Book of Abstracts

The Fifth Serbian International Conference on Applied Artificial Intelligence (SICAAI)

May 20-21, 2026, Kragujevac, Serbia

The Fifth Serbian International Conference on Applied Artificial Intelligence (SICAAI)

May 20-21, 2026, Kragujevac, Serbia

The Fifth Serbian International Conference on Applied Artificial Intelligence, Kragujevac - Book of Abstracts

Editor
Professor Nenad Filipović

Technical assistants
Đorđe Ilić
Ognjen Pavić
Lazar Dašić
Teodora Jeremić
Nemanja Marković
Dajana Jovanović
Anđela Stojadinović

Proofreaders
Neda Vidanović Miletić
Tamara Janevska

Publisher
University of Kragujevac, Serbia

Press
Izdavačko Preduzeće Epoha DOO Požega

Impression
100 copies

Year of publication
2026

ISBN- 978- 86-81037- 93-5

CIP - Каталогизација у публикацији Народна библиотека Србије, Београд

004.8(048)

SERBIAN international conference on applied artificial intelligence (5 ; 2026 ; Kragujevac)
Book of Abstracts / The Fifth serbian international conference on applied artificial intelligence (SICAAI), may 20-21, 2026, Kragujevac, Serbia ; [editor Nenad Filipović]. - Kragujevac : University of Kragujevac, 2026 (Požega : Epoha). - 100 str. : fotogr. ; 30 cm

Tiraž 100. - Str. 5: Welcome message / Nenad Filipović.

ISBN 978-86-81037-93-5

a) Вештачка интелигенција -- Примена -- Апстракти

COBISS.SR-ID 193602057

Organizers

- University of Kragujevac

Sponsors

- Ministry of Science, Technological Development and Innovation of the Republic of Serbia
- Bioengineering Research and Development Center (BIOIRC)
- Data cloud technology
- IDEA DOO
- City of Kragujevac
- Vodéna
- Mind Group
- AA Craft Studio

Welcome Message

Dear colleagues and students,

On behalf of the Organizing Committee, it is a pleasure to welcome you at the Fifth Serbian International Conference on Applied Artificial Intelligence (AAI2026) which takes place at Kragujevac, Serbia, on May 20th-21st, 2026.

AAI2026 provides an exceptional Serbian and international forum to share the state-of-the-art research knowledge and results on the innovative theories, methodology and applications of artificial intelligence and its sub-domain like deep learning, machine learning in different areas such as medicine, economy, education, law, smart city, government, industry etc. Moreover, the conference aims to provide a platform for researchers and practitioners for both academia and industry to share the information about cutting-edge developments in the field of artificial intelligence.

It also aims to:

- provide early-stage researchers with an inspiring event allowing them to connect to relevant experts in related fields;
- provide an exciting venue for researchers to network and establish national and international collaborations;
- bring together leading experts from all relevant scientific domains to enhance the understanding of Artificial Intelligence;

Topics cover the following:

AI IN DOMAIN-SPECIFIC APPLICATIONS

- AI in Computational Biology, Medicine and Biomedical Applications
- AI in WWW, Communication, Social Networking, Recommender Systems, Games and E-Commerce
- AI in Finance and Risk Management

AI IN DATA ANALYTICS AND BIG DATA

- Visual Analytics for Big Data
- Computational Modeling for Big Data
- Large-scale Recommendation and Social Media Systems
- Cloud/Grid/Stream Data Mining for Big Velocity Data
- Semantic-based Big Data Mining

MACHINE LEARNING AND DATA MINING

- Pre-processing, Dimension Reduction and Feature Selection Computing, Bayesian and Neural Networks
- Learning Graphical Models and Complex Networks
- Active, Cost-Sensitive, Semi-Supervised, Multi-Instance, Multi-Label and Multi-Task Learning
- Transfer/Adaptive, Rational and Structured Learning

