University of Debrecen
Faculty of Informatics
Debrecen, Hungary

The 2026 IEEE 4th Conference on Information Technology and Data Science

General Information

COMMITTEE MEMBERS

General Chair

András Hajdu	University of Debrecen, Hungary

Honorary Chair

Dániel András Drexler	IEEE Hungary Section Chair

Program Committee Chair

István Fazekas	University of Debrecen, Hungary

Session Chairs

András Hajdu	Artificial Intelligence and its Applications
Sándor Baran and István Fazekas	Stochastic models in Data Science
György Vaszil	Automata, Logic, and Models of Computation
Ágnes Vathy-Fogarassy	Predictive Analytics
Andrea Pintér-Huszti	Cybersecurity
Márton Ispány	Temporal Data Science with Applications
Tamás Bérczes	Autonomous Vehicles and Embedded Systems
Attila Gilányi and Mohammad Omar Wedyan	Virtual Reality and Its Applications
László Szathmáry	Software Technology
Zoltán Gál and Alex Doboli	Internet of Things, Operational Technology, and Network Intelligence

Program Committee

Attila Adamkó	University of Debrecen, Hungary
László Bakó	Sapientia Hungarian University of Transylvania, Romania
Jaume Baixeries	Universitat Politècnica de Catalunya (Barcelona), Spain
Sándor Baran	University of Debrecen, Hungary
Tamás Bérczes	University of Debrecen, Hungary
Pascal Bondon	Université Paris-Saclay, CentraleSupélec, France
Henning Bordihn	University of Potsdam, Germany
Sándor Bozóki	Corvinus University of Budapest, Hungary

Shinnosuke Seki	University of Electro-Communications, Tokyo, Japan
Désiré Sidibe	Université Evry Val Essonne, France
Nataša Sladoje	Uppsala University, Sweden
László Szathmáry	University of Debrecen, Hungary
Gábor Szederkényi	Pázmány Péter Catholic University, Hungary
János Sztrik	University of Debrecen, Hungary
György Terdik	University of Debrecen, Hungary
Do Van Tien	Budapest University of Technology and Economics, Hungary
AlinTisan	Royal Holloway, University of London
Arban Uka	Epoka University, Albania
Pál Varga	Budapest University of Technology and Economics, Hungary
György Vaszil	University of Debrecen, Hungary
Ágnes Vathy-Fogarassy	University of Pannonia, Hungary
Mohammad Wedyan	Jordan University of Science and Technology in Irbid, Jordan

Local Organizing Committee Chair

Imre Varga	University of Debrecen, Hungary

Local Organizing Committee

Erzsébet Tóth	University of Debrecen, Hungary
István Fazekas	University of Debrecen, Hungary
János Tóth,	University of Debrecen, Hungary
Piroska Biró	University of Debrecen, Hungary
Bence Hegedűs	University of Debrecen, Hungary
László Mészáros,	University of Debrecen, Hungary
Sándor Pecsora	University of Debrecen, Hungary
Anikó Szakál	IEEE Hungary Section
Csaba Benedek	John von Neumann Computer Society, Hungary
Roland Kunkli	University of Debrecen, Hungary

Treasurer

Arnold Pintér	University of Debrecen, Hungary

Editors

Piroska Biró	University of Debrecen, Hungary
Erzsébet Tóth	University of Debrecen, Hungary

General contact e-mail: citds2026@inf.unideb.hu

Contents

PLENARY AND KEYNOTE SPEAKERS

The Next Challenges in Autonomous Driving: Rare Failures

My T. Thai, University of Florida, USA

Abstract: Autonomous vehicle (AV) policies excel in everyday driving, yet they routinely collapse in rare, safety-critical edge cases. The culprit isn't just a lack of data, but it is a critical vulnerability to minor distribution shifts. While current world models attempt to simulate these scenarios, they optimize for visual realism over physical validity and temporal consistency. The result? Generated videos that look plausible to humans but provide little value for rigorously testing or repairing the underlying AV policy.

In this keynote, I dissect why this gap persists, exposing the limitations of recent world models and LLM-driven scenario generation. I will then introduce our simulation-first framework designed to systematically discover and refine AV policy failures. By anchoring our approach in simulators to enforce strict physical constraints, we deploy critics and planners to pinpoint exactly why a policy fails and how to modify a scenario to expose and repair it. Finally, I will present our sim-to-real pipeline, which transports these simulated failures into the real domain while preserving vital physical, temporal, and policy-relevant failure semantics. Ultimately, this approach delivers real-world models capable of genuinely stress-testing and hardening the next generation of autonomous driving policies.

Professional Summary: My T. Thai is a Nelms Endowed Chair Professor and Associate Director of the Nelms Institute for the Connected World at the University of Florida. An IEEE Fellow and globally recognized leader in trustworthy AI and optimization, her research addresses critical challenges in today's interdependent, complex systems and high-impact applications such as network infrastructure, autonomous systems, healthcare, and cyber-security. She has authored 7 books and over 350 publications in premier journals and conferences, earning multiple best paper awards and shaping both the theory and practice of network science and machine learning.

Responding to the global call for responsible and safe AI, Dr. Thai is a pioneer in developing deep explanations for black-box ML models and defending against explanation-guided attacks, work recognized with the 2023 AAAI Distinguished Paper Award. That same year, she received

the Web Science Trust Test-of-Time Award for her landmark work on combating misinformation in social media. Her earlier honors include the IEEE Big Data Security Women of Achievement Award (2022), the DTRA Young Investigator Award (2009), and the NSF CAREER Award (2010).

Dr. Thai has served as Program Chair for several leading IEEE conferences and holds editorial positions on multiple international journals. She is currently Editor-in-Chief of ACM Computing Survey, ranked #1 in Computer Science, Theory & Methods, and of the Springer Journal of Combinatorial Optimization, as well as editor of Springer's Optimization and Its Applications book series.

Resource Orchestration and Scheduling Across the Mist–Fog–Cloud Continuum: Trends and Future Directions

Helen Karatza, Professor Emeritus, Department of Informatics, Aristotle University of Thessaloniki, Greece

Abstract: The rapid proliferation of Internet of Things (IoT) applications is increasingly exposing the limitations of centralized cloud infrastructures in accommodating the scale, velocity, and geographical distribution of data generated by pervasive sensing devices. To address the stringent latency, bandwidth, and scalability requirements of emerging IoT services, the Cloud computing paradigm has evolved towards a hierarchical computing continuum that incorporates Fog and Mist computing. By moving computation progressively closer to the data sources, Fog computing enables low-latency processing at the network edge, while Mist computing extends this capability to resource-constrained devices, supporting ultra-low-latency execution directly at or near the sensing layer. Many contemporary IoT applications, including industrial automation, autonomous systems, intelligent transportation, smart healthcare, and critical infrastructure monitoring, are inherently time-critical and operate under strict Quality-of-Service (QoS) and deadline constraints. Consequently, efficient resource management and scheduling across the cloud–fog–mist continuum have become fundamental challenges for ensuring predictable performance, scalability, and reliable service delivery in highly distributed environments. This keynote explores the state of the art in scheduling methodologies for heterogeneous cloud–fog–mist infrastructures, with particular emphasis on real-time scheduling, workflow scheduling, fault-tolerant execution, energy-aware resource management, and Quality-of-Service optimization. Modern applications exhibit diverse computational structures, including independent Bag-of-Tasks workloads, tightly coupled parallel applications requiring Gang scheduling, and complex workflows represented as Directed Acyclic Graphs (DAGs), where precedence constraints and end-to-end deadlines significantly complicate scheduling decisions. A particular focus is placed on real-time scheduling techniques based on imprecise (approximate or partial) computation. Imprecise computation provides an effective mechanism for addressing overload conditions by allowing applications to trade computational precision for timeliness. Rather than producing an exact result after its deadline—when its value may be substantially diminished—a computation may instead return an intermediate result of acceptable quality before its deadline. This approach

assumes monotonic applications, whose output quality improves progressively as execution advances. Such applications are typically modelled as consisting of a mandatory component, which must complete to ensure functional correctness, followed by an optional component that incrementally improves result quality. Although approximate execution is not universally applicable, as some workloads require exact computation, it offers an effective scheduling mechanism for a wide class of time-critical applications. The keynote further investigates how imprecise computation can be integrated with application-directed checkpointing to provide fault tolerance in large-scale distributed real-time systems. By periodically checkpointing intermediate computation states—particularly upon completion of the mandatory execution phase—applications can recover efficiently from software failures while preserving timing guarantees and reducing recovery overhead. This combination of adaptive execution and fault-tolerant scheduling offers a promising approach for improving the resilience of cloud-native real-time applications. The unprecedented growth of data generated by sensors, mobile devices, cyber-physical systems, social media, and large-scale IoT deployments has also driven the emergence of Big Data analytics, whose computational demands increasingly rely on distributed cloud–fog–mist infrastructures. Many data-intensive analytics applications exhibit real-time characteristics, as delayed results may significantly reduce their operational value. Consequently, scalable scheduling algorithms must simultaneously address computational performance, communication overhead, timing constraints, reliability, and heterogeneous resource availability. Energy efficiency and sustainability constitute another fundamental dimension of scheduling in modern distributed computing infrastructures. As data centers and edge platforms continue to expand, reducing energy consumption and carbon emissions has become a primary design objective. Energy-aware scheduling strategies must therefore balance performance, deadline satisfaction, resource utilization, and Quality-of-Service requirements while minimizing operational cost and environmental impact, contributing to the broader vision of sustainable cloud and edge computing. The continued evolution from centralized cloud infrastructures towards fully integrated cloud–fog–mist ecosystems is fundamentally transforming the way distributed computing platforms are designed, deployed, and managed. Future computing environments will be characterized by highly heterogeneous resources, dynamic workloads, geographically distributed infrastructures, and increasingly autonomous edge devices, all of which introduce new challenges for efficient resource orchestration. In this context, scheduling is emerging as a key enabling technology for achieving scalability, predictability, resilience, and sustainability across the entire computing continuum. Designing scheduling strategies that can simultaneously satisfy real-time constraints, optimize resource

utilization, tolerate failures, minimize energy consumption, and adapt to highly dynamic execution environments remain a major research challenge. At the same time, the growing integration of artificial intelligence, machine learning, digital twins, and data-driven optimization into distributed computing systems is creating new opportunities for intelligent, self-adaptive scheduling and autonomous resource management. These developments are expected to play a central role in supporting next-generation IoT applications, including cyber-physical systems, smart cities, Industry 5.0, connected healthcare, autonomous transportation, and large-scale digital infrastructures. This keynote reviews recent advances in scheduling algorithms and resource management frameworks across the cloud–fog–mist continuum, presents representative research contributions addressing real-time execution, fault tolerance, approximate computing, workflow management, and energy-aware optimization, and discusses the open research challenges that must be addressed to realize scalable, dependable, and sustainable distributed computing platforms for future intelligent IoT ecosystems.

Professional Summary: Helen Karatza (senior member of IEEE, ACM, SCS) is a Professor Emeritus in the Department of Informatics at the Aristotle University of Thessaloniki, Greece. Dr. Karatza's research interests include cloud, fog and mist computing, energy efficiency, fault tolerance, resource allocation, scheduling algorithms and real-time distributed systems. Dr. Karatza has authored or co-authored over 270 technical papers and book chapters. She is included in the list of the world’s top 2% of scientists (Stanford University rankings). She has received the ACM SIGSIM Distinguished Contributions Award 2025. She served as an elected member of the Board of Directors at Large of the Society for Modeling and Simulation International. She also served as chair and keynote speaker at international conferences. Dr. Karatza is Senior Associate Editor of the Elsevier journal “Simulation Modelling Practice and Theory”, an Editor of “Future Generation Computer Systems” of Elsevier, and the Topical Editor for Edge Computing of IEEE Transactions on Services Computing. She was Editor-in-Chief of the Elsevier journal “Simulation Modelling Practice and Theory”, Editor-in-Chief of “Simulation Transactions of The Society for Modeling and Simulation International”, Associate Editor of “ACM Transactions on Modeling and Computer Simulation” and Senior Associate Editor of the “Journal of Systems and Software” of Elsevier. She served as Guest Editor of multiple Special Issues in international journals. More info about her activities and publications can be found at: https://karatza.webpages.auth.gr/

Complementary AI Paradigm: Logical Synthesis Model - Neural Mesh

Péter Baranyi (Corvinus Institute for Advanced Studies (CIAS), Institute of Data Analytics and Information Systems Corvinus University of Budapest, Budapest, Hungary, Hungarian Research Network (HUN-REN))

Abstract: Large Language Models (LLMs) have fundamentally transformed Artificial Intelligence by learning statistical representations of knowledge from massive datasets. Their success has established statistical learning as the dominant paradigm of modern AI. However, future Physical AI systems require considerably more than accurate prediction. They require mathematical representations that directly support modelling, analysis, optimization, verification, and controller synthesis. This paper argues that these requirements call for a complementary representation paradigm in which the objective is not merely to approximate functions but to learn mathematically structured representations suitable for engineering synthesis. We show that the recently proposed Neural Mesh embodies this philosophy by introducing the first neural architecture whose trainable parameters are intentionally aligned with the mathematical degrees of freedom governing Tensor Product (TP) model construction and control-oriented system synthesis. Rather than viewing TP model transformation as an offline mathematical procedure, the Neural Mesh incorporates its governing principles directly into the learning architecture. This design philosophy naturally extends beyond a single architecture and motivates a broader family of synthesis-oriented neural models, referred to as Logical Synthesis Models (LSMs). Consequently, the Neural Mesh should be regarded not simply as a new neural network but as the first realization of a learning paradigm specifically developed for Engineering AI and Physical AI.

Professional Summary: Peter Baranyi has made significant contributions to the fields of non-linear control theory and modeling. Among his key inventions is the TP model transformation, a sophisticated form of higher-order singular value decomposition for continuous functions. This transformation is crucial in the development of nonlinear control design theories and facilitates new optimisation theory. Baranyi's scientific achievements have been recognised with several prestigious awards, including the Investigator Award from Sigma Xi, the Kimura Award, and the International Design Gabor Award. He has published over 100 journal papers and authored four books, resulting in an h-index of 54.

From Data to Decisions: AI, Digital Twins and Integrated Modelling for Climate and Air Quality

Thanh-Binh Nguyen, International Institute for Applied Systems Analysis (IIASA), Austria / University of Debrecen, Hungary

Abstract: Translating regional climate and air-quality strategies into effective urban actions requires consistent integration of regional scenarios with detailed city-level information. This paper presents an AI-enhanced integrated modelling approach for urban climate and air-quality assessment. The methodology combines two main steps. First, regional activity pathways are downscaled to construct internally consistent city scenarios covering activities, technologies, emissions, costs and environmental impacts. Artificial intelligence and heterogeneous urban data support sector-specific regional-to-city transformations and high-resolution assessment. Second, AI-assisted retrieval is used to identify and adapt mitigation measures relevant to the city-specific activity and technology structure. The proposed approach combines established domain models with data-driven methods to support consistent and transparent urban policy assessment.

Professional Summary: Thanh-Binh Nguyen is a researcher at the International Institute for Applied Systems Analysis (IIASA), Austria, and is affiliated with the Faculty of Informatics at the University of Debrecen, Hungary. His research lies at the intersection of information technology, data science and environmental modelling, with a particular focus on the development and application of the GAINS (Greenhouse Gas – Air Pollution Interactions and Synergies) integrated assessment modelling framework.

His research interests include environmental information systems, integrated assessment modelling, data integration and harmonisation, spatial and temporal downscaling, air-quality and climate modelling, artificial intelligence, and data-driven decision-support systems. He has worked extensively on computational approaches for connecting large and heterogeneous environmental, socioeconomic and scenario datasets with models for analysing air pollution, greenhouse-gas emissions and mitigation strategies.

A particular focus of his current research is the development of GAINS-City, which aims to bridge the gap between regional integrated assessment and urban-scale environmental analysis. This includes the integration and downscaling of activity and emission data, connection with high-resolution atmospheric modelling such as uEMEP, incorporation of monitoring

observations, and the use of AI and data-model fusion to improve urban air-quality and population-exposure assessment.

His broader research explores how AI, integrated models and environmental Digital Twins can be combined to transform complex and heterogeneous data into transparent and actionable information for climate and air-quality decision-making.

ABSTRACTS OF CONTRIBUTED TALKS

Real-Time Traffic Classification for Normal and Abnormal Flows in SDN

Safa Abbas (University of Kufa), Mohammed Nsaif (University of Kufa)

Network traffic classification is an important aspect of network management. This paper presents a real-time fine-grained classification for Software Defined Networking (SDN). Therefore, the paper addresses the research gap in encrypted traffic classification of both normal and abnormal traffic by exploring Machine Learning (ML) techniques with real-time testing evaluation. The method employs the Distributed Internet Traffic Generator (D-ITG) to simulate a wide range of normal real-time traffic, hping3 to simulate DoS traffic and UDP flood, and VLC player to simulate video traffic. The results show that deep neural networks (NNs) have the best prediction performance in both offline and online testing for a wide range of real-time traffic such as video, VoIP, DNS, ping, UDP flood, and gaming traffic, compared to other methods such as Logistic Regression (LR) and Support Vector Machine (SVM). The outcome of this paper could be used in future research directions, which in turn helps increase productivity by controlling different traffic flows, providing sufficient bandwidth for high-priority data, improving routing efficiency, and enhancing security.

FreqTextSeg3D: Frequency-Aware Text-Guided 3D Brain Tumour Segmentation via Hierarchical Semantic Decomposition and Dual-Path Cross-Modal Fusion

Biniyam Mulugeta Abuhayi (Faculty of Informatics, University of Debrecen), András Hajdu (Faculty of Informatics, University of Debrecen)

Accurate 3D brain tumour segmentation is critical for diagnosis, treatment planning, and prognosis monitoring. Although recent vision-language methods have begun to incorporate radiology reports into segmentation pipelines, they typically treat clinical text as a monolithic embedding, ignoring its rich sub-semantic structure. We present FreqTextSeg3D, a novel framework that introduces three tightly coupled contributions: (i) Hierarchical Semantic Text Decomposition (HSTD), which disentangles a radiology report into anatomical, morphological,

and diagnostic streams via learned soft gates with a diversity loss; (ii) Frequency-Aware Cross-Modal Fusion (FACMF), a dual-path module that combines spatial cross-attention with 3D FT amplitude/phase modulation to inject structured text priors into visual features; and (iii) Region-Text Contrastive Alignment (RTCA), a symmetric InfoNCE loss that anchors each text stream to its corresponding tumour sub-region. Evaluated on BraTS2020 under 5-fold cross-validation, FreqTextSeg3D achieves mean Dice scores of 91.36% (Whole Tumour), 86.49% (Tumour Core), and 77.74% (Enhancing Tumour), with an overall mean of 85.19% (standard deviations 0.52, 1.63, and 2.44, respectively), matching or surpassing state-of-the-art image-only and text-guided segmentation methods while providing interpretable, semantically grounded predictions.

On Some Extensions of Boolean Grammars

Patrik Adrián (Doctoral School of Informatics, University of Debrecen)

Context-free grammars (CFGs) are a standard formalism for specifying syntactic structures in computer science, with applications ranging from programming language compilers to natural language models. Despite their widespread adoption and computational tractability, CFGs lack the capacity to natively express some logical constraints: they cannot model the intersection of syntactic properties, nor can they directly handle complementation or exclusionary rules. Boolean grammars, initially introduced by Alexander Okhotin, address this limitation by augmenting the standard CFG framework with explicit logical connectives—specifically, conjunction and negation. This provides an expressive formalism capable of describing complex, overlapping, and non-context-free constraints entirely within the grammar itself. However, this increased expressive power comes at a theoretical cost. The combination of self-referential grammatical rules intertwined with explicit negation spawns logical paradoxes and self-contradictions, typified by the pathological rule where a nonterminal derives exactly its own negation ($X \to \neg X$). Under classical two-valued logic, such grammatical equations lack well-defined models. Furthermore, traditional attempts to obtain solutions using "Negation as Failure" suffer from non-monotonicity, where the temporary assumption of falsehood can trigger an infinite, unresolvable feedback loop.

In our previous work, we successfully resolved these paradoxes by adopting Fitting's information-minimal semantics over Kleene's strong three-valued logic (K3). By establishing a monotonic evaluation model that tolerates an "indeterminate" truth value for deductively

unfounded words, we guaranteed a unique least fixed point that naturally gave rise to efficient polynomial-time parsing algorithms. In this talk, we will abstract away from the specific confines of K3 and string-parsing algorithms. Instead, we investigate a generalized algebraic and topological theory of formal language equations. We achieve this by treating a formal language as a characteristic function mapping an arbitrary structural domain to an arbitrary space of truth values, rather than treating a language as a discrete set of strings, or even as a bipartite pair of positive and negative string sets. We generalize the underlying mechanisms of Boolean grammars across multiple, orthogonal mathematical dimensions, transforming the formalism into a universal framework for logic.

