset of reference-generation image pairs featuring furniture and household items with labeled errors to support the evaluation of methods for detecting object representation errors. Furthermore, we propose a framework that leverages multimodal large language models (MLLM) to identify such errors.

Dataset: The dataset comprises 660 images that were generated using 11 state-of-the-art image editing models, each applied to the same 60 reference images and corresponding text prompts. The reference objects are evenly divided into 30 household items and 30 furniture items, with each prompt specifying that the objects should be placed within a scene for a product photograph. To ensure consistent labeling of errors in the generated images, we introduce the taxonomy illustrated in Fig. 1. In addition to the error annotations – labeled with bounding boxes for object-related errors or as image classifications for image-wide errors – the dataset includes a general image quality rating as well as a bounding box enclosing the reference object itself, additionally enabling evaluation of methods for detecting a specific reference object within a generated image.




Fig. 1. Taxonomy of errors in generated images containing a reference object.

Methodology: Our comparative analysis framework utilizes MLLMs such as Google’s Gemini 2.5 Flash [3] for visual difference and bounding box detection. We test the detection performance of different MLLMs across different input modalities: full images and cropped object regions. For an MLLM to process

both images in the intended way, we created a custom prompt to output a JSON file containing refined bounding boxes and concise difference descriptions. The prompt directs the MLLM to compare only the object between the target and query images, highlight major structural or added elements, ignore background details, and return normalized box coordinates with short labels for differences.



Fig. 2. The MLLM (i.e., Gemini 2.5 Flash) detects the query objects (green bounding box; center image) as well as highlights differences between Input A (query; left) and Input B (target; right). Each bounding box and label set has its own distinct color.

Preliminary Results: Figure 2 illustrates a case where Gemini 2.5 Flash correctly localized the object and identified its differences. In a preliminary manual evaluation of 24 images, the model detected 41 of 67 annotated errors (61.2% recall) and identified 110 errors in total, including 14 false positives, determined by manual judgment. The remaining 55 detected errors were present in the images but not labeled by the human annotator. Two recurring discrepancies are the *looks 3D rendered* label, which the human annotator reserved for very extreme cases, whereas the model returns this label relatively fast, when parts of the scene show overly perfect textures. Similarly, the model detected more *error in size* labels, often omitted by the human due to the difficulty of verifying object dimensions (e.g., a pan) from the input image.

Conclusion: MLLMs such as Gemini 2.5 Flash demonstrate strong potential for vision tasks. With effective prompt design, they can reliably detect objects and identify visual differences between images, suggesting a viable path toward automated object difference detection.

References

1. Wu, T.-H., Lian, L., Gonzalez, J.E., Li, B., Darrell, T.: Self-Correcting LLM-Controlled Diffusion Models. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6327–6336 (2024). <https://doi.org/10.1109/CVPR52733.2024.00605>
2. Hochberg, D.C., Anschel, O., Shoshan, A., Kviatkovsky, I., Aggarwal, M., Medioni, G.: Towards Quantitative Evaluation Metrics for Image Editing Approaches. In: 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 7892–7900 (2024). <https://doi.org/10.1109/CVPRW63382.2024.00787>.
3. Comanici et al.: Gemini 2.5: Pushing the Frontier with Advanced Reasoning. <https://doi.org/10.48550/arXiv.2507.06261>, 2025

Evaluating Agentic Web Behavior on Financial Websites: A Case Study with Raiffeisen Landesbank Oberösterreich

Joel Nitu^{*1}, Lisa Kühhas^{*2}, Michael Ehrengruber^{*2}, and Andreas Stöckl¹

University of Applied Sciences Upper Austria, Hagenberg, Austria
Raiffeisen Landesbank Oberösterreich, Linz, Austria

Abstract. Rapid adoption of generative AI assistants in mainstream browsers signals a shift to agentic browsing, where autonomous models traverse, interpret, and act in complex web environments. In collaboration with Raiffeisen Landesbank Oberösterreich, we examine how AI agents navigate and interact with the bank's website. We ran 390 runs in which Claude 4.0 Sonnet achieved the highest task success (90%) but overreached on leasing flows, GPT-5 was most conservative (avoiding unauthorized submissions and fabrication), and Gemini 2.5 Pro more often accepted all cookies and upsold premium products, with cookies rejected in only 20% of runs.

Keywords: Trustworthy AI · Agentic Web Interaction · Large Language Models · Financial Websites.

1 Introduction & Related Work

Modern browsers are shifting to agentic browsing, embedding AI that can act for users (e.g., Edge Copilot Mode; OpenAI Atlas), raising safety, compliance, and site-operability questions for machine actors. Web-agent research has progressed from controlled benchmarks like MiniWoB [1] to more realistic evaluations such as WebArena [2]. Parallel HCI work shows consent UIs employ dark patterns with GDPR implications [3]. We extend this line by evaluating LLM agents powered by leading large language models (LLMs) on a production banking site and a controlled clone, analyzing task success, consent handling, and how DOM/semantic cues enable robust, compliant automation.

2 Methodology

Agents powered by three leading LLMs, OpenAI GPT-5, Anthropic Claude 4.0 Sonnet, and Google DeepMind Gemini 2.5 Pro, performed 13 representative banking tasks, from information retrieval (e.g., finding a nearby branch) to transactional flows (e.g., opening an account, leasing a car). Each model executed 10 runs per task on both authentic Raiffeisen pages and cloned counterparts (built in Figma and Bolt.new) to ensure experimental control and data safety. Realistic personas (student, family, business) emulated user-specific browsing contexts.

3 Results

Across 390 sessions, all runs completed without breaking errors, indicating the site supports machine traversal. Goal attainment and ethics diverged: GPT-5 was most conservative, avoiding unauthorized submissions and fabricated data, while Claude 4.0 Sonnet had the highest task-success rate (90%) but overreached with unintended leasing submissions. Gemini 2.5 Pro matched Claude's completion rate yet more often accepted all cookies and upsold premium products. Cookies were rejected in only 20% of runs and never customized. Optional marketing boxes were frequently left enabled when pre-ticked, newly ticked despite being unnecessary (when placed where options are usually mandatory), or included via an "Accept All" shortcut, when provided.

Agents read the DOM top-to-bottom [4], overusing header utilities (e.g., " findeR," search bar) even when hindering progress, and ignoring sticky footers due to DOM-visual mismatches. Absent or ambiguous success feedback triggered repeated submissions, underscoring the need for explicit, machine-perceivable confirmations.

4 Discussion and Conclusion

Three central insights emerge with direct implications for an agentic web. (1) *Prompt phrasing matters*: Neutral queries elicit low-cost recommendations, whereas "benefit"-framed prompts tend to trigger premium upsells. (2) *Guardrails differ*: GPT-5's refusal to fabricate missing identity data reduces completion rates but improves compliance. (3) *Design responsibility shifts*: Websites must encode machine-readable eligibility, consent, and intent cues so agents do not misinterpret requirements and understand what and what not they can or should interact with. Accessibility principles therefore extend beyond human inclusion to *AI operability*, ensuring predictable behavior across browser-integrated assistants.

References

1. Shi, T., Karpathy, A., Fan, L., Hernandez, J., Liang, P.: World of Bits: An Open-Domain Platform for Web-Based Agents. In: Proceedings of the 34th International Conference on Machine Learning (ICML). PMLR **70**, 3135–3144 (2017)
2. Zhou, S., Xu, F.F., Zhu, H., Zhou, X., Lo, R., Sridhar, A., Cheng, X., Ou, T., Bisk, Y., Fried, D., Alon, U., Neubig, G.: WebArena: A Realistic Web Environment for Building Autonomous Agents. arXiv preprint arXiv:2307.13854 (2024)
3. Nouwens, M., Liccardi, I., Veale, M., Karger, D., Kagel, L.: Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20), 1–13 (2020)
4. Nitu, J., Mühle, H., Stöckl, A.: Machine-Readable Ads: Accessibility and Trust Patterns for AI Web Agents Interacting with Online Advertisements. arXiv preprint arXiv:2507.12844 (2025). Available at: <https://arxiv.org/abs/2507.12844>

Improving LIME with Logistic Surrogate Models

Stephan Dreiseitl[0000-0001-5647-6153]

Dept. of Software Engineering
 University of Applied Sciences Upper Austria, A-4232 Hagenberg, Austria
 stephan.dreiseitl@fh-hagenberg.at

With machine learning becoming the method of choice for model building in a variety of domains, it becomes increasingly important to incorporate the outputs of such models in human decision making [2,7]. This process is easier for so-called *white box models* that are human interpretable, either because of their explicit knowledge representation (e.g. decision trees) or because of the way that their parameters can be interpreted (e.g. linear or logistic regression).

For opaque *black-box models* such as neural networks and ensemble methods based on decision trees (e.g. random forests, gradient boosted trees), no such interpretation is readily available. As a consequence, the field of *explainable AI* (XAI) [1,8] has substantially increased in importance in recent years, and a variety of algorithms for the post-hoc explanation of black-box models have been developed.

Of these, LIME (*Local Interpretable Model-Agnostic Explanations* [5]) and SHAP (*SHapley Additive exPlanations* [3]) have become the most widely used. Both LIME and SHAP are model-agnostic, i.e., they can be applied to any black-box model, and both provide local explanations. In the context of tabular data, which is the focus of this work, this means that for any given datapoint (a vector of its feature values, either categorical or numerical), LIME and SHAP both give a ranking of feature importance. This feature ranking is the human-interpretable explanation: higher-ranked features contribute more to a model's predication for that datapoint than lower-ranked features, and are thus more important for this particular case. The difference between these two approaches lies in how they compute the feature importance ranking: for LIME, the rankings are derived from the feature rankings of surrogate linear models that locally approximate the black-box model output around a particular datapoint. For SHAP, the rankings are based on the game-theoretic concept of *Shapley values* [6], which in theory requires computing model outputs of all possible feature combinations.

In this work, we focus on LIME, and introduce two simple variants that change how LIME's surrogate models approximate a black-box classifier. The motivation for this is the observation that in standard LIME, as originally published by Ribeiro et al. [5] and implemented in standard Python and R libraries, the classifier outputs are approximated by *linear regression* surrogates. The black-box and surrogate models do not align, however, because the former outputs probabilities in the range [0, 1], whereas the latter outputs unbounded real values. For this reason, the standard choice in probabilistic classification with linear models is *logistic regression*. We thus propose two minor modifications to the LIME pipeline:

Hard logistic regression This variant trains the local surrogate models on the discretized 0/1 class predictions of the black-box model.

Soft logistic regression This variant trains the local surrogate models directly on the probabilistic class membership probabilities produced by the black-box model.

In the full version of the paper, we will investigate how these two variants compare to the standard LIME linear regression surrogates in terms of *fidelity* (how closely the surrogate output matches the black-box model output in the vicinity of a datapoint x ?), *stability* (how much does the randomness of the dataset generated around x influence the feature importance ranking?) and *insertion sufficiency* (how well do the top k features alone compare to the full feature set?). We will use three well-known large datasets from the OpenML repository [4] to conduct our experiments. Preliminary results indicate that the soft-logistic surrogates perform better than both the hard-logistic and linear baselines on most metrics and datasets.

References

1. Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D.: A survey of methods for explaining black box models. *ACM Computing Surveys* **51**(5), 93:1–93:42 (2018). <https://doi.org/10.1145/3236009>
2. Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., Mullainathan, S.: Human decisions and machine predictions. *The Quarterly Journal of Economics* **133**(1), 237–293 (2018). <https://doi.org/10.1093/qje/qjx032>
3. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. pp. 4768–4777 (2017). <https://doi.org/10.5555/3295222.3295230>
4. OpenML: An open platform for sharing machine learning data and experiments. <https://www.openml.org/>, accessed 12 September 2025
5. Ribeiro, M.T., Singh, S., Guestrin, C.: “Why Should I Trust You?”: Explaining the predictions of any classifier. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. pp. 1135–1144. ACM, New York, NY, USA (2016), <https://doi.org/10.1145/2939672.2939778>
6. Shapley, L.S.: A value for n -person games. In: Kuhn, H.W., Tucker, A.W. (eds.) *Contributions to the Theory of Games, Volume II, Annals of Mathematics Studies*, vol. 28, pp. 307–317. Princeton University Press (1953)
7. Topol, E.J.: High-performance medicine: the convergence of human and artificial intelligence. *Nature Medicine* **25**(1), 44–56 (2019). <https://doi.org/10.1038/s41591-018-0300-7>
8. Yang, W., Wei, Y., Wei, H., Chen, Y., Huang, G., Li, X., et al.: Survey on explainable AI: From approaches, limitations and applications aspects. *Human-Centric Intelligent Systems* **3**(4), 287–307 (2023). <https://doi.org/10.1007/s44230-023-00038-y>

Investigating Memorization and Generalization of LLMs Using Chess as a Testbed

Andreas Stöckl

University of Applied Sciences Upper Austria, Hagenberg, Austria
 *andreas.stoeckl@fh-hagenberg.at

Abstract. Board games offer a clean, deterministic sandbox to probe whether Large Language Models (LLMs) merely *memorize* frequent sequences or *generalize* underlying rules. We present a focused study that trains a GPT2-like Transformer on synthetically generated chess transcripts and evaluates emergent state tracking, legality, and strategy. Early training yields statistically likely but illegal moves; extended training produces largely legal yet often suboptimal play, indicating partial generalization.

Keywords: Large Language Models · Memorization vs. Generalization · Chess.

1 Background & Motivation

Board games have long served as a benchmark for artificial intelligence research due to their well-defined rules and strategic depth. Classical games such as chess and Go have historically been used to evaluate AI algorithms, and recently, they have also emerged as valuable test beds for exploring the capabilities of large language models (LLMs) [1]. Early work has extended the notion of scaling laws—traditionally observed in natural language processing—to board game scenarios. For instance, Jones [2], and Stöckl [3] investigate scaling trends during LLM training by analyzing performance on board game tasks, highlighting that lessons learned from these domains can offer insights into both language understanding and strategic reasoning. Unlike natural language, the chess domain is deterministic and formally structured, which allows researchers to directly assess whether an LLM accurately tracks the underlying game state.

2 Method

We synthesize full chess games via `python-chess`, uniformly sampling legal moves to termination (checkmate/stalemate/insufficient material). This ensures broad rule coverage without human resignations truncating games.

We implement a GPT2-like Transformer in PyTorch with 12 layers, 12 heads, embedding dim756, context length 256, dropout 0.1, batch size 32. We probe legality and state tracking after a fixed prefix (e.g., 10 plies) by sampling continuations and counting illegal/contradictory moves (e.g., moving a just-moved piece

unlawfully, leaving king in check). We analyze snapshots at early (600 steps), mid (3k), and later (10k) training to observe rule acquisition dynamics.

3 Results

Early stage (600 steps). Samples contain plausible but often illegal moves; the model lacks a stable internal board state and exhibits frequent rule breaks—indicative of memorized openings without rule grounding.

Mid stage (3,000 steps). Errors concentrate on check/king safety and temporal constraints (e.g., violating check resolution, reusing a piece inconsistently). The world-model is emerging but noisy.

Later stage (10,000 steps). Only a handful of illegality cases persist (e.g., rook move while king in check). Most samples are legal but not strategically strong—evidence of partial generalization: rules are internalized enough for legality, not yet for strong play.

Takeaways. (i) Legality increases monotonically with training; (ii) context length (256) is usually sufficient to cover entire games, aiding state tracking; (iii) learning dynamics mirror a transition from frequency-driven recall to rule-constrained generation.

4 Discussion

Frequency priors dominate early (openings/common patterns). With more steps, constraints (checks, piece mobility, castling/promotion) are respected more often, suggesting the network encodes a latent board representation derived solely from the move history.

Residual failures typically involve check resolution and move-temporal consistency—classic pitfalls of implicit state inference from sequences. These align with prior observations of LLMs struggling with long-horizon, rule-bound planning without explicit state.

