BOOK OF ABSTRACTS

The 9th International Conference on the Dynamics of Information Systems (DIS 2026)

May 31- June 4, 2026, Czech Republic

UNIVERZITA J. E. PURKYNĚ V ÚSTÍ NAD LABEM

JAN EVANGELISTA PURKYNĚ UNIVERSITY IN ÚSTÍ NAD LABEM
Faculty of Science

Springer

CAO Center for Applied Optimization

UF UNIVERSITY of FLORIDA

iSoGO

OXEUS CLOUD SOLVER

Conference Organization

General Chairs:

- Hossein Moosaei, Jan Evangelista Purkyně University, Czech Republic
- Zbyšek Posel, Jan Evangelista Purkyně University, Czech Republic
- Panos M. Pardalos, University of Florida, USA

Program Committee Chairs:

- Hossein Moosaei, Jan Evangelista Purkyně University, Czech Republic
- Zbyšek Posel, Jan Evangelista Purkyně University, Czech Republic
- Panos M. Pardalos, University of Florida, USA

Organizing Committee Chairs:

- Natálie Růtová, Jan Evangelista Purkyně University, Czech Republic

Organizing Committee:

- Jiří Škvor, Jan Evangelista Purkyně University, Czech Republic
- Pavel Beránek, Jan Evangelista Purkyně University, Czech Republic
- Jiří Škvára, Jan Evangelista Purkyně University, Czech Republic
- Michaela Tichá, Jan Evangelista Purkyně University, Czech Republic
- Alena Černíková, Jan Evangelista Purkyně University, Czech Republic
- Lenka Bauerová, Jan Evangelista Purkyně University, Czech Republic

Contents

Invited Speakers

Panos M. Pardalos, University of Florida, USA

A New Frontier: From a Single Network to a Network of Networks

This lecture examines the fundamental shift from isolated, monolithic systems to the expansive "Network of Networks" architecture that underpins modern global infrastructure. We move beyond traditional single-layer analysis to explore the intricate interdependencies among critical domains—for example, the Energy–Financial nexus, where real-time market signals influence grid stability, and the Transportation–Digital nexus, where autonomous logistics depend on ubiquitous communication.

Problems in networks of networks are far more complex than those in single networks. For example, in a single network, the propagation of failures can often be predicted and contained. In contrast, within a "Network of Networks," such failures become exponentially more difficult to anticipate due to hidden interdependencies—connections that remain invisible until they trigger cascading and often unpredictable effects.

Biography: Panos Pardalos received his PhD (Computes and Information Sciences) from the University of Minnesota. He is an Emeritus Distinguished Professor in the Department of Industrial and Systems Engineering at the University of Florida, and an affiliated faculty of Biomedical Engineering and Computer Science & Information & Engineering departments.

Panos Pardalos is a world-renowned leader in Global Optimization, Mathematical Modeling, Energy Systems, Financial applications, and Data Sciences. He is a Fellow of AAAS, AAIA, AIMBE, EUROPT, and INFORMS and was awarded the 2013 Constantin Caratheodory Prize of the International Society of Global Optimization. In addition, Panos Pardalos has been awarded the 2013 EURO Gold Medal prize bestowed by the Association for European Operational Research Societies. This medal is the preeminent European award given to Operations Research (OR) professionals for "scientific contributions that stand the test of time."

Panos Pardalos has been awarded a prestigious Humboldt Research Award (2018-2019). The Humboldt Research Award is granted in recognition of a researcher's entire achievements to date – fundamental discoveries, new theories, insights that have had significant impact on their discipline.

Panos Pardalos is also a Member of several Academies of Sciences, and he holds several honorary PhD degrees and affiliations. He is the Founding Editor of Optimization Letters, Energy Systems, and Co-Founder of the International Journal of Global Optimization, Computational Management Science, and Springer Nature Operations Research Forum. He has published over 600 journal papers, and edited/authored over 200 books. He is one of the most cited authors and has graduated 71 PhD students so far. Details can be found in www.ise.ufl.edu/pardalos

Panos Pardalos has lectured and given invited keynote addresses worldwide in countries including Austria, Australia, Azerbaijan, Belgium, Brazil, Canada, Chile, China, Czech Republic, Cyprus, Denmark, Egypt, England, France, Finland, Germany, Greece, Holland, and etc.

Peter Richtárik, KAUST, Saudi Arabia

From the Ball-proximal (Broximal) Point Method to Efficient Training of Large Language Models

Non-smooth and non-convex global optimization poses significant challenges across various applications, where standard gradient-based methods often struggle. We propose the Ball-Proximal Point Method, Broximal Point Method, or Ball Point Method (BPM) for short – a novel algorithmic framework inspired by the classical Proximal Point Method (PPM) [8], which, as we show, sheds new light on several foundational optimization paradigms and phenomena, including non-convex and non-smooth optimization, acceleration, smoothing, adaptive stepsize selection, and trust-region methods. At the core of BPM lies the ball-proximal ("broximal") operator, which arises from the classical proximal operator by replacing the quadratic distance penalty by a ball constraint. Surprisingly, and in sharp contrast with the sublinear rate of PPM in the nonsmooth convex regime, we prove that BPM converges linearly and in a finite number of steps in the same regime. Furthermore, by introducing the concept of ball-convexity, we prove that BPM retains the same global convergence guarantees under weaker assumptions, making it a powerful tool for a broader class of potentially non-convex optimization problems. Just like PPM plays the role of a conceptual method inspiring the development of practically efficient algorithms and algorithmic elements, e.g., gradient descent, adaptive step sizes, acceleration [1], and "W" in AdamW [9], we believe that BPM should be understood in the same manner: as a blueprint and inspiration for further development. Generalization non-Euclidean ball constraints can be found in the follow-up work [3].

The Broximal Point Method (BPM) [2] offers an idealized optimization framework based on iteratively minimizing the objective function over norm balls centered at the current iterate. It enjoys striking global convergence guarantees, converging linearly and in a finite number of steps for proper, closed and convex functions. However, its theoretical analysis has so far been confined to the Euclidean geometry. At the same time, emerging trends in deep learning optimization, exemplified by algorithms such as Muon [4] and Scion [6], demonstrate the practical advantages of minimizing over balls defined via non-Euclidean norms which better align with the underlying geometry of the associated loss landscapes. We ask whether the convergence theory of BPM can be extended to this more general, non-Euclidean setting. We give a positive answer, showing that most of the elegant guarantees of the original method carry over to arbitrary norm geometries. Along the way, we clarify which properties are preserved and which necessarily break down when leaving the Euclidean realm. Our analysis positions Non-Euclidean BPM as a conceptual blueprint for understanding a broad class of geometry-aware optimization algorithms, shedding light on the principles behind their practical effectiveness.