We first demonstrate how to generalize the codomain of truth. This generalized truth space requires a dual-order structure: a truth ordering (forming a lattice to evaluate conjunctions and disjunctions) and a knowledge ordering (a chain-complete poset to model the accumulation of certainty). This dual-order abstraction extends the grammatical framework beyond classical Boolean algebra. For instance, the negation operator need not be an involution—a property where double negation strictly cancels out—to maintain a mathematically well-defined semantics. By relaxing the requirement for involutive negation, our framework naturally accommodates a vast array of non-classical logical systems, including intuitionistic logic, Gödel logic, and abstract modal operators. The only algebraic prerequisite for resolving grammatical paradoxes is that all logical operators remain strictly monotonic with respect to the knowledge ordering. This ensures that the deductive derivation of the language continuously accumulates information, culminating in a stable semantic model regardless of the underlying logic's unique characteristics.

Moving beyond the discrete multi-valued logics of the codomain, this framework also opens the door to continuous and infinitary mathematical spaces. We show how the discrete knowledge ordering can be replaced by a topological metric, thereby generalizing the framework into the realm of quantitative, weighted, and probabilistic Boolean grammars. In this continuous paradigm, the truth space becomes a semiring, and by requiring the generalized language operator to act as a contraction mapping within a complete metric space, paradoxes are resolved topologically rather than algebraically. For example, a paradoxical rule like $X \rightarrow 1 - X$ is no longer assigned an abstract "indeterminate" logical value; instead, it converges to a stable equilibrium via Banach's fixed-point theorem. Furthermore, we extend the framework to encompass infinite structures. By including completely distributive lattices and incorporating greatest fixed-point semantics, we outline how generalized Boolean grammars can evaluate over infinite words (ω-languages) and continuous data streams. This extension allows

grammars to express co-inductive liveness properties—such as guaranteeing that a specific structural pattern occurs infinitely often.

We also generalize the structural domain and the grammatical operations themselves. Classical formal language theory operates almost exclusively over the free monoid of strings, where concatenation relies heavily on existential quantification over finite string partitions. In our generalized theory, we replace standard string concatenation with an arbitrary algebraic combinator operating over arbitrary abstract structures. We demonstrate that as long as this abstract combinator relies on a well-founded decomposition function—meaning that the evaluation of a structural element depends only on a strictly partially ordered, finite set of structural sub-components—the monotonic fixed-point semantics remains entirely intact. Consequently, the exact same mathematical equations and fixed-point theorems that define Boolean string grammars can also define Boolean tree grammars evaluated over non-associative groupoids. Furthermore, this immediately extends to Boolean graph grammars evaluated over trace monoids, offering a powerful new theoretical tool for specifying logical constraints in concurrent systems and distributed network topologies.

Towards AI-Assisted Hardware Engineering: An Overview of LLM-Driven RTL Design Frameworks

Shiva Ahir (Stony Brook University), Alex Doboli (Stony Brook University), Simona Doboli (Hofstra University)

Large Language Models (LLMs) are increasingly being explored for register-transfer-level (RTL) generation, hardware verification, debugging, and design automation tasks. Unlike traditional electronic design automation (EDA) workflows that rely heavily on formal methods and handcrafted heuristics, LLM-based systems introduce natural-language-driven hardware synthesis and reasoning capabilities. Recent work has illustrated promising results in Verilog generation, testbench synthesis, simulation-guided refinement, and retrieval-augmented hardware design. However, existing approaches remain fragmented across different architectures, benchmarks, and evaluation methodologies, while challenges such as functional correctness, silent failures, scalability, and hardware-aware reasoning remain unresolved. This paper presents a focused overview of LLM-assisted RTL generation and verification systems. We organize existing methods into direct generation, retrieval-augmented, verification-aware, and agentic multi-stage architectures, and discuss commonly used benchmarks, evaluation

metrics, and tool flows. Finally, we highlight major limitations and outline future research directions toward reliable and production-ready AI-assisted hardware design systems.

HCGA: A Hybrid Clustering-Genetic Algorithm for Optimal Multi-Constrained Team Formation

Ali Ali (University of Debrecen), Gergely Kovásznai (Eszterházy Károly Catholic University), Mohammed Abdulhadi (University of Technology), Mustafa Mezaal (Imam Jaafar Al-Sadiq University)

Optimal team formation under simultaneous skill and diversity constraints is an NP-hard combinatorial optimization problem. We introduce HCGA (Hybrid Clustering- Genetic Algorithm), a four-stage metaheuristic combining KMeans archetype discovery, multi-strategy population initialization, adaptive genetic annealing, and simulated annealing refinement with incremental evaluation. We evaluate HCGA against seven algorithms — Genetic Algorithm (GA), Simulated Annealing (SA), K-Means, Particle Swarm Optimization (PSO), a neural network scorer (NN), a greedy heuristic, and a Memetic Algorithm (MA) — on four synthetic datasets (50–1000 students), under a strictly compute-matched protocol with 30 seeds per configuration, paired statistical tests, and a component-wise ablation. HCGA significantly outperforms five of the seven baselines at every scale, but a well-budgeted SA initially outperforms all hybrids, including HCGA. A diagnostic decomposition attributes HCGA's deficit almost entirely to its refinement-stage cooling schedule rather than to its design; after a disclosed schedule correction and ablation-guided simplification, the revised method (HCGA-R) attains fitness statistically indistinguishable from the tuned SA baseline on all four datasets (largest difference 0.03%, pairedWilcoxon with Holm correction) while significantly outperforming the six other algorithms and satisfying all diversity constraints. Beyond the algorithm, our results carry a methodological warning: the apparent superiority of hybrid metaheuristics can be an artifact of unmatched evaluation budgets.

Learning Hysteretic Routing Policies for Heterogeneous Queues with Switching Costs

Mostafa Ali (Johannes Kepler University Linz), Dmitry Efrosinin (Johannes Kepler University Linz)

Routing jobs efficiently in heterogeneous server pools requires balancing queue performance against the cost of switching servers between active and idle states. Recent actor–critic methods learn effective threshold routing policies for heterogeneous queues without model knowledge, but optimise queue occupancy alone and treat switching as cost-free. We propose HACHQ, which augments the learning objective with explicit per-event switching penalties and replaces the single-threshold policy with a two-threshold hysteresis parameterisation. The dead zone between the two thresholds suppresses unnecessary switching, while the switching-cost-aware objective adapts the dead-zone width to the cost magnitude. On two-server systems, HACHQ attains near-optimal cost, within one percent of the exact hysteretic optimum computed by relative value iteration, while reducing switching. On multi-server systems the gains grow with scale, reaching up to 37% reduction in total cost and 77% reduction in switching rate as switching cost grows. We analyse the mechanism that produces these gains.

The Invisible Signature: A Stylometric and Probabilistic Approach to Detecting AI-Generated Text in Academic Thesis Introductions

Elyah Frisco Andriantsialo (University of Fianarantsoa), Volatiana Marielle Ratianantitra (University of Fianarantsoa), Thomas Mahatody (University of Fianarantsoa)

Since ChatGPT became publicly available in November 2022, the use of Large Language Models (LLMs) in academic writing has spread well beyond institutions with abundant digital resources. The University of Madagascar is no exception: conversational AI assistants, accessible via a simple smartphone, have become commonplace tools among master's and doctoral students. Yet unlike traditional plagiarism, AI-generated text is, by definition, original, which renders textual comparison software powerless. This study explores an alternative approach grounded in stylometry and probabilistic analysis. Drawing on a corpus of 600 thesis introductions, 300 written by Malagasy students between 2015 and 2019 and 300 generated by

GPT-4 on the same topics, we extract two types of markers: classical lexical indicators (TTR, POS tag distribution) and probabilistic metrics (perplexity, burstiness). Our results confirm that AI-generated texts exhibit significantly higher syntactic regularity (mean PPL = 21.6 vs. 87.3). We also uncover an under-documented signal: the frequent presence of typographic artefacts (asterisks, double dashes, Markdown icons) in 63% of generated texts, entirely absent from human writing. Finally, we discuss the risk of false positives for writers whose language of instruction is a second language.

Explainable Multimodal AI for Crop Health Monitoring: Benchmarking Large Multimodal Models Against Spectral Vegetation Indices

Sumaiya Ashfaque (University of Debrecen), Gergő Bogacsovics (Faculty of Informatics, University of Debrecen)

Satellite-based crop monitoring can support timely agricultural decisions, but vegetation-index outputs often require expert interpretation. This paper investigates whether large multimodal models can serve as explainable interpretive layers on top of Sentinel-2-derived crop-health evidence. In this research, a 60-site benchmark was constructed in Hajdu-Bihar County, Hungary, using five-year Sentinel-2 histories converted into RGB--NDVI benchmark panels and deterministic spectral-temporal reference labels. Gemini 2.5 Flash, GPT-4o-mini, and a consensus ensemble were compared with NDVI-based rules, Random Forest, and lightweight convolutional neural network baselines. According to our experimental results, the Gemini-OpenAI consensus achieved the highest multimodal anomaly accuracy of 82.61%, while the Random Forest baseline achieved the strongest binary anomaly result with 89.55% accuracy and 93.17% F1-score. The results suggest that multimodal models are useful for interpretable crop health analysis, while classical machine learning methods remain strong for structured binary classification.

Development and Implementation of a Blockchain-Integrated ABAC Framework for Secure Multi-Tenant Cloud Environments

Kris C. Bahaya (University of Southeastern Philippines Obrero), Ariel Roy L. Reyes (University of Southeastern Philippines Obrero)

The multi-tenant cloud computing structure can be vulnerable to a range of well-known security problems, including unauthorized cross-tenant access and overprivileged identities. This research introduced a hybrid Attribute-Based Access Control (ABAC) framework that employs the Polygon blockchain to enhance digital civil registry security within multi-tenant cloud computing. It is a seven-layered defense-in-depth framework that uses Twofish encryption and SHA-256 hashing as part of the overall solution. The results of the completed experiments confirm a 100% success rate in achieving tenant isolation and preventing unauthorized access. The use of the blockchain consensus introduced an average latency between 6.09 and 9.88 seconds. By limiting on-chain transactions to security-critical transactions, high throughputs were achieved. Overall, this framework represented an inexpensive and secure means of ensuring data protection in LGU services, with transaction costs.

The Limits of Localized T30 Prediction in Rectangular Rooms: A Residual Architecture Study

Bence Bakos (Eötvös Loránd University), Csaba Huszty (ENTEL Engineering Research & Consulting Ltd.), András Lukács (Eötvös Loránd University)

In architectural acoustic design, estimating room acoustic descriptors such as reverberation time (T30) is crucial for assessing auditory quality. While an increasing number of neural network approaches generally yield promising results in acoustic descriptor estimation, they typically estimate a single, room-averaged value. Predicting location-dependent variations of sound fields remains a significant mathematical challenge. Standard spatial architectures mapping non-sequential global layouts to continuous 3D coordinates frequently fail or suffer from mean-collapse, where the network lazily predicts the global room average across all gridpoints. To address this, we propose a two-stage residual prediction pipeline. We evaluate two coordinate-mapping architectures: a Factorized Hypernetwork and a Dense Cross-Attention Head, both utilizing Fourier Feature Mapping (FFM). Our Oracle analysis reveals that T30 variance in

rectangular rooms is overwhelmingly global, with the true room mean achieving a baseline where ~87% of predictions fall within tight error margins. Despite this physical uniformity, our residual architectures are capable of mapping localized spatial variance, partially overcoming mean-collapse and pushing Oracle accuracy to ~88.4%. However, a residual-map analysis over the full validation set shows this recovered variance to be only partially correct ($r = 0.22$): the spatial head learns almost exclusively the radial, distance-to-source component of the field, leaving its geometry-driven directional structure mainly unmodeled. Escaping mean-collapse is therefore necessary but not sufficient for spatial fidelity. Dynamic spatial mapping is architecturally viable for high-resolution acoustics, but its practical benefit for T30 remains bounded by the near-diffuse behavior of rectangular rooms.

An Energy-Aware Hybrid GA-PSO Routing Protocol for Solar-Harvesting Wireless Sensor Networks in Smart Cities

Mohammed Bamatraf (University of Kufa, College of Computer Science and Mathematics), Mohammed Nsaif (University of Kufa, College of Computer Science and Mathematics)

Energy-harvesting wireless sensor networks (EHWSNs) play a critical role in providing a sustainable infrastructure in smart cities. Nevertheless, ensuring continuous operation is still difficult owing to the intermittent nature of solar irradiance and suboptimal routing structures. Existing schemes like MLEACH make use of probability-based clustering and hence result in large disparities in energy consumption, leading to premature failure of some nodes. To solve this problem, this work proposes a new energy-aware hybrid routing scheme that makes use of the combination of Genetic Algorithms (GA) and Particle Swarm Optimization (PSO). Herein, GA explores the search space to obtain an elite set of diverse routing schemes whereas the PSO is responsible for fine-tuning the obtained routing schemes to ensure minimal delay in communication and maximization of the use of the solar energy. The performance of the proposed scheme was analyzed based on hourly solar irradiance values collected from the NSRDB of 2019 in Iraq for 90 days of winter season. From a set of 50 independent Monte Carlo simulations, the proposed hybrid scheme increases the lifetime of the system by 2.13× over that of Pure GA and by 1.89× over that of MLEACH. In addition, the proposed hybrid protocol ensures nearperfect packet delivery rate along with higher energy efficiency, resulting in more than 95% node survival at the end of the simulation period.

Predictive Modeling of Misinformation Diffusion Patterns on Facebook and X in the 2025 Philippine Midterm Elections

Carmela Lei M. Bañaga (Mapua University Makati), Niña Alexa A. Dinglasan (Mapua University Makati), Ysobella C. Villariba (Mapua University Makati), Richard N. Monreal (Mapua University Makati), Florenzo Isaac E. Romance (Mapua University Makati)

Social media platforms such as Facebook and X play a central role in political information dissemination in the Philippines, particularly during election periods when the rapid spread of misinformation can influence voter perception and democratic processes. This study examines the structural diffusion patterns of verified election-related misinformation during the 2025 Philippine midterm elections and develops predictive models for early classification of amplification behavior. 65 misinformation cases, involving approximately 20,000 posts, were quantified using amplification metrics, including size, depth, width, reach, and speed. Unsupervised clustering methods—K-Means++, Ward's Linkage, and Average Linkage—were applied to standardized structural features, with validation metrics supporting a five-cluster solution representing distinct diffusion archetypes: low-amplification, layered constrained, deep high-intensity, broad shallow, and temporal resurfaced cascades. For predictive modeling, the generated clusters were used to train Logistic Regression, Random Forest, and Support Vector Machine classifiers following feature selection via the Boruta algorithm, which identified size, reach, depth, and width as key predictors. Stratified 5-fold cross-validation demonstrated strong internal consistency, and evaluation on an unseen test set yielded 92.86% accuracy and a macro-averaged F1-score of 0.8939 across models. Logistic Regression was selected as the final model due to its interpretability and comparable performance. The findings indicate that misinformation diffusion patterns exhibit structurally predictable characteristics, suggesting that early detection based on cascade geometry is feasible. This integrated framework, combining network analysis, clustering, and supervised learning, provides actionable insights for election regulators, fact-checking organizations, and social media platforms to support timely intervention during high-stakes democratic events.

From Turbulence to Take-off: Modernising an airline's Digital Platform with Domain-Driven Design

Tamás Bartók (EPAM Systems), Ferenc Héjja (Doctoral School of Informatics, University of Debrecen, EPAM Systems), Attila Szabó (Wizz Air Ltd.)

The digital platform — serving millions of passengers across web and mobile — was held back by a tightly-coupled architecture: page loads exceeding 25 seconds, frequent outages caused by a third-party component, and Backend-for-Frontend components that had grown into sprawling monoliths with hidden cross-dependencies. With the airline's growth strategy demanding a step-change in digital capability, incremental fixes were no longer an option.

This talk tells the story of the airline's Architecture Modernisation Program — a ground-up rethinking of an e-commerce platform that handles flight search, booking, payment, ancillaries, and disruption management at scale. We will walk through how a small architecture team applied Domain-Driven Design to carve eleven clearly bounded domains (Search, Order, Offer, Pricing, Payment, Ancillary, Seating, Customer, Agency, Discount Programs, and Flight Exception Handling) out of a system where everything was tangled with everything else.

Attendees will learn how we modelled availability mathematically on the critical path, chose Azure-primary and AWS-secondary multi-cloud topology, and designed a BFF and API Gateway layer to replace duplicated business logic without a big-bang rewrite. We will be honest about the hard parts: aligning product owners on bounded context boundaries, sequencing domain migrations to avoid circular dependencies, and building a governance model that keeps an agile delivery team anchored to architecture principles.

The session concludes with concrete takeaways: a practical approach to domain readiness scoring, a phased migration roadmap template, and the lessons learned from turning architecture blueprints into funded implementation plans — leaving attendees with a reusable playbook for large-scale platform modernisation in a high-traffic, compliance-heavy environment.

Reaction systems and approximation spaces

Péter Battyányi (Faculty of Informatics, University of Debrecen)

In this talk, we elaborate on the connection between reaction systems and computation with rough sets. Using rough-set approximations, we define the definite and probable results of a reaction system. Among other results, we show that computations based on these

approximations may lead to outcomes that differ significantly from the expected ones, even in the case of computations of length two.

Similarity-Aware BERT Model for Drug-Target Interaction Prediction

Dylan Bock (University of Northern Iowa), Sinisa Grozdanic (Animal Eye Consultants of Iowa), Vladimir Jacimovic (University of Montenegro), Aleksandar Poleksic (University of Northern Iowa)

Modern drug discovery relies on accurate computational methods for predicting drug-target interactions. When trained on hundreds of millions of protein sequences, transformer-based methods such as ProteinBERT demonstrate high accuracy in predicting protein function and structural properties, including disordered regions and secondary structure. However, these models are limited in their ability to accurately predict associations between proteins and other biomedical entities, such as drugs. This is in part due to the lack of explicit regularization that incorporates biological and chemical similarity structures. In this study, we propose a model, SimBERT-DTI, that incorporates sequence similarities among proteins and structural similarities among drugs to enhance predictions of drug-target interactions (DTIs). Our results demonstrate that this general framework can be utilized to develop more accurate and robust DTI prediction methods.

Explainability and Verifiability of Artificial Intelligence in ICU Mortality Prediction

Yuqi Chen (Eszterhazy Karoly Catholic University), Gergely Kovásznai (Eszterhazy Karoly Catholic University)

High area under the receiver operating characteristic curve (AUROC) alone is insufficient for assessing intensive care unit (ICU) mortality models, because calibration, feature behavior, explanation stability, and external generalizability also affect trustworthiness. We present a verification and explainability framework combining an XGBoost–multilayer perceptron (MLP) soft-voting ensemble trained on the MIMIC-IV dataset, five verification experiments, and a Streamlit-based decisionsupport prototype. The framework identifies monotonicity and calibration problems and evaluates segmented-risk A/B/C monotonicity diagnosis, component-

level Platt scaling, and external validation on the eICU-CRD dataset. The ensemble achieved AUROC 0.888 on MIMIC-IV and 0.856 on eICU-CRD, while Platt scaling reduced the expected calibration error (ECE) from 0.090 to 0.006 internally and from 0.167 to 0.045 externally.

Multi-Objective Fine-Tuning Optimization for Pneumonia Detection Using Evolutionary Search and Successive Halving

Radu-Ionuț Costin (Technical University of Cluj-Napoca, North University Center at Baia Mare), Ștefan Vasile Oniga (Technical University of Cluj-Napoca, North University Center at Baia Mare), Ramona-Iulia Babiciu (Technical University of Cluj-Napoca, North University Center at Baia Mare)

This study investigates resource-aware validationstage selection of fine-tuning configurations for chest X-ray pneumonia classification. Rather than proposing a new convolutional neural network (CNN) architecture, it combines the Non-dominated Sorting Genetic Algorithm II (NSGA-II) with Successive Halving (SH) to search over pretrained-backbone configurations using three objectives: validation area under the receiver operating characteristic curve (AUC), search-stage training-and-validation time, and trainable parameter count. On the cleaned balanced chest X-ray dataset, NSGA-II/SH was useful mainly for resource-aware configuration selection during development. In validation logs, it found the strongest observed maximum-budget validation AUC among the evaluated methods and reduced mean training-and-validation time per successful backbone-level evaluation by 68.6% relative to the manually tuned reference configurations; estimated search cost for successful evaluations is reported separately in GPU-hours. In held-out testing across three final seeds, the frozen featureextractor baseline achieved the best predictive performance, whereas the final model selected by NSGA-II/SH achieved the lowest mean final retraining/evaluation runtime. Deploymentoriented cost was reported separately using model size, floatingpoint operations (FLOPs), GPU memory, inference latency, and throughput. Overall, NSGA-II/SH was most useful for selecting compact and faster-to-train fine-tuning configurations, while the strongest held-out predictive performance was obtained by the frozen feature-extractor baseline.