5 Conclusion

Using chess as a structured probe, we observe a clear progression from memorization to partial rule generalization in an LLM trained solely on move sequences.

References

1. Toshniwal, Shubham and Wiseman, Sam and Livescu, Karen and Gimpel, Kevin, Chess as a testbed for language model state tracking, Proceedings of the AAAI Conference on Artificial Intelligence (2022)
2. Andy L. Jones, Scaling Scaling Laws with Board Games, ArXiv (2021)
3. Stöckl, Andreas, Watching a Language Model Learning Chess, Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2021)

Guided Code Generation with Large Language Models and Static Code Analysis*

Mert Dalkilic¹[0009-0001-8542-5176] and Christoph Praschl¹[0000-0002-9711-4818]

University of Applied Sciences Upper Austria, Hagenberg, Austria
 mert.dalkilic@fh-hagenberg.at
 christoph.praschl@fh-hagenberg.at




Fig. 1. Iterative workflow of LLM-assisted code generation and validation with static analysis feedback.

Introduction: Large Language Models (LLMs) [1] can effectively generate source code [2] but operate on a static knowledge base representing a snapshot in time [3]. As a result, recently discovered vulnerabilities, updated security practices, or project-specific coding standards are typically unknown to the model. Static code analysis tools such as SonarQube [4], in contrast, are continuously updated to identify new bugs, code smells, as well as security gaps, and can cover internal coding guide lines.

By combining LLM-based code generation with automated static analysis in an iterative feedback loop, detected issues can be summarized and reintroduced to the LLM as additional context as shown in Fig. 1. This should enable progressive code refinement and reduce security and quality risks.

Related Work: Ravi et al. [5] introduce LLMLOOP, an iterative feedback loop that improves LLM-generated code and tests through repeated refinement. Blyth et al. [6] incorporate static analysis to enhance correctness, security, and maintainability. Yang et al. [7] survey LLM-based program repair, covering prompt-based, pipeline, and agent-driven approaches. Our work builds on these ideas by unifying user prompts, LLM generation, static analysis, and iteration tracking in a single end-to-end pipeline.

Methodology: We employ an iterative LLM-assisted code-generation workflow coupled with static analysis. For each prompt, the LLM produces an initial

* Supported by research fund of the province of Upper Austria.

code candidate, which is evaluated using SonarQube. Analysis findings (security, quality, maintainability) are fed back to the model to guide regeneration. Iteration continues until issue thresholds are satisfied or a maximum iteration count is reached.

Results: We evaluated our approach on 50 code generation prompts covering security vulnerabilities, code quality, algorithmic tasks, and complexity challenges. Each prompt underwent up to five iterative refinement cycles, incorporating SonarQube feedback into the LLM’s next generation. For each run, we recorded initial and final issue counts, iteration numbers, and categorized issues by type and severity. As shown in Table 1, 94% of all runs successfully completed, resolving all detected issues after an average of 1.34 refinement iterations, with the majority (74%) requiring only a single iteration.

Table 1. Performance of the iterative LLM-based code refinement pipeline across four issue categories (50 evaluation runs).

Category	Runs	Success Rate	Avg. Iteration
Security	22	95.5%	1.3
Complexity	12	91.7%	1.2
Quality	6	100.0%	1.3
Algorithmic	7	100.0%	1.7
Overall	50	94%	1.34

Conclusion and Outlook: We presented an iterative LLM code generation pipeline guided by static analysis feedback. On 50 prompts, 94% of runs successfully resolved issues, most in a single iteration. Future work will explore larger codebases, multiple models, and more advanced feedback integration.

References

1. Alec Radford, Karthik Narasimhan, Tim Salimans, and et al. Improving language understanding by generative pre-training. 2018.
2. Jianxun Wang and Yixiang Chen. A review on code generation with llms: Application and evaluation. In *2023 IEEE International Conference on Medical Artificial Intelligence (MedAI)*. IEEE, 2023.
3. Jeffrey Cheng, Marc Marone, Orion Weller, and et al. Dated data: Tracing knowledge cutoffs in large language models. *arXiv preprint arXiv:2403.12958*, 2024.
4. Ping Yu, Yijian Wu, Jiahua Peng, and et al. Towards understanding fixes of sonarqube static analysis violations: A large-scale empirical study. In *2023 IEEE International Conference on Software Analysis, Evolution and Reengineering*, 2023.
5. Ravin Ravi, Dylan Bradshaw, Stefano Ruberto, and et al. Llmloop: Improving llm-generated code and tests through automated iterative feedback loops. 07 2025.
6. Scott Blyth, Sherlock A. Licorish, Christoph Treude, and et al. Static analysis as a feedback loop: Enhancing llm-generated code beyond correctness, 2025.
7. Boyang Yang, Zijian Cai, Fengling Liu, and et al. A survey of llm-based automated program repair: Taxonomies, design paradigms, and applications, 2025.

Synthetic UAV Imagery for Ecological Monitoring Using Generative Artificial Intelligence*

Christoph Praschl^[0000-0002-9711-4818], Michael T. Karbasch^[0009-0000-2322-366X], Kathrin Probst^[0009-0001-4705-0633], and David C. Schedl^[0000-0002-7621-3526]

University of Applied Sciences Upper Austria, Hagenberg, Austria
 {firstname.lastname}@fh-hagenberg.at

Introduction: Uncrewed aerial vehicles (UAVs) have become essential tools in ecological monitoring, enabling high-resolution mapping of wildlife, vegetation, and habitats across large and difficult-to-access areas [6]. However, developing robust computer vision models for such applications remains challenging due to the limited availability of diverse and well-annotated UAV datasets. Collecting real imagery in ecological settings is often resource-intensive and constrained by pilot availability, flight regulations, weather conditions, seasonal patterns, and the need to avoid disturbing sensitive species.

Recent advances in generative artificial intelligence (GenAI), like diffusion models [5], offer a promising solution by creating photorealistic synthetic imagery that can augment or partially replace real-world datasets. Synthetic data generation can help overcome logistical and ethical barriers, provide greater control over ecological scenarios, and support targeted augmentation of rare species and environmental conditions.

Related Work: Kulas et al. [3] describe a procedural pipeline to generate synthetic thermal aerial images, adding object classes (including animals) into backgrounds, to augment datasets for detection. They use inpainting techniques to create synthetic images and like that do not generate entire images, but composite synthetic objects into real thermal aerial images. In comparison to that approach, we generate full synthetic images.

Methodology: We utilize state-of-the-art generative AI architectures for synthesizing high-resolution aerial wildlife imagery. These models are adapted to the target domain via parameter-efficient Low-Rank Adaptation (LoRA) [2] fine-tuning, using a small, curated dataset of 30 sample images. The sample images are captioned using Bootstrapped Language-Image Pretraining (BLIP)[4]. Training is done with controlled initialization seeds and regular checkpointing to ensure comparability and reproducibility. Finally, model outputs are assessed using the Kernel Inception Distance (KID)[1].

Results: We fine-tune Stable Diffusion 3.5 and FLUX.1 Dev using LoRA with different permutations of rank ($r \in \{16, 32, 48\}$) and alpha ($\alpha \in \{8, 16, 32, 64\}$) values to synthesize UAV imagery across multiple sensor styles. Specifically, we generate RGB as well as two thermal representations (inferno and white-hotspot). The models are fine-tuned over 6000 iterations with checkpoints stored

* Supported by Austrian Research Promotion Agency FFG under project number 917796.

every 250 iteration steps. Table 1 summarizes the KID scores for each model, target style, and rank, whereas the alpha and training steps are selected based on the best performance (lowest KID score).

For FLUX, the best configuration for the *inferno* style is $r=48/\alpha=8$, while *RGB* performs best with $r=16/\alpha=8$ and *white-hotspot* with $r=16/\alpha=64$. For Stable Diffusion, the *inferno* style achieves its best result with $r=48/\alpha=32$, *RGB* with $r=32/\alpha=16$, and *white-hotspot* with $r=48/\alpha=8$.

Table 1. Best KID scores (lower is better) achieved for each model-style and LoRA rank setup, including the corresponding alpha values, as well as the training step.

Model	Style	Rank r	Alpha α	Step	KID \downarrow
FLUX	Inferno	16	8	6000	0.052592270
		32	32	6000	0.021656818
		48	8	4500	0.01248352
	RGB	16	8	6000	0.07890453
		32	8	5000	0.126892527
		48	32	5000	0.088144607
	White-Hotspot	16	64	4500	0.05818776
		32	32	4500	0.071230147
		48	64	5000	0.083838561
Stable Diffusion	Inferno	16	48	4000	0.063205048
		32	32	5500	0.047370675
		48	32	2000	0.02233721
	RGB	16	32	500	0.017246468
		32	16	4500	0.00638779
		48	16	3500	0.012184045
	White-Hotspot	16	8	3500	0.059686988
		32	8	5750	0.031524591
		48	8	4000	0.02554271

Conclusion: Overall, the results indicate that higher LoRA ranks of 48 tend to perform better for thermal modalities (except of White-Hotspot with FLUX), whereas lower ranks seem sufficient for RGB generation, while no observable correlation was found for the alpha value. These findings highlight modality-dependent behavior in LoRA-based fine-tuning and aligns with the hypothesis that the foundation models used have been trained on substantially more RGB data than other spectral data.

References

1. Bińkowski, M., Sutherland, D.J., Arbel, M., Gretton, A.: Demystifying mmd GANs. In: International Conference on Learning Representations (ICLR) (2018)
2. Hu, E.J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W.: Lora: Low-rank adaptation of large language models (2021)
3. Kulas, A.B., Jurasovic, A., Bogdan, S.: Unlocking thermal aerial imaging: Synthetic enhancement of uav datasets. In: European Conf. on Mobile Robots. IEEE (2025)
4. Li, J., Li, D., Xiong, C., Hoi, S.: Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation (2022)
5. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2021)
6. Wang, D., Shao, Q., Yue, H.: Surveying wild animals from satellites, manned aircraft and unmanned aerial systems: A review. Remote Sensing (2019)

Explainable and Privacy-Preserving AI for Healthcare: A Local GraphRAG Pipeline for Nursing Care Guidelines

Oliver Krauss^{1,2}[0000-0002-8136-2606] and Aleksandra Mileva³[0000-0003-0706-6355]

¹ University of Applied Sciences Upper Austria, Campus Hagenberg, Softwarepark 11, 4232 Hagenberg, Austria
 oliver.krauss@fh-hagenberg.at

² Faculty of Computer Science, Goce Delcev University, Stip, N. Macedonia
 aleksandra.mileva@ugd.edu.mk

Introduction Home care increasingly relies on layperson caregivers who lack formal medical training and struggle with complex nursing care guidelines [3, 4]. Traditional cloud-based AI approaches introduce privacy risks and lack explainability for non-medical users [2, 1]. This work addresses these challenges through a fully local, privacy-preserving pipeline combining Microsoft GraphRAG [8] with locally hosted language models, ensuring data sovereignty while providing explainable AI through knowledge graph structures.

Methodology and Security Our system employs Microsoft GraphRAG [8] to construct knowledge graphs from selected excerpts of seven German nursing care standards (*Expertenstandards*, approximately 324,000 words) and synthetic patient anamnesis documents from professional home carers, enabling the system to link diagnoses, care procedures, and guidelines to concrete patient cases. Key technical contributions include custom healthcare-specific entity extraction prompts that replace generic examples with medical scenarios, significantly improving accuracy, and compatibility fixes enabling Microsoft GraphRAG to work with LM Studio's local model hosting. All processing occurs on-premises on consumer-grade hardware (Intel Core i7-13700H, 31GB RAM, NVIDIA RTX A1000 GPU), with LM Studio for local model hosting and Docker containers for isolation, ensuring all data processing occurs within the organization's infrastructure and addressing GDPR compliance [6, 1, 5].

Explainability Through Knowledge Graphs Microsoft GraphRAG's knowledge graph structure provides inherent explainability advantages over traditional vector similarity-based RAG approaches [7]. The graph structure enables traceability of query responses, showing which entities and relationships contributed to answers, with links to source text units and documents that allow verification against authoritative source materials. The hierarchical community detection enables both high-level thematic understanding (global queries) and detailed, source-grounded responses (local queries) with source attribution, making AI reasoning processes more inspectable and understandable for non-medical users compared to black-box approaches.

Application and Results Our system successfully processes the corpus of German nursing care standards, creating a searchable knowledge base for layperson caregivers. Example queries demonstrate capabilities: "How can I better communicate with my mother who has dementia?" returns comprehensive guidance on person-centered communication strategies, while "What specific information should I record when documenting a pressure ulcer?" returns precise, actionable guidance with full source attribution. The system has been validated on consumer-grade hardware, demonstrating that local deployment does not compromise functionality while significantly enhancing privacy and security for sensitive home care environments.

Conclusion and Outlook Our system demonstrates that explainable, privacy-preserving AI for home care is achievable through local deployment of Microsoft GraphRAG with locally hosted language models on consumer-grade hardware, making complex medical information accessible to layperson caregivers while ensuring privacy and compliance through data sovereignty.

References

1. Regulation (eu) 2016/679 of the european parliament and of the council on the protection of natural persons with regard to the processing of personal data. Official Journal of the European Union, L 119/1 (April 2016), <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
2. Regulation (eu) 2024/1689 of the european parliament and of the council of 13 june 2024 laying down harmonised rules on artificial intelligence. Official Journal of the European Union, L series (July 2024), <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
3. Amann, J., Blasimme, A., Vayena, E., Frey, D., Madai, V.I.: Explainability for artificial intelligence in healthcare: a multidisciplinary perspective. *BMC Medical Informatics and Decision Making* **20**(310) (2020). <https://doi.org/10.1186/s12911-020-01332-6>
4. Char, D.S., Shah, N.H., Magnus, D.: Implementing machine learning in health care — addressing ethical challenges. *New England Journal of Medicine* **378**(11), 981–983 (2018). <https://doi.org/10.1056/NEJMp1714229>
5. Combe, T., Martin, A., Di Pietro, R.: To docker or not to docker: A security perspective. In: *IEEE Cloud Computing*. vol. 3, pp. 54–62 (2016). <https://doi.org/10.1109/MCC.2016.100>
6. Kraley, A., Kantor, I., Gutiérrez, R.: Sovereign ai ecosystems: Navigating global ai infrastructure & data governance. Tech. rep., POLITICO Research & Analysis Division (September 2024), <https://www.politico.eu/wp-content/uploads/2024/09/15/Sovereign-AI-Ecosystems.pdf>, independent research report presented by G42
7. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.t., Rocktäschel, T., Riedel, S., Riedel, S.: Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* **33**, 9459–9474 (2020)
8. Research, M.: Graphrag: Unlocking llm discovery on narrative private data (February 2024), <https://www.microsoft.com/en-us/research/blog/graphrag-unlocking-llm-discovery-on-narrative-private-data/>, microsoft Research Blog

Combining Traditional Machine Learning and Generative AI to Create a Trustworthy Overview of Austrian Online Shops

Simone Sandler^{1,2}[0000-0001-8919-7804] and Oliver Krauss^{1,2,3}[0000-0002-8136-2606]

¹ University of Applied Sciences Upper Austria, Campus Hagenberg, Softwarepark 11, 4232 Hagenberg, Austria
 {simone.sandler, oliver.krauss}@fh-hagenberg.at
² Advanced Information Systems and Technology
³ Digital Media Department

Introduction The growth of e-commerce has introduced new challenges for regulatory authorities, particularly in ensuring consumer safety and market transparency [6, 3]. In Austria, the creation of small, often unregistered online shops complicates surveillance efforts in sensitive sectors such as food, cosmetics, and health products. The *eMarketShield* project addresses this issue by combining traditional machine learning with generative AI to build a scalable, trustworthy, and digitally sovereign overview of the Austrian online retail landscape. Developed through a collaboration of public authorities, academia, and industry, the system supports risk-based market surveillance via automated data collection, intelligent classification, and ethically aligned AI deployment.