Latest developments in deep learning optimization have brought about radically new algorithms based on the Linear Minimization Oracle (LMO) framework, such as Muon [4] and Scion [6]. After over a decade of Adam's [5] dominance, these LMO-based methods are emerging as viable replacements, offering several practical advantages such as improved memory efficiency, better hyperparameter transferability, and most importantly, superior empirical performance on large-scale tasks, including LLM training. However, a significant gap remains between their practical use and our current theoretical understanding: prior analyses (1) overlook the layer-wise LMO application of these optimizers in practice, and (2) rely on an unrealistic smoothness assumption, leading to impractically small stepsizes. To address both, we propose a new LMO-based method called Gluon, capturing prior theoretically analyzed methods as special cases, and introduce a new refined generalized smoothness model that captures the layer-wise geometry of neural networks, matches the layer-wise practical implementation of Muon and Scion, and leads to con- vergence guarantees with strong practical predictive power. Unlike prior results, our theoretical stepsizes closely match the fine-tuned values reported in [6]. Our experiments with NanoGPT and CNN confirm that our assumption holds along the optimization trajectory, ultimately closing the gap between theory and practice.

Recent optimizers like Muon [4], Scion [6], and Gluon [7] have pushed the frontier of large-scale deep learning by exploiting layer-wise linear minimization oracles (LMOs) over non-Euclidean norm balls, capturing neural network structure in ways traditional algorithms cannot. Yet, no principled distributed framework exists for these methods, and communication bottlenecks remain unaddressed. The very few distributed variants are heuristic, with no convergence guarantees in sight. We introduce EF21-Muon, the first communication-efficient, non-Euclidean LMO-based optimizer with rigorous convergence guarantees. EF21-Muon supports stochastic gradients, momentum, and bidirectional compression with error feedback-marking the first extension of error feedback beyond the Euclidean setting. It recovers Muon/Scion/Gluon when compression is off and specific norms are chosen, providing the first efficient distributed implementation of this powerful family. Our theory covers non-Euclidean smooth and the more general communication savings with no accuracy degradation.

References

[1] Kwangjun Ahn and Suvrit Sra. “Understanding Nesterov’s Acceleration via Proximal Point Method”. Symposium on Simplicity in Algorithms (SOSA), 2022, pp. 117–130.

[2] Kaja Gruntkowska, Hanmin Li, Aadi Rane, and Peter Richtárik. “The ball-proximal (=”broximal”) point method: a new algorithm, convergence theory, and applications”. arXiv preprint arXiv:2502.02002, 2025.

[3] Kaja Gruntkowska and Peter Richtárik. Non-Euclidean broximal point method: a blueprint for geometry-aware optimization. arXiv preprint arXiv:2510.00823, 2025

[4] Keller Jordan, Yuchen Jin, Vlado Boza, Jiacheng You, Franz Cesista, Laker Newhouse, and Jeremy Bernstein. Muon: An optimizer for hidden layers in neural networks, 2024.

[5] Diederik P. Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: arXiv preprint arXiv:1412.6980 (2014).

[6] Thomas Pethick, Wanyun Xie, Kimon Antonakopoulos, Zhenyu Zhu, Antonio Silveti-Falls, and Volkan Cevher. “Training deep learning models with norm-constrained LMOs”. arXiv preprint arXiv:2502.07529 (2025).

[7] Artem Riabinin, Kaja Gruntkowska, Egor Shulgin, and Peter Richtárik. “Gluon: Making Muon and Scion great again! (Bridging theory and practice of LMO-based optimizers for LLMs)”. arXiv preprint arXiv:2505.13416 (2025).

[8] R. T Rockafellar. “Monotone operators and the proximal point algorithm”. In: SIAM Journal on Control and Optimization 14.5 (1976), pp. 877–898.

[9] Z. Zhuang, M. Liu, A. Cutkosky, and F. Orabona. “Understanding AdamW through proximal methods and scale-freeness”. Transactions on Machine Learning Research (2022).

[10] Kaja Gruntkowska, Yassine Maziane, Zheng Qu, and Peter Richtárik. Drop-Muon: Update less, converge faster, arXiv preprint arXiv:2510.02239, 2025.

[11] Kaja Gruntkowska, Alexander Gaponov, Zhirayr Tovmasyan, and Peter Richtárik. Error feedback for Muon and friends, 14th International Conference on Learning Representations (ICLR 2026).

Biography: Peter Richtárik is a professor of Computer Science at the King Abdullah University of Science and Technology (KAUST), Saudi Arabia, where he leads the Optimization and Machine Learning Lab . His research interests lie at the intersection of mathematics, computer science, machine learning, optimization, numerical linear algebra, and high-performance computing. Through his work on randomized and distributed optimization algorithms, he has contributed to the foundations of machine learning, optimization and randomized numerical linear algebra. He is one of the original developers of Federated Learning. Prof Richtárik’s works attracted international awards, including the Charles Broyden Prize, SIAM SIGEST Best Paper Award, Distinguished Speaker Award at the 2019 International Conference on Continuous Optimization, the IMA Leslie Fox Prize (three times), and a Best Paper Award at the NeurIPS 2020 Workshop on Scalability, Privacy, and Security in Federated Learning. Several of his works are among the most read papers published by the SIAM Journal on Optimization and the SIAM Journal on Matrix Analysis and Applications. Prof Richtárik serves as an Area Chair for leading machine learning conferences, including NeurIPS, ICML and ICLR, and is an Action Editor of JMLR, and Associate Editor of Numerische Mathematik and Optimization Methods and Software. In the past, he served as an Action Editor of TMLR and an Area Editor of JOTA.

Roman Belavkin, Middlesex University, UK

Value of Information and Entropic Optimal Transport

The optimal transport theory pioneered by Gaspar Monge and developed in probabilistic setting by Leonid Kantorovich has gained popularity recently in the context of machine learning and statistical applications due to the appearance of new algorithms (e.g. Sinkhorn) that simplify computations using entropic regularization. These methods allow for more efficient estimations of transport metrics (aka the Kantorovich or Wasserstein metric) to compare distributions. It was shown in 2018 that optimal transport with entropic regularization is closely related to the value of Shannon's information, and that the latter is a relaxed version of the former. Therefore, many solutions and examples developed by Stratonovich and his colleagues in the 1960s for the value of information can be readily used as lower bounds on the optimal transport costs. I will discuss further generalizations of these concepts by considering other types of information.

Biography: Roman Belavkin obtained MSc in Physics from the Moscow State University and PhD in Computer Science from the University of Nottingham. His research interests span several areas including geometric analysis of optimal and learning systems, dynamics of information, value of information, quantum information, topology of information, geometry and combinatorics of mutation and recombination of sequences, optimal control of evolutionary algorithms, cognitive modelling. Roman joined Middlesex University in 2002, where he participated in several research projects and organized research seminars of the Artificial Intelligence group. From 2009 Roman has been the Principle Investigator of the EPSRC project `SANDPIT: Evolution as an Information Dynamic System', which was led by Middlesex University in collaboration with Universities of Manchester, Keele and Warwick. In this project, Roman developed a theory of optimal control of mutation rate in evolutionary systems, and the team discovered plastic mutation rates in microbes (http://doi.org/skb, http://doi.org/cb9s). Roman's current work is on geometric and dynamic value of information theory, which has applications in parameter control and optimization of learning, adaptive and evolving systems. Roman has many international collaborations: He has been an associate member of the `Centre of Applied Optimization' in the University of Florida, USA; his collaboration with Tokyo University of Science was recognized in 2014 by the award from the university's president Professor Akira Fujishima. Roman has been a keynote speaker at many international conferences, workshops and research seminars. He also serves on the editorial board of the `Optimization Letters' and `SN Operations Research Forum' journals.