AAI2026 will host 16 keynote researchers**.** We have received more than 100 high-quality research papers. As a result of the strict review process and evaluation, the committee selected over 80 papers as extended abstracts. After the review, full papers from the AAI2026 conference will be published by Springer Verlag in the series "**Applied Artificial Intelligence 5: Medicine, Biology, Chemistry, Financial, Games, Engineering**" and Special issue of IEEE OPEN JOURNAL OF ENGINEERING IN MEDICINE AND BIOLOGY with title: "**AI and computational modelling in cardiovascular disease**". We must also admit that the conference certainly would not have been so successful without the efforts of many people who were actively engaged in organization of such a major academic event. We express gratitude to the members of the program and scientific review committee as well as to all the chairs, organizers and committee members for their dedication and support.
On behalf of the Organizing Committee, we wish you all a pleasant stay at Kragujevac and a productive conference.

Prof. Nenad Filipović, Conference Program Chair

Conference Committees

Conference Chair

- Nenad Filipović, University of Kragujevac

Honorary Advisory Board

- Miloš Đuran
- Miloš Kojić
- Veljko Milutinović

Scientific Committee

- Martin Aleksandrov (TU Berlin, Germany)
- Sandra Avila (University of Campinas (Unicamp), Brazil)
- Christian Blum (Spanish National Research Council (CSIC), Spain)
- Carlos Cardonha (University of Connecticut, United States)
- Vinay Chaudhri (United States)
- John Chinneck (Carleton University, Canada)
- Andy Chun (City University of Hong Kong, Hong Kong)
- Andre Augusto Cire (University of Toronto, Canada)
- Bradley Clement (Jet Propulsion Laboratory, United States)
- Dubravko Ćulibrk (University of Novi Sad, Serbia)
- Veljko Milutinović (University of Kragujevac and University of Belgrade, Serbia)
- Diane Cook (Washington State University, United States)
- Gabriella Cortellessa (CNR-ISTC, National Research Council of Italy, Italy)
- Lizhen Cui (Shandong University, China)
- Akay Metin (University of Houston, USA)
- Allen Robert (University of Southampton, UK)
- Zoran Bosnić (University of Ljubljana, Slovenia)
- Zlatan Car (Univeristy of Rijeka, Croatia)
- Ciaccio Edward (Columbia University, USA)
- Themis Exarchos (University of Ioannina, Greece)
- Dimitrios Fotiadis (University of Ioannina, Greece)
- Nikola Jorgovanović (University of Novi Sad, Serbia)
- Zoran Marković (IIT, Serbia)
- Michalopoulos George (University of Pittsburgh, USA)
- Nikita Konstantina (National Technical University of Athens, Greece)
- Zoran Obradović (Temple University, USA)
- Ouzounis Christos (King's College, UK)
- Pattichis Constantinos (University of Cyprus, Cyprus)
- Sheu Phillio (University of California, USA)
- Stojanović Radovan (University of Montenegro, Montenegro)
- Tsiknakis Manolis (Hellenic Mediterranean University, Greece)
- Yang Guang-Zhong (Imperial College London, UK)
- Zervakis Michalis (University of Crete, Greece)
- Andre de Carvalho (University of São Paulo, Brazil)
- Luca Di Gaspero (DPIA – University of Udine, Italy)
- Matthew Gaston (Carnegie Mellon University, United States)
- Carmen Gervet (Université de Montpellier, France)
- Odd Erik Gundersen (Norwegian University of Science and Technology, Norway)
- Koen Hindriks (Vrije Universiteit Amsterdam, Netherlands)
- Neil Jacobstein (Singularity University, United States)
- Binbin Jia (Southeast University, China)
- Elias Khalil (Georgia Institute of Technology, United States)
- Lars Kotthoff (University of Wyoming, United States)