Cross-Dataset Evaluation of Foundation Model-Based Acoustic-Text Fusion for Dementia-Related Speech Analysis

Seyma Derdiyok (Yıldız Technical University), Imran Gul (Yıldız Technical University), Hamza Osman Ilhan (Yıldız Technical University), Gorkem Serbes (Yıldız Technical University)

This study examines the cross-dataset generalizability of a foundation-model-based acoustic–text framework for dementia-related speech analysis across three datasets: ADReSSo2021, Ivanova, and Delaware. The proposed pipeline combines WavLM-Large acoustic embeddings with Qwen3-ASR-1.7B transcription and Qwen3-Embedding-0.6B text representations, integrated through validation-guided weighted late fusion. Results demonstrate a dataset- and task-dependent contribution of multimodal fusion: fusion yielded the highest F1-scores on ADReSSo2021 (90.18%) and Delaware (62.12%), whereas acoustic representations showed greater prominence across the Ivanova tasks. Label-composition analyses further revealed clearer HC–AD than HC–MCI separability, highlighting the potential role of MCI heterogeneity in classification difficulty. Chunk-length ablation across 5, 10, 15, and 30 s revealed task-specific preferences for temporal context, with the best performance occurring at different segment durations across tasks. These findings emphasize the importance of speech protocol, diagnostic composition, dataset characteristics, and temporal context when evaluating acoustic–text integration for dementia-related speech analysis.

Machine Learning-Based Molecular Classification of Endometrial Cancer Using Clinicopathological and Immunohistochemical Features: A Chinese Population Study

Alshuwaili Dhafer G. Honi (Doctoral School of Informatics, University of Debrecen), Hakim Bahlok Jebur (Institute of Clinical Pathophysiology, Semmelweis University), László Szathmáry (Faculty of Informatics, University of Debrecen)

The Cancer Genome Atlas (TCGA) molecular classification of endometrial cancer (EC) into POLEmut, MMRd, p53abn, and NSMP subtypes has substantially improved prognostic stratification. However, molecular classification typically depends on sequencing, whereas surrogate clinical workflows combine immunohistochemistry (IHC) with targeted POLE

testing, limiting accessibility in many clinical settings. This study developed, evaluated, and interpreted machine learning models for automated EC molecular subtype classification using routinely available clinicopathological and IHC features, and assessed the relative contributions of molecular and clinical predictors through a controlled feature-ablation analysis. A cohort of 229 Chinese patients with EC was analyzed. Twenty-one features were engineered from clinical and molecular variables, and six ML algorithms (Random Forest, Gradient Boosting, Logistic Regression, Support Vector Machine, Naive Bayes, and Multi-Layer Perceptron) were evaluated using both 5-fold and 10-fold stratified cross-validation. Class imbalance was addressed using SMOTE applied exclusively within the training folds. On the full feature set, Gradient Boosting achieved 96.1% accuracy in 5-fold cross-validation and 97.0% in 10-fold cross-validation (macro F1 = 97.0%), exceeding the performance of a transparent rulebased classifier (93.9% accuracy). Feature importance analysis identified POLE (26.7%), TP53 (20.5%), and mismatch-repair (MMR) markers (20.5% combined) as the dominant contributors to classification. Feature-ablation experiments showed that models trained using only molecular features maintained performance comparable to the full-feature models, whereas removing molecular markers reduced accuracy to 52–60% (macro F1 $\leq$ 0.44), approaching the majority-class baseline (63.8%). These findings indicate that the predictive performance is primarily driven by molecular biomarkers, while routine clinicopathological variables alone provide limited discriminatory power. The proposed framework provides an accurate, automated, and interpretable approach for implementing IHC-based molecular classification of endometrial cancer in routine clinical practice.

Deepfake Image Detection Using Spatial Rich Model Features and Deep Learning Fusion

Adnane Essoufi (University of Debrecen), Gergő Bogacsovics (Faculty of Informatics, University of Debrecen)

We present a spatial rich model (SRM)-based fusion approach for accurate and reliable deepfake image detection in a binary real-versus-fake facial image classification setting. The proposed method combines the effectiveness of modern deep learning models with residual-oriented forensic information extracted through SRM-related processing. First, we evaluate multiple baseline architectures, including EfficientNetV2-S, ResNeXt50, ViT-B/16, and a lightweight custom CNN, under the same validation protocol. Then, we select EfficientNetV2-

S as the main RGB backbone and extend it with an additional SRM-related branch, fused into the network at different depths. We show that the proposed SRM-based fusion models outperform the strongest RGB-only baseline. Finally, we further improve the results by combining the strongest fusion variants with selected baseline models using soft voting. The proposed ensemble achieves 97.48% accuracy, 96.46% precision, 98.57% recall, and 97.50% F1-score on the global test set, showing that SRM-related forensic cues can effectively complement deep visual representations for deepfake image detection.

Empirical Analysis of Swarm Intelligence Convergence Curves in WSN Energy Optimization

Levente Filep (Doctoral School of Informatics, University of Debrecen), Zoltán Gál (Faculty of Informatics, University of Debrecen)

The convergence behavior of swarm intelligence algorithms is commonly represented through best-fitness-versus-iteration curves, yet the empirical shape of these curves has received comparatively little attention in the literature. In this paper, we investigate whether such curves can be approximated by two simple parametric decay models: a shifted exponential and a shifted inverse power law. The models are fitted using nonlinear least squares to the convergence curves of several well-known swarm intelligence algorithms applied to a wireless sensor network transmission-power minimization problem, and the resulting fit quality is evaluated using the coefficient of determination. Since raw convergence curves typically contain an initial transient phase, a steady decay region, and a terminal stall region that cannot be accurately described by a single decay law simultaneously, we additionally introduce preprocessing methods for trimming these regions. Two methods are proposed for removing the leading transient: a geometric parameter-free approach and a method based on a running comparison of iteration-to-iteration improvements. A sliding-window procedure is further used to identify and remove the terminal stall region. Four preprocessing configurations are evaluated across all algorithms, providing an empirical characterization of how well their convergence behavior follows either exponential or power-law decay. Across the algorithms tested, both models yield good fits after preprocessing, with the preferred model being algorithm-dependent.

Hyperparameter-Aware Evaluation of Deep Learning Architectures for AI-Generated Image Detection

Ilham Gafarov (University of Debrecen), Zeyad Qasem (University of Debrecen), Kerem Duzenli (Doctoral School of Informatics, University of Debrecen), Balázs Harangi (Faculty of Informatics, University of Debrecen), Mokhaled Al-Hamadani (Doctoral School of Informatics, University of Debrecen)

The rapid advancement of generative artificial intelligence has raised growing concerns regarding the authenticity and reliability of digital visual content. Reliable detection of AI-generated images has therefore become an important research problem in digital forensics, media verification, and content integrity. This study evaluates four deep learning architectures—ResNet-50, EfficientNetV2-S, ViT-B/16, and Xception—for binary classification of real and AI-generated images using the CIFAKE dataset. Each model is assessed under baseline and hyperparameter-optimized configurations to examine the effects of learning rate, batch size, optimizer selection, weight decay, data augmentation, and backbone-freezing strategy. The results show that hyperparameter tuning improves all evaluated models, although the magnitude of improvement is architecture-dependent. EfficientNetV2-S and ViT-B/16 benefit most from tuning, while ResNet-50 and Xception show more moderate gains. Among individual models, the tuned ViT-B/16 achieves the highest accuracy and F1-score, whereas the tuned weighted soft-voting ensemble obtains the best overall performance, reaching 99.07% accuracy and 99.07% F1-score on the test set. These findings provide practical guidance for selecting and tuning deep learning models for synthetic image detection and highlight the need for further evaluation on higher-resolution datasets and diverse generative models.

FAME: From Adversarial Feature Vectors to Malware Evasion via SMT

Endre Glocker (Budapest University of Technology and Economics), József Sándor (Budapest University of Technology and Economics), Roland Nagy (Budapest University of Technology and Economics)

Internet of Things (IoT) malware detection remains a challenging task due to the severe resource constraints of IoT devices, which require lightweight yet effective security solutions.

SIMBIoTA-ML addresses this challenge by leveraging machine learning (ML) classifiers operating on TLSH-based static features for efficient malware detection. However, like other ML-based systems, SIMBIoTA-ML is vulnerable to adversarial evasion attacks. While prior work investigated this problem through ad-hoc binary modifications, it did not provide a systematic or guaranteed approach for generating adversarial examples.

In this paper, we propose FAME (From Adversarial feature vectors to Malware Evasion), a framework for generating adversarial malware samples against SIMBIoTA-ML. FAME applies established ML evasion attacks in the feature space and addresses the inverse feature-mapping problem between continuous adversarial feature vectors and corresponding executable binaries. To this end, we formulate the discretization of continuous adversarial feature vectors as a Satisfiability Modulo Theories (SMT) problem, enforcing both TLSH validity constraints and SIMBIoTA-ML evasion constraints. Using the target TLSH hash determined by the SMT solver, the original malware binary is then modified accordingly. By combining adversarial ML, constraint solving, and controlled binary modification, FAME enables the scalable and reliable generation of executable adversarial malware samples.

Precedence-Constrained Treatment Sequence Scheduling in Outpatient Healthcare Using a Heuristic Genetic Algorithm

Veronika Gombás (Department of Computer Science and Systems Technology, University of Pannonia), Ágnes Vathy-Fogarassy (Department of Computer Science and Systems Technology, University of Pannonia)

Background: Outpatient care often involves ordered sequences of clinical tasks rather than a single appointment. Scheduling such sequences is challenging because resource availability, occupied calendars, treatment durations, and precedence constraints must be considered simultaneously. Objectives: This paper proposes a heuristic genetic algorithm for precedence-constrained outpatient treatment sequence scheduling. Two healthcare-relevant objectives are considered: minimizing the completion time of the whole sequence and minimizing the time span between the first and last clinical tasks. Methods: Each individual encodes a complete clinical task sequence as calendar–time-slot assignments. The baseline genetic algorithm is extended with precedence-aware penalties and heuristics for preserving or repairing clinical task order during crossover and mutation. The method was evaluated on four synthetic outpatient calendar scenarios with different clinical task sets, calendar horizons, occupancy

levels, and task orders. Results: The results show that feasibility strongly depends on scenario complexity, clinical task order, and the applied precedence-guided heuristic. In the more complex scenarios, precedence-guided weighted mutation-based heuristics produced a higher proportion of feasible schedules than the baseline genetic algorithm. Conclusion: The proposed method provides a suitable framework for outpatient clinical task sequence scheduling with ordering constraints and shows that targeted heuristics are essential for handling more complex calendar structures.

Generative AI in Higher Education: A Systematic Literature Review of Stakeholder Perspectives, Uses, Risks, and Governance

Deddy Gusman (Doctoral School of Informatics, University of Debrecen), David Aditya Sulistiawan (University of Debrecen), Nurmalina (University of Pahlawan Tuanku Tambusai), Gyöngyi Bujdosó (Faculty of Informatics, University of Debrecen)

Generative artificial intelligence (GenAI) has become a major teaching, assessment, and governance issue in higher education. This systematic literature review synthesizes 17 peer-reviewed studies published between 2023 and 2025 and screened within a review window from January 1, 2023 to March 8, 2026. Guided by PRISMA 2020 principles, the review examines educator perspectives on GenAI use, perceived benefits, assessment risks, AI literacy needs, and governance responses, while using student and institutional evidence as contextual signals. The evidence base is intentionally educator-centered and dominated by early empirical, cross-sectional, and self-report studies. The synthesis shows that GenAI supports resource generation, feedback, explanations, planning, and workload reduction, but also raises concerns about assessment validity, academic integrity, unreliable outputs, over-reliance, uneven preparedness, and limited policy support. The review contributes an educator centered readiness and governance response framework linking pedagogical judgement, assessment redesign, AI literacy, disclosure practices, professional development, and policy coherence. It concludes that responsible GenAI adoption should be understood as supervised pedagogical augmentation supported by educator readiness and institutional governance alignment.

A Recency-Weighted Monte Carlo Engine for Possession-Level Basketball Forecasting

Lóránd Heidrich (Eszterházy Károly Catholic University, Faculty of Informatics, University of Debrecen), Ádám Kovács (Institute of Mathematics and Informatics of Eszterházy Károly Catholic University)

This paper presents a modular Monte Carlo (MC) engine for retrospective forecasting experiments on historical National Basketball Association games. Instead of fitting a single discriminative winner classifier, the system estimates player event rates from one-year individual box-score histories and simulates games as sequences of possession-level events. The study states the simulation assumptions formally and evaluates the engine with a season-wide benchmark. The framework is mechanistic and modular: a rolling box-score estimator, position-aware role sampling, explicit possession transitions, and separable calibration produce full score distributions and permit controlled what-if experiments. On a matched set of 1106 games from the 1991– 1992 regular season, the MC estimate averaged over three runs reaches an area under the receiver operating characteristic curve (AUC) of 0.677 after tie normalization. Averaging per-game probabilities across three independent 1000-epoch exports and applying a training-season home-prior logit calibration yields 0.666 accuracy, Brier score 0.216, log loss 0.623, and AUC 0.677. Training-tuned Elo and regularized logistic regression provide stronger binary forecasting metrics, whereas the MC engine provides complementary process-level transparency and distributional outputs, including score distributions, modal scores, tie rates, plausible score ranges, and uncertainty summaries.

A Unified Agentic Development Lifecycle Framework for the Software Industry

Ferenc Héjja (Doctoral School of Informatics, University of Debrecen), Tamás Bartók (EPAM Systems, Institute of Informatics, University of Szeged), Gergely Kocsis (Faculty of Informatics, University of Debrecen)

The adoption of generative and agentic artificial intelligence is reshaping software engineering by enabling autonomous or semi-autonomous agents to plan, generate code, use tools, test software, and interact with development environments. However, traditional Software Development Lifecycle models and AI-assisted software development practices do not fully

address the lifecycle needs of agentic software systems. This paper proposes a Unified Agentic Development Lifecycle framework for the software industry. The framework organizes agentic software engineering into Design, Build, and Operate phases, supported by a continuous control plane for evaluation, monitoring, governance and security. The paper provides an industry-oriented taxonomy, lifecycle meta-model, phase mapping and maturity alignment for software organizations adopting agentic AI. The proposed framework is evaluated conceptually through comparative mapping and is positioned as a foundation for future empirical validation in industrial software engineering settings.

EnergyEquity: A Multi-Agent Conversational AI for Personalized Household Energy Poverty Assessment in Hungary

Dat Tan Ho (University of Debrecen), András Hajdu (Faculty of Informatics, University of Debrecen), Tamás Márton Bérczes (Faculty of Informatics, University of Debrecen)

Energy poverty affects 6–18 % of Hungarian households depending on the indicator used, yet existing tools target municipalities and utilities rather than homeowners. We present EnergyEquity, a four-agent (Profiler, Estimator, Explainer, Advisor) conversational system whose central reliability mechanism ties every quantitative span in the agent's prose to a tool output within ±0.5 %: 98.4 % of 68,956 spans across 300 generated advisories meet this bound. Demand is predicted by an XGBoost regressor over 2,596 deduplicated TABULA archetypes (R2 = 0.940 held-out) with TreeSHAP attribution. On 100 synthetic scenarios at three seeds, Qwen3.5-27B reaches 98.8 % / 94.5 % on a self-designed 18+10-criterion rubric. A paired Qwen3.5-9B trades 7.8 pp accuracy for a 3.6× latency reduction (55.7 vs 199.5 s/cell). EnergyEquity shows that an open-weight multi-agent pipeline with auditable tool grounding can deliver decree-compliant household-level energy advisory under fully local inference.

AI-Enhanced Grouped Time-Series Forecasting with Weather and Holiday Effects

Dat Tan Ho (University of Debrecen), Man Duc Nguyen (Duy Tan University), Tamás Márton Bérczes (Faculty of Informatics, University of Debrecen)

Operational wholesale forecasting depends on two levers: external drivers such as weather and holidays, and hierarchical coherence across customer groups. Their interaction with modern deep sequence models on multi-year business-to-business data is rarely measured with formal uncertainty quantification. We benchmark four model families (error–trend–seasonal exponential smoothing (ETS), SARIMA, the Temporal Fusion Transformer (TFT), and the Time-series Dense Encoder (TiDE)) on 179 weeks from a Hungarian wholesale distributor (~5.1 M records, 93,287 modelled customer–item series), under a matched rolling-origin protocol with eight non-overlapping 12-week windows, and compare Bottom-Up against Top-Down reconciliation. Pooled across the windows, TFT significantly outperforms ETS, SARIMA, and TiDE at every horizon (group-clustered paired tests, $p \leq 0.007$) and attains the best macro-mean MAE and WAPE, though individual-window means can favour the statistical baselines. TFT also beats a covariate-equipped LightGBM reference ($p \leq 0.003$), while the same covariates significantly worsen a per-series SARIMAX, and all families clear a seasonal-naive floor by a wide margin. Covariate effects are marginal on average (at most ≈0.6 % of MAE, even under perfect prognosis), yet regime-conditional: corrected holiday features cut group-total error by 5 % on major-holiday peak weeks. Top-Down reconciliation degrades series MAE by 35–60 %, so Bottom-Up is the default, under which the statistical baselines give the most accurate group totals in most windows. Per-group difficulty spans an order of magnitude, so TFT is the strongest global default but monitoring stays segment-aware.

Reliability and Failure Assessment of 2D Pose Estimation Methods for Drone-Based Swimmer Video Analytics

Dominik Homoki (Pázmány Péter Catholic University, Faculty of Information Technology and Bionics), Tamás Zsedrovits (Pázmány Péter Catholic University, Faculty of Information Technology and Bionics), György Cserey (Pázmány Péter Catholic University, Faculty of Information Technology and Bionics), András Horváth (Pázmány Péter Catholic University, Faculty of Information Technology and Bionics), László Grand (Pázmány Péter Catholic University, Faculty of Information Technology and Bionics)

We evaluate four off-the-shelf 2D pose estimation models (ViTPose-H, RTMPose-L, YOLO11x-Pose, and MediaPipe Pose) on drone-recorded indoor swimming footage annotated with COCO-17 keypoints across five motion phases. Using three swimmer sequences totalling 300 ground-truth-annotated frames, we measure accuracy via PCK@0.05 on keypoints marked as visible (v=2) and OKS over annotated keypoints under the COCO-style visibility convention and reliability via keypoint completeness and detection failure rate. ViTPose-H leads overall (PCK@0.05: 0.780, OKS: 0.648, completeness at τ =0.5: 0.613), while RTMPose-L achieves near-equivalent accuracy (PCK@0.05: 0.754) with less than half the inference time. Turn execution is a consistent failure mode across all four models: pooled PCK@0.05 falls by 51–85 % relative to normal strokes, and completeness falls near zero in every case. A joint-level analysis identifies lower-body keypoints (knees and ankles) as the primary source of reliability loss, driven by submersion during the flip turn. These results quantify the domain gap between COCO-trained models and aquatic drone deployment, and indicate the need for targeted fine-tuning for swimming-specific applications.

A Descriptive Framework for Archival-Source-Based 3D Architectural Modeling

András Horkai (Obuda University Ybl Miklós Faculty of Architecture and Civil Engineering)

In the digital reconstruction of historic buildings, the quality and completeness of available archival sources significantly influence the accuracy and reliability of the resulting 3D models. In archive-based modeling, the data used is heterogeneous: architectural drawings, historical photographs, textual documents, analog examples, and schematic representations may all be

present. The diversity of sources and the building's physical status lack a uniform classification framework that is practical. This study aims to introduce a five-category classification system that categorizes digital reconstruction scenarios based on the type of archival sources and the physical status of the building. The study proposes a structured project description framework for documenting and comparing projects. The results highlight that systematizing reconstruction scenarios enhances transparency in modeling decisions and supports data-driven interpretation and comparability of 3D models. The approach aligns with data-driven and information-oriented modeling perspectives.