Zbyšek Posel, Jan Evangelista Purkyně University, Czech Republic

AI-Driven Labeling and Data Augmentation for ICU Monitoring

Continuous monitoring of patients' vital functions in intensive care units (ICUs) generates high-frequency physiological signals. Among these, electrocardiography (ECG), arterial blood pressure (ABP), and intracranial pressure (ICP) are routinely recorded, particularly in patients with traumatic brain injury. In addition, clinically relevant indices such as the pressure reactivity index and cerebral perfusion pressure are derived from these signals to assess patient status.

However, raw recordings are frequently contaminated by anomalies arising from patient movement, sensor manipulation, device malfunction, and routine clinical interventions. Robust detection and management of such artifacts are essential to ensure data integrity, enable reliable derivation of physiological parameters, and reduce the risk of misclassification of clinically significant events. Recent approaches increasingly rely on machine learning to automate anomaly detection and handling, reflecting the recognition that artifacts are ubiquitous in continuous neuromonitoring and cannot be reliably addressed through manual or ad hoc methods alone. Despite this progress, human annotation remains the ground truth for both model training and evaluation, and together with data availability, defines the upper bound of achievable performance. While several physiological signal databases are publicly available, none provide sufficiently reliable anomaly labels, limiting the development and benchmarking of new methods. To address this limitation, we propose an AI-assisted framework for rapid waveform pre-labeling, enabling efficient human validation and curation. The framework is designed to facilitate the creation of large-scale, multicentric labeled datasets for anomaly detection in physiological signals.

A further challenge lies in the high inter- and intra-patient variability of ICU physiological signals, which complicates the generalization of anomaly detection models. Signal characteristics may vary not only across patients but also due to differences in acquisition methods and treatment conditions over time. To mitigate these challenges and alleviate data scarcity, we also introduce a method for extending real ABP datasets through synthetic signal generation. The proposed approach embeds hemodynamic envelope dynamics and links them with fast signal components such as heart rate variability and waveform morphology. By coupling envelope properties across nearest neighbors in the embedded space, the method produces realistic synthetic ABP signals that closely resemble true physiological recordings.

Biography: Zbyšek Posel is the Deputy Head of the Department and Head of the Data Analysis and Simulation Unit at the Faculty of Science, Jan Evangelista Purkyně University in Ústí nad Labem. His research focuses on computational modeling, computer analysis of electro-physiological signals (EEG, EKG, EMG), image analysis of biological and material systems, and mesoscale simulations of polymers and complex materials. He combines classical algorithms with modern machine learning approaches to address challenges in biomedical data analysis, materials science, and complex system simulations. Dr. Posel has led and contributed to numerous international and national projects, including FP6 MULTIPRO, H2020 VIMMP, TA ČR METAMORPH, and OP JAK DIGITECH, and has conducted extended research stays at the University of Trieste. He teaches a wide range of courses in data analysis, signal and image processing, parallel programming, and system simulation, and supervises research on anomaly detection in physiological signals, polymer self-organization, and artificial signal generation. His work has been published in leading journals such as *Polymers*, *ACS Nano*, *Langmuir*, and *Soft Matter*, and he maintains scientific collaborations with hospitals, research centers, and universities across Europe.

Regular Papers

A Reinforcement-learning-based Column Generation Algorithm for Integrated Operating Room Planning and Scheduling

Hossein Hashemi Doulabi, Department of Mechanical, Industrial and Aerospace Engineering, Concordia University, Montréal H3G 1M8, Québec, Canada CIRRELT, Université de Montréal, Montréal H3C 3J7, Québec, Canada

Operating room planning and scheduling are vital components of hospital management, contributing to improved efficiency, patient satisfaction, staff well-being, and overall quality of care delivery. In this talk, we propose a novel mixed integer programming model to formulate integrated operating room planning and scheduling problems, where several mandatory and elective surgeries are to be assigned and scheduled in operating rooms on different days. Aside from the standard working hours in each operating room, we also take into account the potential for performing surgeries in overtime periods. In addition, our approach also takes into account the availability of surgeons by considering their allowed surgical time on each day. We propose a column generation (CG) algorithm to solve large-scale instances. In order to enhance the CG, we integrate the Reinforcement Learning Algorithm and the Genetic Algorithm and develop a hybrid algorithm to generate initial columns for the CG algorithm. For our analysis, we employed two sets of test instances: one consisting of synthetic data and the other based on real-world cases from a local hospital in Naples, Italy. Computational experiments demonstrate that our proposed model and methodology yields an average optimality gap of 1.23% for synthetic instances and 1.49% on real-world scenarios, significantly outperforming previous solution methodologies in the literature. Additionally, we demonstrate that the developed CG algorithm provides a high-quality solution for large-scale instances where other models and methods fail to obtain even a feasible solution. To further evaluate robustness under uncertainty, we examined scenarios with 20% variability in surgery durations. The results indicate that incorporating a 120-minute buffer time minimizes the overall cost. Moreover, we investigated the impact of emergency surgeries by either introducing additional cases or escalating surgical priorities. For synthetic instances, the inclusion of emergency surgeries increased the total rescheduling cost by 4.13%, whereas in the real-world Naples cases, priority escalation led to only a 0.11% increase, highlighting the resilience of our proposed model in practical hospital settings.

Automated Information Extraction from Historical Archives via OCR and NER

Zbyněk Zajíc, Department of Cybernetics and New Technologies for the Information Society, Faculty of Applied Sciences, University of West Bohemia in Pilsen, Technická 8, 301 00 Plzeň, Czech Republic
Ivan Gruber, Department of Cybernetics and New Technologies for the Information Society, Faculty of Applied Sciences, University of West Bohemia in Pilsen, Technická 8, 301 00 Plzeň, Czech Republic
Adam Frémund, Department of Cybernetics and New Technologies for the Information Society, Faculty of Applied Sciences, University of West Bohemia in Pilsen, Technická 8, 301 00 Plzeň, Czech Republic
Luděk Müller, Department of Cybernetics and New Technologies for the Information Society, Faculty of Applied Sciences, University of West Bohemia in Pilsen, Technická 8, 301 00 Plzeň, Czech Republic

This paper presents a comparative study of publicly available Visual Language Models in the task of digitizing historical documents, specifically in the tasks of Optical Character Recognition (OCR) and Named Entity Recognition (NER). To benchmark the chosen approaches, we utilize a challenging dataset composed of different historical archives, provided by the Czech Institute for the Study of Totalitarian Regimes. In the OCR task, the Qwen2.5VL-7B model achieved the best performance, with a Character Error Rate (CER) of 4.84% and Word Error Rate (WER) of 17.30% for typewritten texts, and 10.88% / 27.42% for mixed texts. In the NER task, we compared two different testing protocols and found that the fine-tuned Qwen2.5VL-7B model again performed best when using the original image input directly.