- Hoong Chuin Lau (Singapore Management University, Singapore)
- Jimmy Lee (The Chinese University of Hong Kong, Hong Kong)
- Lee Mccluskey (University of Huddersfield, United Kingdom)
- Felipe Meneguzzi (Pontifical Catholic University of Rio Grande do Sul, Brazil)
- Mitra Nasri (Delft University of Technology, Netherlands)
- Barry O'Sullivan (University College Cork, Ireland)
- Michael Orosz (University of Southern California Information Sciences Institute, United States)
- Simon Parsons (University of Lincoln, United Kingdom)
- Andrew Perrault (Harvard University, United States)
- David Pynadath (University of Southern California, United States)
- Claude-Guy Quimper (Laval University, Canada)
- Howard Shrobe (Massachusetts Institute of Technology, United States)
- Madhav Sigdel (University of Alabama in Huntsville, United States)
- David Stracuzzi (Sandia National Laboratories, United States)
- Dimitris Stripelis (University of Southern California, United States)
- Nirmalya Thakur (University of Cincinnati, United States)
- Kevin Tierney (Bielefeld University, Germany)
- Michael Trick (Carnegie Mellon University, United States)
- Pradeep Varakantham (Singapore Management University, Singapore)
- Deng-Bao Wang (Southeast University, China)
- Shinjae Yoo (Brookhaven National Laboratory, United States)
- Yingqian Zhang (Eindhoven University of Technology, Netherlands)
- Jovan Stojanović (Serbian AI Society, Serbia)

Organizing Committee

- Vladimir Ranković, University of Kragujevac
- Tijana Geroski, University of Kragujevac
- Bogdan Milićević, University of Kragujevac
- Aleksandra Vulović, University of Kragujevac
- Igor Saveljić, University of Kragujevac
- Miljan Milošević, University of Kragujevac
- Boban Stojanović, University of Kragujevac
- Miloš Ivanović, University of Kragujevac
- Vesna Ranković, University of Kragujevac

Local Organizing Committee

- Ognjen Pavić, University of Kragujevac
- Lazar Dašić, University of Kragujevac
- Đorđe Ilić, University of Kragujevac
- Đorđe Dimitrijević, University of Kragujevac
- Anđela Blagojević, Univeristy of Kragujevac
- Teodora Jeremić, University of Kragujevac
- Nemanja Marković, University of Kragujevac
- Filip Filipović, University of Kragujevac
- Jovana Marković, University of Kragujevac
- Milica Kaplarević, University of Kragujevac
- Milena Đorđević, University of Kragujevac
- Marija Gačić, University of Kragujevac
- Neda Vidanović Miletić, University of Kragujevac

Table of Contents

A Hybrid Deep Learning and Gradient Boosting Model for Fully Automated Non-Invasive Fractional Flow Reserve Assessment 21

Comparative Analysis of LSTM and BiLSTM Models for Patient Blood Glucose Prediction 22

AIoT Integration in Healthcare: Wearable Cardiac Monitoring Case Study 23

From Inbox to Knowledge Graph: An Agentic Pipeline for Email Entity and Relationship Extraction ..... **Error! Bookmark not defined.**

Applications of Artificial Intelligence in Balance Disorders and Rehabilitation 25

Graph Algorithms as First-Class Reasoning Primitives for Small Language Models: A Graph-of-Thoughts Workflow with Embeded Graph Algorithms for Financial Fraud Detection ..... **Error! Bookmark not defined.**

Specialized Agents over Production Knowledge Graphs: A Verified Graph-Rag Architecture for Manufacturing Operations ..... **Error! Bookmark not defined.**

Next Level of Auditing, Risk Management and Compliance in the Telecommunications Industry - AI as a Facilitator 28

Energy Management Challenges and Solutions for Hyperscale AI Data Center 29

Impact of Image Compression Levels on Wildfire Detection Ability using YOLO Models .. 30

AIoT for Small Farms and Households: A Feasible Approach 31

A Data-Driven Machine Learning Framework for Defining Ecological Patterns and Detecting Climate-Driven Responses in Freshwater Ecosystems 32

Scalable Skyscape Analysis using the Barnes-Hut Adaptive Clustering Algorithm 33

AI-Assisted Authoring for Gamified Microlearning: Embedding Generation in the Content Creator's Workflow ..... **Error! Bookmark not defined.**

Monitoring the Impact of Urbanization on Land Cover Change in Serbian Cities using Sentinel-2 Data and Machine Learning 35