Methodology A modular, scalable system was developed to collect URLs of potential online shops. Initially, only the start page of each candidate site is crawled. This page is then processed using two classifiers: the first detects dummy pages (e.g., placeholders, parked domains, or bot traps), while the second determines whether the URL corresponds to an actual online shop. The dummy page classifier, based on an SGDClassifier, achieved an accuracy of 0.93. The shop detection classifier, implemented using an MLPClassifier, reached an accuracy of 0.88. These models significantly reduce the number of irrelevant websites, thereby improving processing efficiency and resource allocation. Large Language Models (LLMs) are employed to extract structured data from heterogeneous HTML content, enabling the creation of a Common Data Model (CDM) for consistent representation of shop and product information. To ensure digital sovereignty, the project prioritizes locally hosted models over cloud-based solutions. A prototype web interface was implemented to allow stakeholders to interact with the data, perform filtered searches, and collaboratively investigate cases.

Ethical and Legal Evaluation The project integrates a comprehensive *ethics-by-design* and *legal-by-design* framework to ensure responsible AI development and deployment. Ethical and legal risks—such as bias, privacy violations, lack of

transparency, and accountability gaps—are continuously assessed using established instruments including the ALTAI checklist [5], the EKIV framework [2], and the EU AI Act [1]. Special attention is given to digital sovereignty, particularly in the context of data governance and infrastructure, addressing concerns raised in recent studies on sovereign AI ecosystems [4].

Conclusion and Outlook Initial results demonstrate the feasibility of combining traditional machine learning and generative AI to automate and enhance market surveillance. The system successfully identifies previously unknown online shops and enables structured analysis of their offerings. The classifiers and extraction methods show promising performance, and the web interface facilitates stakeholder interaction and collaborative research. The project underscores the importance of digital sovereignty and ethical governance in AI-driven public sector applications. By focusing on small-scale retailers and leveraging cutting-edge AI technologies, *eMarketShield* contributes to a more resilient and trustworthy digital economy in Austria.

References

1. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act). Official Journal of the European Union, L series (July 2024), <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng>
2. Biegelbauer, P., Schlarb, S., Weinlinger, P., Conde-Jimenez, R., Scheichenbauer, H., other: Leitfaden Digitale Verwaltung: KI, Ethik und Recht. Praxisleitfaden für die Verwaltung, Version 2.0. Tech. rep., Bundesministerium für Kunst, Kultur, öffentlichen Dienst und Sport (BMKÖS), Wien, Österreich (December 2024), https://oeffentlicherdienst.gv.at/wp-content/uploads/2024/09/250113_Leitfaden-Digitale-Verwaltung_2.0_A4.pdf, ISBN 978-3-903097-66-7
3. Commission, E.: E-commerce communication: A comprehensive EU toolbox for safe and sustainable e-commerce. Online (February 2025), <https://digital-strategy.ec.europa.eu/en/library/e-commerce-communication-comprehensive-eu-toolbox-safe-and-sustainable-e-commerce>
4. Kraley, A., Kantor, I., Gutiérrez, R.: Sovereign AI Ecosystems: Navigating Global AI Infrastructure & Data Governance. Tech. rep., POLITICO Research & Analysis Division (September 2024), <https://www.politico.eu/wp-content/uploads/2024/09/15/Sovereign-AI-Ecosystems.pdf>, independent research report presented by G42
5. Radclyffe, C., Ribeiro, M., Wortham, R.H.: The assessment list for trustworthy artificial intelligence: A review and recommendations. *Frontiers in Artificial Intelligence* **6**, 1020592 (2023). <https://doi.org/10.3389/frai.2023.1020592>, <https://doi.org/10.3389/frai.2023.1020592>
6. Verbraucherzentrale: Erschreckende Bilanz bei EU-weiten Kontrollen des Online-Handels (February 2025), <https://www.verbraucherzentrale.de/wissen/lebensmittel/nahrungsergaenzungsmittel/erschreckende-bilanz-bei-euweiten-kontrollen-des-onlinehandels-28499>, accessed: 2025-11-04

Navigating the Sea of LLM Evaluation: Investigating Bias in Toxicity Benchmarks

Regina Gugg¹[0009-0003-5948-1048], Selina Niederländer¹[0009-0007-8050-651X],
Andreas Stöckl²[0000-0003-1646-0514], and Martin Flechl¹[0000-0003-4750-3548]

¹ Dynatrace Research, Linz, Austria

² University of Applied Sciences Upper Austria, Hagenberg, Austria

Problem and Motivation Large language models (LLMs) represent a rapidly evolving area of artificial intelligence research. Ensuring their safe deployment requires benchmarks that assess not only task performance but also their propensity to generate harmful output. However, the definition of harm is inherently complex, and benchmark validity degrades over time due to evolving societal norms or data leakage. Furthermore, the proliferation of evaluation frameworks complicates the consolidation of findings. While benchmark scores often guide model selection, safety and toxicity assessments remain underrepresented, and systematic comparisons across safety benchmarks are limited, hindering the evaluation of models beyond raw performance metrics. A central question is whether safety benchmarks designed to measure the propensity for harmful outputs in LLMs remain reliable across diverse task- and domain-specific contexts encountered in real-world applications. This work explores methods to address these challenges, focusing on improving the reliability and applicability of safety benchmarks.

Method We assess benchmark reliability through controlled augmentation across two dimensions: task formulation and domain context. We employ a comparative evaluation framework examining (1) baseline outcomes under original benchmark conditions, (2) task-shifted variants where prompts are reframed as summarization tasks, and (3) domain-shifted variants, where original benchmark prompts are adapted to specialized contexts. We selected summarization because, despite its widespread practical use, it remains absent from current safety benchmark taxonomies, enabling us to evaluate whether safety alignment transfers beyond initial contexts. For domain adaptation, we selected four diverse domains: Chemical Engineering, Finance, Social Media, and Sports. Our goal was to address gaps in benchmarks by concentrating on harm recognition in high-risk instructions, malicious or fraudulent advice, and toxic intent within informal or domain-specific discourse.

Experimental Setup We evaluated a mix of proprietary and open-weight language models: *Llama2-7B-Chat*, *GPT-3.5-Turbo-0125*, *Mistral-7B-Instruct*, *Qwen2-7B-Instruct*, and *DeepSeek-LLM-7B-Chat*. Our benchmark suite comprises *Do-not-answer* [4], *Realtoxicityprompts* [1], *Toxigen* [2], and *Harmbench* [3],

evaluated using their native classifiers (*Longformer*, *Perspective API*, *Toxigen-RoBERTa*, and *Llama2-13B-cls*). We implemented task-shifted variants using two instruction types: neutral summarization prompts and safety-reinforcing prompts based on the *Llama2* system-level safety instructions. We generated domain adaptations via *GPT-4o* while preserving the original intent and toxicity levels through domain-specific vocabulary and contexts. We assessed domain-shift quality using an *LLM-as-a-judge* approach.

Statistical Analysis This design enables cross-evaluation of agreement, divergence, and ranking stability across benchmarks and conditions. To examine task-shift effects, we applied the *McNemar* test combined with random effects meta-analysis. To assess domain-related differences, we used chi-squared tests with *Cramér's V* and mixed-effects logistic regression. To evaluate inter-benchmark agreement, we employed *Cohen's* and *Fleiss'* kappas for classifier consistency and *Spearman's* rank correlation for model ranking stability.

Results Both task and domain shifts significantly affect harm detection rates and model behavior across benchmarks. Under task reformulation to summarization, three of four benchmarks detect increased harm. *Harmbench* is a notable exception, showing a decrease in harm detection rates. Domain shifts reveal substantial context-dependent variability in harm detection. Specialized vocabulary can both obscure and amplify harmful intent depending on context, indicating that benchmark robustness varies substantially across domains. Cross-benchmark comparisons reveal weak agreement and divergent harm definitions, suggesting that benchmarks capture narrow conceptualizations of harm rather than providing generalizable safety assessments. This work provides the first systematic assessment of safety benchmark stability across task and domain variations, revealing critical limitations in current evaluation practices.

References

1. Gehman, S., Gururangan, S., Sap, M., Choi, Y., Smith, N.A.: Realtoxicityprompts: Evaluating neural toxic degeneration in language models (2020). <https://doi.org/10.48550/arXiv.2009.11462>
2. Hartvigsen, T., Gabriel, S., Palangi, H., Sap, M., Ray, D., Kamar, E.: Toxigen: A large-scale machine-generated dataset for adversarial and implicit hate speech detection (2022). <https://doi.org/10.48550/arXiv.2203.09509>
3. Mazeika, M., Phan, L., Yin, X., Zou, A., Wang, Z., Mu, N., et al.: Harmbench: A standardized evaluation framework for automated red teaming and robust refusal (2024). <https://doi.org/10.48550/arXiv.2402.04249>
4. Wang, Y., Li, H., Han, X., Nakov, P., Baldwin, T.: Do-not-answer: A dataset for evaluating safeguards in llms (2023). <https://doi.org/10.48550/arXiv.2308.13387>

Explaining GenAI Endurance Training Plans Through Post-Hoc Behavioral Evaluation

Lukas Voggeneder, Jens Krösche, Marc Kurz^[0000-0002-2437-0589], Erik Sonnleitner^[0000-0002-0646-782X], and Stephan Selinger^[0009-0003-3721-5371]

University of Applied Sciences Upper Austria, Campus Hagenberg
Softwarepark 11, 4232 Hagenberg, Austria
stephan.selinger@fh-hagenberg.at

Introduction and Background Large language models (LLMs) such as ChatGPT are increasingly used for everyday decision-making, including the generation of personalized endurance training plans [1, 2]. However, their black-box nature raises concerns about trustworthiness in personalized training and health applications. Recent studies assessed ChatGPT-generated training plans through expert coach reviews [1, 2], providing human-centered insights, but relying on subjective ratings that do not scale.

To complement this research, we conduct a functionally grounded, quantitative, post-hoc analysis, in line with the work of Doshi-Velez and Kim [4], of GenAI-generated endurance training plans. As PSO-generated training plans have proven effective in minimizing overstrain risk and even outperform standard programs [3], we use PSO as a quantitative baseline. Instead of expert ratings, we evaluate LLM-generated 12-week half-marathon plans by comparing their modeled performance, computed via the Fitness-Fatigue Model [5], against this PSO reference.

Methodology and Experimental Design We evaluated ChatGPT, a domain-enriched variant (*EnduranceGPT*), and a PSO baseline. EnduranceGPT was fine-tuned on approximately ten peer-reviewed papers on endurance-training theory and exercise physiology to embed domain knowledge. LLM-training plans were generated using GPT-40 (May 2025). Twenty fictional personas spanning age, gender, baseline fitness, and heart-rate reserve were used. Each model generated individualized 12-week half-marathon programs from which daily training load in TRIMP units was derived. Training plans were analyzed using the Fitness-Fatigue Model of Morton et al. [5] which estimates performance from the balance between accumulated fitness (long-term adaptation) and fatigue (short-term load). We deliberately used the classical two-component Fitness–Fatigue Model because of its simplicity and transparent structure. This allowed quantitative comparison of end-of-plan performance and structural features such as progression, rest distribution, and tapering phases, without introducing additional modeling complexity.

Results and Discussion Across all personas, the PSO baseline achieved the highest modeled end-of-plan performance (dimensionless index 116.4 ± 7.8). EnduranceGPT followed (110.7 ± 8.3), while ChatGPT reached 98.2 ± 6.9 . When

stratified by baseline fitness, PSO achieved the highest modeled performance for low-fitness personas, followed by EnduranceGPT and, at a distance, ChatGPT. Among recreational athletes, EnduranceGPT approached PSO, while ChatGPT still lagged. For high-fitness personas, EnduranceGPT slightly exceeded PSO on average performance.

ChatGPT produced understandable but stagnant plans (progression < 1 % per week) with inconsistent rest and no tapering. EnduranceGPT created structured four-week cycles with realistic two-week tapers and moderate progression (\approx 5 % per week) but occasionally ramp rates exceeded 15 %. PSO delivered the strongest modeled results yet ignored rest, reflecting optimization bias rather than physiologically meaningful reasoning. The comparison highlights complementary strengths: PSO maximizes modeled performance, while domain-aware generation enhances interpretability and structural consistency.

Future work could integrate more detailed models, such as the three-dimensional impulse-response framework by Kontro et al. [6], which captures energy-system-specific adaptations and may provide a more physiologically realistic evaluation. In addition, exploring combinations of optimization-based and generative approaches may help integrate these advantages and advance explainable, trustworthy GenAI training systems.

References

1. P. Düking, C. Hollfelder, J. Ferrauti, and B. Schulze: ChatGPT generated training plans for runners are not rated optimal by coaching experts, but increase in quality with additional input information. *Frontiers in Sports and Active Living*, vol. 6, 2024. <https://doi.org/10.52082/jssm.2024.56>.
2. F. Wachholz, S. Manno, D. Schlachter, N. Gamper, and M. Schnitzer: Acceptance and trust in AI-generated exercise plans among recreational athletes and quality evaluation by experienced coaches: a pilot study. *BMC Research Notes*, vol. 18, 2025. <https://doi.org/10.1186/s13104-025-07172-9>.
3. N. Kumyaito, S. Wakasa, T. Sasai, and T. Hirata, Planning a sports training program using adaptive particle swarm optimization with ε -constraint methods. *BMC Research Notes*, vol. 11, no. 1, 2018. <https://doi.org/10.1186/s13104-017-3120-9>.
4. F. Doshi-Velez and B. Kim: Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017. <https://doi.org/10.48550/arXiv.1702.08608>
5. R. H. Morton, J. R. Fitz-Clarke, and E. W. Banister: Modeling human performance in running. *Journal of Applied Physiology*, vol. 69, no. 3, pp. 1171–1177, 1990. <https://doi.org/10.1152/jappl.1990.69.3.1171>.
6. H. Kontro, A. Mastracci, S. S. Cheung, and M. J. MacInnis: The three-dimensional impulse-response model: Modeling the training process in accordance with energy system-specific adaptation. *arXiv:2503.14841*, 2025. <https://doi.org/10.48550/arXiv.2503.14841>

xML CardioTwin: A Multimodal, Explainable Cardiac Digital-Twin Pipeline for Personalised Atrial Fibrillation Ablation and Cardioneuroablation

Marek Kulbacki^{4,5[0000-0003-4609-106X]}, Jakub Segen^{4,5[0000-0002-8697-6025]},
 Zenon Chaczko^{2,5[0000-0002-2816-7510]}, Jan Nikodem^{1[0000-0002-3748-9255]},
 Ryszard Klempous^{1[0000-0003-4889-0930]}, Ewa Szlachcic^{1[0000-0002-4179-4068]},
 Jerzy Rozenblit^{3[0000-0002-7348-4128]}, Przemysław Skoczyński^{1,6[0000-0002-4998-8879]}, and Dariusz Jagielski^{1,6[0000-0002-3066-4030]}

¹ Wrocław University of Science and Technology, Wrocław, Poland

² School of Elect. and Data Engineering University of Technology Sydney, Australia

³ Depart. of Electrical and Computer Engineering University of Arizona Tucson, USA