Dynamic Knowledge Lifecycle Management for Academic Advising RAG Systems

Anuraj Sule, Wilfrid Laurier University, Waterloo, ON, Canada
Abdul-Rahman Mawlood-Yunis, Wilfrid Laurier University, Waterloo, ON, Canada

Retrieval-Augmented Generation (RAG) is the standard architecture for grounding Large Language Models (LLMs) in domainspecific facts. However, current literature largely treats the underlying vector store as a static artifact, which poses significant risks in highvelocity domains like academic advising. In this paper, we present a Dynamic Knowledge Base Management System (DKBMS) evaluated on a corpus of 10,000 university documents. We formalize an end-to-end pipeline that leverages metadata filtering to ingest new documents and surgically remove outdated embeddings. While underlying vector engines support CRUD operations, we demonstrate a unified workflow that reduces knowledge update latency by 99.8% compared to the static snapshot approach often used in academic deployments. Crucially, we analyze retrieval stability to prove that these operations do not degrade the semantic integrity of the remaining knowledge base.

Iterative improvement GA based algorithm for a car production process specified by an extended task graph

Adam M. Górski, Jagiellonian University, Dept. of Information Technologies, Prof. Stanisława Łojasiewicza 11, Cracow, Poland

In this paper, a genetic iterative improvement methodology for a car production process is proposed. The process is specified using an extended task graph. Such a representation allows distinguishing not only resources but also tasks. Thus, it is possible to indicate which tasks can be executed by more than one resource at once and which can be executed only by specialized resources. Unlike existing solutions, in this approach each individual is built starting from a suboptimal solution. Then next solutions are generated by making local changes like allocating new resources, deallocating resources or reschedule tasks between resources. Using an iterative improvement algorithm in connection with genetic algorithms allows escaping local minima of optimization parameters and generating better-quality results.

From Words to Games: Evaluating Large Language Models as Game-Theoretic Modelers

Michaela Tichá, Jan Evangelista Purkyně University, Ústí nad Labem, Czech Republic

Large language models (LLMs) are increasingly evaluated as strategic agents that play games, yet their ability to model real-world strategic situations as formal game-theoretic objects remains largely untested. In this paper we shift the perspective from LLM-as-player to LLM-as-modeler and ask whether contemporary frontier models can translate a natural-language description of a strategic interaction into a faithful formal representation: identifying the players, enumerating strategies, constructing a payoff structure, classifying the game type, and locating its equilibria. We introduce a benchmark of fifteen scenarios spanning three difficulty levels, ranging from canonical two-player matrix games to multi-stage interactions with incomplete information, and evaluate three frontier LLMs against an expert-constructed gold standard using a six-criterion scoring rubric. Our results show that all three models perform strongly on canonical games ($M \geq 14.7/15$) but diverge increasingly on complex multi-feature scenarios, with payoff construction (K3) and equilibrium justification (K5) being the most discriminating criteria. Gemini 3.1 Pro exhibits notably lower performance and higher variance on Tier 3 scenarios ($M = 12.1$, $SD = 2.0$) and is the least stable across replications. We argue that game-theoretic formalization is a tractable, informative, and so far overlooked test of LLM reasoning that complements existing strategic-reasoning benchmarks.

Rethinking Cloud Elasticity: AutoML-Driven Predictive Scaling for Bursty Workloads

Mohib Magbool, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Dariush Ebrahimi, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Daniel Ramirez Bumaguin, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Mustafa Daraghmeh, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada

Cloud computing platforms depend on elasticity to satisfy Service Level Agreements (SLAs) while controlling operational costs. However, widely deployed reactive autoscaling mechanisms respond only after demand has already increased, which leaves bursty workloads exposed to provisioning latency, transient under-provisioning, and avoidable SLA violations. This paper addresses that limitation by proposing a multi-objective hybrid scaling framework in which Automated Machine Learning (AutoML) forecasts near-future workload demand and feeds a tightly coupled decision module for proactive capacity provisioning, while a reactive fail-safe protects against extreme spikes. The framework is evaluated using burst-dominant traces from the Alibaba Cluster Trace and implemented inside the published ECLYPSE cloud-edge simulation framework. The results show that the most statistically accurate forecasting model is not necessarily the most operationally effective. In particular, an ensemble-based Random Forest model outperforms lower-error linear alternatives at the system level by preserving a safety margin that reduces SLA violations from 83 to 5, cuts total operational cost by 87.5%, and enables approximately 84% of the infrastructure to remain in idle or low-power states. These findings demonstrate that publication-grade predictive elasticity must be optimized against system-level objectives, not prediction error alone, and that hybrid AutoML-guided scaling can deliver reliable and cost-efficient cloud operation under highly volatile workloads.

When Does Forecasting Improve Cloud Auto-Scaling? A Trace-Driven Evaluation

Muneeb Saeed, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Sulaiman Sultan, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Dariush Ebrahimi, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Mustafa Daraghmeh, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada

Cloud auto-scaling remains difficult to configure because reactive threshold policies must balance service reliability against infrastructure cost under volatile demand. This paper investigates when lightweight forecasting actually improves auto-scaling outcomes under a controlled trace-driven setting by comparing four one-step-ahead predictors: Reactive, Shifted Load, Moving Average, and ARIMA. The framework operates on a real CPU-usage trace, applies a threshold-based scaling policy with a cool-down constraint, and evaluates cost, scale-out events, scale-in events, oscillations, and SLA violations. The baseline results show clear behavioral differences across predictors: ARIMA reduces cost by about 50% relative to the reactive baseline (1000 vs. 2016 instanceunits) while eliminating oscillations entirely (0 vs. 192). Moving Average provides an intermediate outcome, while Reactive and Shifted Load remain the least stable strategies. Scenario-based stress testing further shows that forecasting is most useful under persistent bursty demand, offers limited benefit under stable high-load conditions, and has minimal effect under isolated spikes. Overall, the results show that forecasting improves cloud auto-scaling when workload variability is sufficiently high to reward anticipation but not so abrupt that smoothing suppresses timely reactions. The paper contributes an interpretable evaluation framework, a quantitative comparison of lightweight forecasting strategies, and practical guidance on when prediction-driven scaling is most beneficial.