A Multimodal AI-Based Framework for Detection of Muscle Fatigue and Knee Biomechanics using sEMG and CNN-LSTM Models 36

Deep Learning for Alzheimer's Disease Classification: A Comparative Analysis of Convolutional Neural Network Architectures 37

LLM-Based Master Thesis Report Generation 38

A Robust Postprocessing Grounding Framework for Monocular 3D Motion Recovery: Benchmarks on MLB Baseball Footage 39

Generating BDD Scenarios from Source Code using Large Language Models 40

Comparative Analysis of Different LLM Models in Detecting Security Vulnerabilities in Source Code 41

Hierarchical Feature Fusion for Improved Optical Character Recognition using CNN-Transformer Architecture 42

Stabilizing Service Operations Through Applied Artificial Intelligence under Workforce Pressure: Evidence from Hospitality Organizations 44
The Role of Artificial Intelligence and Digital Marketing in Shaping Green Investment Trends and Organizational Competitiveness: Evidence from Serbia and the European Union..... 45
AI-Supported Analysis of Organizational, Investment and Digital Marketing Determinants of Sustainable Competitiveness in Serbia and the European Union..... 46
Synergy Between Corporate Economic Analysis and Banking Management in the Credit Management Process..... 47
From Predictive to Prescriptive AI: Value Optimization in Complex Decision Systems 48
An Informer-Based Deep Learning Model for Stock Market Index Prediction..... 49
Recovering Demographics from Survey Responses: An LLM-Based Approach..... 50
Conversational AI Agent for Intelligent Search over Open Access Repositories of Research Outputs 51
A Local Retrieval Augmented Generation Approach for Internal Medicine Question Answering 52
Hybrid Coding Assistance: Evaluating Local-Plus-Paid LLM Workflows Against Paid-Only Coding 53
Neural Algorithmic Reasoning for Zero-Sum Game Theory: Mastering Minimax Adversarial Search 54
LLM-Based Optimization of X-Ray Spectral Filtering for Improved CT Dose Efficiency 55
The Deep-Micro-Core Project: Machine Learning Algorithms for the Identification of the Core Microbiome from Integrated Data on Multiple Species..... 56
Three-Domain EEG Feature Selection for Distinguishing Neutral and Incongruent Stroop States 57
Physics-Informed Neural Networks for Vancomycin Pharmacokinetic Parameter Estimation from Sparse Concentration Data 58
Hepatotoxicity Classification upon Drug Chemical Composition and Molecular Structure - Deep Learning Approach 59
Deep Learning for Zebrafish Embryo Phenotype Classification using the Embryonet Dataset 61
Physics-Informed Neural Networks for Modeling the Hodgkin-Huxley Electrophysiological System 63
Digitalization and Digital Transformation of the Economy: Two Sides of the Same Coin..... 64
Transforming Machine Learning Models to Decision Support Systems 65
Perception of Generative AI Among Serbian Students in a Global Context 66
Entrepreneurship in Twelve Steps: An LLM-Oriented Case Study..... 67
Twelve Entrepreneurship-Oriented Paradigms Supported with AI 68
Mapping a Specific Science-Oriented Software onto 12 Paradigms of Interest for Entrepreneurship 69