Paradata in Virtual Reconstruction: Towards a Conceptual Model of Transparent Reconstruction

András Horkai (Obuda University Ybl Miklós Faculty of Architecture and Civil Engineering)

Virtual reconstruction has become a central method in digital heritage research, particularly for documenting destroyed, altered, or unrealized architectural heritage. However, reconstruction workflows remain inherently interpretative, especially in archive-based modeling, where incomplete or contradictory sources require inferential decisions. This review examines how recent literature conceptualizes and operationalizes paradata as a mechanism for documenting interpretative reasoning, provenance, and uncertainty in digital heritage workflows. Based on the thematic synthesis, the paper further proposes a conceptual model that integrates the identified research directions into a unified representation linking conceptual principles, reconstruction workflows, and implementation technologies. Using a focused qualitative thematic literature review, the paper analyses recurring themes, including the distinction between metadata and paradata, semantic provenance modeling, CIDOC-CRMbased documentation, HBIM interoperability, FAIR data principles, explainable reconstruction workflows, and the longterm preservation of born-digital heritage assets. The review identifies a consistent shift from purely visual reconstruction toward machine-readable transparency frameworks that connect sources, interpretative assumptions, workflows, and resulting visualizations. At the same time, the literature reveals persistent methodological fragmentation, limited standardization, and insufficient integration of uncertainty documentation into operational reconstruction pipelines. Recent studies increasingly emphasize semantic interoperability, knowledge graphs, provenance ontologies, and lightweight paradata schemas

that support reuse, scholarly validation, and digital preservation. The proposed three-layer conceptual model extends the literature by integrating previously fragmented perspectives into a coherent structure that may support future archive-based virtual reconstruction projects and interoperable paradata implementations.

A Canonical Rigid Direct-System Representation of Finite Local-Unit-Aligned Totally Ordered Monoids

Sándor Jenei (Eszterházy Károly Catholic University and University of Pécs)

We study finite local-unit-aligned totally ordered monoids, meaning finite totally ordered monoids in which the greatest right and left local units of each element coincide. The paper proves that every such monoid has a canonical rigid representation as a finite chain-indexed direct system of component monoids, and that every rigid direct system of the corresponding kind reconstructs such a monoid. The representation is induced intrinsically by the local-unit map τ, through the finest τ-multiplication-coherent partition of the positive idempotent skeleton. The resulting component monoids, transition maps, directed lexicographic order, and product formula recover both the order-theoretic and algebraic structure.

Efficient Scheduling of Mixed-Criticality Pipeline Workflows with Imprecise Computations in Distributed Environments

Helen D. Karatza (Aristotle University of Thessaloniki)

In distributed systems, workloads consist of jobs with heterogeneous criticality levels. Higher-criticality jobs are prioritized over lower-criticality ones and impose stricter constraints on execution accuracy, tolerating less imprecision. This heterogeneity complicates scheduling, particularly in ensuring low latency for critical jobs while maintaining the required Quality of Service (QoS). To address this, controlled imprecise computation is employed, allowing bounded relaxation of accuracy under specific conditions. Jobs are modeled as pipeline workflows executed on distributed processors. We adopt periodic scheduling, where both the acceptable imprecision and waiting-time tolerance are adapted according to job criticality.

Simulation results demonstrate that enabling partial computation for PWs within the periodic scheduling framework improves system performance across diverse load conditions.

A Void-Aware Depth-Based Routing Protocol for Deep-Ocean Life Extension in Underwater WSNs

Wajih Khelifi (Doctoral School of Informatics, University of Debrecen), Zoltán Gál (Faculty of Informatics, University of Debrecen)

Underwater Wireless Sensor Networks (UWSNs) are essential for deep-sea monitoring, but their performance is limited by severe energy constraints, high acoustic attenuation, and dynamic 3D topologies. Existing routing solutions often fail to sustain bottom-layer source nodes and lack adaptive mechanisms to bypass communication "dead zones", leading to premature network failure and catastrophic packet loss. To address these fundamental gaps, this paper proposes the Void-aware Deep-Ocean Life Extension Depth-Based Routing (VDOLE_DBR) algorithm. The algorithm introduces density-adaptive volume partitioning and an energy-balancing cost function to dynamically optimize cluster head selection. Furthermore, VDOLE_DBR employs a multi-hop vertical pipeline strategy that actively circumvents communication voids to ensure continuous data delivery to the surface. Comprehensive simulations validate that VDOLE_DBR significantly outperforms existing protocols, including LEACH, EEUMC, and GSA, by substantially delaying node mortality metrics (FND, HND) and extending the functional monitoring lifespan. Additionally, the proposed algorithm maintains a dominant and stable Cumulative Packet Delivery Ratio (PDR) across varying depths and node densities, demonstrating robust scalability in highly constrained subsea environments.

Short-Term Traffic Congestion Prediction Using Public Transportation Data

Robert-Lucian Kis (University of Medicine, Pharmacy, Science, and Technology of Targu-Mures), Adrian-Silviu Roman (University of Medicine, Pharmacy, Science, and Technology of Targu-Mures), Teri Lenard (University of Medicine, Pharmacy, Science, and Technology of Targu-Mures)

Urban traffic congestion imposes significant costs on commuters, the local economy, and the environment, and its prediction remains a challenging problem due to the complex spatiotemporal dynamics of urban traffic flow. In this paper, we present a congestion prediction approach for short term forecasting that leverages location data from public bus data in the city of Târgu Mureș, Romania. The proposed pipeline integrates Automatic Vehicle Location (AVL) data collection, map matching, Travel Time Index (TTI) estimation, and machine learning-based forecasting to deliver 15-minute-ahead congestion predictions. Ten prediction architectures of increasing complexity were evaluated, ranging from naive baselines to recurrent and graph-based deep learning models. The results demonstrate that temporal traffic history is essential for accurate forecasting, and that Graph WaveNet achieves the lowest Mean Absolute Error (MAE=0.546), outperforming models that rely on the fixed road graph structure. The Spatial-Temporal Identity (STID) model achieves superior results on RMSE, MAPE, and R2, representing an efficient alternative. The findings confirm that publicly available AVL data can serve as a reliable and cost-effective source for urban traffic monitoring in cities lacking dedicated sensor infrastructure.

Applicability of a Spectral-Geometric Transformer for Resolution-Agnostic Estimation of Global Horizontal Irradiance from Geostationary Satellite Observations

Attila Knolmajer (University of Pannonia Department of Electrical Engineering and Information Systems), Attila Fodor (University of Pannonia Department of Electrical Engineering and Information Systems), Ágnes Vathy-Fogarassy (University of Pannonia Department of Computer Science and Systems Technology)

This paper investigates the applicability and transferability of a resolution-agnostic spectral–geometric transformer model for estimating surface global horizontal irradiance from

geostationary satellite observations. The model represents geolocation-specific original SEVIRI channel values, derived spectral features, and auxiliary solar-geometric, seasonal, and topographic descriptors as scalar tokens, which are processed by a transformer encoder without spatial positional encoding. Instead of learning image patterns tied to a fixed satellite resolution, the model learns nonlinear relationships among spectral and geometric variables and predicts the clear-sky index, from which global horizontal irradiance is reconstructed. The model was trained using four years of Meteosat Second Generation (MSG2) Spinning Enhanced Visible and Infrared Imager (SEVIRI) observations and ground measurements from Hungarian stations, and was evaluated on independent stations, different temporal aggregation levels, sky conditions, and geographically distant station groups. Regional validation showed high accuracy, especially after hourly and daily aggregation, while partly cloudy conditions remained the most challenging because of the scale mismatch between satellite pixels and point measurements. Transfer experiments indicated good applicability in several European stations, but also revealed systematic domain shift in South American and African or Atlantic African regions. SHapley Additive exPlanations-based (SHAP-based) interpretation confirmed the dominant role of visible spectral information and the complementary contribution of infrared and solar-geometric variables. The proposed model should therefore be interpreted as a modular radiativeattenuation estimator rather than a standalone future forecasting system; in forecasting applications, it can be coupled with an upstream model that predicts future satellite spectral states.

Comparing evolutionary algorithms and reinforcement learning algorithms in "mate-in-N" puzzles

Kristóf Kondor (University of Debrecen), Géza Horváth (Faculty of Informatics, University of Debrecen)

This article presents an analysis of two learning-based artificial intelligence algorithms, Deep Q-Network (DQN) and Genetic Algorithm (GA), applied to solving mate-in-N chess puzzles. Both methods were implemented in Python and integrated with the Stockfish engine to ensure reliable reward calculation. Neither algorithm retains knowledge between puzzles, each task is solved with an untrained initial state. The study analyzes learning behavior, success rate, stability and the advantages of the two methods across multiple difficulty levels. Multiple performance statistics were analyzed, including win rates, reward progression, runtime and

solution accuracy. The results show that reinforcement-learning based DQN achieved higher win rates with lower overall runtime, while the evolutionary GA demonstrated the ability for higher quality solutions and greater creativity. These results show the strengths of the two approaches in structured tactical decision-making tasks.

LLM SQL Cross-Dialect Translation: An Exploratory Study of Form-Content Dialectical Thinking

Arsenij Kroptya (Admiral Makarov National University of Shipbuilding Mykolayiv, Ukraine), Ihor Mykheliev (Admiral Makarov National University of Shipbuilding Mykolayiv, Ukraine)

Recent studies show that large language models (LLMs) translate SQL across dialects with accuracy below 40% on average, which is unexpectedly low for what looks like a language manipulation task. SQL dialect translation is not only a syntactic rewriting task where syntax can be manipulated independently of content. Correct cross-dialect translation requires recognizing when identical syntax carries different meaning across systems. This is a dialectical situation in the sense of interdependence of form and content: the query remains itself and keeps semantics unchanged even when syntax changes are required. We analyze the chain-of-thought traces of reasoning LLMs on three deliberately small translation cases in which syntactic rewriting alone is likely to fail, define explicit criteria for recognizing dialectical thought forms in the traces, and add four experiments built on thought forms: scenarios in which session configuration rather than syntax or schema fixes the meaning, executable witnesses that models must construct for their own translations, a structured antithesis pass, and piecewise translation, evaluated on nine reasoning models. The study is exploratory and tests whether LLMs apply dialectical thought forms, not a new SQL translation benchmark. Newergeneration models solved every configuration scenario that older models failed. Yet only 15 of 43 model-generated witnesses could demonstrate a semantic divergence, and a post-hoc antithesis step corrected eleven translations. The safety-aligned claudefable-5 refused some translation requests outright but translated the same queries correctly when the task was framed as self-critique. We frequently observed reasoning consistent with the interdependence-of-form-and-content pattern in the models' traces, including in cases where the produced query ignored it, and therefore observe a gap between declarative recognition and operational use of knowledge.

Lightweight Visual Question Answering in Radiology Domain

Pang Lang Kueh (Universiti Malaya), Wai Lam Hoo (Universiti Malaya)

The increasing volume of imaging studies has motivated research on radiology Visual Question Answering (VQA), where systems combine image understanding and natural language processing to answer questions about medical images. This paper proposes a lightweight radiology VQA prototype based on MobileNetV3-Small and DistilBERT, connected through an attention-based fusion module and a dual-head classifier for binary and open-ended questions. The system is trained and evaluated on VQA-RAD, a clinician-curated benchmark of 315 radiological images and 2,248 question–answer pairs that span X-ray, CT, and MRI modalities. Using the official VQA-RAD test split, the lightweight model achieves an overall accuracy of 69.4%, compared with 66.3% for a ResNet-34 with BERT-base baseline, while reducing model size by 47% (269.59 MB vs. 511.69 MB). The measured mean inference time improves modestly from 16.357 ms to 15.751 ms on the tested device, indicating that the principal efficiency gain is model compactness rather than latency. The lightweight model also achieves higher recall on binary questions (78.81% vs. 65.25%), although these results remain preliminary because VQA-RAD is small and no external clinical validation is performed. A public Streamlit research demo and GitHub repository are provided to support reproducibility.

Requirements Analysis for the Development of a Mobile Application for Emergency Medication Dosage Calculation

Csaba Sándor Kuki (University of Debrecen), Balázs Palcsik (University of Debrecen), Piroska Biró (Faculty of Informatics, University of Debrecen)

Nowadays, the use of digital devices is becoming increasingly widespread and accepted. Methods based on modern technology have gained ground in many areas of life. In fact, we could say that information technology complements many fields, thereby meeting the expectations of the 21st century. One of the most important areas in this regard is healthcare, which has become deeply intertwined with IT developments in recent years. It is unimaginable that we would not provide the most up-to-date solutions to the problems arising in healthcare. Emergency patient care is one of the most critical areas in medicine. Providing care in this field presents numerous challenges. Of particular note is the calculation of medication dosages,

which typically occurs under time pressure, and calculation errors can also occur during the process. A mobile app can be a great tool for reducing cognitive load and supporting rapid decision-making. Before beginning development, it is essential to consider the perspectives of the target users so that the mobile app aligns with their expectations and needs even before it is created. We wanted to gain insight into the users' perspective through a preliminary questionnaire survey. In this article, we would like to present the direction emerging from the survey results regarding the medication dosage calculator mobile application.

LLM-Assisted SAT Solver Engineering: A Case Study on Enhancing CSFLOC

Gábor Kusper (Eszterházy Károly Catholic University)

This paper presents a case study on LLM-assisted SAT solver engineering. The central question is whether an LLM-assisted, human-supervised design process can be useful when improving a specialized SAT solver implementation. The case study focuses on CSFLOC, a SAT solver based on counting subsumed full-length ordered clauses. Starting from CSFLOC19, we used ChatGPT as an interactive engineering assistant to analyze the source code, identify promising modification points, generate instrumentation, and propose an alternative learned-clause cache policy. The LLM-generated modification, implemented in CSFLOC20, was not the best final solution, but it was better than the original baseline on larger random unsatisfiable 3-SAT instances and, more importantly, it exposed a simple engineering opportunity that had not been noticed in our previous CSFLOC development. Based on this insight, we designed CSFLOC21, a stronger human-generated cache policy inspired by the LLM-generated version. Experiments on selected SATLIB instances show that CSFLOC21 improves over CSFLOC19 on six of nine benchmarks, with the largest reduction from 51.913 seconds to 23.929 seconds on uuf150-01.cnf. The results do not show that CSFLOC competes with modern CDCL solvers such as CaDiCaL. Rather, they support the conclusion that LLM-assisted engineering can act as a useful second reader for specialized solver code.

Error-Driven Prompt Optimization for Sustainable Tabular Reasoning

Róbert Lakatos (Faculty of Informatics, University of Debrecen), Árpád Pándy (Faculty of Informatics, University of Debrecen), András Hajdu (Faculty of Informatics, University of Debrecen)

The environmental cost of adapting Large Language Models (LLMs) is a growing concern in artificial intelligence research. We present a systematic comparison between gradientbased fine-tuning and a rule-based prompt optimization framework for financial tabular reasoning. Using the TAT-QA dataset, we demonstrate that our Error-Driven Code Generation Agent (CGA) achieves 66.45% accuracy on the Qwen3 4B model, surpassing GPT-3.5 Turbo. Crucially, this performance is attained with a fraction of the energy required for LoRA or full fine-tuning. Specifically, our proposed solution consumes, on average, >14x less electricity than fine-tuning approaches, significantly reducing operational overhead. Our results indicate that systematic prompt engineering offers a sustainable, privacycompliant alternative for industrial AI deployment, achieving high precision without the carbon footprint of iterative training cycles. This work provides evidence for an energy-efficient (Green AI) alternative in regulated sectors, proving that strategic model adaptation can be both high-performing and energy-efficient.

From Flat Storage to Structured Graphs: Adaptive Memory for Domain-Specific LLM Reasoning

Erick Machungo (University of Debrecen), Gergő Bogacsovics (Faculty of Informatics, University of Debrecen)

Large language models (LLMs) achieve strong general performance but struggle in specialized domains requiring structured, evolving knowledge. Existing context adaptation frameworks, such as Dynamic Cheatsheet and Agentic Context Engineering (ACE), accumulate domain knowledge as flat bullet-point lists, ignoring the relational structure inherent to many knowledge-intensive fields. We present Graph-Structured Adaptive Memory (GSAM), a framework that replaces flat storage with an ontology-grounded knowledge graph, enabling principled cross-concept transfer, explicit failure-cascade modeling, and structurally-aware retrieval. GSAM extends the ACE architecture with two new components, a Graph Constructor

and a Graph Retriever, operating over a domain ontology layer pre-populated from the XBRL financial taxonomy. Evaluated on financial numeric entity recognition (FiNER) and formula reasoning tasks, GSAM achieves 64.0% and 65.0% accuracy, respectively, in the online setting, surpassing ACE by 11.0% on both benchmarks, and reaches 71.0% and 81.0% in the offline setting. On the newly introduced FiNER-Transfer benchmark, GSAM achieves a near-transfer rate of 64.3% compared to 26.2% for ACE, with negative transfer reduced from 19.0% to 7.1%. Retrieval precision improves nearly fourfold, from 6.5% to 25.2%, and the repeated failure rate is reduced by 54.9%.

Creating EMBeD, the Embedded Malware Benchmark Dataset

Dávid Maliga (Budapest University of Technology and Economics), Dorottya Papp (Budapest University of Technology and Economics), Levente Buttyán (Budapest University of Technology and Economics)

IoT devices, such as wireless routers, IP cameras, and smart home gadgets, are widely adopted in everyday life. Unfortunately, it is both technically feasible and economically viable for attackers to infect them with malware. In response to this threat, the research community has made significant efforts to develop various malware detection methods for IoT devices. However, the published solutions are rarely accompanied by the dataset of malware samples on which they were evaluated by their authors. As a matter of fact, currently there does not exist any publicly available benchmark dataset of IoT malware binaries that would make independent validation and comparison of IoT malware detection methods possible in a meaningful way. To address this issue, in this paper, we make the first steps towards an Embedded Malware Benchmark Dataset (EMBeD) by proposing a methodology to create such a dataset. We are currently using this methodology to build EMBeD, which we intend to make publicly available for the benefit of the research community.

Numerically Stable CUDA Implementation of the High-Order Multivariate Moment Computations

Tamás Menyhárt (Unversity of Debrecen), András Hajdu (Faculty of Informatics, Unversity of Debrecen), Róbert Lakatos (Faculty of Informatics, Unversity of Debrecen)

The accurate computation of higher-order central moments (variance, skewness, and kurtosis) is numerically chal- lenging on massively parallel architectures due to floating-point cancellation and non-deterministic reduction order. Classical single-pass updates are less stable for higher-order moments. We present a CUDA implementation that combines hierar- chical pairwise reduction with corrected two-pass accumula- tion to improve numerical robustness while maintaining high throughput. The method uses column-wise partitioning and multi-level reductions across threads, warps, and blocks, with partial statistics accumulated in registers and shared memory before a final merge. On a 2,097,152 × 512 dataset (1.07 billion float32), the kernel completes in 91 ms versus 8.71 s (PyTorch CPU), 6.44 s (PyTorch CUDA), 32.23 s (NumPy), and 46.98 s (SciPy), corresponding to 96×, 71×, 355×, and 517× speedups, respectively. When each GPU thread processes multiple elements, performance is limited primarily by global memory bandwidth.

Quantization-Induced Error Propagation: Activation Analysis within the PEP-AI Framework for Sustainable High-Efficiency and Explainable AI

Tamás Menyhárt (Unversity of Debrecen), András Hajdu (Faculty of Informatics, Unversity of Debrecen), Róbert Lakatos (Faculty of Informatics, Unversity of Debrecen)

Post-training quantization to INT8 enables energy- efficient deployment of deep neural networks on edge devices, yet standard metrics such as mean Average Precision (mAP) fail to capture internal instabilities. We propose a Precise, Explainable, and Provable (PEP-AI) framework for validating quantized convolutional neural networks beyond output-level evaluation. We analyze intermediate activations across Faster R-CNN, YOLO, EfficientNet, and DenseNet using ONNX-based feature extraction. Our results show that despite stable detection perfor- mance, internal feature maps exhibit median relative errors of 4–6% in deeper

layers, with cumulative deviations reaching up to 20%. We identify an attenuation effect where Max-Pooling and Batch Normalization suppress quantization noise, particularly in object regions, where errors correlate with detection failures. In addition, smaller activation deviations (~6%) may lead to missed detections in safety-critical scenarios such as autonomous driving. At the same time, INT8 quantization reduces model size by 67% and inference time by 50.2%, significantly improving energy efficiency and sustainability. We introduce a visual-statistical framework based on sym- metric logarithmic and hexbin analysis to reveal hidden error distributions. Our findings highlight that feature-level validation is useful for energy-efficient AI deployment.