Reliability-Aware Hybrid Autoscaling: Integrating Forecasting and AIOps with LLM Guardrails

Sufiya Rahemtulla, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Mustafa Daraghmeh, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada
Dariush Ebrahimi, Wilfrid Laurier University, Waterloo, Ontario, N2L 3C5, Canada

Cloud autoscaling remains a difficult control problem because modern containerized services must absorb abrupt workload surges while also limiting the cost of unused capacity. Conventional reactive autoscalers, including the Kubernetes Horizontal Pod Autoscaler (HPA), respond only after utilization rises and therefore expose a reaction gap during which Service Level Agreement (SLA) violations accumulate. Predictive approaches reduce this delay, but many are evaluated primarily with average forecasting accuracy and provide limited protection against rare but operationally expensive bursts. This paper presents a reliability-aware hybrid autoscaling framework that combines a lightweight linear-regression forecaster with an anomalytriggered AIOps supervisory agent. The forecasting layer provides shorthorizon workload estimates from lagged demand features. The AIOps layer revises fallback scaling thresholds only when anomalous residual patterns suggest that the forecast is becoming unreliable. To prevent unsafe or economically inefficient actions, all AIOps recommendations are f iltered through deterministic Python guardrails that validate schema, whitelist actions, and bound control parameters before they are applied. Theframework is evaluated in a trace-driven simulation built from Google Cluster-Usage Traces v3 under steady, intermediate, and bursty demand regimes. The proposed hybrid configuration reduces the SLA violation rate from 14.2% for the reactive baseline to 2.1%, while also lowering average provisioning latency and limiting resource waste to 8.4%. An ablation shows that removing the guardrails improves violation rate only marginally, to 1.2%, but increases resource waste to 48.7%, demonstrating that bounded AI assistance is essential for practical deployment. The findings show that reliability-aware autoscaling benefits from combining proactive forecasting with carefully constrained LLM-assisted operations. Rather than replacing deterministic control, the LLM is most useful as a bounded advisory mechanism for tail events that are poorly handled by reactive thresholds and average-case predictors alone.

Deep learning framework based on multi-source data fusion for real-time monitoring and prediction of air quality index in smart cities using CNN–LSTM hybrid architecture

Fatemeh Ashtari Gargari, Faculty of Electrical and Computer Engineering, University of Tabriz, Tabriz, Iran
Kimia Shirini, Faculty of Multimedia, Tabriz Islamic Art University, Tabriz, Iran
Sina Samadi Gharehveran, Faculty of Electrical and Computer Engineering, University of Tabriz, Tabriz, Iran
Siamak Pedrammehr, Faculty of Design, Tabriz Islamic Art University, Tabriz, Iran
Roohallah Alizadehsani, School of Medicine, Queensland University of Technology, Brisbane, Australia

This paper proposes an intelligent artificial intelligence–based system for monitoring, predicting, and managing air pollution in smart cities. By leveraging data from Internet of Things (IoT) sensors, satellite imagery, and meteorological data, a hybrid deep learning model incorporating Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) models has been developed for accurate prediction of major pollutant concentrations such as PM2.5, NO2, and CO in the upcoming hours. The main objective is to reduce the environmental and health impacts of air pollution through real-time alerts and optimization of traffic and energy consumption in smart cities. The system was evaluated on real data from Tehran city and improved prediction accuracy up to 95%, reduced computational resource consumption by 30%, and enabled intelligent decision-making by urban managers. The results demonstrate the high efficiency of the proposed approach in real-world scenarios.

Recent Advances in Clustering and Classification-based Algorithms for Big Data management

Emrah Arslan, Departmant of Computer Engineering, Faculty of Engineering, KTO Karatay University, Türkiye
Javad Hassannataj Joloudari, Department of Computer Engineering, Technical and Vocational University (TVU), Tehran, Iran
Javad Hassannataj Joloudari, Faculty of Electrical and Computer Engineering, University of Birjand, Birjand, Iran
Javad Hassannataj Joloudari, Department of Computer Engineering, Islamic Azad University, Babol, Iran
Hanieh Khosravi, Department of Computer Engineering, Technical and Engineering Faculty, Islamic Azad University, South Tehran Branch, Tehran, Iran
Bita Mesbahzadeh, Department of Industrial Engineering, K. N. Toosi University of Technology, Tehran, Iran
Amir Mashmool, Faculty of Mathematics and Computer Science, University of Bremen, Bremen, Germany
Roohallah Alizadehsani, School of Medicine, Queensland University of Technology, Brisbane, Australia

The exponential growth of heterogeneous data generated across domains such as social media, Internet of Things (IoT), e-commerce, healthcare, and financial systems has intensified the demand for scalable and intelligent Big Data analytics. Despite significant advances, effectively extracting actionable knowledge from high-volume, high-velocity, and high-variety data remains a fundamental challenge due to limitations in scalability, interpretability, and computational efficiency. This paper presents a systematic and comprehensive review of clustering and classification algorithms, positioning them as the core analytical paradigms underpinning modern Big Data analytics. A structured comparative analysis is conducted across diverse benchmark datasets, encompassing text, graph, and image modalities, to evaluate algorithmic performance, scalability, and domain suitability. The study further examines the evolution from classical machine learning techniques to deep learning architectures, highlighting key trade-offs between accuracy, efficiency, and explainability. In addition, hybrid and ensemble frameworks that integrate supervised and unsupervised learning are critically analyzed, demonstrating their effectiveness in addressing the complexity of real-world data environments. The paper also identifies persistent challenges, including data quality, bias, privacy risks, and the lack of model transparency. To address these issues, emerging research directions such as explainable AI, self-supervised learning, and decentralized analytics are explored. Overall, this work provides a unified and forward-looking perspective on the synergistic role of clustering and classification in enabling robust, scalable, and interpretable Big Data solutions.

Network distance based on graphlets

David Hartman, Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 271/2, 182 07 Prague, Czech Republic
Jaroslav Hlinka, Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 271/2, 182 07 Prague, Czech Republic
Anna Pidnebesna, Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 271/2, 182 07 Prague, Czech Republic
František Szczepanik, Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 271/2, 182 07 Prague, Czech Republic. Faculty of Mathematics and Physics, Charles University, Malostranské nám. 25, Prague 1 11800, Czech Republic

Comparing networks is fundamental across disciplines ranging from computational biology to econometrics. There exist two approaches: alignment-free methods, which compare aggregate structural summaries, and alignment-based methods, which derive distances from explicit vertex mappings. We study graphlet-based methods that are powerful, mostly alignment-free approaches, which offer a unique blend of local and global structural insights. We propose using graphlet-based alignment to support a network distance based on graphlets. We compare this approach to existing graphlet-based methods to evaluate the effect of including alignment in the distance metric. To evaluate this, we introduce the α-BA model—a Barabási-Albert variant that interpolates between preferential attachment and deterministic growth while preserving order and density. Using PERMANOVA to assess discriminative power across α-classes, we find that an alignment-based approach, utilizing the Hungarian algorithm on graphlet-derived weights, significantly outperforms standard alignment-free methods. Finer-grained tests confirm that our alignment-based metric reliably distinguishes adjacent α-classes where alignment-free methods fail. These findings highlight the potential of alignment-based graphlet comparison for broader network analysis tasks.