Constitution and Artificial Intelligence 70
AI Memorization Represents Copyright Infringement: Regional Court in Munich Judgment *Gema vs. Open AI* 71
Artificial Intelligence: Lawmaking and the Application of Law 72
The Compliance Function in an AI Environment: Between AI Governance, Data Protection and Cybersecurity 73
Trust under Pressure: Explainable AI for Crisis Communications - A Literature Review 74
Legal Aspects of Liability of AI Systems with Special Reference to the EU AI Act 75
A Cybersecurity Trust Check Module for Risk Triage of Suspicious Messages 76
Holistic Secure Platform for Intelligent Threat Awareness and Business-Continuity in Trusted Adaptive Healthcare Environments 77
Metaheuristic Optimization of LGBM for Software Defect Prediction 78
Visual Analysis of User Interactions over Time for the Detection of Malicious Activities using Machine Learning 79
AI Risk Management: A Structured Approach to Identifying, Assessing, and Mitigating AI-Driven Threats 80
Design of STRATIFYHF: AI-Driven and Computational Modelling Cloud-Based Decision Support System for Heart Failure 81
Echocardiography View Classification as an Important Step to Heart Failure Diagnosis: A Case Study of the ECHOjepa Foundation Model 82
Differentiating Suspected and Confirmed Heart Failure using Machine Learning and Refined Vocal Features 83
Comparative Analysis of U-Net-Based Architectures for Coronary Artery Segmentation using X-Ray Angiography Images 84
SoC-Based Implementation of CNN Model for End-Diastolic Volume Classification from Echocardiogram via HLS4ML 85
Evaluation of Latest ChatGPT Models for Disc-Based Personality-Type Recognition from Interview Responses 86
Acceleration with Twelve Paradigms: An LLM-Oriented Case Study 87
Analysis of Attention Mechanisms in the Context of Computational Paradigms 88
Comparison of Acceleration Paradigms for Selected AI Algorithms 89
A Semantic JSON Schema for Knowledge Graphs of Electrospun PVDF Processing, Properties, and Applications 90
A Two-Level Cloud-Edge Framework for Personalized Glycemic Prediction 91
Discovering Physical Laws from Experimental Data using Regression Analysis 92
What Does Unsupervised ML Say? Clustering Virtual Reality Behavior for Typical and Physically Disabled Persons 93
Mapping AI Applications in Education - A Conceptual Synthesis of the Literature 94

Automated License Plate Recognition in Constrained Environments: A Case Study on AI-Assisted Development, Security Implications and Sustainability..... 95

Towards Intelligent Automated Essay Grading using Large Language Models..... 96

RGB-D Based 4DOF Grasp Prediction via CNN Segmentation for Assistive Robotics..... 97

Advances in Artificial Intelligence and Deep Learning for Neurosonographic and Chest X-Ray Analysis of Premature Infants..... 98

Mapping Ai In Osteochondral Tissue Engineering: A Reproducible Python Framework…...99

Keynote Speakers

Prof. Dan Shechtman
Technion, Haifa, Israel

Technological Entrepreneurship - Key to World Prosperity and Peace

Over the past several decades we witness a shift toward national policies that encourage innovation and technological entrepreneurship. The call for more investment in entrepreneurship echoes around the globe as it becomes clear that, except for a few countries, natural resources like oil and minerals are not enough to sustain economies, while human ingenuity is indeed the most important, sustainable natural resource.
So, is there hope for everybody in the world to improve their lives? Can technological entrepreneurship be motivated and taught so that generations of determined entrepreneurs will build up thriving economies? The clear answer to both questions is yes and it all starts with education in general and scientific-technical education in particular. This is a long process, but there is a way to expedite it – start with the already educated engineers and scientists. These are the first candidates to open entrepreneurial endeavors. They can make the difference, but need motivation, instruction and encouraging economic environment that fosters creation of successful start-ups. These pioneering entrepreneurs can then serve as role models to others. The name of the game is motivation. If this nucleus of capable people is motivated toward entrepreneurship, a process can start that will make a huge difference in a life of a country. Living examples to countries that underwent this process are China, Israel and Singapore whose societies shifted from agrarian to industrial within several decades thanks to the spirit of entrepreneurship and the motivation to create high-tech industries led and guided by individual engineers and scientists.
In my talk I will explain the need for technological entrepreneurship and describe my involvement in turning Israel into a startup nation.

Prof. Hermann Maurer
Fakultät für Informatik und Biomedizinisch Technik, Institut für Human-Centered Computing (HCC), Graz University of Technology, Austria

Combating polarization through credible information exchange: a trust-centered social media architecture

Social media has become one of the primary sources of information, especially for younger generation. However, the absence of credibility signals in this medium has enabled widespread misinformation, propaganda, and polarization in society. This paper presents ATSNET (A Trustworthy Social Media Network), a hybrid framework that integrates agentic AI analytics coupled with domain expert oversight to assess the credibility and authenticity of publicly shared viewpoints. ATSNET employs layered NLP processing techniques, source analysis, and knowledge-driven inference using large language models to generate post content and author specific credibility rankings. These rankings are prepared as data matrix and communicated through multi-level trust indicators. By combining scalable AI with transparent, pluralistic expert validation ATSNET aims to foster an accountable trusted digital public discourse.