Seasonal Effects in Anomaly Detection for Urban Transit Buses

László Mészáros (Faculty of Informatics, University of Debrecen), László Kovács (Faculty of Informatics, University of Debrecen), András Hajdu (Faculty of Informatics, University of Debrecen)

Predictive maintenance in public transportation increasingly depends on deep anomaly detection models; however, seasonal environmental variations can significantly alter the underlying sensor data distribution, leading to elevated false alarm rates during seasons of the year. Using telematics data collected from ten urban transit buses over a 19-month period, we demonstrate that the latent space of a 1D Convolutional Variational Autoencoder follows a predictable sinusoidal trajectory at the fleet level ($R^2 = 0.89$). We further show that the latent distributions observed in January and July differ significantly (Mahalanobis distance $D = 1.95$), resulting in an empirical false positive rate of 16.7% when a static anomaly threshold calibrated on January data is applied to healthy July samples. Model selection using AIC and BIC reveals a dual-component drift structure consisting of a fleet-wide synchronized seasonal cycle combined with a vehicle-specific linear degradation trends.

3D Retinal Reconstruction and Deep Learning for Diagnosis of Dry and Wet Age-related Macular Degeneration from Optical Coherence Tomography Images

Lilla Mezei (University of Debrecen), Beáta Bajdik (Department of Ophthalmology, University of Debrecen Clinical Centre), Balázs Harangi (Faculty of Informatics, University of Debrecen)

This study presents a 3D Optical Coherence Tomography (OCT) image reconstruction workflow, developed primarily for OCT systems that do not natively export volumetric data, together with a deep learning-based framework for classifying age-related macular degeneration (AMD) into its two primary subtypes, dry and wet AMD. The reconstruction module aligns and interpolates the gaps between sequential B-scans to produce high-quality volumetric visualizations of retinal structures and is used purely for visualization, independently of the classification pipeline. Classification is performed by a binary convolutional neural network trained directly on stacks of the original, non-interpolated 2D B-scans, treating each set of 25 aligned slices per eye as a single input tensor. Across five-fold cross-validation, the proposed model achieved an average F1-score of 0.84 ± 0.15, an accuracy of 0.88 ± 0.10, and an AUC of 0.92 ± 0.09, with a peak fold-level validation accuracy of 96.00%. These results indicate that preserving cross-slice spatial context substantially improves classification performance over single-slice or unaligned baselines. By supporting both visualization and automated subtype classification, the presented approach has the potential to assist ophthalmologists in the timely and accurate diagnosis and management of AMD.

Towards Robust Arabic Stance Detection: Cross-Domain Generalization and Joint Training for Misinformation Mitigation

Ali Mhrez (Doctoral School of Informatics, University of Debrecen), Tibor Tajti (Faculty of Informatics, University of Debrecen, Eszterházy Károly Catholic University)

Stance detection is a pivotal component of automated fact-checking and misinformation mitigation. Machine learning research on Arabic stance detection largely focuses on in-domain evaluations, which leaves a significant gap in understanding how models generalize across the

diverse linguistic patterns found in the real-world news ecosystem. In this paper, we present a systematic evaluation of ten Arabic Transformer models across two publicly available benchmarks: AraStance and UnifiedFC. We investigate cross-domain transferability and the impact of training distributions on model robustness. Our results reveal a significant performance degradation during domain shifts and uncover a generalization disparity, where training on a high-signal dataset can outperform in-domain training for minority classes in noisier domains. Furthermore, we demonstrate that joint training significantly improves model resilience without incurring performance costs on the source domain. Finally, we identify the Discuss class as a bottleneck in Arabic stance detection. This work offers a leaderboard for Arabic stance detection and provides critical insights for the development of robust, universal misinformation detection systems.

Incorporating Weather Data Into Machine Learning Models for Estimating Tertiary-sector Activity

Naoto Minakawa (Keio University), Yoshiyuki Suimon (Keio University)

This paper presents a machine-learning-based model for estimating economic activity in Japan’s tertiary sector, incorporating weather variables as input features. The experimental results demonstrate that weather-related data contributes effectively to the estimation of economic activity in certain industries. Furthermore, an analysis based on SHAP values reveals non-linear relationships between weather conditions and economic performance. In particular, sectors such as road passenger transport, railway passenger transport, and daily life–related personal and amusement services exhibit negative SHAP values when the average temperature falls below approximately 15 °C. In the case of road passenger transport, SHAP values become positive when total precipitation ranges from 100 mm to 200 mm, while values outside this range are negative. These findings indicate that weather variables carry predictive information for specific service sectors and reveal the non-linear patterns used by the models. While this link may be intuitive, this study contributes by quantifying these relationships and showing how it varies across sector.

Secure Data Exchange Protocol for IPFS

Csaba Nagy (University of Debrecen), Andrea Huszti (Faculty of Informatics, University of Debrecen), Norbert Oláh (Faculty of Informatics, University of Debrecen)

The adoption of decentralized Web3 architectures has led to the growing use of IPFS for distributed storage. IPFS nevertheless raises privacy concerns, as network participants may monitor content requests and associate them with specific users or data objects. While Triple Hashing has been proposed to obscure Content Identifiers (CIDs), additional protection mechanisms are still required against profiling and active network-level attacks. This paper proposes an enhanced privacy framework for IPFS that integrates digital signatures and asymmetric encryption to secure the Triple Hash exchange. Furthermore, we introduce decoy requests to mitigate user profiling and a novel encrypted caching mechanism that allows resource-rich nodes to serve content without compromising data privacy. Formal verification using ProVerif was used to evaluate the security properties of the proposed protocol, including content authenticity and key secrecy.

Recognition of GNSS Spoofing in V2I Environments with Physics-Informed Gaussian Process

Tamás Ormándi (Budapest University of Technology and Economics, Department of Control for Transportation and Vehicle Systems), Balázs Varga (Budapest University of Technology and Economics, Department of Control for Transportation and Vehicle Systems)

V2X reliability is paramount for intelligent transportation systems, especially adaptive applications like adaptive traffic lights which are highly sensitive to positioning errors. This abstract addresses V2X vulnerability to GNSS spoofing by proposing an infrastructure-centric detection framework. A physics-informed Gaussian Process Regression model synthesizes Received Signal Strength Indicator metrics and road topology to establish a probabilistic wireless baseline. By cross-referencing claimed coordinates against expected signal behavior, the framework identifies anomalies via physical residuals and epistemic uncertainty, achieving a 95% detection probability at a 2.1-meter offset in severe Non-Line-of-Sight urban environments.

Enhancing Transformer-Based RGB-T Image Fusion via Hybrid Pixel Upsampling

Bilgesu Örten (Middle East Technical University), Cansel Fıçıcı (Ankara University), Mehmet Cem Çatalbaş (Ankara University)

Guided thermal super-resolution (GTSR) is a cross-modal deep learning task that reconstructs high resolution thermal images from low resolution inputs using structural guidance from high resolution RGB images. This problem is relevant to a wide range of intelligent systems and data driven applications where high resolution thermal sensors are limited by hardware constraints. Recent transformer-based fusion models achieve strong cross modal performance but rely on fixed bicubic interpolation during early upsampling, which can degrade high frequency spatial details prior to feature extraction. In this work, a hybrid upsampling module is proposed that combines learnable PixelShuffle based feature expansion with a bicubic residual baseline, integrated into a SwinFuSR based pipeline without modifying the backbone architecture. Experiments on the Cross spectral Image Dataset for Image Super resolution (CIDIS) at scale factor ×2 demonstrate that the proposed approach outperforms classical interpolation methods including Lanczos in terms of PSNR and SSIM, while maintaining low computational overhead. The results demonstrate the effectiveness of hybrid early upsampling design in improving both the accuracy and efficiency of reconstruction in transformer-based thermal image fusion systems.

Accelerated Keystroke Dynamics: An XGBoost Approach to Continuous Biometric Authentication

Erik Pál (University of Debrecen), András Hajdu (Faculty of Informatics, University of Debrecen), Róbert Lakatos (Faculty of Informatics, University of Debrecen)

In the digital age, safeguarding data is paramount, yet user convenience often compromises security, leading to inadequate protection. Traditional authentication methods struggle to keep up with rising unauthorized access. This research addresses this gap by developing a passive, continuous user identification system based on keystroke dynamics. Our approach analyzes low-dimensional features, including key-hold times, transition times, typing speed, and error rates. We implemented a custom Python keylogger and trained an XGBoost machine learning

model, achieving 99.8% accuracy and an Equal Error Rate of 0.0017% on our benchmark dataset. This performance outperformed the selected TypeNet baseline on our dataset by 1.2% in accuracy and demonstrates significantly faster inference times (63 ms vs. 180 ms) with minimal CPU load. This work formalizes typing dynamics as a time-dependent stochastic process and offers a robust, efficient, and user-friendly solution for continuous biometric authentication, supporting proactive security systems.

Deep Learning-Based Prediction of Potential Drug Candidates for Alzheimer's Disease Among Philippine Natural Products

Anne Kelsea Palo (Mapua University), Gielome Martinez (Mapua University), Karylle Cate Tuazon (Mapua University)

Alzheimer's disease (AD) remains one of the most prevalent neurodegenerative disorders worldwide. This study explored the potential of Philippine medicinal plants as sources of compounds for AD-related drug discovery through a deep learning-based compound–protein interaction (CPI) prediction framework. Four model architectures, ChemBERTa+ESM, Graphormer+ESM, TransformerCPI2.0+ProteinBERT, and MMCL-CPI, were trained on a curated dataset of 15,169 compound–protein interactions derived from BindingDB, targeting acetylcholinesterase, butyrylcholinesterase, and NMDA receptor subunits. The trained models were then applied to a screening dataset of 1,503 unique compounds sourced from ten Department of Health-approved medicinal plants. Among the four models, ChemBERTa+ESM achieved the best predictive performance (ROC-AUC = 0.895, PR-AUC = 0.751), though it exhibited a wider train-validation gap indicative of moderate overfitting, while Graphormer+ESM showed more stable, consistent learning behavior. Screening of the natural product library yielded 8,034 compound–target predictions, with 1-hexanol from Allium sativum obtaining the highest prediction score (0.7116) against GRIN2B, and several other Allium sativum-derived compounds ranking highly against cholinesterase targets. These findings suggest that Philippine medicinal plants, particularly Allium sativum, may harbor compounds with potential relevance to AD-related protein targets. However, the results represent computational prioritization rather than confirmed biological activity, showing the need for further structural and experimental validation.

A Hybrid Sim-to-Real Framework for Industrial Object Detection Using Synthetic Data Augmentation and Inference-Time Fusion

János Polgár (Doctoral School of Informatics, University of Debrecen), Gergő Bogacsovics (Faculty of Informatics, University of Debrecen)

We present a hybrid sim-to-real framework for the accurate detection of delivery boxes in a warehouse environment, without depending on large-scale, annotated real-world data. The framework combines realistic synthetic data generation using NVIDIA Isaac Sim to supplement limited real-world annotations, and an inference-time Weighted Box Fusion (WBF) strategy to mitigate distribution mismatch during inference between synthetically generated and real data. We systematically scale synthetic data augmentation across varying dataset proportions while maintaining strict train/test separation and evaluating models only on real images to prevent bias. Our experimental results using YOLO26 as an object detection model demonstrate that while performance gains saturate beyond 50% synthetic augmentation for the standalone model, inference-time fusion consistently delivers superior results across all tested configurations, even when the standalone YOLO26 model using the synthetic data plateaus. The fused approach achieves robust detection metrics, achieving 95% mAP@50, 95% precision, and 90% recall.

Temporal-Aware Graph Attention Network With XGBoost Ensemble for Blockchain Transaction Fraud Detection

Nithish Kumar Ragupathi (University of Debrecen), Norbert Oláh (Faculty of Informatics, University of Debrecen)

Blockchain transaction fraud detection is a critical challenge as cryptocurrency adoption grows rapidly. Graph neural networks (GNNs) and temporal models show potential, but key questions persist about their advantages over tabular baselines, stability in realistic class distributions, and the true predictive value of temporal patterns. This paper presents a systematic empirical study across two real-world blockchain datasets—the Bitcoin Elliptic dataset and an Ethereum phishing dataset—comparing ten models including classical machine learning (ML) methods, five GNN architectures, a Temporal-Aware Graph Attention Network (ATGAT), and a Hybrid ATGAT+XGBoost ensemble. Experiments are conducted under both balanced (50:50) and

imbalanced settings using strict temporal splits to prevent future leakage. Our results reveal that XGBoost remains a formidable baseline, achieving F1 = 0.6889 and ROC-AUC = 0.8731 on Bitcoin, while GNN-only approaches consistently underperform under temporal evaluation. The Hybrid model achieves competitive performance (F1 = 0.6816 on Bitcoin; ROC-AUC = 0.9874 on Ethereum) and provides more robust generalisation than pure GNN models, though it does not clearly surpass standalone XGBoost on Bitcoin. Temporal fraud rate analysis confirms that fraud dynamics evolve over time, motivating temporal-aware architectures even when their measured gains are modest.

Surrogate model of a manufacturing system using random forest method

Gábor Ruzicska (Faculty of Engineering, University of Debrecen), Levente Czégé (Faculty of Engineering, University of Debrecen), Tamás Mankovits (Faculty of Engineering, University of Debrecen)

This paper investigates the development of a surrogate model for a simulation model of a manufacturing system made in Tecnomatix Plant Simulation. Discrete-event simulations have a need of significant computational resources, which limits the scope for analysis and optimization. As an alternative, the use of machine learning methods to estimate system behavior is proposed. Based on a generated dataset (over 70,000 observations), a random forest model is trained to predict key performance indicators: throughput, failure time and congestion. It is shown that the random forest provides high accuracy ($R^2 > 0.86$) and allows for the identification of key factors influencing the system's behavior.

Efficient Local Retrieval-Augmented Generation via Neighboring Document Slice Contextualization

Martin Tibor Sándor (University of Debrecen), Róbert Lakatos (Faculty of Informatics, University of Debrecen)

State-of-the-art AI architectures like Retrieval-Augmented Generation (RAG) using LLMs often rely on "Red AI" infrastructure, creating significant barriers to privacy-centric, localized applications. Anthropic's Contextual Retrieval addresses chunk-level context loss but introduces linear computational scaling that exceeds the limits of consumer-grade hardware.

This study proposes a novel "Document Slice" architecture that optimizes context generation by utilizing a sliding window of neighboring chunks rather than the entire document. By constraining the context window to a fixed radius, we enable high-performance RAG on a single GPU with 8GB of VRAM. Our empirical evaluation, conducted on a "Silver Standard" dataset of 250 complex queries, demonstrates that the proposed method achieves a 2.5x increase in preprocessing speed and a 51% reduction in carbon emissions and energy consumption. Crucially, these efficiency gains are achieved with a statistically insignificant relative decrease in Recall@20 of only 0.9% compared to the full-context baseline. We note that the current evaluation is limited to the English-language scientific literature; generalization to other domains, languages, and document types remains an open question. Yet, this work bridges the gap between industrial-scale performance and accessible "Green AI" deployment, providing a scalable framework for local research assistants in resource-constrained environments.

Open-Ended LLM Problem Solving: Structured Scaffolding for Concept Combination, Agentic Control, and Cross-Coupled Symbolic Tasks

Hashmath Shaik (Stony Brook University), Gnaneswar Villuri (Stony Brook University), Alex Doboli (Stony Brook University)

Open-ended problem solving by Large Language Models (LLMs) requires three coupled capabilities: (i) combining concepts beyond the LLM's training distribution, (ii) performing this combination through agentic architectures whose internal reasoning can be inspected and adaptively intervened upon, and (iii) applying the result to cross-coupled symbolic combinatorial problems whose correctness depends on runtime semantics rather than text alone. We address these capabilities through structured external scaffolding (knowledge graphs, validator agents, mechanistic interpretability probes, symbolic representations, and parameterized statistical models). A foundational LLM-based Cognitive Architecture and its team-problem-solving extension establish the agentic substrate. A graph-augmented generator-validator framework with sparse-autoencoder and circuit-discovery feedback delivers concept combination with first-generation causal probes. A static and a real-time validation method instantiate two regimes of agent intervention. Three application threads (LLM-based code generation under coupled semantics, LLM and Symbolic Knowledge Graph integration for analog circuit topology synthesis, and a granularity-aligned model-based generalization to

graph-transformer maze navigation) demonstrate cross-coupled symbolic combinatorial scope. Across all three capabilities, structured external scaffolding consistently outperforms end-to-end neural generation. Open directions include moving from feedback-style mediation to representation-level inference-time interventions, scaling agentic architectures beyond small teams, and automating the construction of symbolic and statistical scaffolds.

Analysis Based Initial Tuning Method for Steering Systems

Dávid Somogyi (Faculty of Information Technology and Bionics at Pázmány Péter Catholic University), Csaba Fazekas (thyssenkrupp Components Technology Hungary Ltd.), Gábor Szederkényi (Faculty of Information Technology and Bionics at Pázmány Péter Catholic University)

This paper investigates the influence of EPS system nonlinearities when a linear boost curve controller is used as an initial tuning approach for a nonlinear boost characteristic system. Electric Power-Assisted Steering (EPS) systems use nonlinear boost curves to generate assist torque based on driver input and vehicle speed, strongly affecting steering feel, but their calibration is time-consuming due to many tuning parameters. The proposed method uses a system-level analysis of EPS dynamics, combining operating-point fitting, closed-loop linearization, and dominant pole analysis to determine initial boost parameters and controller gains that ensure stable steering response. A linear boost controller is applied to a nonlinear EPS model to evaluate its effect on assist torque, transmission behavior, and steering feel. Simulation results show good agreement with measurements obtained on a real EPS test bench, confirming that the proposed initialization provides a reliable tuning baseline while accurately capturing the dominant steering dynamics.

Chaotic Systems Modeling with Federated Learning Approaches and Recurrent Neural Networks

Georgiana-Stefania Stoie (George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Targu Mures), Roland Bolboaca (George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Targu Mures)

This paper presents a federated learning framework for forecasting chaotic Lorenz system dynamics using LSTM, GRU, and Temporal Fusion Transformer (TFT) architectures. The

study reveals a performance inversion between training environments: in a centralized setup, the GRU and LSTM models excel with scores of 0.0260 and 0.0428 while the TFT performs poorly at 0.4216. Conversely, in a federated setup, the recurrent models struggle, whereas the TFT demonstrates superior stability with a leading score of 0.0633, significantly outperforming the GRU result of 0.1949 and the LSTM score of 0.1995.

Pattern-Oriented Fake News Detection via Five-Dimensional Knowledge Graph and Retrieval-Augmented Generation: A Study on Chinese-Language Media

Yi-An Sung (Soochow University), Chang-Yi Kao (Soochow University)

Fake news increasingly exploits discourse-level manipulation rather than outright fabrication, rendering traditional fact-checking approaches insufficient. This paper proposes a pattern-oriented framework that combines a five-dimension knowledge graph (KG) with retrieval-augmented generation (RAG) to detect fake news by identifying recurring discourse patterns. The five dimensions—Technique (14 propaganda devices), Category (9 news topics), Entity (open-set named entities), DeceptionSignal (8 semantic cues), and TargetEmotion (8 affective states)—are organized into a Neo4j knowledge graph: the Category dimension is assigned during preprocessing by a multilingual zero-shot classifier, whereas the remaining four dimensions are extracted by a large language model (LLM). During inference, a novel discrete-matching retrieval mechanism queries the KG for historically similar cases from both fake and real news pools, weighted by inverse document frequency and pivoted category-scale normalization, to provide contextual evidence for the LLM's classification decision. Experiments on the Cofacts dataset (42,632 Traditional Chinese articles) using a locally deployed 8B-parameter model (Llama-3-Taiwan-8B-Instruct) demonstrate that the RAG-enhanced system achieves an accuracy of 78.78% on the full test set of 8,527 articles, outperforming the LLM-only baseline by 12.29 percentage points. Fake news recall improves substantially, from 63.36% to 86.02%. However, the system's 55.73% false-positive rate (FPR) limits deployment to human-reviewed triage rather than autonomous decision making. On a 10% stratified subset of 853 articles, statistical testing confirms a significant improvement with a medium effect size. An ablation study on the 1% tuning subset showed an accuracy difference of 10.47 percentage points relative to the no-retrieval configuration. The entire system runs

locally with zero API cost, and the retrieved historical cases provide traceable justification for each classification decision.