Initial Evaluation of Stacked Hourglass Networks for Automated Cervical Vertebrae Keypoint Detection

Adéla Svítilová, Jan Evangelista Purkyně University in Ústí nad Labem, Pasteurova 3544/1, Czech Republic
Jaroslav Radimský, Jan Evangelista Purkyně University in Ústí nad Labem, Pasteurova 3544/1, Czech Republic
Petr Kubera, Jan Evangelista Purkyně University in Ústí nad Labem, Pasteurova 3544/1, Czech Republic. Faculty of Nuclear Sciences and Physical Engineering, Czech Technical University, Prague, Czech Republic
Jiří Škvára, Jan Evangelista Purkyně University in Ústí nad Labem, Pasteurova 3544/1, Czech Republic. Institute of Chemical Process Fundamentals of the Czech Academy of Sciences, v.v.i., Prague, Czech Republic

Precise assessment of cervical spine alignment from X-ray images is crucial for diagnosing spinal deformities and evaluating postoperative outcomes. Important radiographic features, such as vertebral edges, provide clinically relevant information about spinal anomalies. However, manual detection of these edges is time-consuming and prone to inter-observer variability. This work explores a deep learning-based approach using keypoint detection for the automatic identification of C2–C7 vertebral corners on lateral cervical spine X-ray images. We employ a convolutional neural network inspired by the Stacked Hourglass architecture, which directly detects keypoints of vertebral corners from the entire image. The model is trained on an annotated dataset of nearly 5000 X-ray images, where individual keypoints (i.e., corners of the vertebral bodies) were manually labeled by experts. The network leverages multi-scale feature extraction to capture both fine details and the global context of the spine, thereby improving detection accuracy even with overlapping structures or variable imaging angles. The procedure we propose demonstrates that keypoint detection enables reliable automated analysis of cervical spine X-rays with high accuracy. Such systems can facilitate faster and more objective spine alignment assessment for clinicians, reduce subjectivity, and accelerate diagnosis.

Applications of federated learning in optical coherence tomography image analysis

Ali Jafarizadeh, Nikookari Eye Center, Tabriz University of Medical Sciences, Tabriz, Iran
Hamidreza Ashayeri, Research Center for Evidence-based Medicine, Iranian EBM Center: A Joanna-affiliated group, Tabriz University of Medicine Science, Tabriz, Iran
Fatemeh Razmjooei, Poostchi Ophthalmology Research Center, Department of Ophthalmology, Medical School, Shiraz University of Medical Sciences, Shiraz, Iran
Mahdi Mohammadkhani, Nikookari Eye Center, Tabriz University of Medical Sciences, Tabriz, Iran
Roohallah Alizadehsani, School of Medicine, Queensland University of Technology, Brisbane, Australia

Artificial intelligence (AI) can handle multiple tasks, process many parameters, and increase task speed and accuracy. The training stage is a crucial step for building a capable model. Medical data is among a person`s most sensitive data, which can be used to train AI models. This makes data privacy a crucial factor to consider. Federated learning (FL) is one proposed method for increasing data privacy. In FL, as opposed to centralized learning, a global model is trained on local servers, and then the learnt parameters are sent back to the central server. FL has been applied in the segmentation and diagnosis of ophthalmologic diseases such as diabetic eye disease, age-related macular degeneration, and glaucoma. FL has shown performance comparable to centralized learning and, in some cases, outperformed the centralized training model. In this article, we aim to summarize the recent developments and applications of FL in optical coherence tomography.

Comparing the Effectiveness of Influence Maximization Algorithms Across Changing Network Parameters

Anna Kmentová, Faculty of Mathematics and Physics, Charles University, Malostranské nám. 25, Prague 1 11800, Czech Republic
Dávid Kubek, Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 271/2, Prague 8 18200, Czech Republic. Faculty of Mathematics and Physics, Charles University, Malostranské nám. 25, Prague 1 11800, Czech Republic
David Hartman, Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 271/2, Prague 8 18200, Czech Republic. National Institute of Mental Health, Topolová 748, Klecany 25067, Czech Republic. Unicorn University, Kolbenova 942, 190 00 Praha 9, Czech Republic
Jaroslav Hlinka, Institute of Computer Science of the Czech Academy of Sciences, Pod Vodárenskou věží 271/2, Prague 8 18200, Czech Republic. National Institute of Mental Health, Topolová 748, Klecany 25067, Czech Republic

Influence maximization (IM) is a fundamental problem in network science. While researchers have built many algorithms to improve upon the original greedy approach, existing surveys typically compare these tools on fixed test instances. Consequently, the field lacks a systematic understanding of how changing network properties alter algorithm performance. To close this gap, we compare the solution quality of various IM algorithms across changing network parameters to capture performance trends and determine algorithm suitability. We generate diverse random network models, integrating features like preferential attachment and community structure. By systematically varying quantitative and qualitative network parameters, we test multiple algorithms and rank their relative performance using the Friedman–Nemenyi statistical test. We find that changing network configurations can significantly shift algorithm rankings. Specifically, network directedness impacts algorithm choice the most, whereas community structure layout shows little effect. While some algorithms maintain consistent results, others evolve as network parameters change; however, the Inuence Maximization via Martingales (IMM) algorithm performs well overall. Ultimately, these f indings show that underlying network properties must be considered rather than relying on default choices when selecting an IM algorithm.

A Scalable VNS-Based Heuristic for Multi-Depot Drone–Truck Arc Routing

SHROUQ GAMAL, Department of Industrial and Systems Engineering, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia
Mujahid N. Syed, Department of Industrial and Systems Engineering, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia. Interdisciplinary Research Center for Intelligent Secure Systems, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia

Multi-depot drone–truck arc routing with endurance and deployment coupling constraints quickly becomes computationally intractable for large-scale instances. This paper concentrates on the heuristic acceleration of the Phase I patrolling model introduced in [1], which ensures full coverage of required network segments while minimizing deployment and sustainability-related costs. The underlying formulation extends classical arc routing problems by incorporating multiple mobile depots, multiple coordinated drones, endurance limits, and deployment-dependent routing decisions within a mixed-integer programming framework. Although exact methods provide modeling flexibility, they struggle to scale for dense networks and realistic operational sizes. To address this limitation, we propose a structured VNS-based heuristic specifically designed for the problem's multi-depot and endurance-constrained structure. The approach begins with a decomposition-based initialization: a Minmax assignment model determines strategic truck positioning, followed by an m-TSP formulation to construct feasible drone routes. Building on this initialization, a tailored Greedy m-TSP warm-start is embedded within a Variable Neighborhood Search (VNS) framework. The proposed VNS introduces five problem-specific neighborhood operators—route swapping, sub-route reversal, depot relocation, crossover recombination, and balanced route merging—carefully designed to maintain feasibility under endurance and depot coupling constraints while enabling effective exploration of the solution space. Unlike generic metaheuristics, the proposed method explicitly exploits structural properties of cooperative drone–truck arc routing. Computational results demonstrate substantial improvements in convergence speed and optimality gaps compared to direct MILP approaches, particularly for large and dense instances. The findings highlight the effectiveness of structured initialization combined with customized neighborhood design for scalable cooperative aerial–ground routing.