Prof. Milan Đorđević
School of Computer Science, UOWD Building, Dubai Knowledge Park, Dubai, UAE

How Computers See: Convolutional Neural Networks for Medical Image Understanding

This presentation begins with how computers "see" by converting images into numerical pixel grids and using Convolutional Neural Networks (CNNs) to learn visual patterns at multiple levels, from edges and textures to complex structures. CNNs apply convolution, non-linear activation, and pooling to build increasingly abstract feature representations suitable for classification and detection tasks. Within this framework, we highlight pneumonia detection from pediatric chest X-ray images as an example of how learned features can distinguish healthy from diseased lungs and support early, consistent clinical decision-making in real-world healthcare environments.

Prof. Aleksandar Kartelj
The Faculty of Mathematics, University of Belgrade, Serbia

Navigating Large Combinatorial Search Spaces: From Classical Optimization to Topological and Neural Methods

Large combinatorial search spaces arise in many NP-hard optimization problems, where classical metaheuristics such as variable neighborhood search and genetic algorithms have shown long-term effectiveness. These methods rely on carefully designed neighborhoods and perturbations but often ignore deeper structural properties of the search space. This talk presents two complementary extensions of classical optimization. The first enforces topological regularities to guide exploration and stabilize search trajectories. The second integrates neural models to learn search patterns and bias optimization dynamics. Neural guidance naturally complements classical heuristics and hybrid frameworks, including integer linear programming based methods, leading to more adaptive and scalable optimization strategies.

Prof. Uroš Marković
School of Electrical Engineering, University of Belgrade, Serbia

A Comparison of Selected Emerging Computing Paradigmsfor the Floating Point Matrix Multiplication

General Matrix Multiplication (GEMM) is the cornerstone of modern computational workloads, from scientific simulations with big data, all the way to generative AI. Understanding GEMM's cross-platform performance is critical, as global investment in AI-driven architectures is essential for further developments in the field. This work analyzes floating-point GEMM kernels across a set of diverse computing paradigms, addressing the unique constraints of each one. We evaluate single- and multi-core CPU implementations, CUDA-based GPU kernels for many-core processing, and data-flow architectures utilizing Google's TPU. Additionally, we present a sketch of a stream-based kernel for the FPGA approach of Maxeler. This comparative study provides insights into optimizing this deceptively complex algorithm, concluding with an overview of emerging hardware trends and future work improvements.

Prof. Yun Li
Industrial Artificial Intelligence Centre (i4AI Centre); Shenzhen Institute for Advanced Study, UESTC, Shenzhen 518110, China

Embodied Intelligence and World Models

Today, the artificial intelligence (AI) large model is forging towards the physically-aware world-based model, spawning embodied intelligence. This talk covers the evolution of World Models and their role in Embodied Intelligence, as well as "AI for Engineering" for industrial innovation. The talk starts from AI Models based on the Artificial Neural Network (ANN), the Kolmogorov-Arnold Network (KAN), the large language model (LLM), the vision-language-action (VLA) model, and the state-space model (SSM). Together with nature-inspired computing, they help elevate "Computer-Aided Design" (CAD) in the third paradigm of science to "Computer-Automated Design" (CAutoD) in the fourth. This facilitates breaking through the intelligence limits of human engineers, whereby enlightening original creativity, enhancing design performance, and shortening development cycles. While the LLM-based AI is shifting to the world-based AI just now, a question arises as to how the next generation of AI and EI models may address real-world generalisation in a changing condition, a topic that is also to be discussed.