SNART: Real-Time Translation via ECAPA-TDNN-Based Biometric Speaker Filtering

Gergő Szabó (University of Debrecen)

Real-time simultaneous translation systems deployed in lecture halls and conference rooms face three intertwined challenges: latency, recurring cloud API cost from continuous audio streaming, and translation quality degradation triggered by ambient noise and off-target speech, which produce transcription hallucinations that propagate into the translation layer. This paper presents SNART, a licensed translator software addressing these issues through a two-stage architecture. A locally executed biometric gatekeeper, built on an ECAPA-TDNN speaker embedding network fine-tuned on a domain-specific Hungarian speech corpus with batch-hard triplet loss, validates incoming audio in three-second sliding windows. Only segments whose 192-dimensional embedding exceeds a cosine-similarity threshold against a pre-enrolled reference voiceprint are forwarded to cloud speech-to-text (STT) services (optional). Transcripts pass through a large language model (LLM) translation step augmented with a retrieval-augmented glossary that handles proper-noun phonetic correction and terminology consistency. The same LLM pass returns a structured side-channel signal that supports runtime zero-shot speaker enrollment: when an utterance contains a self-introduction, the extracted name is bound to the live voiceprint in a local JSON registry. Visual delivery is handled by a transparent PyQt6 overlay whose typography is sized by a trigonometric calculator driven by viewer distance. Demonstrating its efficiency, SNART reduces word error rate from 48.4% to 6.5% and improves BLEU from 43.4 to 84.5 on a five-minute Hungarian benchmark relative to the native captioning of Microsoft Teams.

Data-Efficient Chest X-Ray Classification Under Severe Label Scarcity with Lightweight Backbones

Péter Szilágyi (Faculty of Engineering, University of Debrecen), Gergely Sógor (University of Debrecen), János Tóth (Faculty of Informatics, University of Debrecen)

This study investigates data-efficient chest X-ray classification with lightweight convolutional backbones under severe label scarcity. We compare linear probing, partial fine-tuning, full fine-tuning, and prototypical networks across three binary pathology detection tasks using balanced training subsets with 20, 50, 100, or 500 positive and negative examples per task, repeated across random seeds. An external evaluation examines performance under cross-dataset shift. Prototypical networks were competitive and achieved the strongest aggregate results for EfficientNet-B0, whereas full fine-tuning performed best for MobileNetV3-Small in terms of macro-F1 and AUROC. However, the advantage of prototypical networks over full fine-tuning was modest and not uniform. Backbone selection further exposed a trade-off between predictive performance, inference latency, and model size. These findings suggest that metric learning may be useful in selected low-label settings, but full fine-tuning remains a strong and simpler practical baseline.

Automated Generation of Attack Trees from System Models

Dániel Zoltán Szilvágyi (Budapest University of Technology and Economics), Levente Buttyán (Budapest University of Technology and Economics)

Attack trees support threat modeling by decomposing high-level attacker goals into smaller subgoals connected by AND/OR logic. Manual construction is time-consuming and does not scale to complex systems. This paper presents a rule-based generator that builds attack trees automatically from a system model represented as typed entities and their relations. Starting from a goal predicate, the generator performs backward chaining over precondition rules, producing an AND/OR dependency structure whose leaves are either model facts or basic attack steps that represent concrete attacker actions. To avoid combinatorial blow-ups, the generator reuses identical substructures via memoization, turning the output into a compact directed graph. Instead of pruning recursive dependencies during generation, cycles are recorded explicitly and may become the subject of a later analysis phase. We demonstrate the proposed

solution on an industrial control system and evaluate the effect of the main scalability optimizations.

A Hybrid Entropy-Aware Method for Hyperparameter Optimization

Muhammad Talal Bin Tayyab (University of Debrecen), Gergo Bogacsovics (Faculty of Informatics, University of Debrecen)

Hyperparameter optimization is an important but computationally expensive part of machine learning model development, especially when the search space contains continuous, integer-valued, and categorical hyperparameters. Conventional approaches such as random search, Gaussian-process Bayesian optimization, and tree-structured Parzen estimators are widely used, but they may become less effective in mixed, irregular, and discontinuous hyperparameter spaces. This paper proposes HOpt, a hybrid entropy-aware hyperparameter optimization method designed for budget-constrained optimization in heterogeneous search spaces, combining probabilistic modeling, rank-based softmax weighting, entropy-aware bandwidth adaptation, Spearman-based importance guidance, and proposal-pool softmax selection. Our method learns weighted density models from good-performing hyperparameter configurations, monitors search diversity via an entropy score, and adaptively balances exploration and exploitation by broadening the sampling distribution when diversity decreases. The proposed method is evaluated on three MedMNISTv2 medical image classification datasets and compared against random search, Optuna-TPE, scikit-optimize Gaussian Process Bayesian optimization, and BoTorch Gaussian Process Bayesian optimization under a controlled cross-validation protocol. Experimental results show that HOpt achieves competitive final test performance while maintaining favorable runtime behavior. In particular, HOpt obtains the highest mean test F1-macro on OCTMNIST, matches the highest mean test F1-macro on PathMNIST, and achieves the second-highest mean test F1-macro on BloodMNIST. Across all three datasets, HOpt is consistently faster than the Gaussian Process-based baselines while producing stronger or comparable predictive performance.

RL DDPG-Based Adaptive Access Point Selection for Cell-Free Massive MIMO Systems in 6G Networks

Djamila Talbi (Faculty of Informatics, University of Debrecen), Zoltán Gál (Faculty of Informatics, University of Debrecen)

This paper proposes a Deep Deterministic Policy Gradient (DDPG)-based adaptive access point (AP) selection framework for cell-free massive multiple-input multiple-output (CF-mMIMO) systems in beyond-5G and 6G networks. Unlike conventional full-cooperation or heuristic clustering methods, the proposed approach learns a continuous AP weighting policy that dynamically selects a subset of APs based on channel conditions and spatial distribution. A reward function jointly optimizing spectral efficiency, AP locality, and utilization efficiency is introduced. Simulation results demonstrate that the proposed framework achieves high spectral efficiency with significantly reduced AP activation, leading to improved energy efficiency and reduced fronthaul overhead. The learned policy exhibits spatially localized cooperation, confirming the effectiveness of reinforcement learning for scalable CF-mMIMO optimization.

Variational Causal Learning for Dynamic Treatment Effects in Time Series Data with an Application to Wholesale Electricity Markets

Masahiro Tanaka (Faculty of Economics, Fukuoka University).

This paper develops a scalable variational method for estimating dynamic treatment effects in time-series data. The proposed approach builds on a moment-based quasi-posterior that identifies causal response paths through instrumental-variable estimating equations rather than a fully specified likelihood. To stabilize horizon-specific estimates, the method uses a smoothness prior over the treatment-effect path, while a coordinate-ascent variational approximation replaces simulation-based posterior computation. The variational distribution keeps a full Gaussian factor for the stacked coefficient vector, preserving dependence across horizons and covariates, and yields closed-form updates for efficient implementation. We apply the method to daily Danish wholesale electricity markets to estimate the dynamic effects of wind and solar generation on day-ahead prices, using weather-driven renewable potentials as instruments. The variational estimates broadly agree with those from a Gibbs sampler while

reaching stable values substantially faster. The empirical results show negative short-run price effects of wind generation and weaker effects of solar generation.

MM-TTM-C: Multimodal Topic Tracking Model with Confidence for Twitter Sentiment Analysis

Tomoe Taniguchi (Ochanomizu University), Ichiro Kobayashi (Ochanomizu University)

We propose the Multimodal Topic Tracking Model with Confidence (MM-TTM-C), a probabilistic generative framework extending the Topic Tracking Model by jointly modeling word--topic and sentiment--topic distributions through sequential Bayesian updating, with a label confidence parameter for robustness to annotation noise in sentiment labels.
The resulting topic-level sentiment drift series supports two capabilities: Granger causality analysis for identifying directional dependencies among topics, and one-step-ahead forecasting in which anomalous deviations trigger early-warning alerts.
Experiments on a cryptocurrency tweet dataset spanning 26 weekly time steps show that MM-TTM-C reveals inter-topic sentiment dynamics not captured by existing approaches.

Robustness of Retinal Vessel Segmentation to Image Quality Degradations

János Tóth (Faculty of Informatics, University of Debrecen)

Retinal vessel segmentation is essential for automated fundus image analysis, but its reliability depends strongly on image quality. Because benchmark datasets contain mostly high-quality images, the robustness of vessel segmentation methods under realistic degradations remains insufficiently characterized. We investigate this issue using controlled synthetic degradations and two U-Net-based models, one trained on mixed-quality images and the other exclusively on high-quality images. The effects of blur, contrast reduction, and non-uniform illumination are evaluated across multiple severity levels, together with no-reference image quality assessment (NR-IQA) metrics as candidate indicators of segmentation reliability. The results reveal degradation-specific failure modes. Blur and severe compound degradations substantially reduce performance, mainly through recall loss, whereas contrast reduction causes a sharp performance drop in both models. Although several NR-IQA metrics are associated

with performance under specific degradations, no metric shows consistent alignment with segmentation reliability across conditions. These results motivate task-aware quality control for reliable clinical deployment.

L1-aware grammatical error correction for ESL writing: performance-disparity trade-offs

Anita Valieva (University of Debrecen), János Tóth (Faculty of Informatics, University of Debrecen)

Grammatical error correction (GEC) systems for learner English are usually trained as general-purpose models and rarely account for the writer's first language (L1), although L1 influences error patterns and may affect correction quality. To address this, this paper investigates whether lightweight L1 conditioning improves GEC for ESL writing. Specifically, using the FCE learner corpus, we focus on essays written by learners whose L1 is Chinese, Russian, or Spanish and compare BERT-based GECToR models in four settings: baseline, continued training, L1-aware input, and L1-aware training on a downsampled L1-balanced training set. For evaluation, we use ERRANT to compute precision, recall, and the precision-weighted F-score (F0.5) on both the original and a balanced test set. Our findings show that the L1-aware model obtains the best overall performance, whereas the balanced model yields the smallest performance gap across learner groups but lower overall precision and F0.5. Thus, the results suggest a performance-disparity trade-off in this setting: lightweight L1 conditioning improves average correction quality, while balanced training reduces cross-group score dispersion at the cost of overall F0.5.

Hierarchical Structured Abstraction as an Inductive Bias for Out-of-Distribution Generalisation

Gnaneswar Villuri (Stony Brook University), Hashmath Shaik (Stony Brook University), Alex Doboli (Stony Brook University)

Flat machine learning models fit training-distribution statistics and degrade when that distribution shifts at test time. We present a cross-domain research arc demonstrating that hierarchical structured abstraction, which decomposes a prediction task into fewer, topologically-grounded decisions over domain- specific primitives, acts as an effective

inductive bias that mitigates out-of-distribution (OOD) degradation. Four domains are examined: (i) dialog relation extraction, where a two-layer MLP–LLM architecture exploiting closed-class word structure improves over flat LLM extraction by 13.6%; (ii) conceptual combination, where graph augmented generator–validator optimisation improves relational linkage by 8.8 points over zero-shot prompting; (iii) code generation with coupled semantics, where a knowledge-graph intermediate layer with simulated-annealing selection outperforms Reflexion and OpenEvolve on 7 of 8 benchmark problems at 300–600 evaluations; and (iv) graph navigation, where a pointer network restricted to topological landmarks with exact BFS bridging achieves 95.3% vs. 91.0% OOD success on maze navigation and 85.1% vs. 78.0% on zero-shot road-network transfer, both at a 20× compute reduction. A cross-domain theoretical analysis shows that the error-compounding bound from the navigation work (characterising when a hierarchical policy dominates a flat one as a function of decision-count ratio T /L) is consistent with the observed gains across all four settings, and probing experiments confirm the mechanism is representational, not merely inferential.

Fusing Semantics and LLM-Driven Emotions by Gated Attention for Implicit Hate Speech Detection

Yuan-Chih Yu (Chinese Culture University, Taipei, Taiwan)

Detecting implicit hate speech continues to be a significant challenge for natural language processing systems, as harmful intent is often concealed within polite, neutral, or coded language. Traditional text-based classifiers often struggle to identify these subtle forms of hostility because they lack the emotional context needed to interpret the underlying aggression. To address this issue, we propose the Emotion Gated Attention Fusion Network (EGAF-Net), an architecture that integrates textual semantics with psychological emotions. Unlike conventional approaches that rely on static lexicons, our method utilizes a Large Language Model (LLM) to extract a 9-dimensional affective vector directly from the text. This vector represents Plutchik's eight basic emotions along with an overall measure of arousal. Simultaneously, we employ a frozen pre-trained Transformer model (RoBERTa) to capture deep linguistic semantics. To combine these two modalities, we introduce a feature-wise Gated Attention Layer that utilizes a sigmoid-driven masking mechanism to reduce modality-specific noise and dynamically adjust the scale of semantic-emotional features. In evaluations on the Implicit Hate Corpus (IHC), our model demonstrates improved discriminative capability,

achieving a competitive ROC-AUC score of 0.7714 and outperforming pure-semantic baselines. However, further architectural refinements are necessary to address the low recall rates observed across all tested baselines. Additionally, qualitative error analysis indicates that our gating mechanism successfully identifies several high-risk instances of implicit hate that traditional models misclassify as safe. These findings suggest that incorporating LLM-driven emotional intelligence is a promising approach to uncovering implicit toxicity.

Transformation-Invariant Latent Consistency for Robust 3D Point Cloud Watermarking

Aya Zairi (Doctoral School of Informatics, University of Debrecen), Róbert Tornai (Faculty of Informatics, University of Debrecen)

This document quantifies the robustness limits of convolutional 3D watermarking to global affine deformations in order to quantify its resilience limits. The necessity for strong intellectual property protection has increased as 3D digital assets become essential to industrial digital twins and metaverse platforms. The performance of architectures like PointNet under global structural alterations continues to be a crucial security bottleneck, despite their notable resilience to local perturbations like Gaussian jitter and point cloud simplification. We present a parameterised ”Global Twist” attack, which is a smooth, non-linear transformation that maintains local manifold continuity while changing long-range structural relationships. Our empirical analysis identifies a basic architectural constraint known as ”Geometric Myopia,” in which convolutional neural network (CNN) encoders’ locally limited receptive fields are unable to preserve latent-space stability even when the geome try is perceptually unaltered. We demonstrate that, even with minimal geometric distortion, watermark integrity in typical autoencoders fails above a critical threshold of dL ≈ 4.8 for real-world assets such as the Stanford Bunny. We suggest a novel Latent Consistency Constraint (LCC) to mitigate this issue. We explicitly impose transformation invariance within the embedding space by incorporating a contrastive Mean Squared Error (MSE) loss between the original and distorted latent vectors during training. The Pareto frontier of watermark integrity is substan tially stabilised by the suggested LCC, according to the results, providing a strong foundation for safe 3D asset authentication in dynamic digital contexts.

Graph Theory–Driven Methods for Early Optimization of Neural Networks

Dániel Márk Zombori (University of Debrecen), Imre Varga (Faculty of Informatics, University of Debrecen)

This paper investigates how underlying network properties, including centrality, plasticity, and supra-strength, can be leveraged to guide weight pruning. The proposed approach aims to identify and remove redundant parameters at an early stage while minimizing the resulting degradation in model performance. The method resulted in only a slight decrease in performance on the evaluated datasets, demonstrating its ability of improving overparametrized neural networks.

EXTENDED ABSTRACTS

From Turbulence to Take-off: Modernising an airline's Digital Platform with Domain-Driven Design

Tamás Bartók
i.) EPAM Systems
ii.) Institute of Informatics
University of Szeged
Szeged, Hungary
ORCID: 0009-0005-3384-1919
Email: tamas_bartok@epam.com

Ferenc Héjja
i.) Doctoral School of Informatics
University of Debrecen
ii.) EPAM Systems
Debrecen, Hungary
ORCID: 0009-0008-5770-5577
Email: hejja.ferenc@inf.unideb.hu

Attila Szabó
Wizz Air Ltd.
Budapest, Hungary
ORCID: 0009-0002-2196-8371
Email: attila.szabo@wizzair.com

***Abstract*—The digital platform — serving millions of passengers across web and mobile — was held back by a tightly-coupled architecture: page loads exceeding 25 seconds, frequent outages caused by a third-party component, and Backend-for-Frontend components that had grown into sprawling monoliths with hidden cross-dependencies. With the airline's growth strategy demanding a step-change in digital capability, incremental fixes were no longer an option.**

This talk tells the story of the airline's Architecture Modernisation Program — a ground-up rethinking of an e-commerce platform that handles flight search, booking, payment, ancillaries, and disruption management at scale. We will walk through how a small architecture team applied Domain-Driven Design to carve eleven clearly bounded domains (Search, Order, Offer, Pricing, Payment, Ancillary, Seating, Customer, Agency, Discount Programs, and Flight Exception Handling) out of a system where everything was tangled with everything else.

Attendees will learn how we modelled availability mathematically on the critical path, chose Azure-primary and AWS-secondary multi-cloud topology, and designed a BFF and API Gateway layer to replace duplicated business logic without a big-bang rewrite. We will be honest about the hard parts: aligning product owners on bounded context boundaries, sequencing domain migrations to avoid circular dependencies, and building a governance model that keeps an agile delivery team anchored to architecture principles.

The session concludes with concrete takeaways: a practical approach to domain readiness scoring, a phased migration roadmap template, and the lessons learned from turning architecture blueprints into funded implementation plans — leaving attendees with a reusable playbook for large-scale platform modernisation in a high-traffic,compliance-heavy environment.

***Index Terms*—Solution Architecture, Non-Functional Requirements, Quality Attributes, Availability, fault tolerance, Key Performance Indicator (KPI)**

REFERENCES

[1] L. Bass, P. Clements, and R. Kazman, Software Architecture in Practice, 4th ed., Addison-Wesley, 2022.
[2] E. Evans, Domain Driven Design, Tackling Complexity in the Heart of Business Software, Addison-Wesley, 2003.
[3] T. Bartok, F. Hejja, and G. Kocsis, "Availability Modeling in Software Architecture," in *Proc. 2024 IEEE 3rd Conf. on Information Technology and Data Science (CITDS)*, Debrecen, Hungary, Aug. 2024, pp. 1–6, doi: 10.1109/CITDS62610.2024.10791390.
[4] C. Ericson, Fault Tree Analysis - A History, Proceedings of the 17th International Systems Safety Conference, 1999.
[5] H.A. Watson, Launch control safety study, Bell Telephone Laboratories, Murray Hill, N.J. USA, 1961.
[6] M. Glinz, "On Non-Functional Requirements," The 15th IEEE International, 15-19 October, 2007, pp. 21-26.
[7] ISO/IEC 25010, Systems and software engineering — Systems and software Quality Requirements and Evaluation (SQuaRE) — System and software quality models
[8] Nurbojatmiko, E. K. Budiardjo, and W. C. Wibowo, Slr on Identification & Classification of Non-Functional Requirements Attributes, and Its Representation in Functional Requirements, In Proceedings of the 2018 2nd International Conference on Computer Science and Artificial Intelligence (CSAI '18). Association for Computing Machinery, New York, NY, USA, 151–157, 2018.
[9] J. Yuan, A Conditional Probability Approach to Reliability with Common-Cause Failures, IEEE Transactions on Reliability, vol. R-34, no. 1, pp. 38–42, April 1985.
[10] AWS Service Level Agreements (SLAs). [Online]. Available: https://aws.amazon.com/legal/service-level-agreements/?aws-sla-cards (Accessed Jul. 2026.)

979-8-3195-0700-6/26/$31.00

On properties of feasibility in non-standard Heyting arithmetic

Péter Battyányi
Departement of Computer Science
Faculty of Informatics, University of Debrecen
Debrecen, Hungary
battyanyi.peter@inf.unideb.hu

***Abstract*—We examine nonstandard Heyting arithmetic extended with a feasibility predicate. Feasibility is defined as a downward closed property containing all numerals and closed under applications with primitive recursive functions. Making use of Kleene's arithmetical realizability, we demonstrate for a theory slightly weaker than Heyting arithmetic that provably feasible terms coincide with the set of numerals. Moreover, we show that disjunction and existential properties are preserved in our arithmetical theory.**

***Index Terms*—Heyting arithmetic, realizability, non-standard arithmetic**

I. Introduction

We are concerned with non-standard Heyting arithmetic extended with a feasibility predicate. We consider intuitionistic arithmetic, called Heyting aritmetic, and we augment it with a non-standard element. Then we add to the theory axioms expressing a feasibility relation. In more detail, Heyting arithmetic with non-standard elements is obtained by adding the axioms $\underline{n} < c$ for all natural numbers n to our arithmetical theory. It was as early as in 1934, when Skolem proved with an argumentation based on the compactness theorem for first-order logic that first-order arithmetic is consistent with these axioms [7]. Since then, non-standard models of arithmetic were intensively studied, see, for example, the monographs [3] and [4]. Following the idea of Dragalin [2], we introduce a feasibility predicate which is a formalized counterpart of the notion of a proper cut in a nonstandard model. That is, we assume that feasibility is downward closed and no primitive recursive function can lead out of the set of feasible elements. Moreover, the infinite element c is not feasible.