Learning Optimal Decision Trees with MILP-Based Splitting and Greedy Decomposition

MOHAMED AHMED, Department of Industrial and Systems Engineering, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia
Mujahid N. Syed, Department of Industrial and Systems Engineering, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia. Interdisciplinary Research Center for Intelligent Secure Systems, King Fahd University of Petroleum and Minerals, Dhahran, Saudi Arabia

Learning optimal decision trees is a fundamental problem in interpretable machine learning and is known to be NP-hard. This work presents a compact Mixed-Integer Linear Programming (MILP) formulation for training univariate binary classification trees to global optimality for a prespecified depth. The model adopts a node-based representation of a full binary tree, where each internal node selects exactly one feature along with a corresponding continuous threshold. Binary variables govern the routing of each training instance from the root to a single leaf, ensuring that left and right branches are followed according to feature–threshold conditions. Leaf predictions are determined majority class voting, implemented through linear constraints with slack variables that penalize classification impurity. Finally, we propose a hybrid approach that combines greedy tree construction with Optimal Classification Trees (OCT). The method decomposes the full tree into a main structure and smaller subtrees grown from selected leaves where further optimization is required. This approach not only improves predictive accuracy but also significantly improves interpretability compared to training a single OCT of equivalent depth.

Adiabatic Quantum Computing in Urban Waste Collection Logistics Optimization

Mahdia Ahmadi, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Catarina Baldaia, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
P. J. Rodrigues, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Getúlio Peixoto Igrejas, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Rui P. Lopes, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal

This work investigates the application of QUBO-based optimization concepts, motivated by adiabatic quantum computing and quantum annealing, to combinatorial optimization in urban waste collection logistics, using the Travelling Salesman Problem (TSP) as a representative routing framework. First, the mathematical foundations of adiabatic quantum computing and quantum annealing are reviewed, with particular emphasis on the Quadratic Unconstrained Binary Optimization (QUBO) formulation and its interpretation in terms of Hamiltonians. The 0–1 knapsack problem is introduced as an illustrative example of QUBO encoding, followed by the formulation of the TSP in both classical and QUBO-compatible forms. A computational study is then conducted to compare strong classical routing methods with QUBObased approaches on routing instances of varying size, including ORTools Guided Local Search, 2-opt, and Ocean-neal classical simulated annealing applied to the QUBO model. Performance is assessed in terms of solution quality, relative gap, execution time, and cost–time trade-offs. The results show that strong classical approaches remain clearly superior in practical performance, whereas QUBO-based methods solved with a classical simulated annealing engine exhibit limitations in scalability and solution quality. To extend the study toward a more realistic urban waste collection setting, a multi-vehicle Vehicle Routing Problem (VRP) was also examined through a hybrid decomposed workflow. In this extension, the allocation of waste containers to vehicles was handled classically, while the sequencing of each vehicle route was optimized independently through a TSP QUBO solved with Ocean-neal classical simulated annealing. The results show that this decomposed hybrid approach yields an intermediate level of solution quality, clearly improving over a simple greedy multi-vehicle baseline, but remaining consistently inferior to OR-Tools VRP in terms of total route cost. A simplified estimate of execution time in the quantum paradigm is also discussed only as a theoretical reference associated with the annealing stage, without being interpreted as a directly achievable end-to-end runtime under the present experimental conditions. The study is framed within the RAICYCLE project and is intended as a methodological contribution toward smart urban waste management in a Portuguese municipal context.

Global QUBO Formulation of the Vehicle Routing Problem for Urban Waste Collection: Feasibility, Cost, and Scalability

Mahdia Ahmadi, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Catarina Baldaia, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
P. J. Rodrigues, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Getúlio Peixoto Igrejas, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal
Rui P. Lopes, CeDRI, SusTEC, Instituto Politécnico de Bragança, Campus de Santa Apolónia, 5300-253 Bragança, Portugal

Quadratic Unconstrained Binary Optimization (QUBO) has become one of the main formulation layers for mapping combinatorial optimization problems to adiabatic quantum computing and quantumannealing-oriented solvers. In routing, however, most practical studies avoid fully global formulations and instead rely on decomposition or hybridization. This raises a fundamental question: how far can a monolithic global QUBO model of the Vehicle Routing Problem (VRP) be pushed before feasibility or scalability breaks down? This paper addresses that question in the context of urban waste collection. We formulate a global VRP-QUBO model in which waste-container assignment and route ordering are encoded simultaneously across multiple vehicles departing from a common depot. The formulation is evaluated on synthetic Euclidean instances with up to 48 waste containers and 6 vehicles using Ocean-neal simulated annealing as an exploratory classical QUBO solver. The results show that the global formulation is structurally more robust than initially expected. All tested configurations remained fully feasible except one, and the feasibility frontier was not reached within the explored range under a threshold of 0.8. At the same time, solve time, logical-variable count, and QUBO-term count increase substantially with problem size, indicating that the dominant limitation in the explored regime is computational growth rather than immediate structural collapse. These findings position the global formulation as both a valid proof of concept and a strong baseline for future comparison with decomposed and hybrid routing workflows. The study is framed within the RAICYCLE project and is intended as a methodological contribution toward smart urban waste management in a Portuguese municipal context.

Extreme points and faces in the moment problem

Stephan Weis, Department of Control Engineering, Faculty of Electrical Engineering, Czech Technical University in Prague, Czechia

We shed new light on the moment problem of minimizing an integral functional over a set of probability measures under finitely many moment constraints. We address whether the optimal value coincides with the optimal values restricted to the extreme points of the constraint set and to certain finitely atomic measures.

We approach this question in two different ways. The first relies on the Weizsäcker–Winkler integral representation for Radon measures [1], while the second uses the Richter–Tchakaloff theorem [2,3]. In both approaches, the characterization of extreme points by Weizsäcker and Winkler [4,5] leads to an affirmative answer, though under different sets of assumptions.

Our main contribution is an improved characterization of extreme points. We provide a novel description of extreme points for arbitrary affine constraints on any convex set. This characterization is based on the notion of the smallest face of a convex set containing a given point, a concept studied in convex analysis and optimization by Alfsen [6] and further developed in recent works [7–10].

Using this framework, we recover the Weizsäcker–Winkler characterization of extreme points in the moment problem while weakening their assumptions and significantly simplifying their proof.

Our motivation arises from numerical methods for the polyconvex envelope in elasticity theory [11]. Pointwise, this envelope can be interpreted as the optimal value of a moment problem over finitely atomic measures [12]. Our results show that the polyconvex envelope can equivalently be evaluated over Borel measures, making it amenable to numerical approximation via the moment sum-of-squares hierarchy, a standard method in polynomial optimization [13].

This is joint work with Didier Henrion and Martin Kružík.

References

[1] H. v. Weizsäcker and G. Winkler, *Integral representation in the set of solutions of a generalized moment problem*, Math. Ann. 246(1):23–32, 1979.
[2] K. Schmüdgen, *The Moment Problem*, Graduate Texts in Mathematics, vol. 277, Springer, 2017.
[3] A. Shapiro, *On duality theory of conic linear problems*, in M.Á. Goberna and M.A. López (eds.), *Semi-Infinite Programming*, Nonconvex Optimization and Its Applications, vol. 57, Springer, 2001.

