

Editorial

Ivan Luković¹, Mirjana Ivanović², Sonja Ristić³, and Pavle Milošević¹

¹ University of Belgrade, Faculty of Organizational Sciences,
Belgrade, Serbia

{ivan.lukovic,pavle.milosevic}@fon.bg.ac.rs

² University of Novi Sad, Faculty of Sciences,
Novi Sad, Serbia

mira@dmf.uns.ac.rs

³ University of Novi Sad, Faculty of Technical Sciences,
Novi Sad, Serbia, Retired Professor
sdristic@uns.ac.rs

Vol. 23, No. 4, regular issue of ComSIS is in front of you. This issue of ComSIS contains 11 regular research papers, covering several computer science disciplines, from the applications of AI and LLMs, through complex and collaborative networks, cloud computing, to ontological modeling.

The authors of the paper: Quantifying Multi-Level Effects of AI Adoption through an End-to-End Machine Learning Pipeline, Na Tang and Yang Yang, propose a construction of an end-to-end machine learning pipeline — from raw heterogeneous data input, through three cascaded modules (micro: Random Forest + SHAP; meso: XGBoost; macro: K-means++ clustering), to actionable policy outputs — to quantify the complex socio-economic effects of AI adoption across countries and industries. The authors present a study that moves beyond isolated econometric models and provides quantifiable, actionable insights for balancing AI innovation with social welfare, by coupling machine learning methods in a tripartite analytical pipeline.

In the paper: PocketPsyc: A Multi-Model AI System for Real-Time Emotion-Aware Mental Health Support, Ayesha Maqsood, Altaf Hussain, Marran Al Qwaid, Zara Haider, and Muazzam A. Khan Khattak present PocketPsyc, a mobile-based system designed to deliver scalable, evidence-based mental health support through Cognitive Behavioral Therapy (CBT). The system integrates three specialized AI models: a fine-tuned BART model for therapeutic response generation, a RoBERTa-base classifier for real-time emotion recognition, and a TinyLlama-1.1B model for personalized mindfulness guidance. To ensure user privacy, the platform employs client-side AES-256 encryption and row-level security mechanisms. Additionally, a hybrid crisis detection module combines clinical threshold monitoring with linguistic cue analysis to identify high-risk scenarios.

The authors of the paper SAMS-M: Explicit Sentiment-Guided Alignment and Multi-dimensional Mutual Supervision for Multimodal Sentiment Analysis, Shishu Qi, Yulei Zhang, Siyang Zhang, Jia Xu, Kuan-Ching Li, and Guangli Zhu, advocate that Multimodal Sentiment Analysis leverages the fusion of heterogeneous data to achieve fine-grained emotional understanding, which finds extensive application in large-scale public opinion monitoring and data mining. The authors propose an explicit sentiment-guided alignment and multi-dimensional cross-modal mutual supervision-based model for multimodal sentiment analysis. The model primarily employs a fine-grained sentiment-saliency directed alignment mechanism, which leverages bidirectional cross-attention to couple textual sentiment cues with visual saliency, enabling precise localization of sentiment-

relevant regions. Furthermore, they introduce a tripartite strong contrastive learning strategy to mitigate distribution discrepancies between heterogeneous modalities within a shared latent space, thereby enhancing cross-modal coherence and complementarity. Finally, the authors design a noise-robust gating-based fusion module, which, together with text augmentation and deep supervision, facilitates effective joint optimization.

Nonlinear and regime-shifting price behavior of cryptocurrencies challenges traditional econometric and deep learning models, creating a need for more adaptive forecasting frameworks. Addressing this research gap, in the paper: Hybrid LSTM–HMM Framework for Regime-Aware Cryptocurrency Price Forecasting, Tran Kim Toai and Nguyen Phan Nhat Minh present a study that proposes a hybrid Hidden Markov Model–Long Short-Term Memory (HMM–LSTM) architecture that jointly captures latent market regimes and temporal dependencies in cryptocurrency price dynamics. The HMM identifies latent market regimes, while the LSTM captures temporal dependencies. Technical indicators, including Bollinger Bands, EMA, and MACD, support feature benchmarking.

Lijun Xiao, Xinwu Jian, Zijian Wang, Sisi Zhou, and Kuanching Li, the authors of the research presented in the paper: Hybrid Path Planning for Industrial Autonomous Vehicles: Integrating Neural Heuristics with Energy-Constrained Search advocate that energy-aware path planning for industrial autonomous vehicles requires more than a geometric shortest path: terrain, payload, drivetrain efficiency, turning, and a hard task-level energy budget jointly affect feasibility. They present Neuro-PathNet-EC, a physics-aware constrained planner with learned multimodal guidance. The planner operates on heading-aware grid states and evaluates directed edge energy from friction and elevation. Exact Pareto labels enforce the hard energy budget, while a goal-conditioned ResNet-18 network fuses four raster channels with vehicle semantics and predicts an eight-channel dense cost-to-go map. The learned heuristic is used only for guidance; it is not assumed admissible or consistent.