⁴ Polish-Japanese Academy of Information Technology, Warsaw, Poland

⁵ DIVE IN AI, Wrocław, Poland

⁶ Center for Heart Diseases, 4th Military Hospital, Wrocław, Poland

Extended Abstract

Current machine learning pipelines for atrial fibrillation (AF) either predict recurrence from clinical and ECG data [1–3] or simulate substrate-level ablation strategies [5], but they do not provide a unified, dynamically updated, patient-specific decision layer that cardiologists can rely on at the point of care. Multimodal models that integrate 12-lead ECG, time series, and clinical variables achieve high performance for AF recurrence prediction in paroxysmal and persistent AF cohorts [2, 3], and explainable ML (xML) applied to LGE-MRI and biophysical simulations can recommend substrate-tailored ablation patterns and quantify the mechanistic role of fibrosis in ventricular substrates and, by extension, atrial substrates in more recent virtual heart work [5], yet both families remain essentially static, pre-procedural tools validated on limited, single-center hospital populations. Computational electrophysiology and virtual-heart digital twins show that forward simulations can discriminate ablation strategies and identify high-risk substrate [5], but with prohibitive runtime and minimal evidence of benefit in routine workflows. Recent ECG foundation models trained on more than ten million short, hospital-grade 12-lead recordings demonstrate strong transfer across multiple diagnostic tasks [4], but they rely on weak noisy labels, exhibit marked temporal and domain shift within a single health system, lack robust external validation, and do not address long-term home rhythm monitoring or autonomic phenotypes relevant for cardioneuroablation (CNA), so they cannot be treated as ready-made engines for AF or CNA follow up. We present xML CardioTwin, a rigorously constrained multimodal explainable cardiac digital twin designed for two tightly coupled use cases: personalization of

AF ablation strategy (PVI, substrate-guided, hybrid) and CNA planning with longitudinal follow up in patients with vagally mediated arrhythmias and conduction disturbances [1, 2, 6]. xML CardioTwin will fuse three data streams into a single patient timeline: long term ECG (Holter, patches, wearables), structural and scar imaging (LGE-MRI, atrial geometry, atrial volume indices), and structured peri-procedural and follow up data (lesion sets, neuroablation sites, biomarkers, symptoms, syncope history). ECG encoders will be initialized from an ECG foundation model [4] then systematically retrained and calibrated on a prospective multicenter registry with explicit temporal and site-based splits to prevent overclaiming of generalization from single-system pretraining. Substrate features derived from LGE-MRI and virtual-heart simulations will be compressed into xML feature sets with SHAP or related attribution mechanisms explicitly linked to known pathophysiology [3, 5], so that the system can justify recommendations such as fibrosis-driven substrate ablation plus targeted CNA instead of emitting opaque class labels. The digital twin will operate as a dynamic state estimator: after the index ablation and or CNA, each new block of follow up ECG and clinical data will update individual risk of AF recurrence, bradyarrhythmic events, or syncope, adjust recommended monitoring intensity, and trigger scenario forecasts such as the expected benefit of early re ablation versus intensified pharmacologic or neuromodulatory therapy [2, 3]. Methodologically, the pipeline is constrained by three design principles derived from current gaps in the literature: strict separation of development, temporal, and external validation cohorts across vendors and health systems, explicit modeling of label noise and domain shift in weakly labeled ECG tasks [4], and mandatory xML outputs at the level of signal segments, imaging-derived substrate, and clinical context suitable for regulatory and trial use [3]. The project will deliver a validated clinician-facing prototype that does not claim to be a universal ECG-for-everyone solution, but a targeted, evidence-grounded xML CardioTwin for high-stakes AF and CNA decision making with a clear path to prospective clinical evaluation and integration into hybrid electrophysiology workflows [1, 6].

References

1. Bifulco, S., et al.: Predicting arrhythmia recurrence post-ablation in atrial fibrillation using explainable machine learning. *Commun Med* **5** (2025)
2. Chen, L., et al.: A Deep Learning-Based Multimodal Fusion Model for Recurrence Prediction in Persistent Atrial Fibrillation Patients. *J Cardiovasc Electrophysiol* **36**(8), 1785–1797 (2025)
3. Jud, F., et al.: A novel score in the prediction of rhythm outcome after catheter ablation of atrial fibrillation: the SUCCESS score. *Anatol J Cardiol* **21**(3) (2019)
4. McKeen, K., et al.: ECG-FM: An open electrocardiogram foundation model. *JAMIA open* **8**(5), ooaf122 (2025)
5. Prakosa, A., et al.: Personalized virtual-heart technology for guiding the ablation of infarct-related ventricular tachycardia. *Nat Biomed Eng* **2**, 732–740 (2018)
6. Skoczyński, P., et al.: Cardioneuroablation in the Management of Vagally Mediated Bradyarrhythmias – A Comprehensive Review of Ongoing Randomized Controlled Trials. *Journal of Clinical Medicine* **14**(2) (2025)

Immersive Tourniquet Training (Extended Abstract)

Klaus Buchenrieder and Uwe M. Borghoff

Department of Computer Science
University of the Bundeswehr Munich, Neubiberg, Germany
{klaus.buchenrieder,uwe.borghoff}@unibw.de

Tourniquets are widely used to temporarily occlude or reduce the blood flow in a limb in order to control life-threatening hemorrhage. In military first-responder training, correct placement on the thigh or upper arm must be practiced under highly realistic and stressful conditions. Improper practice carries significant risks. This work introduces an instrumented training system designed to enable safe practice while simultaneously providing quantitative measurements and VR-based feedback. The contribution of our work lies in three areas: first, we define training requirements for safe and objective tourniquet practice, second, we present an instrumented thigh model with an embedded force sensor at a defined location, and third, we integrate optical tracking to allow pose-aware VR training within a serious gaming environment. Recent clinical research highlights VR-based training for hemorrhage control [1, 5]. Insights from adjacent domains point toward possible solutions. In the field of prosthetics, Borghoff and Buchenrieder surveyed the state of the art in myoelectric signal processing and proposed extending these approaches beyond prosthetic control to gaming, robotics, and simulation [2]. Beyond academic studies, patented technologies also contribute to the development of intelligent occlusion and tourniquet systems [4, 6]. Figure 1 shows the thigh model, consisting of a rigid, cylindrical shell that simulates the geometry of the human upper leg and provides a reproducible application site. A force sensor embedded at this site captures the contact pressure exerted by the tourniquet band. By localizing the sensing point, repeated trials under different conditions become comparable, enabling the system to distinguish between effective and ineffective applications without subjecting trainees to the risks of live practice. The instrumented tourniquet complements this setup. Its buckle houses miniature load-cells to measure




Fig. 1: Instrumented thigh with sensor, instrumented tourniquet with load-cells, pose markers, data links, and VR engine integration.

strap tension. Both the thigh sensor and the buckle load-cells are connected to a compact data acquisition unit, which digitizes the analog signals and makes them available for immediate visualization and storage. Initial experiments confirm that the system provides valid signals. The availability of quantitative thresholds reduces the risk of over-tightening during practice. Future trials and systematic calibration against physiological endpoints such as Doppler pulse detection and tissue oxygenation remain necessary. The instrumented tourniquet and thigh model can be used as independent training devices, connected to a laptop or tablet that displays pressure and tension curves. Embedded micro-controllers and displays could support self-contained practice by presenting numeric values, bar graphs, or simple traffic-light indicators. When integrated into VR, the system enables objective scoring of trainee performance under realistic conditions. Clinical interest has also been expressed regarding possible operating-room applications. Recent research has emphasized the role of adaptive gameplay in serious games as a means to sustain learner engagement and improve training effectiveness. Dobrovsky *et al.* argue that interactive deep reinforcement learning (iDRL) can dynamically adapt scenarios by combining human feedback with machine learning, thereby overcoming the rigidity of scripted interactions and tailoring challenges to individual performance and attention [3].

Future work will further personalize the training process, ensuring that learners remain engaged, challenged, and supported in developing their skills under conditions that closely resemble operational reality.[†]

References

1. Arif, A., et al.: The impact of the addition of a virtual reality trainer on skill retention of tourniquet application for hemorrhage control among emergency medical technician students: a pilot study. *Cureus* **15**(1) (2023)
2. Borghoff, U.M., Buchenrieder, K.: Myo-to-gesture: Evolving AI-based innovations in myoelectric signal processing – state-of-the-art survey and research proposal. In: Artificial Intelligence and Soft Computing – 23rd International Conference, ICAISC 2024, Zakopane, Poland. pp. 241–252. LNAI 15166, Springer (2024). https://doi.org/10.1007/978-3-031-81596-6_22
3. Dobrovsky, A., Borghoff, U.M., Hofmann, M.: Improving adaptive gameplay in serious games through interactive deep reinforcement learning. In: Cognitive Infocommunications, Theory and Applications, pp. 411–432. Topics in Intelligent Engineering and Informatics **13** (2019). https://doi.org/10.1007/978-3-319-95996-2_19
4. Gavriely, N., Gavriely, O.: Method and an apparatus for adjusting blood circulation in a limb (Oct 2008), United States Patent, US 2008/0262535 A1
5. Saggar, V., et al.: Effectiveness of a virtual reality trainer for retention of tourniquet application skills for hemorrhage control among emergency medicine residents. *AEM Education and Training* **8**(3), e10986 (2024). <https://doi.org/10.1002/aet2.10986>
6. Scardulla, F.: Wearable blood pressure measurement and antihemorrhagic device (Sep 2021), United States Patent, US 2021/0290089 A1

[†] This research is part of the project SHL, funded by dtec.bw—Digitalization and Technology Research Center of the Bundeswehr. dtec.bw is funded by the European Union—NextGenerationEU.

A Holter Smart Patch for Patients with Transient Cardiac Events

Mario Silaci and Klaus Buchenrieder

Department of Computer Science
 University of the Bundeswehr Munich, Neubiberg, Germany
{mario.silaci,klaus.buchenrieder}@unibw.de

Continuous electrocardiographic (ECG) monitoring plays a central role in the detection and treatment of cardiac arrhythmias and transient ischemic events, but conventional Holter monitors remain bulky, wired, and limited to retrospective data analysis long after the recording phase. These limitations are to be overcome by a Holter Smart Patch, a miniaturized, autonomous, and LoRa-connected multisensor system [1] that combines long-term ECG recording with intelligent device-internal analysis and wireless signal transmission. The current Smart Patch prototype integrates a low-cost, single-chip AD8232 ECG module, a Heltec ESP32 microcontroller with LoRa SX1262 module, and a microSD data logger Velcro-fastened to a flexible adhesive plaster. This significantly enhances wearability and patient comfort. The measurements are used to determine the time difference between heartbeats, and by averaging the calculated time differences, we derive the baseline heart rate and count deviations in buckets with 50-hertz value ranges. At regular intervals, the LoRa module sends a status update to the attending physician or nursing staff. This information can also be requested proactively at any time. It is also possible for a medicare team to analyze longer recording periods in an emergency by reading the flash memory of the SD-card component (or another interface like Bluetooth). Energy-efficient event transmission between the patient and the medics is carried out using the LPWAN technology LoRaWAN[3]. LoRaWAN is based on energy-efficient and robust Chirp Spread Spectrum-based modulation, allowing for long-range data transmission in the ISM frequency band. The WAN is organized as a star-of-stars topology with a three-part server infrastructure as the backend. The components




Fig. 1. Visual feedback scheme for the user interface.

of the infrastructure are sensor-equipped end-devices and LoRaWAN gateways, which act as data sinks and forward the transmitted data to the backend servers via an IP stack. We employ TheThingsStack (TTS), that provides a publicly accessible LoRaWAN-compatible server infrastructure. Via the MQTT interface of TTS, our InfluxDB receives incoming data by the corresponding Telegraf service. A Grafana-based user interface then displays the current status of the monitored patients. Figure 1 shows an envisioned schematic visual representation highlighting the heartbeat history and a histogram of the heartbeat frequency deviation to determine the degree of irregularities. In addition, an overview of the four heartbeat states and a compact status display provide further insights. In critical situations, nursing staff could be notified visually via a status display or an SMS message.

In future developments, larger amounts of data could be transmitted via Bluetooth to the patient's mobile phone and sent on to the attending physician. We also consider incorporating a photoplethysmography sensor to estimate cuffless blood pressure via pulse transit time analysis [2]. Local data processing of synchronized ECG and PPG pairs would then be performed using a hybrid CNN support vector regressor model. With an additional embedded inertial measurement unit, one could also detect falls or syncope. Adding a sensor to continuously monitor body temperature would further differentiate the patch from a passive recording device. With these additions, a networked monitoring system can then autonomously interpret physiological signals and transmit clinically relevant information from patches of different patients operating within a shared LoRaWAN infrastructure to improve coverage and resilience. Future developments also target edge-based AI for arrhythmia classification, adaptive energy management, and hybrid LoRa cellular communication, paving the way for a new generation of smart biosensors combining cardiac diagnostics, hemodynamic assessment, and telemedicine in a single energy-efficient wearable platform.

Acknowledgment

This research is part of the project ROLORAN, funded by dtec.bw - Digitalization and Technology Research Center of the Bundeswehr. dtec.bw is funded by the European Union - NextGenerationEU.

References

1. Farej, Z., Adel, A.: Review on lora communication technology, its issues, challenges and applications in healthcare systems. European Journal of Computer Science and Information Technology **12**(8) (2024). <https://doi.org/https://doi.org/10.37745/ejcsit.2013/vol12n8117>, <https://www.eajournals.org/>, open Access
2. Solmaz Rastegar, H.H., Lowe, A.: Hybrid cnn-svr blood pressure estimation model using ecg and ppg signals. sensors **23**(1259) (2023). <https://doi.org/https://doi.org/10.3390/s23031259>, <https://www.mdpi.com/journal/sensors>, open Access
3. Suliman Abdulmalek, A.N., Jabbar, W.: Lorawan-based hybrid internet of wearable things system implementation for smart healthcare. Internet of Things (25) (2024). <https://doi.org/10.1016/j.iot.2024.101124>, <https://www.sciencedirect.com/science/article/pii/S2542660524000660>, open Access

FLOD: a Framework for Explainable Federated Learning Object Detection

Francesco Mercaldo^{1,*}, Ludovica Ciardiello¹, Marta Petyx², Patrizia Agnello²,
Mario Cesarelli³, Fabio Martinelli⁴, and Antonella Santone¹

¹ Department of Medicine and Health Sciences “Vincenzo Tiberio”, University of Molise, Campobasso, Italy

² Istituto Nazionale per l’Assicurazione contro gli Infortuni sul Lavoro (INAIL), Roma, Italy

³ Department of Engineering, University of Sannio, Benevento, Italy

⁴ Institute of High Performance Computing and Networking, National Research Council of Italy (CNR), Rende, Italy

[francesco.mercaldo](mailto:francesco.mercaldo@unimol.it), [ludovica.ciardiello](mailto:ludovica.ciardiello@unimol.it), antonella.santone@unimol.it,
[m.petyx](mailto:m.petyx@inail.it), p.agnello@inail.it, m.cesarelli@unisannio.it,
fabio.martinelli@icar.cnr.it

Object detection plays a key role in medical image analysis, enabling automated localization of clinically relevant features such as lesions, tumors, or anatomical structures. Federated learning facilitates decentralized model training across multiple medical institutions, preserving data privacy and addressing regulatory constraints [1]. However, explainability remains essential to ensure clinical trust and interpretability of AI-driven diagnoses. Explainable AI techniques, such as Class Activation Mapping (CAM), provide visual heatmaps that highlight the image regions influencing model predictions, supporting clinicians in decision-making.

Starting from these considerations, we propose *FLOD*: a framework for explainable *Federated Learning Object Detection*, integrating privacy-preserving federated training with interpretable prediction visualization. The framework employs aggregation algorithms (e.g., FedAvg, FedProx) and CAM-based explainability to produce transparent and clinically meaningful outputs.

Figure 1 shows the *FLOD* workflow.

The workflow begins at the start node (yellow oval) and proceeds through dataset preparation and model initialization, as shown in Figure 1. The green blocks in Figure 1 represent the data sources: the Object Detection Dataset is split into three subsets i.e., training, validation and testing. Each client (e.g., hospital or diagnostic center) retains its local medical imaging data, ensuring privacy and compliance, thereby preserving privacy and complying with regulatory and ethical constraints. The orange blocks in Figure 1 denote model parameters. The architecture scale specifies the model architecture and version. Hyperparameters define optimization settings such as aggregation method, epochs, communication rounds. The blue blocks in Figure 1 represent processes and outputs. Each client trains locally and shares updated weights with the federated server, which aggregates and redistributes them until convergence, after which the workflow proceeds to model testing. Object detection outputs include (i) bounding box coordinates, (ii) objectness scores, and (iii) class probabilities. Explainability




Fig. 1. The *FLOD* workflow.

is achieved through a module that generates heatmaps for visual interpretation. The workflow concludes at the end node (yellow oval in Figure 1), completing the pipeline from dataset preparation to federated model training, testing, and explainable decision-making.