Prof. Branislav Bajat
Department of Geodesy and Geoinformatics, Faculty of Civil Engineering, University of Belgrade, Belgrade, Serbia

Application of Artificial Intelligence in Spatial Data Science

The rapid expansion of spatial data, driven by advances in satellite technologies, sensor networks, and location-based services, has created new opportunities for understanding
complex spatial phenomena across domains such as urban systems, environmental monitoring, and infrastructure management. At the same time, the volume, velocity, and heterogeneity of such data exceed the capabilities of traditional analytical approaches. In this context, Artificial Intelligence (AI) has emerged as a key enabler for transforming raw spatial data into actionable knowledge.
This work explores the application of AI techniques within Spatial Data Science, with an emphasis on real-world applications and integrative methodologies aligned with
contemporary AI research and practice. The paper examines how machine learning methods (e.g., Random Forest, Support Vector Machines) and deep learning approaches (e.g., Convolutional Neural Networks) can be effectively utilized for spatial prediction, classification, and pattern recognition tasks, particularly in the analysis of remote sensing data and geospatial imagery. Furthermore, the concept of Geospatial Artificial Intelligence (GeoAI) is discussed as a unifying framework that integrates AI, Geographic Information Systems (GIS), and large-scale data infrastructures, enabling scalable and automated solutions.
However, the application of AI in spatial domains also introduces specific challenges, including spatial autocorrelation, data scarcity and quality issues, model interpretability, and transfer learning. Addressing these challenges is essential for ensuring the reliability and responsible deployment of AI-based spatial systems. Special attention is given to application domains that reflect current priorities in applied AI research, including urban planning, environmental sustainability, disaster risk management, real estate assessment, and precision agriculture. In these contexts, AI-driven spatial models support decision-making processes through tasks such as change detection, resource optimization, and real-time monitoring. These contributions are consistent with the growing need for AI systems that address practical challenges across industry, government, and society.

Prof. Kyoung Mu Lee
Dept. of ECE, Seoul National University (SNU), Seoul, South Korea

Lessons Learned from the 2025 AAI Keynote

This presentation opens up the discussion about "Lessons Learned from the 2025 AAI Keynote of Mu Lee, Editor in Chief of the AI Journal with the Highest Impact Factor (IF) - IEEE/PAMI". It also gives the "guidelines" for publishing in currently the best-rated WoS journal in Computing - ACM/CSUR. With IF=28, it is the third best rated journal on the Planet, after Nature with IF=65 and Science with IF=45. In recent years, authors from Serbia have published successfully in ACM/CSUR.

Prof. Emil Jovanov
The University of Alabama in Huntsville, the USA

Seamless Health Intelligence: AI-Driven Internet of Medical Things for Ubiquitous Physiological Monitoring

Advances in artificial intelligence (AI) and the Internet of Medical Things (IoMT) are enabling a paradigm shift from episodic clinical measurements to continuous, real-time health monitoring in everyday environments. This talk presents an integrated framework for seamless physiological monitoring using multimodal sensors embedded in objects of daily use—an approach referred to as "Smart Stuff." By combining photoplethysmography (PPG), electrocardiography (ECG), bioimpedance, and video-based sensing, the system enables unobtrusive assessment of key health indicators, including blood pressure, blood glucose trends, cardiac activity, and autonomic nervous system dynamics. AI-driven signal processing and machine learning models fuse heterogeneous data streams to extract clinically relevant features, detect anomalies, and provide personalized feedback. As a representative example, the talk highlights the development of a smart water bottle equipped with integrated physiological sensors, enabling continuous monitoring during routine daily activities. This approach demonstrates the potential of ubiquitous, seamless monitoring to transform preventive healthcare and support scalable, patient-centered medical systems.