The dynamics of artificial intelligence in healthcare: From data silos and external validity issues to privacypreserving and explainable solutions

Faezeh Darvishi, Liver and Gastrointestinal Diseases Research Center, Tabriz University of Medical Sciences, Tabriz, Iran
Hamidreza Ashayeri, Research Center for Evidence-based Medicine, Iranian EBM Center: A Joanna-affiliated group, Tabriz University of Medicine Science, Tabriz, Iran
Parisa Rostami, Research Center for Evidence-based Medicine, Iranian EBM Center: A Joanna-affiliated group, Tabriz University of Medicine Science, Tabriz, Iran
Hadi Vahedi, Research Center for Evidence-Based Health Management, Maragheh University of Medical Sciences, Maragheh, Iran
Ali Jafarizadeh, Nikookari Eye Center, Tabriz University of Medical Sciences, Tabriz, Iran
Hossein Moosaei, Jan Evangelista Purkyně University, 400 96 Ústí nad Labem-město, Czech Republic
Roohallah Alizadehsani, School of Medicine, Queensland University of Technology, Brisbane, Australia

Artificial intelligence (AI) promises faster, more accurate screening, diagnosis and treatment monitoring, yet clinical gains remain limited by fragmented data, opaque models, poor external validity and complex implementation hurdles. In this review, English-language original articles published in PubMed and Web of Science were searched using combinations of “artificial intelligence,” “machine learning,” “deep learning,” “healthcare,” and challenge-related terms. Several results were found in this study; Limited, biased or siloed datasets hinder reliable training, and multimodal fusion demands rigorous harmonization. Black-box architectures reduce clinician trust, driving demand for explainable AI. Algorithms often falter on external cohorts, underscoring the need for multicenter validation and robust transfer or federated learning. Clinical uptake is slowed by regulatory ambiguity, workflow disruption, high costs, scant physician training, variable patient acceptance and absent long-term monitoring. In conclusion, AI can measurably improve every stage of patient care, but real-world impact depends on richer data, transparent models, external validation and secure, clinician-centered implementation under adaptive regulation. Coordinated technical, clinical and policy innovation is essential to translate rapid laboratory advances into safe, equitable and sustainable healthcare benefit.

Integrating Standards and Agile Workflows for a Multilayered Quality Assurance Approach in the Virtual Museum Framework

Olena Shlyakhetko, Comenius University Bratislava, Odbojárov 10, 82005 Bratislava, Slovakia
Yuriy Syerov, Comenius University Bratislava, Odbojárov 10, 82005 Bratislava, Slovakia. Lviv Polytechnic National University, 79000 Lviv, Ukraine
Natalia Kryvinska, Comenius University Bratislava, Odbojárov 10, 82005 Bratislava, Slovakia

This paper presents a comprehensive approach to software quality assurance by integrating the rigorous requirements of the ISO/IEC 25010 software quality model with agile development workflows. Specific software sub-characteristics are evaluated according to their relative importance to achieve the required quality levels. Furthermore, the study discusses both the technical and social approaches necessary to ensure a high standard of software. A multilayered quality architecture is proposed, encompassing a structured testing pyramid comprising unit, integration, and end-to-end tests. To enforce reliability, a minimum code coverage requirement of 70% is established using automated testing libraries. Additionally, static code analysis tools are implemented to ensure strict code consistency, readability, and security. To further optimize the software development life cycle, the proposed framework incorporates an agile organizational structure with a three-tier development environment (Development, Testing, Production), robust peer-review protocols via pull requests and spot checks, and Kanban-based task management. Ultimately, the study demonstrates that achieving high software quality requires a synchronized integration of technical testing methodologies, standardized quality metrics, and structured team collaboration.

Environmental Risk and Antimicrobial Resistance Potential of Antibiotics in Aquatic Systems: A Modeling Study of Ciprofloxacin and Doxycycline

Eva K. Lee, NSF-Whitaker Center for Operations Research in Medicine and Healthcare, the Data and Analytics Innovation Institute, Atlanta, Georgia, USA. AccuHealth Technologies, Atlanta, Georgia, USA. Georgia Institute of Technology, Atlanta, Georgia, USA
Ferdinand H. Pietz, Strategic National Stockpile, the Administration for Strategic Preparedness and Response (ASPR/HHS), Washington D.C., USA
Peijue Zhang, Georgia Institute of Technology, Atlanta, Georgia, USA
Cheng Chen, Georgia Institute of Technology, Atlanta, Georgia, USA

The work is partially supported by a grant from the National Science Foundation, and the Centers for Disease Control and Preven on. The findings and conclusions in this paper are those of the authors and do not necessarily represent the official posi on of the funding agencies. The widespread use of antibiotics has raised increasing concern regarding their environmental persistence and contribution to antimicrobial resistance (AMR). In this study, we develop a dynamic modeling framework to assess the environmental risk of antibiotics in aquatic systems using predicted environmental concentrations (PECs) and predicted no-effect concentrations (PNECs). The framework integrates pharmaceutical usage, excretion, wastewater treatment removal, environmental dilution, and degradation processes, and is applied to two commonly used antibiotics: Ciprofloxacin and Doxycycline. Under baseline conditions, modeled PECs in downstream environments remained below conventional ecotoxicological thresholds. However, concentrations in sewage retention systems approached or exceeded no observed effect concentrations (NOECs), particularly during periods of elevated antibiotic use. Simulated mass dispensing scenarios further amplified environmental exposure, resulting in moderate to high ecological risk for ciprofloxacin. When antimicrobial resistance-based thresholds (PNECAMR_{AMR}AMR) were incorporated, ciprofloxacin concentrations exceeded resistance selection thresholds across all scenarios, including baseline conditions. In contrast, doxycycline exhibited substantially lower environmental risk and did not exceed ecotoxicological thresholds. Monte Carlo simulations demonstrated that these findings are robust to parameter uncertainty. These results suggest that conventional ecotoxicological risk assessments may underestimate the environmental impact of antibiotics by failing to account for resistance selection at sub-inhibitory concentrations. Integrating AMR-based thresholds into environmental risk frameworks provides a more comprehensive assessment and highlights the importance of considering environmental consequences in both routine antibiotic use and large-scale public health interventions.

Shadow Technical Debt: The Dark Matter of Engineered Systems

Jason Pyeron, PD Inc International, Baltimore, MD 21218, USA

Shadow technical debt is not defined by who or what created it. It is defined by the fact that it exists before an organization has acknowledged, measured, or budgeted for it. Although the phenomenon is not new, recent discussion has highlighted the role of distributed and agentic development environments in accelerating the creation of hidden liabilities within complex engineered systems.

Over the past two decades, project management for systems has emphasized risk management and resource-constrained decision making. It stresses focusing limited resources where they can most effectively reduce uncertainty and system risk. This paper frames shadow technical debt as a long-standing condition within engineered systems and gives it socioeconomic context, enabling management within existing frameworks.

Building on established risk and resource allocation methodologies, including cost-managed discovery and prioritization approaches, the paper outlines practical techniques for identifying shadow technical debt within tightly constrained budgets. As contemporary agentic and automated engineering processes increase the rate at which systems are created and modified, the volume of hidden technical debt is growing by orders of magnitude and may exhibit a superlinear relationship. Efficient methods for discovering and classifying these liabilities are therefore essential for converting shadow debt into actionable technical debt and bringing it onto the engineering balance sheet.