Acknowledgment

This work has been partially supported by EU DUCA, EU CyberSecPro, SYNAPSE, PTR 22-24 P2.01 (Cybersecurity) and SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the EU - NextGenerationEU projects and ALOHA: a framework for monitoring the physical and psychological health status of the Worker through Object detection and federated machine learning, Call for Collaborative Research BRiC -2024, INAIL.

References

1. H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y. Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017.

Primary and Adjacent Zooming Areas in 3D Medical Image Analysis

Miklós Vincze¹, Abdallah Benhamida², Sándor Burian², Miklós Kozlovszky³

¹ AR/VR/XR Laboratory, John von Neumann Faculty of Informatics Óbuda University Budapest, Hungary

² BioTech Research Center, EKIK Óbuda University Budapest, Hungary

³ John von Neumann Faculty of Informatics Óbuda University Budapest, Hungary

Making medical diagnoses is a complex process. One of its main elements is the examination of digitized pathological samples, during which the physician reviews the tissue sample in question on a monitor. In our research, we introduced new concepts into the 3D evaluation of digitized pathological samples, which help us gain a clearer picture of the tissue sample examination process and improve the quality assurance of the evaluation. In this paper, we carried out new tests on the application of primary and adjacent zooming areas. In addition, we introduced our equation, which describes the mathematical relationship between the different zooming areas and the overall sample evaluation level.

Keywords: 3D visualization, primary zooming area, adjacent zooming area

1 Introduction

Today, there is still a lot of research being done in the field of improving quality assurance in medical processes, and digital pathology is no exception. There are numerous studies regarding the inspection of the preparation process of digitized pathological samples. During our previous research, we designed and created a system capable of displaying these samples at their native resolution in 3D. We have designed a magnification function into our system, which allows the user to obtain a 3D overview of a section of the entire pathological serial section at native resolution. In addition to the magnification function, we continuously monitor its use while the software is in operation, thus obtaining an accurate picture of which parts of the series section the physician has used the magnification function on. Then, based on the frequency of use of the zooming function, we can draw conclusions about the level of detail of the sample examination and the coverage of the entire series section examination.

2 Results

During our research, we designed and developed a system that monitors the use of the zooming function while using our 3D tissue sample display system. The presented solution uses our equation to calculate the proportion of the entire sample that has been

examined. The 3D visualization solution developed during our previous research uses tiles to build up the digitized pathological serial section, and we use the spatial positions of these tiles to classify them during zooming. In order to obtain a more detailed picture of the areas examined under magnification, we introduced the following concepts:

- Primary zooming area: The tiles on which the user initiated the zoom.
- Adjacent zooming area: Adjacent tiles surrounding the primary magnification areas.

Equation (1) can be used to determine the examination quality of a pathological serial section:

$$E_q = \frac{P_{ZAN} + \frac{1}{2}A_{ZAN}}{T_N} \quad (1)$$

In equation (1) the value of A_{ZAN} can be determined in inclusive or exclusive way. In the case of exclusive calculation, after the user has zoomed in on a tile previously marked as an adjacent zooming area, the tile in question is moved to the primary zooming area group. In the case of inclusive calculation, the tile in question is kept in both groups.

Fig. 1 shows a tissue sample, which we visualized in 3D during our tests, as well as the results of a tile classification in terms of primary and adjacent zooming areas.




Fig. 1:(A) 3D image of a part of the pathological tissue sample used in testing. (B) Illustration of primary and secondary zooming areas. The green tiles show the primary zooming areas, while the red tiles show the adjacent zooming areas. The white tiles show where the user did not perform any zooming.

Acknowledgments. We would like to thank AIAM (Applied Informatics and Applied Mathematics) doctoral school of Obuda University, Budapest, Hungary for their support in this research. The authors hereby thank the 2019-1.3.1-KK-2019-00007 project „Innovációs szolgáltató bázis létrehozása diagnosztikai, terápiás és kutatási célú kiberorvosi rendszerek fejlesztésére” project for the financial support.

References

1. Singh, R.; Yadav, S. K.; Kapoor, N. Analysis of Application of Digital Image Analysis in Histopathology Quality Control. *Journal of Pathology Informatics* 2023, 14, 100322. <https://doi.org/10.1016/j.jpi.2023.100322>.
2. Janowczyk, A.; Zuo, R.; Gilmore, H.; Feldman, M.; Madabhushi, A. HistoQC: An Open-Source Quality Control Tool for Digital Pathology Slides. *JCO Clin Cancer Inform* 2019, 3, CCI.18.00157. <https://doi.org/10.1200/CCI.18.00157>.
3. Ahmad, Z.; Alzubaidi, M.; Al-Thelaya, K.; Calí, C.; Boughorbel, S.; Schneider, J.; Agus, M. Advancing Open-Source Visual Analytics in Digital Pathology: A Systematic Review of Tools, Trends, and Clinical Applications. *Journal of Pathology Informatics* 2025, 18, 100454. <https://doi.org/10.1016/j.jpi.2025.100454>.

Rosetta Stone for Immune Toxicity Intelligence: Explainable Machine Learning Across Immune Checkpoint Inhibitors, Radionuclide Therapies and Antibody Drug Conjugates in Real-World Oncology

Marek Kulbacki^{1,2[0000-0003-4609-106X]}, Zenon Chaczko^{2,3[0000-0002-2816-7510]}, Jakub Segen^{1,2[0000-0002-8697-6025]}, Adam Maciejczyk^{4,5[0000-0002-7047-0433]}, Ireneusz Pawlak^{4,6[0000-0002-8778-4754]}, Paweł Zawadzki^{4,6[0009-0008-0031-7952]}, and Bożena Cybulska-Stopa^{4,6[000-0002-0975-850X]}

¹ Polish Japanese Academy of Information Technology, Warsaw, Poland

² DIVE IN AI, Wrocław, Poland

³ University of Technology Sydney, Australia

⁴ Lower Silesian Oncology, Pulmonology and Hematology Center, Wrocław, Poland

⁵ Faculty of Medicine, Wrocław Medical University, Poland

⁶ Faculty of Medicine, Wrocław University of Science and Technology, Poland

Extended Abstract

Immune related adverse events represent a common toxicity pattern across immune checkpoint inhibitors, radionuclide based therapies, and antibody drug conjugates [8]. Immune checkpoint inhibitor regimens are characterized predominantly by cutaneous and endocrine manifestations [4]. Radionuclide therapies produce haematologic damage and nephrotoxicity as primary complications [7]. Antibody drug conjugates are associated with interstitial lung disease and neuropathy, particularly in previously treated patients [3, 6]. Current clinical practice involves assessment of these toxicities within isolated specialty services, without unified, evidence based comparison across therapeutic classes. This organizational fragmentation limits comprehensive patient stratification and constrains treatment sequencing optimization. We propose the Rosetta Stone for Immune Toxicity Intelligence as an integrated analytical framework addressing this critical clinical gap through unified representation of immune related adverse events across three major immunotherapeutic drug families. Retrospective analysis from DCOPiH Center in Wrocław, Poland comprises three cohorts: immune checkpoint inhibitor treated patients, approximately 200 individuals; radionuclide therapy recipients, approximately 70 patients receiving Lu177 PSMA and Ra223; and antibody drug conjugate treated patients, approximately 130 individuals with primary HER2 positive breast cancer diagnosis. Clinical and laboratory variables, treatment exposure parameters, and adverse event outcomes were encoded according to unified scheme incorporating organ system involved, CTCAE grade severity, timing of symptom onset measured in days from treatment initiation, requirement for immunosuppressive intervention, and treatment

modification status. This standardized encoding enables direct integration with contemporary hospital information systems.

The analytical methodology comprises four sequential stages: first, descriptive analysis characterizes toxicity distributions and inter-therapeutic class relationships; second, ensemble machine-learning models (Random Forest, XG-Boost and logistic regression with L1/L2 regularization) identify cross-class predictors of clinically significant toxicity with explicit correction for class imbalance using cost-weighted loss functions and SMOTE oversampling [1]; third, unsupervised clustering with hierarchical methods and Gaussian mixture models yields three clinically interpretable phenotypes, with Phenotype A denoting early-onset multi-organ high-grade toxicity requiring immunosuppressive management, Phenotype B delayed-onset isolated-organ low-grade toxicity amenable to dose modification, and Phenotype C predominantly hematologic toxicity associated with radionuclide therapy [2]; fourth, explainable machine-learning analysis using SHAP and LIME with Individual Conditional Expectation visualizations, cross-validated by Spearman rank correlations, identifies baseline and treatment-derived features driving phenotype assignment and risk stratification [5]. Expected results are distinct yet overlapping toxicity spectra across therapeutic classes, stable cross-class predictors including baseline organ function and systemic inflammatory markers, and phenotype-specific risk hierarchies that provide an analytic basis for treatment optimization and risk-adapted sequencing in comprehensive oncology practice.

References

1. Bai, C., Wu, L., Li, R., Cao, Y., He, S., Bo, X.: Machine learning-enabled drug-induced toxicity prediction. *Advanced Science* **12**(16), 241340 (2025)
2. Bellavia, A., Ran, X., Zimerman, A., Antman, E.M., Giugliano, R.P., Morrow, D.A., Murphy, S.A.: Unsupervised clustering approach to assess heterogeneity of treatment effects across patient phenotypes in randomized clinical trials. *Contemporary Clinical Trials* **148** (2025)
3. Belloni, S., et al.: Prevalence of treatment-related adverse events (traes) with antibody-drug conjugates in metastatic breast cancer patients: A systematic review and meta-analysis. *Critical Reviews in Oncology/Hematology* **204**, 104527 (2024)
4. Da, L., et al.: Organ-Specific Immune-Related Adverse Events Associated With Immune Checkpoint Inhibitor Monotherapy Versus Combination Therapy in Cancer: A Meta-Analysis of Randomized Controlled Trials. *Frontiers in Pharmacology* **10**, 1671 (2020)
5. Johannsson, A., Chukhrova, N.: The crucial role of explainable artificial intelligence (xai) in improving health care management. *Health Care Management Science* **28**(5), 565–570 (2025)
6. Markides, D.M., Hita, A.G., Merlin, J., Reyes-Gibby, C., ching J. Yeung, S.: Antibody-Drug Conjugates: The Toxicities and Adverse Effects That Emergency Physicians Must Know. *Annals of Emergency Medicine* **85**(3), 214–229 (2025)
7. Parihar, A.S., Chopra, S., Prasad, V.: Nephrotoxicity after radionuclide therapies. *Translational Oncology* **15**(1), 101295 (2022)
8. Siddiqui, M.A., et al.: Immune related adverse events in immunotherapy. Challenges in diagnosis, monitoring, and management. *Toxicology Reports* **14**, 102036 (2025)

Applying Machine Learning and Automation to Enhance Student Support in Higher Education

Kate Han^{1[0000-0002-0800-1644]}, Yun Chen^{1[0000-0001-5123-1624]}, and Charlie Seager^{1[0009-0001-5733-2730]}

Salford Business School, University of Salford, Greater Manchester, UK
 k.han3@salford.ac.uk

Workshop: Innovative Teaching Approaches in Computer Science and Automation: Methods, Tools, and Experiences

Keywords: Educational Data Pipeline · Machine Learning · Student Support · UK Higher Education

Extended Abstract

Universities today gather large amounts of student data from digital learning platforms, assessments, and administrative records. These datasets create opportunities to improve student support, but institutions often struggle to combine and use them effectively [2]. This research develops an automated system, called the APT pipeline, that collects, organizes, and analyzes student data to provide early warnings and targeted support [6].

Traditional support methods are reactive, relying on teacher observation or end-of-term results. Machine learning techniques, including Random Forests [7] and Gradient Boosting Machines [4], allow early identification of students who may underperform [8].

The APT modular pipeline integrates multiple data sources, including Module Report CSVs, Programme Board Excel files, and registration databases. It anonymizes student identifiers and names consistently across nested and irregular Excel structures to maintain privacy [5]. Assessment results are linked to students, modules, and programmes. The system calculates metrics such as engagement patterns, submission timing, and grades, which are fed into predictive models to forecast performance and identify at-risk students. This approach supports scalable monitoring of 2,500+ students across 12 undergraduate programmes [1, ?, ?].

Predicted alert levels generate automated alerts for students and advisors, forming a feedback loop that improves outcomes while reducing staff workload. Model features highlight key predictors of student performance, guiding teaching strategies and administrative decisions.

Challenges in educational analytics include heterogeneous and inconsistent data, privacy concerns, and differences in institutional readiness [3, ?]. The APT pipeline addresses these through rigorous preprocessing, anonymization, and a

modular architecture that supports gradual deployment, continuous improvement, and integration with existing systems. Combining historical and real-time data allows proactive, reliable student monitoring.

In conclusion, integrating automated data pipelines, predictive models, and actionable alerts transforms student support in higher education. By moving from reactive to proactive monitoring, institutions can enhance engagement, retention, and performance. Future work will focus on real-time analytics, adaptive interventions, and longitudinal evaluation to improve predictive accuracy and scalability.

Acknowledgment

The authors gratefully acknowledge the support of Salford Business School and the pump-priming funding 2024/25 from the University of Salford that enabled this research.

References

1. Chen, Y., Han, K.: Driving automated student support with modular data pipelines. *International Journal of Educational Technology in Higher Education* **21**, 50–72 (2024). <https://doi.org/10.1186/s41239-024-00341-2>
2. Han, K., Chen, Y., Seager, C.: Utilizing data pipelines for automated student performance monitoring in higher education. *Journal of Learning Analytics* **10**(2), 45–62 (2023). <https://doi.org/10.1234/jla.2023.1023>
3. HESA: Higher education data quality and analytics report. Higher Education Statistics Agency (2020)
4. Natekin, A., Knoll, A.: Gradient Boosting Machines: Principles and Algorithms. Springer (2013)
5. Prinsloo, P., Slade, S.: Student privacy and data ethics in learning analytics. *British Journal of Educational Technology* **49**, 167–179 (2018). <https://doi.org/10.1111/bjet.12515>
6. Rahman, M., Smith, A.: Challenges in implementing learning analytics in higher education. *Computers Education* **130**, 123–137 (2019). <https://doi.org/10.1016/j.compedu.2018.12.012>
7. Rigatti, L.: Random Forests: Fundamentals and Applications. Springer (2017)
8. Schapire, R.E.: Explaining AdaBoost. Springer (2013)

Digital Learning Factory As A Workshop Format: Low-Threshold Introduction To Industry 4.0

Nikolaus Angel¹[0009-0005-5474-7462] and Horst Orsolits²[0000-0003-1915-1063]

UAS Technikum Vienna, Höchstädtplatz 6, 1200 Vienna, Austria
 info@technikum-wien.at
<https://www.technikum-wien.at>

Abstract. Industry 4.0 increases the need for interdisciplinary competences as workers must integrate mechanical, electrical, digital, and organizational knowledge in cyber-physical production environments. While research shows that such skills require intentionally designed learning processes, current learning-factory formats mainly target students. At the same time, non-academic workers—who apply engineering knowledge daily—lack accessible training opportunities. This study addresses this gap by designing and implementing a learning-factory workshop specifically for this target group. Using a concurrent embedded mixed-methods design (QUAN(qual)), a Likert-scale survey with embedded open-ended questions was used to validate and enrich the quantitative findings. The results contribute to adapting learning factories for lifelong learning and interdisciplinary skill development in the Industry 4.0 context

Keywords: Learning Factory · Industry 4.0

1 Introduction

Even after more than 10 years, there is still no uniform definition of Industry 4.0 in the literature which is widely accepted [3]. The authors of one of the first scientific papers [2] to use the term Industry 4.0 describe the basic idea as follows: " The main idea was to merge real and virtual spaces in so-called cyber-physical production systems, building on progress that German industry had already made with the lighthouse projects on the Internet of Things (IoT) and the Internet of Services (IoS)". In view of the changing skill requirements of Industry 4.0, adapted training and further education is necessary. Learning factories offer a suitable, practical infrastructure for systematically developing the relevant skills.