Accurately identifying vulnerable nodes in complex networks is a critical challenge for preventing cascading failures and ensuring system robustness. The research presented in the paper: Vulnerable Node Detection in Complex Networks Based on Dynamic Comparator and Trusted Chain Propagation by Dongsheng Jing, Baozhen Qi, Yu Yang, Jingsong Xue, and Renjun Feng states that existing methods generally suffer from limitations such as static metrics that ignore dynamic loads, high parameter dependency, and a lack of unified fusion mechanisms. To address these issues, the authors propose an approach that is based on dynamic comparator selection and trusted chain propagation, referred as Optimal Parameter (OP). Specifically, the proposed OP approach adaptively adjusts strategies based on node degree: low-degree nodes enable a lightweight mode to reduce overhead, while high-degree nodes employ full sampling to enhance reliability. The method integrates temporal consistency checks, faulty comparator isolation, and cross-layer comparison mechanisms to effectively handle complex scenarios such as malicious camouflage and low connectivity.

Ganlong Zhou, Dezhi Han, Xiang Shen, and Xiaoqi Xin in the paper: Decoupled Global-Local Collaborative Network for Visual Question Answering propose a Decoupled Global-Local Collaborative Network (DGL-Net) that structurally partitions the cross-modal reasoning process into two complementary branches: global semantic understanding and local fine-grained modeling. The global branch employs hierarchical attention fusion to model long-range dependencies and holistic scene semantics, providing contex-

tual constraints for reasoning. The local branch introduces semantic relation enhancement and relevant region screening to explicitly model intra-modal structural relationships, strengthening fine-grained region-word alignment by injecting structural information. The authors state that the proposed method exhibits enhanced robustness and generalization in complex semantic and compositional reasoning tasks, validating the efficacy of the decoupled collaborative strategy.

Cloud computing has become a dominant paradigm for delivering scalable and flexible on-demand resources. However, efficiently executing high-performance computing (HPC) workloads remains challenging, particularly in heterogeneous environments. In the paper: SLA-DE-RALBA: Cost-efficient dynamic enhanced resource-aware load balancing algorithm for cloud computing, Mohsin Nawaz, Altaf Hussain, Marran Al Qwaid, and Muazzam A. Khan Khattak advocate that conventional static scheduling methods often lead to poor resource utilization and increased makespan, while dynamic approaches improve load distribution but introduce significant overhead due to continuous monitoring and real-time decision-making. To address these challenges, they propose an SLA-aware Dynamic Enhanced Resource-Aware Load Balancing Algorithm (SLA-DE-RALBA). The algorithm minimizes load imbalance by considering the computational capacities of virtual machines and ensures the Service Level Agreement compliance through a three-tier priority-based workflow.

Data validation in ETL workflows is most often performed after transformation, even though many data quality issues originate earlier in the pipeline. The authors in the paper: VR-ETL: Leveraging Virtual Reality for Pre-Transformation ETL Data Validation, Maja Cerjan, Kornelije Rabuzin, and Leo Mršić state that in practice errors are detected late, when tracing their source and applying corrections already requires additional effort. The authors in the paper explore whether virtual reality (VR) could be used as an alternative interface for addressing this problem earlier, at the stage of raw data inspection. They propose VR-ETL, a proof-of-concept system that allows users to inspect and correct tabular data before transformation by interacting with it inside an immersive 3D environment. The system imports CSV datasets and applies a set of rule-based validation checks targeting common data quality problems. Detected issues are visualized directly within the 3D representation of the data, using color-coded cells and simple spatial summaries that reflect error distribution across columns and records.

In the paper: Methodology for Rapid Development of PAIS Based on Ontological Modeling and Category Theory, Borivoj Bogdanović, Vidan Marković, Milan Segedinac, and Zora Konjović propose a methodology for the rapid development of Process-Aware Information Systems (PAIS) based on ontological modeling and category theory. Ontology-Driven Code Generation (ODCG) is formalized as a deterministic, multi-phase transformation from the BAB Core ontology, through a platform-specific ontology and library, to source code. Model levels are represented as categories, while transformations are represented as functors. The main contribution is a categorical formalization of stack-parametrized multi-target generation, together with a conditional guarantee of schema-level structural preservation under explicitly stated verification assumptions, including mapping completeness, injectivity, quotient compatibility, and absence of unintended collapse. The methodology distinguishes the categorical specification, in which functoriality follows by construction, from its operational implementation, whose conformance is established through explicit engineering checks.

In their paper: *Leveraging Content-based Metadata in Catalogs to Improve Discoverability in Data Spaces*, Adriana Morejón, Alberto Berenguer, Lucía de Espona, David Tomás, and Jose-Norberto Mazón state that data spaces facilitate sovereignty-preserving data sharing within federated ecosystems, while dataset discovery is often constrained by catalogs that rely primarily on high-level metadata standards such as DCAT-AP. These models typically lack structural and content-level details, which impedes fine-grained relevance assessment and complicates both internal dataset discovery and the federation of external open data into data spaces. The authors propose cbmDCAT-AP, a backward-compatible extension of DCAT-AP that integrates content-based meta-data as primary catalog elements. The proposed extension augments dataset descriptions with field-level semantics, and representative sample values, thereby enabling content-informed discovery without disclosing entire datasets.

On behalf of the ComSIS Consortium, we would like to take this opportunity to give our great thanks to the reviewers and all the authors for their high-quality work, great efforts, and remarkable enthusiasm.