The overspill principle, due to Robinson, for non-standard Peano arithmetic states that we are not able to define a proper cut of a model of Peano arithmetic with an arithmetic formula [4]. In our case, we prove that F, however, defines the proper cut of standard elements. Of course, F does not belong to the original language. A theory admits the disjunction property, if, whenever $A \vee B$ is derivable for the closed formula $A \vee B$, then either A or B is derivable. Similarly, we say that for a theory the existential property fulfills, if, whenever $\exists x A(x)$ is derivable for the closed formula $\exists x A(x)$, then $A(t)$ is derivable for some closed term t. Constructive, non-standard theories may refute both properties [1], [6]. One of the reasons is that they usually require the fulfillment of axioms other than the ones considered in this paper, for example transfer principles [1]. However, we have deliberately chosen the simplest axioms that could possibly be more defendable from a constructive standpoint. As a result of this, our theory preserves both the disjunction and the existential properties.

979-8-3503- 8788-9/24/$31.00 © 2024 IEEE

II. Preliminaries

We consider the constant c and we add to the existing theory the axioms $\underline{n} < c$ for all natural numbers n, $\underline{n}$ denoting $S(\dots S(0))$, where S occurs n times in the expression.

We further extend our theory with a predicate F expressing feasibility. Let F be a one-place predicate symbol. The following axioms are from the paper of Dragalin [2]:

1) $F(0)$,
2) $\forall x \forall y (F(x) \wedge y < x \supset F(y))$,
3) $\forall x (F(x) \supset x < c)$,
4) $\forall x_1 ... \forall x_n (F(x_1) \wedge \dots \wedge F(x_n) \supset F(g(x_1, ..., x_n)))$, for each symbol g standing for a primitive recursive function.

In addition, we accomplish slight modifications to the predicate calculus. In effect, the new axioms state that no fresh term can be introduced in a derivation unless they are feasible.

Finally, in place of the usual induction axiom in arithmetic, we introduce the axiom

$$A(0) \wedge \forall^f x (A(x) \supset A(Sx)) \supset \forall^f x A(x) \quad (Ind^f),$$

where $A(x)$ does not contain F and the notation $\forall^f x A(x)$ stands for $\forall x (F(x) \supset A(x))$. We denote the theory obtained in this way by HAF^*.

III. Results

We state the following results. In what follows, let $\vdash$ denote provability in HAF^*.

Theorem 3.1: Let $\vdash F(t)$ for a closed term t. Then there exists an n such that $\vdash t = \underline{n}$.

The proof of the theorem relies on a realizability technique due to Kleene [5].

As a consequence of the Theorem, we obtain the disjunction and existential properties.

Corollary 3.2: Let A be a closed formula such that $\vdash A$. Then the following assertions are valid.

1) If $A = (B \vee C)$ then either $\vdash B$ or $\vdash C$.
2) If $A = \exists x C(x)$ then $\vdash C(\underline{k})$ for some natural number k.

IV. Future Work

In his paper [2], Dragalin raises some questions concerning the theory HAF, which is Heyting arithmetic with the usual axioms for predicate logic together with a non-standard element and a feasibility relation. For example, if we adopt the formal thesis of Church, that is, $\forall x \exists y A(x, y) \supset \exists e \forall x \exists y(\{e\}(x) = y \wedge A(x, y))$, then it follows from Dragalin's result that $HAF + CT$ is conservative over HAF provided CT contains no occurrence of F. The question emerges what happens when we consider CT with the premise $\forall x \exists y(F(y) \wedge A(x, y))$. Does this mean that we are able to find a y for every x by means of a feasible algorithm? In [2], Dragalin establishes an algebraic model to study provability in arithmetical theories. As a future undertaking, it could be interesting to make use of the characteristics of this model in order to examine the metamathematical properties of these theories.

References

[1] J. Avigad, J. Helzner, "Transfer principles in nonstandard intuitionistic arithmetic," Archive for Mathematical Logic 41(6), 581–602, 2002.

[2] A. G. Dragalin, "Explicit algebraic models for constructive and classical theories with non-standard elements," Studia Logica 55(1), 33–61, 1995.

[3] P. Pudlák, P. Hájek, Metamathematics of First-Order Arithmetic. Springer-Verlag, Berlin Heidelberg, 1993.

[4] R. Kaye. Models of Peano Arithmetic. Clarendon Press, Oxford, 1991.

[5] S. C. Kleene. Introduction to Metamathematics. North Holland, Amsterdam, 1952.

[6] E. Palmgren, "A note on mathematics of infinity," The Journal of Symbolic Logic, 58(4), 1195—1200, 1993.

[7] Th. A. Skolem, "Über die nichtcharakterisierbarkeit der zahlenreihe mittels endlich oder abzählbar unendlich vielen aussagen mit ausschliesslich zahlvariablen. Fundamenta Mathematicae 23, 150–161, 1934.

A Canonical Rigid Direct-System Representation of Finite Local-Unit-Aligned Totally Ordered Monoids

Sándor Jenei
Institute of Mathematics and Informatics
Eszterházy Károly Catholic University, Hungary
Institute of Mathematics and Informatics
University of Pécs, Hungary
jenei.sandor@uni-eszterhazy.hu; jenei@ttk.pte.hu

***Abstract*—We study finite local-unit-aligned totally ordered monoids, meaning finite totally ordered monoids in which the greatest right and left local units of each element coincide. The paper proves that every such monoid has a canonical rigid representation as a finite chain-indexed direct system of component monoids, and that every rigid direct system of the corresponding kind reconstructs such a monoid. The representation is induced intrinsically by the local-unit map τ, through the finest τ-multiplication-coherent partition of the positive idempotent skeleton. The resulting component monoids, transition maps, directed lexicographic order, and product formula recover both the order-theoretic and algebraic structure.**

***Index Terms*—totally ordered monoid, local unit, direct system, canonical decomposition, representation theorem, ordered algebra, algebraic logic**

I. Introduction

The present contribution is a theoretical information-technology result in ordered algebra and algebraic logic, with direct relevance to formal models of computation. We develop a precise representation theorem for finite ordered algebraic structures that can serve as algebraic semantics for logical calculi and finite computational models. Our methods are structural and constructive: they identify canonical components, transition maps, and reconstruction formulas, thereby transforming an initially opaque finite ordered monoid into a finite indexed system whose data are easier to analyze and compare.

The result facilitates the exchanging of theoretical and practical advances in data and information technologies. Although the results are purely mathematical, finite ordered monoids and their homomorphisms are natural objects in semantics, classification of finite algebraic models, and logic-oriented computation. The categorical equivalences make the representation stable under structure-preserving maps, which is important for transferring questions between monoids and their direct-system presentations.

Thus the paper contributes to the foundational side of information technology by clarifying how finite ordered algebraic systems can be decomposed, reconstructed, and functorially compared.

979-8-3195-0700-6/26/$31.00 © 2026 IEEE

II. Problem and Motivation

Finite totally ordered monoids provide a compact setting in which algebraic multiplication and a linear order constrain one another. They occur in ordered algebra, semigroup theory, valuation-type constructions, and algebraic logic; in the commutative literature a common simplifying convention is to put the identity at one end of the chain. The present work does not impose this convention. The identity may lie anywhere in the chain, and no commutativity assumption is made.

The central hypothesis is local-unit alignment. For an element x of a finite totally ordered monoid $\mathbf{M} = \langle M, \leq, \cdot, e\rangle$, consider the greatest right local unit and the greatest left local unit available for x. The monoid is local-unit-aligned when these two greatest local units coincide for every x. Their common value is denoted $\tau(x)$ and belongs to the positive idempotent skeleton $\mathrm{Id}^+(\mathbf{M})$. Thus τ records the intrinsic local-unit geometry of the ordered monoid.

The motivating question is whether this local-unit geometry is merely descriptive or whether it determines a canonical representation of the whole structure. Earlier representation results in related residuated settings used a strong global identity of the form $\tau(xy) = \tau(x)\tau(y)$, equivalently a maximum rule on the idempotent skeleton. In the class of residuated lattices and even so in the wider class of finite totally ordered monoids this identity is exceptional. The main point of the paper is that a representation theorem still survives, but the correct object is not the pointwise τ-layer decomposition. It is a coarser canonical decomposition obtained by asking exactly how much multiplicative information τ retains at block level.

III. Canonical Stratification

The construction begins with a partition of $\mathrm{Id}^+(\mathbf{M})$. It is the τ-multiplication-coherent partition, denoted $\mathcal{C}_{\mathrm{m}}(\mathbf{M})$, which is the finest partition for which the block of $\tau(xy)$ depends only on the blocks of $\tau(x)$ and $\tau(y)$. This distinction is crucial: exact local-unit layers can be too fine to support multiplication, while coherent blocks are exactly coarse enough for a direct-system reconstruction.

For each block $A \in \mathcal{C}_{\mathrm{m}}(\mathbf{M})$ one forms a component layer consisting of all elements whose τ-value belongs to A. We show that each such layer carries a natural component monoid $\mathbf{M}_A$, again finite, totally ordered, and local-unit-aligned. If

$A \leq B$ in the finite chain of coherent blocks, canonical transition maps $\rho_{A,B}\colon \mathbf{M}_A \to \mathbf{M}_B$ are induced by multiplication with suitable block units. These maps are isotone unital homomorphisms and are compatible in the direct-system sense. The resulting data

$$\mathfrak{D}(\mathbf{M}) = \big((\mathbf{M}_A)_{A\in\mathcal{C}_{\mathrm{m}}(\mathbf{M})}, (\rho_{A,B})_{A\leq B}\big)$$

form a finite chain-indexed direct system recovered intrinsically from τ, $\mathrm{Id}^+(\mathbf{M})$, and the coherent block structure. In the finite ordered setting the proper canonical transitions are unit-constant: if $A < B$, then $\rho_{A,B}(x) = e_B$ for all $x \in M_A$. Hence the system attached to $\mathbf{M}$ is rigid.

IV. Reconstruction

Multiplication and order are recovered from this rigid system by common-upper transport. For $x \in M_A$ and $y \in M_B$, with $C = \max\{A, B\}$, one has $xy = \rho_{A,C}(x) \cdot_C \rho_{B,C}(y)$; the order is directed lexicographic, comparing after transport to C and using the index direction to break ties. Thus $\mathfrak{D}(\mathbf{M})$ determines $\mathbf{M}$ up to canonical isomorphism.

Conversely, every rigid finite chain-indexed direct system of finite local-unit-aligned totally ordered monoids reconstructs a monoid on the disjoint union of the components, with this product and order. This gives the canonical rigid direct-system representation theorem: every finite local-unit-aligned totally ordered monoid decomposes canonically into τ-multiplication-cohesive components connected by unit-constant transition maps, and every such rigid system reconstructs a monoid of the same class.

This theorem has a Clifford-type character. Clifford's classical representation theory reconstructs suitable inverse semigroups from simpler components connected over a semilattice, and Plonka-type sums extend this philosophy to universal algebra. The present result is not a special case of those theorems: there is no inverse operation, the structures are totally ordered monoids, and the reconstruction must recover the linear order as well as multiplication. Nevertheless, the organizing principle is analogous. A global algebraic object is recovered from canonical components and connecting morphisms, now determined by local units rather than by inverses or externally imposed semilattice data.

V. Functorial Form

The representation is also functorial in two complementary categories. The first categorical formulation uses strict block-preserving homomorphisms. Such a homomorphism sends each canonical component of the source into a canonical component of the target and preserves the corresponding block unit. These maps match the usual componentwise morphisms of rigid direct systems, giving a categorical equivalence between finite local-unit-aligned totally ordered monoids with block-unit-preserving morphisms and rigid finite direct systems. The second formulation is more intrinsic. A homomorphism $f\colon \mathbf{M} \to \mathbf{N}$ is τ-compatible when it is isotone, unital, multiplicative, and satisfies $\tau_{\mathbf{N}} \circ f = f \circ \tau_{\mathbf{M}}$. Such a map need not be injective on positive idempotents. It may collapse

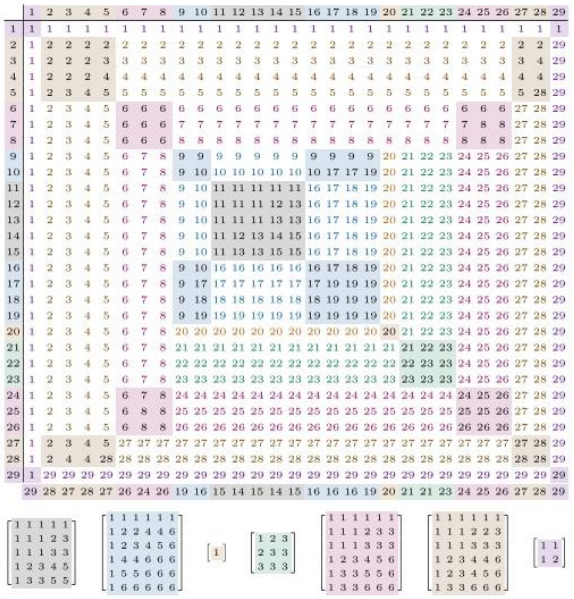

Fig. 1. An example of the canonical τ-multiplication-coherent decomposition. The upper table is the ambient monoid together with the τ-row, and the lower row displays the component monoids. The cells without background coloring illustrate the rigidity phenomenon.

several canonical components of the source into a single component of the target. On the direct-system side this is represented by non-injective isotone index maps together with compatible component morphisms. Rigidity is what makes this collapse well-defined: it prevents a source component from being split by pullback along a target decomposition and supplies the absorption behavior required when several source components are identified.

VI. Conclusion

The paper identifies the canonical τ-multiplication-coherent decomposition as the intrinsic structure governing finite local-unit-aligned totally ordered monoids. This decomposition produces component monoids and transition maps forming a rigid finite direct system from which the original order and multiplication are recovered. Conversely, rigid systems of this type reconstruct exactly the same class of monoids. The resulting theory is transparent, canonical, and functorial, culminating in two categorical equivalences: one for strict block-unit-preserving homomorphisms and one for intrinsic τ-compatible homomorphisms.

References

[1] L. Fuchs, *Partially Ordered Algebraic Systems*. Pergamon Press, Oxford, 1963.

[2] A. H. Clifford, "Semigroups admitting relative inverses," *Annals of Mathematics*, vol. 42, pp. 1037–1049, 1941.

[3] J. Płonka, "On a method of construction of abstract algebras," *Fundamenta Mathematicae*, vol. 61, pp. 183–189, 1967.

[4] S. Jenei, "Group representation for even and odd involutive commutative residuated chains," *Studia Logica*, vol. 110, pp. 881–922, 2022.

Resource Orchestration and Scheduling Across the Mist–Fog–Cloud Continuum: Trends and Future Directions

Helen D. Karatza
Department of Informatics
Aristotle University of Thessaloniki
Thessaloniki, Greece
karatza@csd.auth.gr

***Abstract*—The rapid proliferation of Internet of Things (IoT) applications is increasingly exposing the limitations of centralized cloud infrastructures in accommodating the scale, velocity, and geographical distribution of data generated by pervasive sensing devices. To address the stringent latency, bandwidth, and scalability requirements of emerging IoT services, the Cloud computing paradigm has evolved towards a hierarchical computing continuum that incorporates Fog and Mist computing. By moving computation progressively closer to the data sources, Fog computing enables low-latency processing at the network edge, while Mist computing extends this capability to resource-constrained devices, supporting ultra-low-latency execution directly at or near the sensing layer.**

Many contemporary IoT applications, including industrial automation, autonomous systems, intelligent transportation, smart healthcare, and critical infrastructure monitoring, are inherently time-critical and operate under strict Quality-of-Service (QoS) and deadline constraints. Consequently, efficient resource management and scheduling across the cloud–fog–mist continuum have become fundamental challenges for ensuring predictable performance, scalability, and reliable service delivery in highly distributed environments.

This keynote explores the state of the art in scheduling methodologies for heterogeneous cloud–fog–mist infrastructures, with particular emphasis on real-time scheduling, workflow scheduling, fault-tolerant execution, energy-aware resource management, and Quality-of-Service optimization. Modern applications exhibit diverse computational structures, including independent Bag-of-Tasks workloads, tightly coupled parallel applications requiring Gang scheduling, and complex workflows represented as Directed Acyclic Graphs (DAGs), where precedence constraints and end-to-end deadlines significantly complicate scheduling decisions.

A particular focus is placed on real-time scheduling techniques based on imprecise (approximate or partial) computation. Imprecise computation provides an effective mechanism for addressing overload conditions by allowing applications to trade computational precision for timeliness. Rather than producing an exact result after its deadline—when its value may be substantially diminished—a computation may instead return an intermediate result of acceptable quality before its deadline. This approach assumes monotonic applications, whose output quality improves progressively as execution advances. Such applications are typically modelled as consisting of a mandatory component, which must complete to ensure functional correctness, followed by an optional component that incrementally improves result quality. Although approximate execution is not universally applicable, as some workloads require exact computation, it offers an effective scheduling mechanism for a wide class of time-critical applications.

The keynote further investigates how imprecise computation can be integrated with application-directed checkpointing to provide fault tolerance in large-scale distributed real-time systems. By periodically checkpointing intermediate computation states—particularly upon completion of the mandatory execution phase—applications can recover efficiently from software failures while preserving timing guarantees and reducing recovery overhead. This combination of adaptive execution and fault-tolerant scheduling offers a promising approach for improving the resilience of cloud-native real-time applications.

The unprecedented growth of data generated by sensors, mobile devices, cyber-physical systems, social media, and large-scale IoT deployments has also driven the emergence of Big Data analytics, whose computational demands increasingly rely on distributed cloud–fog–mist infrastructures. Many data-intensive analytics applications exhibit real-time characteristics, as delayed results may significantly reduce their operational value. Consequently, scalable scheduling algorithms must simultaneously address computational performance, communication overhead, timing constraints, reliability, and heterogeneous resource availability.

Energy efficiency and sustainability constitute another fundamental dimension of scheduling in modern distributed computing infrastructures. As data centers and edge platforms continue to expand, reducing energy consumption and carbon emissions has become a primary design objective. Energy-aware scheduling strategies must therefore balance performance, deadline satisfaction, resource utilization, and Quality-of-Service requirements while minimizing operational cost and environmental impact, contributing to the broader vision of sustainable cloud and edge computing.

The continued evolution from centralized cloud infrastructures towards fully integrated cloud–fog–mist ecosystems is fundamentally transforming the way distributed computing platforms are designed, deployed, and managed. Future computing environments will be characterized by highly heterogeneous resources, dynamic workloads, geographically distributed infrastructures, and increasingly autonomous edge devices, all of which introduce new challenges for efficient resource orchestration. In this context, scheduling is emerging as a key enabling technology for achieving scalability, predictability, resilience, and sustainability across the entire computing continuum.

Designing scheduling strategies that can simultaneously satisfy real-time constraints, optimize resource utilization,

979-8-3195-0700-6/26/$31.00 ©2026 IEEE

tolerate failures, minimize energy consumption, and adapt to highly dynamic execution environments remain a major research challenge. At the same time, the growing integration of artificial intelligence, machine learning, digital twins, and data-driven optimization into distributed computing systems is creating new opportunities for intelligent, self-adaptive scheduling and autonomous resource management. These developments are expected to play a central role in supporting next-generation IoT applications, including cyber-physical systems, smart cities, Industry 5.0, connected healthcare, autonomous transportation, and large-scale digital infrastructures.

This keynote reviews recent advances in scheduling algorithms and resource management frameworks across the cloud–fog–mist continuum, presents representative research contributions addressing real-time execution, fault tolerance, approximate computing, workflow management, and energy-aware optimization, and discusses the open research challenges that must be addressed to realize scalable, dependable, and sustainable distributed computing platforms for future intelligent IoT ecosystems.

REFERENCES

[1] G. L. Stavrinides and H. D. Karatza, "A hybrid approach to scheduling real-time IoT workflows in fog and cloud environments," *Multimedia Tools and Applications*, Springer, vol. 78, no. 17, pp. 24639-24655, Sep. 2019.

[2] G. L. Stavrinides and H. D. Karatza, "An energy-efficient, QoS-aware and cost-effective scheduling approach for real-time workflow applications in cloud computing systems utilizing DVFS and approximate computations," *Future Generation Computer Systems*, Elsevier, vol. 96, pp. 216-226, Jul. 2019.

[3] G. L. Stavrinides and H. D. Karatza, "Security, cost and energy aware scheduling of real-Time IoT workflows in a mist computing environment," *Information Systems Frontiers*, Springer, vol. 26, no. 4, Jul. 2024.

[4] H. D. Karatza and G. L. Stavrinides, "Resource allocation and aging priority-based scheduling of linear workflow applications with transient failures and selective imprecise computations," *Cluster Computing*, Springer, vol. 27, iss. 4, pp. 5473–5488, Jul. 2024.