However, Cazeri et al. show that the training formats of learning factories focus primarily on students [4]. According to Liu and Brunhaver [1], it is important to provide lifelong learning opportunities for workers who apply engineering knowledge daily but do not hold engineering titles or degrees. This paper has set itself the task of designing and implementing a learning factory as a workshop format for workers without formal academic education.

2 Methods

2.1 Evaluation design

A concurrent embedded mixed-methods design (QUAN(qual)) was employed: a Likert-scale survey with three embedded open-ended questions was used to complement/validate the quantitative findings—consistent with the Embedded Design and the triangulation variant ‘Validating Quantitative Data Model’[?].

3 Results




Fig. 1. The attitude of the participants motivation towards the learning factory workshop.

References

1. Liu, J.; Brunhaver, S. From the margins to the center: Repositioning workforce development in engineering education. *JEE* 2025, 4, 25. <https://doi.org/10.1002/jee.70022>
2. Kagermann, H.; Wahlster, W. Ten Years of Industrie 4.0. *Sci* 2022, 4, 26. <https://doi.org/10.3390/sci4030026>
3. Culot, G., Nassimbeni, G., Orzes, G., Sartor, M.: Behind the definition of Industry 4.0: Analysis and open questions. *International Journal of Production Economics* **226**, 107617 (2020). <https://doi.org/10.1016/j.ijpe.2020.107617>
4. Cazeri, G. T., Santa-Eulália, L. A. de, Serafim, M. P., & Anholon, R. (2022). Training for Industry 4.0: a systematic literature review and directions for future research. *Brazilian Journal of Operations & Production Management*, 19(3), 1–19. <https://doi.org/10.14488/BJOPM.2022.007>
5. Poszytek, P. (2021). Interdisciplinarity as a key competence on industry 4.0 labor market. *Journal of Systemics, Cybernetics and Informatics*, 19(8), 1-7.
6. Kirkpatrick, Donald L., R. L. Craig, and L. R. Bittel. "Evaluation of training." Evaluation of short-term training in rehabilitation (1970): 35.

Integrating Makeathons and Problem-Based Learning in MINT Degree Programs: Curriculum Design and Fair Assessment Strategies

Simone Arnold¹[0009–0008–0905–1599], Andreas Kneißler²[0000–0002–4072–5044], and Katja Zeume³[0009–0000–3946–8514]

¹ University of Applied Sciences Dortmund simone.arnold@fh-dortmund.de

² Westphalian University of Applied Sciences Gelsenkirchen
andreas.kneissler@w-hs.de

³ Westphalian University of Applied Sciences Gelsenkirchen
katja.zeume@w-hs.de

Abstract. Universities face a growing imbalance between the rising demand for qualified STEM graduates and declining student numbers. To address this, innovative teaching formats are needed that strengthen engagement, emphasize practical relevance, and develop transversal skills such as critical thinking, collaboration, and adaptability. Experience-based approaches like Problem-Based Learning (PBL) and Makeathons foster interdisciplinary teamwork and creative problem-solving, yet their curricular integration remains challenging. This paper explores how both formats can be sustainably embedded in accredited MINT programs while ensuring fair individual assessment within collaborative learning environments.

1 Introduction

While demand for STEM professionals continues to grow, enrollment in these fields has steadily decreased. This imbalance highlights the need for new learning environments that make the relevance and applicability of practical competencies more visible to students. Universities of applied sciences already integrate practice-oriented teaching, but experience-based formats such as PBL and Makeathons can further enhance motivation and deepen understanding by linking theory with real-world challenges. Employers increasingly expect graduates to demonstrate transversal skills — critical thinking, collaboration, adaptability, and the ability to deal with complex problems—that traditional lecture-based teaching rarely promotes. Strengthening these competencies requires learning formats that combine academic rigor with creativity, teamwork, and reflective practice. Integrating such approaches sustainably into existing, accredited curricula raises important questions regarding workload transparency, module coherence, and especially fair, individual assessment in team-centered learning contexts. Addressing these challenges is essential to modernize STEM education and align it with evolving professional expectations.

1.1 Problem-Based Learning

Problem-Based Learning (PBL) is a student-centered approach in which learners collaboratively work on open-ended, real-world problems to acquire both disciplinary and transversal competencies such as self-directed learning and teamwork. Rather than passively receiving content, students engage in inquiry cycles where the problem drives the learning process. Research shows that PBL enhances autonomy and problem-solving skills, particularly when supported by authentic problem contexts and structured facilitation [1, 2].

1.2 Makeathons

A Makeathon (a blend of make and marathon) is an intensive innovation event where interdisciplinary teams design and prototype physical or digital solutions under time constraints. Unlike hackathons, Makeathons emphasize applied engineering and hands-on creation. They bring together students, professionals, and industry partners to develop solutions in areas such as smart energy, mobility, or sustainability (ITQ GmbH, n.d.; MTF Labs, n.d.; Hochschulallianz Ruhr, n.d.).

2 Research Questions

To harness the potential of experience-based learning without compromising academic rigor, this work investigates two core questions:

- How can Makeathon and PBL learning outcomes be sustainably integrated into degree programs while maintaining curriculum quality and accreditation standards?
- How can fair and transparent individual assessment be ensured in collaborative, PBL-based modules that combine personal accountability with team performance?

Acknowledgments. The authors thank the Applied Excellence Department for supporting the development of a new master's degree program with a strong focus on problem-based learning.

Disclosure of Interests. The authors declare no competing interests, the research was conducted within the context of publicly funded academic curriculum development.

References

1. Wagino, W. et al. (2024). *Enhancing Learning Outcomes and Student Engagement: Integrating E-Learning Innovations into Problem-Based Higher Education*. Int. J. of Interactive Mobile Technologies (iJIM), 18(10), 106–124. <https://doi.org/10.3991/ijim.v18i10.47649>
2. Folomieieva, N. et al. (2024). *Problem-based learning in higher education: a path towards training innovative and competent specialists*. Amazonia Investiga, 13(77), 201–217. <https://doi.org/10.34069/AI/2024.77.05.15>

Embedded Engineering Experiences: Impact on Learning Efficiency – Smart Packaging Technologies Pilot

Markus Glück ^[0000-0002-7631-5389], Bernhard Höfig ^[0009-0007-6646-9089], Peter Eichinger

University of Applied Sciences Aalen, Faculty of Optics and Mechatronics
Beethovenstr. 1, 73430 Aalen, Germany

Abstract. This paper discusses an innovative approach to increase the course efficiency in education by “*embedded engineering and learning experiences*” in close interaction with industrial professionals as R&D coaches. In this pilot initiative, addressing smart packaging technologies, we observed substantial improvements on the development of systems engineering skills by on-site project-based learning under agile conditions and ambitious target setting, discovering and integrating new technologies - e. g. AI-based vision and robotics, handling and additive manufacturing - into packaging systems.

Keywords: Embedded Learning · Embedded Engineering · Mechatronic Systems Engineering · AI · Robotics · Vision · Agile Innovation.

1 Motivation

Today, young engineers are extremely challenged as pacemakers of digital transformation, designers of advanced robotics and automation, pioneers in technical revolutions such as artificial intelligence (AI), vision, robotics, and as enablers of smart and sustainable packaging technologies, respectively [1, 2]. As a consequence, our graduates will face complex challenges in mechatronic systems engineering in an industrial region like the “Packaging Valley” in our close neighborhood and need to be prepared accordingly in special classes.

Problem-based learning as student-centered teaching methodology where learners acquire knowledge by working collaboratively to solve complex, real-world problems has proven to be an attractive methodology to enhance the knowledge transfer by practical experience [3, 4]. In combination with agile and human-centered development as well as direct testing in a professional industrial working environment, significant efficiency gains are expected [5].

2 Embedded Engineering and Learning Experiences

Unlike traditional knowledge transfer in classrooms, problem-based learning starts with the presentation of a realistic challenge to the students as a “trigger” to identify what they need to learn, guide their research, and apply their knowledge to find an appropriate solution in dedicated team efforts. The focus is on developing critical thinking, problem-solving, and learning skills together with professional coaches acting as instructors and facilitators. To improve the learning efficiency, we observed, that embedding student teams fulltime into an industrial working environment for a dedicated period in time and stimulating direct interaction between students and skilled R&D workers under agile frame conditions and ambitious target setting significantly leverages the skill development and reduces reservations to new technologies.

3 Smart Packaging Technologies Pilot

In a 2-week seminar on smart packaging technologies, we embedded a student team of mechatronic systems engineering at Gerhard Schubert GmbH, a supplier of automated packaging systems. Agile postings of theoretical baseline were given on a daily basis, starting with a special baseline setting in control technology at the regional Beckhoff Automation Training Center.

Several challenges were presented to the student teams, addressing current R&D challenges to be solved. Daily stand-up meetings on the progress of the sprint cycles were held on the shopfloor. Functional tests at the robots had also to be absolved using their own minimum viable products (MVPs).




Fig. 1. Embedded learning experiences on AI-based vision & control on the shopfloor

All the challenges were devoted to advanced mechatronic systems engineering. Several package types were presented to be handled in a vision-based robot application [6]. Grasping tools were designed, fabricated by 3D printing and integrated into the robot control environment. Interfaces were adapted, cameras calibrated, control parameters adjusted, workpieces identified. The latest AI methods and algorithms from concurrent R&D in the partner company were used. Training data were generated synthetically. Automated labeling was introduced. Finally, the correct function of the models and systems were tested and the results had to be presented on the shopfloor as an exam.

4 Conclusions

Embedded learning and engineering experiences proved to be a powerful method to boost the learning efficiency in student classes. We have observed a significant increase in self-confidence to master complex challenges due to lessons learned in the experience phases. Our student teams developed their skills significantly from a basic understanding of automation to an advanced level of systems engineering, being now able to integrate robotics, vision, and latest AI features.

5 Acknowledgements

The authors gratefully acknowledge the generous support of our collaboration partners Gerhard Schubert GmbH and Beckhoff Automation GmbH & Co. KG.

References

1. Eichinger, P., Höfig, B., Richter, C.: *Education 4.0 for Mechatronics – Agile and Smart*. In: Proc. 18th Intern. Conf. on Research and Education in Mechatronics (REM), Sept. 14-15, Wolfsbüttel (2017) <https://ieeexplore.ieee.org/document/8075250>
2. Nafz, N., Höfig, B., Glück, M., Mauß, N.-M.: *Introduction to Circular System Design and First Use Cases for Sustainable Product Development in Smart Meter Remanufacturing and Robotics*, Computer Aided Systems Theory – EUROCAST 2024: 19th Intern. Conf., Las Palmas de Gran Canaria, Feb. 25 – Mar. 1, 2024, Part III, p. 132-146 (2025). https://doi.org/10.1007/978-3-031-83885-9_13
3. Barrows, H. S.: *Problem-based learning in medicine and beyond*, in Wilkerson, L., Gijselears, W. H. (Eds.): *New directions for teaching and learning*, Vol. 68, p. 3-13, Jossey-Bass, San Francisco, (1996)
4. Strittmatter-Haubold, V.: *Problem Based Learning. Eine Methode für wirksame Lernsettings in der wissenschaftlichen Weiterbildung?*, in: Hochschule und Weiterbildung, Schwerpunktthema: Hochschulweiterbildung und Beruf, 1/2016, DGWF, S. 50-55 (2016)
5. Glück, M.: *Agile Innovation: Mit neuem Schwung zum Erfolg*, Springer Vieweg, Wiesbaden (2025). <https://doi.org/10.1007/978-3-658-46584-1>
6. Glück, M.: *Mensch-Roboter-Kooperation erfolgreich einführen*, Springer Vieweg, Wiesbaden (2025). <https://doi.org/10.1007/978-3-658-49743-9>

Understanding Modern Data Architectures Through Visualized NoSQL Cluster Simulations

Katja Zeume¹, Mert Ayas¹, and Dennis Stroz¹

Westphalian UAS, Gelsenkirchen, Germany

katja.zeume@w-hs.de, mert.ayas@w-hs.de, dennis.stroz@w-hs.de

1 Motivation

NoSQL databases are experiencing significant growth in popularity [5]. However, incorporating NoSQL into STEM education curricula presents challenges. We identify two primary reasons for this challenge: the abundance of varied and complex data models and structures [1, 3], and the recognition that teaching NoSQL is not synonymous with learning these data models [2, 4].

The majority of NoSQL data models have been developed around very large data sets and the necessity to scale architecture horizontally. Each new data model, such as document-based, graph-based, or key-value stores, introduces unique challenges and algorithms that are crucial for understanding the efficiency of these models in practice in addition to learning a new query language. While many NoSQL databases offer easy-to-access (desktop) tools for visualizing and learning their respective query language, these tools do not support teaching their behavior for cluster computing. We present a platform that facilitates learning NoSQL databases in connection with their behavior on a larger distributed architecture. Our implementation consists of a modular framework which can be used to build a personal cluster simulation. The running simulation can then be analyzed and manipulated into erroneous states rather than harming a real cluster.

2 The Cluster Simulation

We need to address the following two problems. First, it is necessary to have a platform for learning the different data models with their respective query languages. Undergraduate STEM students usually start by learning databases on a relational system and learning SQL. The motivation is a simple server-client application, which is sufficient for most foundational cases. Many NoSQL databases are optimized to work with large data sets and provide easily accessible (desktop) tools for visualizing and learning their respective query language. Unlike the standard relational database application, many NoSQL applications work with large data sets that differ not only in size, but also in structure and content.

For our platform the selection of suitable raw sample data is crucial to our objective to learn concepts of cluster computing rather than to perform extensive

computations. It is essential to identify data sets that are both small and efficient, but sufficiently diverse to encompass all properties.

Secondly, students are motivated to learn NoSQL databases for applying them on clusters rather than monolithic systems. Therefore, it is essential to build teaching cases which include the concepts of scalability, redundancy, fault tolerance, and more. Our teaching platform facilitates these concepts using visualizations and simulations of distributed algorithms to assist students in learning data modeling and query writing. The platform demonstrates the capacity to effectively manage these algorithms and emulate diverse machines within a cluster configuration. The foundational concept must possess the capacity to execute on a personal computer or server in a manner that conserves resources. It is also essential that it exhibit sufficient flexibility to accommodate the desired modifications in the communication between two processes. It is essential that this is done in a way that conserves resources in order to give all students the best access.

Acknowledgments. This work is funded by a FREIRAUM grant from *Stiftung Innovation in der Hochschullehre (StIL)*.

References

1. Guo, M., Qian, K., Yang, L.: Hands-on labs for learning mobile and nosql database security. In: 2016 IEEE 40th Annual Computer Software and Applications Conference (COMPSAC). vol. 2, pp. 606–607. Institute of Electrical and Electronics Engineers, Atlanta, GA, USA (2016). <https://doi.org/10.1109/COMPSAC.2016.126>
2. Miedema, D., Taipalus, T., Ajanovski, V.V., Alawini, A., Goodfellow, M., Liut, M., Peltsverger, S., Young, T.: Data systems education: Curriculum recommendations, course syllabi, and industry needs. In: 2024 Working Group Reports on Innovation and Technology in Computer Science Education. p. 95–123. ITiCSE 2024, Association for Computing Machinery, New York, NY, USA (2025). <https://doi.org/10.1145/3689187.3709609>, <https://doi.org/10.1145/3689187.3709609>
3. Mohan, S.: Teaching nosql databases to undergraduate students: A novel approach. In: Proceedings of the 49th ACM Technical Symposium on Computer Science Education. p. 314–319. SIGCSE '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3159450.3159554>, <https://doi.org/10.1145/3159450.3159554>
4. Tripathi, N.: Nosql database education: A review of models, tools and teaching methods. Journal of Systems and Software **226**, 112391 (2025). <https://doi.org/https://doi.org/10.1016/j.jss.2025.112391>
5. Tu, H.: Cassandra vs. mongodb: A systematic review of two nosql data stores in their industry uses. In: 2024 IEEE 7th International Conference on Big Data and Artificial Intelligence (BDAI). pp. 81–86. Institute of Electrical and Electronics Engineers, Beijing, China (2024)

Design Thinking and Project Based Learning as a synergetic approach to learning

Hanna C. Trenkler¹, Simon Kranzer¹, Matthäus Horn¹, Manuel Taferner¹, and Reinhard Portenkirchner¹

Salzburg University of Applied Sciences, Austria, Department Information Technology

1 Extended Abstract

The many new developments in computer science and automation reshape not only industry, but also what students need to learn, and how they should learn it. To prepare for this change, teaching should encourage experimentation, prototyping, and iteration rather than one-directional transfer of knowledge. Design Thinking (DT) and Project-Based Learning (PBL) complement each other. PBL structures complex, real-life challenges, while DT adds a creative mindset with human-centred focus. Together, they build the skills that innovation depends on. Prior work [1] shows that combining the two strengthens creativity and problem-solving.

This paper builds on experiences from a two-day DT workshop for small and medium-sized enterprises (SMEs), organised as part of a bootcamp at Salzburg University of Applied Sciences. The goal was to give participants practical tools they could use to work on challenges from their own work. Practitioners from retail, hospitality, and innovation sectors worked on issues ranging from rural mobility to staff-scheduling challenges. Through personas, empathy maps, and “How Might We” prompts, participants adopted a user perspective. Although time restricted prototyping, engagement remained strong. Learners stayed reflective and developed a clearer sense of the root causes behind their challenges, a learning effect similar to what [3] describe. The guiding research question is: How can insights from a Design Thinking SME bootcamp inform the development of project-based learning settings in technical higher education? The topic is important because technical programmes emphasise efficiency over creativity and user perspective, although these are essential for successful projects.

DT provides an iterative user-centered process that balances desirability, feasibility, and sustainability in its flexible phases (Empathise, Define, Ideate, Prototype, Test) [3]. PBL anchors that process in structured project milestones. Meta-analytic evidence [5] links PBL to higher achievement, motivation, and thinking skills. For educators, combining both means letting go of full control, therefore embracing ambiguity and guiding learning rather than directing it [2].

The bootcamp itself was made up of short theoretical inputs and hands-on teamwork. After exploring user needs, participants reframed their challenges, brainstormed ideas, and outlined realistic next steps based on the resources they actually had. Given the limited time, the group stayed lively and focused

throughout. Several participants later mentioned that they had continued using some methods on their own, which was one of the most rewarding results. This experience reflects the collaborative and creative gains described in [1] for DT-PBL environments.

What worked in the bootcamp could translate into higher education. In automation courses, for instance, students might analyse and redesign real industrial or service processes using a Design Thinking perspective. Computer science students could take a similar approach when developing software, starting from user research, testing ideas early, and refining them through iteration. These types of projects connect technical skills with empathy and creativity, and align closely with industry-integrated PBL models [4] and meta-analytic evidence [5] showing how such approaches strengthen teamwork, adaptability, and motivation, skills that education should cultivate.

The application of DT-PBL in university settings is not without challenges: limited time for iterative cycles [1], assessment systems that undervalue collaboration and reflection [2], and firm curricula. These challenges echo what was observed in the workshop. Attendees initially jumped to solutions, but after they familiarised themselves with reflection and iteration, their solutions improved. The alignment of industry collaboration with course design remains a promising area for further exploration [4].

This work contributes a practice-based analysis of how a DT-oriented SME workshop can inform the design of structured PBL environments in technical higher education. It highlights concrete behavioural changes in problem framing and collaboration, and derives transferable design principles for DT-PBL teaching in computer science and automation. Future work should focus on developing scalable assessment models, integrating extended DT cycles into technical curricula, and strengthening cross-disciplinary collaboration.

Acknowledgments. We gratefully acknowledge the active participation of our partners. This work was funded by the Austrian Promotion Agency and the Federal Ministry for Innovation, Technology and Telecommunications. This research is part of <https://retailization.at>.

References

1. Jia, L., et al.: Design thinking and project-based learning (dt-pbl): A review of the literature. *International Journal of Learning, Teaching and Educational Research* **22**(8) (2023)
2. Lake, D., et al.: Design thinking in higher education: Opportunities and challenges for decolonized learning. *Teaching and Learning Inquiry* **12** (2024)
3. McLaughlin, J., et al.: Design thinking teaching and learning in higher education: Experiences across four universities. *PLoS ONE* **17** (2022)
4. Naseer, F., et al.: Project based learning framework integrating industry collaboration to enhance student future readiness in higher education. *Scientific Reports* **15** (2025)
5. Zhang, L., Ma, Y.: A study of the impact of project-based learning on student learning effects: a meta-analysis study. *Frontiers in Psychology* **14** (2023)

Teaching Computational Thinking and Coding in an Extra-curricular Context: An Assessment of the Status Quo

Lukas Amadeus Schachner [0009-0002-9834-8525]

Aalen University of Applied Sciences, Germany
lukas.schachner@hs-aalen.de

Extended Abstract

Teaching computational thinking and coding is an urgent educational imperative that transcends mere programming. It encompasses “the thought processes involved in formulating a problem and expressing its solution(s) so that a computer – human or machine – can effectively carry out” [1]. Closely associated with information literacy, computational thinking fosters both declarative and procedural skills and analytical capacities. In the context of digital transformation and rapid advances in artificial intelligence, these skills are increasingly vital for participation in a democratic society and for success in STEM fields [2]. Despite its significance, the International Computer and Information Literacy Study showed that in 2018, one third of German students lacked basic ICT knowledge. Since then, governments have launched ICT initiatives and invested in educational infrastructure to address these deficits [3-4].

This article critically examines extracurricular approaches to computational thinking and coding from preschool to secondary level II. Informal ICT learning is highly heterogeneous and presents unique pedagogical challenges. Three main settings are discussed:

Firstly, the global maker and coding movement promotes hands-on, creative learning and attracts young participants through experimentation. Early coding is often introduced via games, robots, and block-based programming (e.g., Scratch), while microcontrollers (Calliope, Arduino, Raspberry Pi) and Python are popular among older students [5-6]. These activities are mostly informal, but increasingly address social expectations and inclusion for disadvantaged groups.

Secondly, German “school laboratories” offer a wide range of extracurricular learning opportunities, with over 400 nationwide and 231 focusing on ICT learning. These labs encourage interest in STE(A)M topics through interactive formats and help mitigate gender disparities [7]. Many labs participate in initiatives such as Code Week, which in 2025 involved 1,458 participants and organizations.

Thirdly, industry-supported training courses like ‘Technik *macht* Spaß!’ provide free-format, practical learning experiences. These programs focus on modern technologies as solutions to societal challenges such as mobility and climate change, allowing methodological innovation. However, systematic research on their effectiveness remains limited.

References

1. Wing, J.M.: Computational thinking’s influence on research and education for all, p. 8. *Italian Journal of Educational Technology* 25 (2), 7-14 (2017). doi: 10.17471/2499-4324/922
2. Mills, K.A., Cope, J., Scholes, L., Rowe, L.: Coding and computational thinking across the curriculum: A review of educational outcomes, p. 611. *Review of Educational Research* 95 (3), 581-618 (2025). doi: 10.3102/00346543241241327
3. Fraillon, J., Ainley, J., Schulz, W., Friedman, T., Duckworth, D. (eds.): IEA International Computer and Information Literacy Study 2018: Preparing for Life in a Digital World. Springer Open, Cham (2020). doi: 10.1007/978-3-030-38781-5
4. Fraillon, J., Rožman M. (eds.): IEA International Computer and Information Literacy Study 2023: Assessment Framework, pp. 6-10. Springer Open, Cham (2025). doi: 10.1007/978-3-031-61194-0
5. Soomro, S.A., Casakin, H., Nanjappan, V., Georgiev, G.V.: Makerspaces fostering creativity: A systematic literature review. *Journal of Science Education and Technology* 32, 530-548 (2023). doi:10.1007/s10956-023-10041-4
6. Marín-Marín, J.-A., Carcía-Tudela, P.A., Duo-Terrón, P.: Computational thinking and programming with Arduino in education: A systematic review for secondary education. *Heliyon* 10, 1-15 (2024). doi: 10.1016/j.heliyon.2024.e29177
7. Pawek, C.: Schülerlabore als interessenfördernde außerschulische Lernumgebungen. In: Brovelli, D., Fuchs, K., von Niederhäusern R., Rempfler, A. (eds.) *Kompetenzentwicklung an außerschulischen Lernorten*, pp. 69-94. LIT-Verlag, Münster/Wien/Zürich (2012)

AI Augmented Virtual Student Exchange for 24.7 Global Engineering Education

Zenon Chaczko¹[0000-0002-2816-7510], Wojciech Bozejko³[0000-0002-1868-8603], Imre J. Rudas²[0000-0002-2067-8578], Atsushi Ito²[0000-0002-9686-4019], Ryszard Klempous³[0000-0003-4889-0930], Jan Nikodem³[0000-0002-3748-9255], Jerzy Rozenblit²[0000-0002-7348-4128], Marek Kulbacki^{4,5}[0000-0003-4609-106X], and Jakub Segen^{4,5}[0000-0002-8697-6025]

¹ University of Technology, Sydney, Australia

² Chuo University, Tokyo, Japan

³ Wroclaw University of Science and Technology, Wroclaw, Poland

⁴ Depart. of Electrical and Computer Engineering University of Arizona Tucson, USA

⁵ Faculty of Computer Science, Polish-Japanese Academy of Information Technology, Warsaw, Poland

⁶ DIVE IN AI, Wroclaw, Poland

Extended Abstract

A Virtual Student Exchange (VSX) conducted in a 24/7, time zone distributed environment was originally conceived as a way to align engineering education with the globally distributed software industry, using the same collaboration technologies as modern enterprises [1–3]. Students in Europe, the United States and Australia work on joint projects across three 8 hour time zones, handing over code, documentation and project ownership in a continuous cycle. Earlier VSX deployments have demonstrated feasibility and clear pedagogical benefits in terms of international teamwork skills, project discipline and exposure to real world constraints [4, 5], but orchestration still depends heavily on human coordinators and manual processes. In parallel, recent work on artificial intelligence in education highlights both the potential of AI for personalisation and orchestration and the need for explicit governance in real learning environments [6, 7]. This paper revisits the VSX concept in the era of AI assisted engineering and proposes an upgraded framework in which multi agent AI systems support planning, execution and assessment of globally distributed student projects. The revised VSX architecture preserves the original Globally Distributed Work Environment paradigm, in which students analyse problems, define system requirements and implement solutions in project based courses, but augments it with several classes of AI agents. A project orchestration agent monitors repositories, issue trackers and communication channels across time zones, generates structured handover reports and flags schedule or dependency risks for academic mentors [1, 2]. Personal learning agents are attached to individual students and adapt to their learning curve, providing tailored explanations, micro exercises and formative feedback while respecting course constraints and assessment integrity in line

with emerging AI in education principles [6]. Human–agent research cells combine students with domain aware agents for literature triage, hypothesis generation and rapid prototyping, accelerating early stage research without replacing human judgement. Language and culture aware communication agents support meetings in which each participant speaks their native language, by providing real time speech transcription, translation and synthetic voice rendering into the meeting language, which reduces friction in cross cultural collaboration and makes participation more equitable. Analytics agents compute longitudinal indicators of participation, workload balance and deadline risk at team and cohort level, enabling earlier intervention than was possible in previous VSX iterations. Coding and design assistants embedded in the development environment provide just in time support for software engineering tasks while enforcing agreed standards for style, documentation and testing. All agents operate under an explicit human in the loop governance model that specifies which decisions must remain under human control and how AI suggestions are logged and audited for educational and ethical purposes [7]. The methodology is evaluated in joint courses linking Wroclaw University of Science and Technology, the University of Technology Sydney, the University of Arizona and partner institutions in Japan, by comparing cohorts using the original VSX model with cohorts using the AI augmented framework on outcomes including project quality, on time delivery, robustness of handovers and student self reported confidence in global teamwork, with log data from development and collaboration tools providing an objective basis for analysing how AI agents reshape the temporal structure of work, the distribution of cognitive effort and the role of human coordinators.

References

1. Chaczko, Z., Davis, J. D., Scott, C.: New Perspectives on Teaching and Learning Software Systems Development in Large Groups. Telecollaboration. In: IADIS International Conference WWW/Internet 2004, Madrid, Spain (2004).
2. Doerry, E., Klempous, R., Nikodem, J., Paetzold, K.: Virtual Student Exchange. Lessons Learned in Virtual International Teaming in Interdisciplinary Design Education. In: Proceedings of the 5th International Conference on IT Based Higher Education and Training (ITHET 2004), pp. 1–6 (2004).
3. Gupta, A., Seshasai, S.: Toward the 24 Hour Knowledge Factory. Massachusetts Institute of Technology Working Paper (2004).
4. Gupta, A.: Expanding the 24 Hour Workplace. Round the Clock Operations Are Not Just for Call Centers Anymore. Wall Street Journal, September 27 (2007).
5. Klempous, R., Maciejewski, H., Nikodem, J., Gonzalez Rodriguez, M., Suarez Araujo, C. Paz: Informatics Engineering Curricula in the European Space Higher Education Area. Towards a BAMA Model. In: ICSE 2003, Coventry, UK, 9–11 September 2003.
6. Holmes, W., Bialik, M., Fadel, C.: Artificial Intelligence in Education. Promises and Implications for Teaching and Learning. Center for Curriculum Redesign, Boston (2019).
7. UNESCO: AI and Education. Guidance for Policy Makers. UNESCO Publishing, Paris (2021).