[5] H. D. Karatza, "Scheduling linear workflows with multicriteria based partial computations in a cloud–fog–mist environment," in *Proceedings of the 9th International Conference on Fog and Mobile Edge Computing (FMEC'24)*, Malmö, Sweden. Sep. 2024, pp. 107-113.

[6] H. D. Karatza, "Scheduling bag-of-task jobs with security requirements and dual-criteria partial computations in a fog–cloud system," *Simulation Modeling Practice and Theory*, Elsevier, vol. 140, 103082, Apr. 2025.

[7] H. D. Karatza, "Scheduling gang jobs with security requirements on a hybrid cloud," in *Proceedings of the 2025 International Conference on Computer, Information and Telecommunication Systems (CITS'25)*, Colmar, France, Jul. 2025, pp. 1-8.

[8] H. D. Karatza, "Scheduling mixed workloads with security requirements in a cloud-fog-mist computing environment," *Simulation Modelling Practice and Theory*, Elsevier, vol. 146, 103231, Jan. 2026.

AI-Enhanced Integrated Modelling for Urban Climate and Air Quality

Binh Nguyen[1,2], Fabian Wagner[1], Gregor Kiesewetter[1], Zbigniew Klimont[1], Susanna Dedring[1]
[1] International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria
[2] Faculty of Informatics, University of Debrecen, Debrecen, Hungary

Abstract. Translating regional climate and air-quality strategies into effective urban actions requires consistent integration of regional scenarios with detailed city-level information. This paper presents an AI-enhanced integrated modelling approach for urban climate and air-quality assessment. The methodology combines two main steps. First, regional activity pathways are downscaled to construct internally consistent city scenarios covering activities, technologies, emissions, costs and environmental impacts. Artificial intelligence and heterogeneous urban data support sector-specific regional-to-city transformations and high-resolution assessment. Second, AI-assisted retrieval is used to identify and adapt mitigation measures relevant to the city-specific activity and technology structure. The proposed approach combines established domain models with data-driven methods to support consistent and transparent urban policy assessment.

Keywords— integrated assessment modelling, artificial intelligence, GAINS City, air quality, urban modelling, downscaling, decision support.

I. INTRODUCTION

Climate change and air pollution are closely interconnected challenges requiring coordinated action across energy, transport, buildings and industry. Integrated assessment models provide a framework for linking socioeconomic activities, technologies and emission-control measures with emissions, environmental impacts and mitigation costs. The GAINS model developed at the International Institute for Applied Systems Analysis has been widely used to assess air-pollution and greenhouse-gas mitigation strategies at global, regional and national scales [1].

Many mitigation measures, however, are implemented at the city level, where emission sources, population exposure and policy priorities can differ substantially from regional averages. The GAINS City concept was developed to bridge this scale gap by extending the GAINS integrated-assessment approach to urban policy analysis. The methodology combines baseline development with the representation of policy interventions and assessment of their effects on emissions, air quality, health and costs [2]. The earlier PMEH implementation demonstrated this approach through an interactive urban decision-support framework.

Since this methodology was developed, the availability and resolution of urban data have increased substantially. Traffic networks, building and heating information, industrial facilities, energy systems, population distributions and environmental observations can provide a much richer description of urban activities and emission sources. At the same time, high-resolution atmospheric modelling, including uEMEP, provides new opportunities to connect regional air-quality modelling with local concentrations and population exposure [3]. These developments create both an opportunity and a challenge: regional GAINS scenarios need to be translated consistently to the city scale without losing the relationships among activities, technologies, emissions, costs and impacts that constitute the integrated-assessment framework.

This paper proposes an AI-enhanced extension of the GAINS City methodology addressing this challenge in two steps. First, regional GAINS activity pathways and associated model information are transformed into a consistent city baseline using local data and AI-assisted regional-to-city downscaling, while uEMEP provides high-resolution air-quality and exposure assessment. Second, the resulting city baseline is used to retrieve and adapt relevant mitigation measures from the regional GAINS system. AI supports the matching of regional information and policy options to city-specific conditions, while established GAINS and atmospheric models retain the underlying scenario, physical and policy relationships.

II. AI-ENHANCED GAINS CITY METHODOLOGY

The proposed methodology extends the GAINS City concept through two main steps: (i) downscaling a regional GAINS scenario to the city level, and (ii) retrieving and adapting relevant mitigation measures from regional GAINS for urban policy assessment. AI supports both steps by exploiting heterogeneous city data, while GAINS and atmospheric models retain the underlying integrated-assessment and physical relationships.

A. Data Integration and Downscaling

GAINS City extends the regional GAINS structure to the city level. Let denote the parent GAINS region r, the city c a GAINS sector s, an activity a, and a scenario year y. City activity is derived as

$$A_{c,s,a,y} = A_{r,s,a,y} SF_{c,s,a,y} \quad (1)$$

with the sector- and activity-specific scaling factor

$$SF_{c,s,a,y} = \frac{I_{c,s,a,y}}{I_{r,s,a,y}} \quad (2)$$

The indicators I depend on the activity being downscaled. Transport may use traffic volumes, road networks and vehicle fleets; buildings may use population, floor area, heating systems and energy demand; and industry may use facility production, capacity and technology information.

Where local indicators are incomplete, AI/ML methods support their estimation from heterogeneous urban data. Regression and tree-based algorithms such as Random Forest or Gradient Boosting can combine available local features to estimate missing activity information. ML-based imputation can fill data gaps, while anomaly-detection methods can identify inconsistent observations. Reliable local statistics are used directly where available and provide reference data for training and validation. The resulting indicators are then used to derive the scaling factors applied to the corresponding GAINS activity pathways.

The downscaled activities retain the corresponding GAINS fuels, technologies and emission controls, adjusted where city-specific information is available. Emissions and

mitigation costs are recalculated within GAINS rather than independently downscaled, maintaining consistency among activities, technologies, emissions and costs.

For air-quality assessment, city emissions are further allocated spatially and temporally using sector-specific local data and linked to uEMEP [3], [5]. uEMEP combines high-resolution urban emissions with regional atmospheric information to calculate concentrations and population exposure, which provide the basis for health-impact assessment.

The resulting GAINS City baseline therefore provides an internally consistent chain from activities and technologies to emissions, costs, concentrations, exposure and impacts, with AI/ML supporting the data-intensive region-to-city downscaling process.

B. Retrieval and Adaptation of Policy Measures

The second step identifies policy measures represented in regional GAINS that are applicable to the city baseline established in Step 1. The original GAINS City methodology represents measures through changes in activities, fuels, technologies or emission controls and accounts for interactions among measures [2]. Rather than defining a separate set of urban measures, the proposed approach exploits the existing GAINS measure database and adapts relevant options to the city-specific activity and technology structure.

For a regional set of GAINS measures M^{reg}, the selection of city-relevant measures can be represented as

$$M^{city} = \mathcal{R}\left(M^{reg}, S^{city}_{baseline}\right), \quad (2)$$

where $S^{city}_{baseline}$ represents the GAINS City baseline developed in Step 1, including its sectoral activities, fuels, technologies, emissions, costs and impacts, and $\mathcal{R}$ denotes the retrieval and matching process.

AI can support this process by matching the characteristics and applicability conditions of regional measures with the corresponding GAINS City sectors, activities, fuels and technologies. It can therefore help identify which measures are technically relevant, rank them according to their potential importance for the city, and detect possible interactions among measures. Local policy information and expert knowledge can subsequently be used to validate the retrieved options and define feasible implementation levels.

Selected measures are then implemented through the established GAINS relationships. Depending on the type of intervention, a measure may modify an activity pathway, shift activity between fuels or technologies, or change the penetration of an emission-control technology. Emissions and costs are consequently recalculated rather than independently scaled, preserving consistency with the GAINS City baseline and the underlying regional scenario.

The environmental consequences of each measure or combination of measures are evaluated using the same modelling procedure as for the baseline. Changes in city emissions are spatially and temporally allocated and processed by uEMEP to estimate changes in concentrations and population exposure, from which health and other impacts can be derived. This provides a consistent framework for comparing policy options in terms of their activity changes, emission reductions, costs and environmental and health benefits.

III. CURRENT DEVELOPMENT AND OUTLOOK

The proposed methodology builds on ongoing development of GAINS City within the Horizon Europe ATMOPOLIS project [4], which provides a practical setting for linking regional integrated assessment with higher-resolution urban emissions, air quality, population exposure, health impacts and costs. The approach will support the construction of consistent city baselines from regional GAINS scenarios and the retrieval and adaptation of policy measures to different urban contexts.

Implementation also requires the integration of data-driven methods with the existing GAINS modelling environment. Current development explores Python-based analytical and AI functionality within the Java-based application to support data integration, regional-to-city downscaling and measure retrieval while retaining the established GAINS calculations and scenario structure.

A proposed future application is CITY-INSIGHT, developed in cooperation with the University of Debrecen, with Debrecen as a case-study city. The concept aims to combine GAINS City with local activity and emission data, high-resolution air-quality modelling and urban monitoring. This would provide a practical environment for evaluating AI-assisted regional-to-city downscaling and measure retrieval using city-specific data and local expertise, and for assessing alternative strategies in terms of emissions, population exposure, health and climate-related impacts.

In the longer term, closer integration of regional scenarios, urban data, atmospheric modelling and observations could provide building blocks for environmental Digital Twins supporting urban climate and air-quality decision-making.

IV. CONCLUSIONS

This paper proposes an AI-enhanced extension of GAINS City for translating regional integrated assessment into urban-scale policy analysis. The methodology follows two main steps: first, regional GAINS activity pathways are downscaled to construct internally consistent city scenarios, from which technologies, emissions and costs are represented within the GAINS framework and high-resolution air quality, exposure and impacts are assessed using uEMEP; second, relevant mitigation measures are retrieved from the regional GAINS system and adapted to the resulting city-specific activity and technology structure. AI supports both transformations by exploiting heterogeneous urban data for regional-to-city scaling and by assisting the identification and matching of policy measures, while established GAINS and atmospheric models retain the underlying integrated-assessment and physical relationships. Ongoing and future applications will provide opportunities to test this approach across different urban contexts and evaluate its potential for more data-informed climate and air-quality decision support.

REFERENCES

[1] M. Amann *et al.*, "Cost-effective control of air quality and greenhouse gases in Europe: Modeling and policy applications," *Environmental Modelling & Software*, vol. 26, no. 12, pp. 1489–1501, 2011.

[2] F. Wagner, J. Borken-Kleefeld, G. Kiesewetter, Z. Klimont, B. Nguyen, P. Rafaj, and W. Schöpp, *The GAINS PMEH-Methodology – Version 2.0*, IIASA Working Paper

WP-18-016, International Institute for Applied Systems Analysis, Laxenburg, Austria, 2018.
[3] B. R. Denby, M. Gauss, P. Wind, Q. Mu, E. G. Wærsted, H. Fagerli, A. Valdebenito, and H. Klein, "Description of the uEMEP_v5 downscaling approach for the EMEP MSC-W chemistry transport model," *Geoscientific Model Development*, vol. 13, pp. 6303–6323, 2020.
[4] ATMOPOLIS Consortium, "Air and Noise Pollution Mitigation Through Integrated, Society-Centred Actions," Horizon Europe project, Grant Agreement No. 101238690, 2026–2029.
[5] Q. Mu, B. R. Denby, E. G. Wærsted, and H. Fagerli, "Downscaling of air pollutants in Europe using uEMEP_v6," *Geoscientific Model Development*, vol. 15, pp. 449–465, 2022.

Recognition of GNSS Spoofing in V2I Environments with Physics-Informed Gaussian Process

Tamás Ormándi
Dept. of Control for Transportation and Vehicle Systems
Budapest University of Technology and Economics
Budapest, Hungary
ormandi.tamas@edu.bme.hu

Balázs Varga
Dept. of Control for Transportation and Vehicle Systems
Budapest University of Technology and Economics
Budapest, Hungary
varga.balazs@kjk.bme.hu

***Abstract*—V2X reliability is paramount for intelligent transportation systems, especially adaptive applications like adaptive traffic lights which are highly sensitive to positioning errors. This abstract addresses V2X vulnerability to GNSS spoofing by proposing an infrastructure-centric detection framework. A physics-informed Gaussian Process Regression model synthesizes Received Signal Strength Indicator metrics and road topology to establish a probabilistic wireless baseline. By cross-referencing claimed coordinates against expected signal behavior, the framework identifies anomalies via physical residuals and epistemic uncertainty, achieving a 95% detection probability at a 2.1-meter offset in severe Non-Line-of-Sight urban environments.**

***Index Terms*—Vehicular communication, GPS Spoofing, Cyber security, Intelligent transportation, Gaussian Process**

I. Introduction

As smart cities integrate connected vehicles, Vehicle-to-Everything (V2X) technology becomes foundational for safety. However, location-dependent services are fragile against GNSS spoofing [1]. Manipulated position errors can induce phantom braking or cause severe rear-end collisions [2].

While vehicle-centric verification is frequently employed [3], it remains susceptible to environmental degradation or jamming. This work introduces a robust secondary defense: a Vehicle-to-Infrastructure (V2I) GNSS spoofing recognition framework. By shifting verification to Roadside Units (RSUs), our system leverages physical layer metrics from Cooperative Awareness Messages (CAMs) [4] and Spatially-Correlated GPR [5] to validate position claims independently. Standard RSSI-based localization often relies on complex sensor fusion (e.g., Kalman filtering) or multi-vehicle triangulation to achieve a $2.2 - 3.2\,\mathrm{m}$ tracking accuracy [7]. The key contribution of this work is the combination of Matérn and RBF kernels to explicitly decouple spatial topology from radio-physical propagation. Unlike these state-of-the-art architectures, the proposed framework requires neither multi-node geometric collaboration nor temporal filtering; a single RSU independently validates a single vehicle's position claim via epistemic uncertainty evaluation.

Project no. EKÖP-25-3-BME-43 has been implemented with the support provided by the Ministry of Culture and Innovation of Hungary from the National Research, Development and Innovation Fund.

II. Methodology

A. System and Threat Model

A trusted, closed RSU infrastructure is assumed. The adversary is considered capable of performing arbitrary spatial spoofing via CAM position injection, but lacks the capability to macroscopically manipulate the physical radio channel characteristics (RSSI) measured at the RSU. Complex attack vectors are outside the scope of this physical-layer validation framework.

B. Data Acquisition and Feature Extraction

Measurements utilized Cohda Wireless MK5 modules in an urban NLOS environment. The passive RSU extracted geospatial coordinates (ϕ, λ) and dual-antenna RSSI values (R_1, R_2) from 10 Hz CAMs. To mitigate channel fluctuations, the feature vector $\mathbf{x} = [\max(R_1, R_2), |R_1 - R_2|, \phi, \lambda]^\top$ synthesizes signal power and spatial geometry. The objective is to approximate the latent function $f : \mathbf{x} \to y \in \mathbb{R}^+$, where $y = \log(d)$ represents logarithmic Euclidean distance, stabilizing path loss variance.

C. Physics-Informed Gaussian Process

Gaussian Process Regression (GPR) maps the input vector to the target distance using a composite additive kernel designed to decouple radio-physical propagation from spatial topology, formulated as $k(\mathbf{x}, \mathbf{x}') = \sigma^2_{phy} k_{RBF}(\mathbf{x}_{phy}, \mathbf{x}'_{phy}) + \sigma^2_{geo} k_{Mat}(\mathbf{x}_{geo}, \mathbf{x}'_{geo})$. Here, $\mathbf{x}_{phy}$ uses a Radial Basis Function (RBF) kernel for smooth signal decay, while $\mathbf{x}_{geo}$ uses a Matérn ($\nu = 3/2$) kernel for abrupt shadowing variations. Training utilized GPyTorch [6].

D. Anomaly Detection Formulation

The GPR yields a predictive mean (μ_{GP}) and variance (σ^2_{GP}). Predictions map back to physical distance ($\hat{d}_*$). A dual-threshold logic evaluates Physics Mismatch and Epistemic

Uncertainty. Thresholds ($\tau_{res} = 0.62$m, $\tau_{unc} = 0.1$) are calibrated to the training set's 99th percentile. Spoofing is flagged if the following condition is met: $\mathbb{I}(|d_{claim} - \hat{d}_*| > \tau_{res}) \vee \mathbb{I}(\sigma_{GP}(\mathbf{x}'_*) > \tau_{unc})$.

III. Results and Performance

A. Detection Sensitivity and Boundaries

Simulating a continuous spectrum of spoofing attacks (Fig. 1) demonstrates rapid convergence. A 95% detection probability occurs at a 2.1-meter deviation, reaching 100% beyond 2.5 meters.

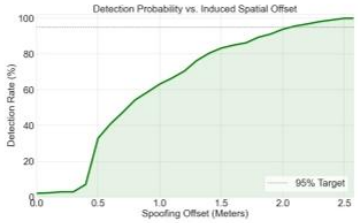

Fig. 1. Detection rate relative to the induced spatial spoofing offset.

Robustness stems from evaluating the residual rather than absolute distance. Figure 2 visualizes this logic. Minor spoofing ($1 - 5$m) drives vertical migration, triggering the residual threshold, while larger deviations (> 10m) shift into high-uncertainty regions.

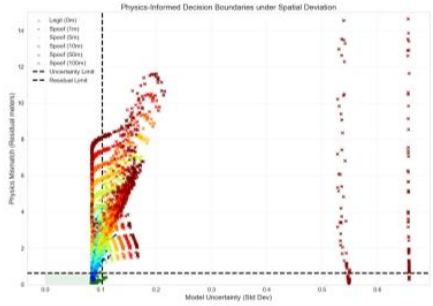

Fig. 2. Physics-Informed decision boundaries under spatial deviation.

B. False Alarm Rate (FAR) Analysis

An audit across a 100m zone (Fig. 3) revealed near-zero FAR in well-mapped areas. Sporadic false detections occurred exclusively in narrow "blind spots" lacking training data, where epistemic uncertainty (σ_{GP}) naturally expands, underscoring the need for uniform spatial calibration.

C. Adversarial Robustness (White-Box)

To evaluate worst-case boundaries, the detector was subjected to a simulated white-box optimization attack. The adversary possesses complete parametric knowledge of the system and seeks to maximize the longitudinal position offset without triggering an anomaly alert. While theoretical evasion under these optimized conditions is high (99.33%), the framework severely bounds the *magnitude* of undetected spoofing.

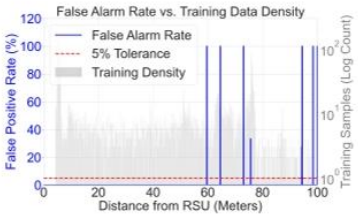

Fig. 3. False Alarm Rates strongly correlate with training data sparsity.

Globally, the mean stealth offset is restricted to 1.33 m. Crucially, in pure Line-of-Sight (LOS) sub-domains—typical during intersection approaches—this mean stealth offset collapses to just 0.79 m (with the 95^{th} percentile reaching only 2.18 m). This demonstrates that even under perfect adversarial knowledge, the spoofing impact is strictly constrained to sub-meter geometries.

IV. Conclusion

This framework utilizes a composite GPR to detect spoofed spatial offsets as small as 2.1 m in NLOS environments by evaluating physical mismatch and epistemic uncertainty. The system operates on a standalone, single-RSU basis without requiring cooperative vehicle triangulation or multi-sensor fusion. Evaluation under a worst-case white-box optimization attack demonstrates that the mean undetected stealth offset is restricted to 1.33 m globally, and collapses to a sub-meter scale of 0.79 m in pure LOS sub-domains. Consequently, spatial manipulation is strictly bounded even under ideal adversarial knowledge. Future work will explore the virtual pre-training of the detector system.

References

[1] Z. Tzoannos et al., "The Impact of Spoofing Attacks in Connected Autonomous Vehicles under Traffic Congestion Conditions," *Telecom*, vol. 5, no. 3, pp. 747–759, 2024.

[2] A. T. Sheik et al., "A Comprehensive Survey of Threats in Platooning—A Cloud-Assisted Connected and Autonomous Vehicle Application," *Information*, vol. 15, no. 1, 2024.

[3] S. Dasgupta et al., "A Sensor Fusion-Based GNSS Spoofing Attack Detection Framework for Autonomous Vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 12, pp. 23559–23572, 2022.

[4] ETSI, "ITS; Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service," ETSI EN 302 637-2, 2019.

[5] S. Chen et al., "Gaussian process regression for enhancing the positional accuracy of collaborative robot," *Measurement*, vol. 256, p. 118142, 2025.

[6] J. R. Gardner et al., "GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration," in *NeurIPS*, 2018.

[7] R. Parker and S. Valaee, "Vehicular Node Localization Using Received-Signal-Strength Indicator," *IEEE Trans. Veh. Technol.*, vol. 56, no. 6, pp. 3371–3380, 2007.