Author Index

Adlberger, S. 225
Affenzeller, M. 40, 42, 44, 48
..... 50, 56, 249
Aghanouri, A. 150
Agnello, P. 287, 339
Aguasca-Colomo, R. 311
Albano, G. 277
Alberich, R. 30
Alirio Cruz, N. 30
Al-Kaff, A. 148, 156, 158
..... 160, 162
Alonso Afonso, H. 62
Alshenoudy, A. 154
Altmann, J. 191
Álvarez-Rodríguez, L. 215
Anderle, L. 233
Angel, N. 347
Araz, M. 136
Armingol, J.M. 138
Arnold, S. 259, 349
Arthofer, K. 225
Attenberger, A. 140
Ayas, M. 353
Bachinger, F. 42, 44, 50
Balaban, N. 70
Balaji, Y. 76
Banaszak, Z. 68
Barreto-Ramallo, A. 291
Bartlett, D.J. 52
Bauer, J. 269
Bauernfeind, S. 217
Bawiec, M.A. 197
Bayaraa, B. 102
Becker, M. 245
Benhamida, A. 341
Berndorfer, J. 144
Bernroider, D. 239
Blum, C. 70
Bocewicz, G. 68
Boehm, S. 207
Bögl, M. 243
Borghoff, U.M. 32, 335
Botttoni, P. 32
Bożejko, W. 112, 359
Bratukhin, A. 255
Braune, R. 243
Bresich, M. 54, 58
Brida, P. 307, 309
Bridova, I. 309
Brito Santana, J.A. 62
Buchenrieder, K. 335, 337
Buhl, C. 263
Burian, S. 341
Caballero-García-Alcaide, D.
..... 138
Caballero-Gil, C. 297
Caballero-Gil, P. 289, 291
..... 293, 297, 299, 301, 303
Cabrera, Y. 303
Cabrera-Aparicio, Y. 303
Caputo, L. 281
Carfora, M.F. 281
Celik, K. 259
Černý, M. 134
Cesarelli, M. 339
Chaczko, Z. 112, 333
..... 343, 359
Chakraborty, S. 187
Chen, Y. 345
Christie, L.A. 142
Ciardiello, L. 339
Comminiello, D. 124
Cordón, B. 215
Cordova, G. 148
Csanák, E. 169
Cybulska-Stopa, B. 343
Czemplik, A. 96, 98, 100
Dalkilic, M. 321
Dalpiaz, C. 46, 189, 257
De Blasio, G.S. 168
de Moura, J. 215, 231
Desmond, H. 52
Di Crescenzo, A. 273, 275
Díaz Peralta, G. 171
Díaz-Santos, S. 297
Dino, V. 136
Djukanović, M. 70
Dobrača, M. 76
Döller, M. 269
Doujak, D. 114
Dreiseitl, S. 317
Duda, A. 197

<i>Dzeroski, S.</i>	70	<i>Gugg, R.</i>	329
<i>Ebner, F.</i>	249	<i>Gül-Ficici, S.</i>	164
<i>Ehrengreuber, M.</i>	315	<i>Günther, H.</i>	233
<i>Eibensteiner, F.</i>	46, 134	<i>Habringier, S.</i>	44, 235, 241
.....	189, 257	<i>Haider, C.</i>	50, 56, 66
<i>Eichinger, P.</i>	351	<i>Haider, M.</i>	217
<i>Escanez-Exposito, D.</i>	297	<i>Halmerbauer, G.</i>	227
.....	299, 301	<i>Hambrock, C.</i>	56
<i>Espelosín, J.</i>	162	<i>Hammer, F.</i>	84
<i>Evangelatos, A.</i>	122	<i>Han, K.</i>	142, 223, 345
<i>Falkner, D.</i>	243, 249	<i>Hanreich, M.</i>	191
<i>Fallast, A.</i>	144	<i>Happich, G.</i>	88
<i>Fernández-Moreira, B.</i>	231	<i>Haraz, H.</i>	76
<i>Ferreira, P.G.</i>	52	<i>Hardi, L.</i>	140
<i>Finck, S.</i>	269	<i>Hartner-Tiefenthaler, M.</i>	207
<i>Flatscher, J.</i>	219	<i>Hasegawa, M.</i>	102
<i>Flechl, M.</i>	329	<i>Hecktor, R.</i>	217
<i>Fleck, P.</i>	56, 66	<i>Hehenberger, P.</i>	86
<i>Fruhmann, M.</i>	56, 66	<i>Heindler, L.</i>	122
<i>Fukumori, R.</i>	102	<i>Heinzl, B.</i>	269
<i>García, F.</i>	148, 156, 158	<i>Hellwig, M.</i>	237, 269
.....	160, 162	<i>Heras, J.</i>	221
<i>García, R.</i>	168	<i>Hernández, C.</i>	285, 293, 295
<i>García Carrasco, A.</i>	295	<i>Hernández-Goya, C.</i>	285
<i>Garcia-Diaz, J.</i>	301, 303	293, 295
<i>Garcia-Martin, E.</i>	215	<i>Herzmanek, D.</i>	217
<i>García-Rodríguez, C.R.</i>	168	<i>Hidalgo, J.I.</i>	229
<i>Garnica, O.</i>	229	<i>Hinterbichler, H.</i>	136
<i>Giedenbacher, J.</i>	239	<i>Hiramatsu, Y.</i>	102
<i>Gil-Cuadrado, A.</i>	158	<i>Hirlaender, S.</i>	181, 261
<i>Gingo, D.</i>	90	<i>Hirsch, T.</i>	144
<i>Giordano, N.</i>	271	<i>Hofer, C.</i>	84
<i>Giorno, V.</i>	279	<i>Höfig, B.</i>	171, 351
<i>Giretzlehner, M.</i>	213	<i>Höfig, K.</i>	267, 269
<i>Glück, M.</i>	171, 351	<i>Holmes, R.</i>	72
<i>Gnatowski, A.</i>	112	<i>Holom, R.M.</i>	136, 152, 154
<i>Godoy, J.</i>	158, 162	<i>Holzinger, F.</i>	235, 247
<i>Göhner, U.</i>	164, 177, 193	<i>Hölzl, P.A.</i>	132
<i>Golombek, M.</i>	74	<i>Horn, M.</i>	183, 355
<i>Gómez, J.</i>	160	<i>Hoyer, A.</i>	140
<i>Gomez-Tornero, J.L.</i>	307	<i>Hrubý, D.</i>	134
<i>González Hernández, A.</i>	168	<i>Hu, B.</i>	144
<i>Goyanes, E.</i>	231	<i>Huber, S.</i>	181, 185, 195
<i>Grinninger, J.</i>	187	261, 265, 269
<i>Groening, H.</i>	239	<i>Huemer, M.</i>	124, 126, 128
<i>Groesser, S.N.</i>	199, 201	130, 132

<i>Hussein, S.</i>	78	<i>Kronberger, G.</i>	52
<i>Hutar, M.</i>	307	<i>Krösche, J.</i>	331
<i>Hutter, D.</i>	237	<i>Kühhas, L.</i>	315
<i>Illarramendi, M.</i>	146	<i>Kulbacki, M.</i>	112, 166, 333
<i>Inés, A.</i>	221	343, 359
<i>Iqbal, H.</i>	148	<i>Kumar Veeranki, S.P.</i>	211
<i>Ito, A.</i>	92, 94, 102, 112, 359	<i>Kurz, M.</i>	331
<i>Iurlano, E.</i>	64	<i>Kwitt, R.</i>	265
<i>Jacob, D.</i>	251, 253	<i>Lang, O.</i>	126, 128, 130, 132
<i>Jagielski, D.</i>	166, 333	<i>Langer, J.</i>	46, 134, 189, 257
<i>Jäkel, T.</i>	173	<i>Larrinaga, F.</i>	146
<i>Janout, H.</i>	219	<i>Leitner, S.</i>	241
<i>Jaroszczuk, M.</i>	36	<i>Leleniewski, N.</i>	106
<i>Jiménez-Suárez, M.A.</i>	311	<i>Lesak, S.</i>	144
<i>Jungwirth, M.</i>	80	<i>Limmer, S.</i>	54, 58
<i>Kammerer, L.</i>	52	<i>Lischka, F.</i>	193
<i>Kammerhofer, T.</i>	118, 120	<i>Liu, Y.</i>	150
<i>Karbasch, M.T.</i>	323	<i>López Boada, B.</i>	148
<i>Karcz-Duleba, I.</i>	98	<i>López-Pastor, J.A.</i>	307
<i>Karder, J.</i>	241	<i>Lunglmayr, M.</i>	116
<i>Kargl, M.</i>	46, 134, 189	<i>Luo, J.</i>	92
<i>Kartelj, A.</i>	70	<i>Maarouf Bassaidi, M.</i>	62
<i>Kastenhofer, S.</i>	217	<i>Machaj, J.</i>	309
<i>Kastner, P.</i>	46, 189	<i>Maciejczyk, A.</i>	343
<i>Kerschbaumer, B.</i>	38	<i>Maletzky, A.</i>	152, 154, 213
<i>Kerschbaumer, S.</i>	207	<i>Maqueda, E.</i>	229
<i>Kinast, A.</i>	243	<i>Marco-Pérez, I.</i>	146
<i>Kirchsteiger, D.</i>	80	<i>Martinelli, F.</i>	287, 339
<i>Klempous, R.</i>	112, 166	<i>Martinucci, B.</i>	271
.....	333, 359	<i>Maschek, A.</i>	313
<i>Kloiber, K.</i>	191	<i>Matzinger, S.</i>	132
<i>Kluwak, K.</i>	166	<i>Mayrhuber, E.</i>	211
<i>Kneißler, A.</i>	259, 349	<i>Meingassner-Lang, A.</i>	128
<i>Kobayashi, K.</i>	281	<i>Mercaldo, F.</i>	287, 339
<i>Kobelrausch, M.</i>	255	<i>Merschak, S.</i>	86
<i>Koleszár, A.</i>	169	<i>Messineo, S.</i>	265
<i>Koller, M.</i>	132	<i>Mileva, A.</i>	325
<i>Königswieser, T.</i>	225, 227	<i>Miró-Julià, M.</i>	30
<i>Korinek, H.</i>	60	<i>Molina-Gil, J.M.</i>	295, 311
<i>Kovács, L.</i>	169	<i>Moreno Pérez, J.A.</i>	62
<i>Kozlovszky, M.</i>	169, 341	<i>Moser, B.</i>	116
<i>Kranzer, S.</i>	183, 355	<i>Mustaro, V.</i>	275
<i>Krauss, O.</i>	217, 325, 327	<i>Muttenthaler, D.</i>	189
<i>Krenn, A.</i>	136	<i>Nagao, Y.</i>	102
<i>Kressel, J.</i>	187	<i>Nakano, T.</i>	102
<i>Kristof, P.</i>	169	<i>Navarro-Amador, F.J.</i>	311

Némethy, V. 169
 Neuhauser, P. 247
 Neussl, D. 120
 Nezamdoust, A. 124
 Nguyen Duy, D. 42, 44
 Niederländer, S. 329
 Nikodem, J. 166, 333, 359
 Nikodem, M. 305
 Nitu, J. 315
 Nobile, A.G. 279
 Nogami, K. 102
 Novo, J. 215, 231
 Oberhauser, F. 152, 154, 249
 Ocansey, E.D. 136, 154
 Och, A. 132
 Oko, J. 82
 Okochi, J. 82
 Olaverri-Monreal, C. 150
 Olesiak, M. 108
 Onda, M. 102
 Orsolits, H. 347
 Ortega, M. 215, 231
 Ortiz de Murua, J. 221
 Otsuka, A. 92, 94, 102
 Ott, S. 203
 Padilla-Pio, J. 301
 Pammer, L. 175
 Papamarkou, T. 223
 Paráda, E. 255
 Paraggio, P. 271
 Pareschi, R. 32
 Parra, D. 229
 Pawlak, I. 343
 Pereira-Domínguez, I. 285
 Pérez, A. 146
 Pérez, B. 146, 156
 Petyx, M. 287, 339
 Pickl, S.W. 209
 Pietrucha, J. 305
 Pirozzi, E. 281, 283
 Plaimer, B. 126, 128
 Pletz, M. 213
 Podann, G. 259
 Porębski, A. 112
 Portenkirchner, R. 183, 355
 Pradeep, V. 90
 Praschl, C. 217, 321, 323
 Preiner, K. 46, 257
 Prielinger, A. 257
 Probst, K. 323
 Putz, V. 84
 Quesada-Arencibia, A. 168
 Quesada-Ojeda, M. 72
 Raidl, G.R. 54, 58, 60, 64
 Ralfs, N. 245
 Rechenmacher, M. 217
 Rehrl, J. 181, 261
 Reich, E. 185
 Reisinger, J. 132
 Resino, M. 156, 162
 Riß, F. 267, 269
 Rodemann, T. 60
 Rodriguez-Vega, M. 289, 291
 Roiss, M. 217
 Rosenstatter, T. 195
 Roth, D. 217
 Rozenblit, J. 166, 333, 359
 Rubio, A.L. 146
 Rudas, I.J. 359
 Sabrowsky-Hirsch, B. 154
 Salvenmoser, M. 40
 Sanchez, L. 162
 Sanchis, A. 138
 Sandler, S. 217, 327
 Santone, A. 287, 339
 Sasaki, A. 92
 Sassinick, O. 195
 Scarpiniti, M. 124
 Schachner, L. 357
 Schäfer, G. 181, 261, 269
 Schaffenberger, A. 144
 Schaffernicht, M. 201
 Schebek, M. 175
 Schedl, D.C. 313, 323
 Scheiblhofer, S. 132
 Schirmeier, F. 173, 179, 193
 Schlechter, T. 74, 76, 78
 82, 90
 Schmidt, M. 140
 Scholz, M. 269
 Schönböck, J. 191
 Schulenburg, P. 164, 263

<i>Schuster, S.</i>	132	<i>Trajkovska, N.</i>	217
<i>Schwaninger, M.</i>	203	<i>Trenkler, H.C.</i>	183, 355
<i>Seager, C.</i>	345	<i>Tsuruta, N.</i>	102
<i>Sebek, M.</i>	34	<i>Uchroński, M.</i>	36
<i>Seco, T.</i>	162	<i>Ulrich, D.</i>	179
<i>Segen, J.</i>	112, 333, 343, 359	<i>Unterweger, A.</i>	195
<i>Seliger, R.</i>	177, 179	<i>Uray, M.</i>	265
<i>Selinger, S.</i>	331	<i>Varga, J.</i>	60
<i>Sengl, D.</i>	255	<i>Velasco, J.M.</i>	229
<i>Serban, R.</i>	313	<i>Ventresque, A.</i>	72
<i>Serrano, E.</i>	156	<i>Vetter, J.</i>	227
<i>Sesmero, M.P.</i>	138	<i>Villani, P.</i>	273
<i>Seyser, D.</i>	175	<i>Vincze, M.</i>	341
<i>Siddiqui, M.K.</i>	187	<i>Voggeneder, L.</i>	331
<i>Silaci, M.</i>	337	<i>Voigt, J.L.</i>	88
<i>Skoczynski, P.</i>	333	<i>Waclawek, H.</i>	269
<i>Slezak, C.</i>	219	<i>Wagner, M.</i>	126, 128, 130, 132
<i>Slezak, P.</i>	219	<i>Wagner, S.</i>	235, 247
<i>Smutnicki, C.</i>	68	<i>Waheed, F.</i>	267, 269
<i>Sonnleitner, E.</i>	331	<i>Walcher, C.</i>	225
<i>Spina, S.</i>	273, 275	<i>Weber, N.</i>	253
<i>Spitzer, F.</i>	239	<i>Weiser, T.</i>	88, 251
<i>Stachno, A.</i>	96, 104, 106	<i>Weiss, A.</i>	64
<i>Stampfer, P.</i>	211	<i>Weissenfels, A.</i>	64
<i>Starzyńska, N.</i>	110	<i>Werth, B.</i>	48, 50
<i>Staudinger, C.</i>	44	<i>Werzi, A.</i>	116
<i>Stein, A.M.</i>	118	<i>Weyermann, A.</i>	251
<i>Steinberger, T.</i>	237	<i>Windner, A.</i>	132
<i>Steindl, M.</i>	154	<i>Winkler, S.</i>	52, 211, 219, 227
<i>Stern, P.</i>	144	<i>Wollitzer, E.</i>	134
<i>Stöckl, A.</i>	313, 315, 319, 329	<i>Wundsam, S.</i>	90
<i>Strehar, A.</i>	90	<i>Yagüe, D.</i>	138
<i>Strobl, M.</i>	227	<i>Zawadzki, P.</i>	343
<i>Stroz, D.</i>	353	<i>Zebec, Z.</i>	70
<i>Suárez Santiago, B.</i>	168	<i>Zeinzinger, M.</i>	134
<i>Sun, W.</i>	229	<i>Zenisek, J.</i>	56, 66
<i>Suproniuk, M.</i>	106	<i>Zeume, K.</i>	233, 259, 349, 353
<i>Syne, L.</i>	293	<i>Zheng, L.</i>	205
<i>Szabó Nagyné, O.</i>	169	<i>Zwettler, G.A.</i>	191
<i>Szlachcic, E.</i>	333	<i>Zyśko, D.</i>	166
<i>Taferner, M.</i>	183, 355		
<i>Thurner, T.</i>	114, 118, 120, 122		
<i>Tilebein, M.</i>	205		
<i>Tobar Nicolau, A.</i>	30		
<i>Tockner, M.</i>	128		
<i>Tomandl, L.</i>	54		

Universidad de Las Palmas de Gran Canaria



Johannes Kepler University Linz



University of Applied Sciences Upper Austria



Museo Elder de la Ciencia y la Tecnología



Fundación Universitaria de Las Palmas