Prof. Rade Hajdin
Faculty of Civil Engineering, University of Belgrade, Serbia

Application of AI in Management of Transportation Infrastructure

Management of transport infrastructure refers to a systematic set of activities that support decision-making for the maintenance and improvement of infrastructure assets, ensuring their long-term, uninterrupted operation. It includes condition monitoring and the planning of optimal maintenance strategies. Transport infrastructure management is a highly data-intensive field, and the need for digitalization was recognized as early as the 1970s. Over the last decade, the potential of artificial intelligence to support and enhance the management of transportation infrastructure has been increasingly recognized and progressively realized. Three major areas of application are either close practical implementation or already in use: inventory acquisition, with a particular focus on BIM; damage detection and the processing of inspection findings and monitoring results; and the extraction of information from unstructured documents, such as reports, testing protocols, and related records. In the area of inventory acquisition, neural networks are used to identify and extract infrastructure objects from point clouds and technical drawings. Inspections and condition surveys can also be greatly facilitated by neural networks, which support the detection of damage in terms of both severity and extent. In addition, large language models are increasingly used to extract relevant information on the condition of infrastructure assets from testing protocols and reports. Finally, recent research on the application of physics-informed neural networks to the reassessment of bridges will also be presented.

Prof. Borko Furht

Florida Atlantic University, Boca Raton, Florida, USA

Industry-Sponsored AI Research: Recent Projects from the NSF IUCRC CAKE Center

Advances in artificial intelligence (AI) and the Internet of Medical Things (IoMT) are enabling a paradigm shift from episodic clinical measurements to continuous, real-time health monitoring in everyday environments. This talk presents an integrated framework for seamless physiological monitoring using multimodal sensors embedded in objects of daily use—an approach referred to as "Smart Stuff." By combining photoplethysmography (PPG), electrocardiography (ECG), bioimpedance, and video-based sensing, the system enables unobtrusive assessment of key health indicators, including blood pressure, blood glucose trends, cardiac activity, and autonomic nervous system dynamics. AI-driven signal processing and machine learning models fuse heterogeneous data streams to extract clinically relevant features, detect anomalies, and provide personalized feedback. As a representative example, the talk highlights the development of a smart water bottle equipped with integrated physiological sensors, enabling continuous monitoring during routine daily activities. This approach demonstrates the potential of ubiquitous, seamless monitoring to transform preventive healthcare and support scalable, patient-centered medical systems.

Prof. Zoran Bosnić

University of Ljubljana, Ljubljana, Slovenia

Educating humans in the age of artificial intelligence

Artificial intelligence is changing not only how students learn, but also what society expects from future graduates. As AI systems increasingly automate cognitive tasks once associated with highly educated professionals, universities are being forced to rethink the relationship between knowledge, employability, creativity, and human expertise. This keynote will explore how generative AI is reshaping higher education through changes in student learning habits, teaching practices, assessment methods, academic integrity, and curriculum design, while also raising broader questions about the future role of universities in society. The lecture will discuss which human competencies may become more valuable in the age of AI, including critical thinking, interdisciplinary reasoning, communication, adaptability, and responsible use of intelligent systems, and how educational institutions can respond to rapid technological and societal transformation. The keynote will conclude with the presentation of the joint European master's programme EMAI as an example of an international, interdisciplinary, and forward-looking educational model designed to prepare students for the emerging AI-driven world.

Prof. Đorđe Jakovljević
Coventry University, United Kingdom

Artificial Intelligence in Heart Failure: Prevention, Diagnosis, and Cure

Heart failure (HF) is a global pandemic currently affecting up to 15 million people in Europe. It is a complex clinical syndrome associated with impaired heart function, poor quality of life for patients and high healthcare costs. Accurate early diagnosis of HF allows implementation of evidence-based prevention and treatment strategies which reduce HF morbidity and mortality and its burden on healthcare. Early diagnosis of cardiac dysfunction and heart failure is challenging, often inaccurate as initial signs and symptoms are non-specific. Following development of a novel, cardiac output response to stress (CORS) test to facilitate early diagnosis and monitoring of heart failure in primary care, we have worked on development of artificial intelligence (AI) based tools to predict cardiac dysfunction and heart failure development in patients living with hypertrophic cardiomyopathy as part of Horizon2020 SILICOFCM project. We are conducting clinical retrospective and prospective studies to facilitate development of AI-based decision support system for risk stratification, early diagnosis, and disease progression in heart failure as part of ongoing STRATIFYHF project funded by Horizon Europe. Preliminary AI models demonstrate high accuracy in differentiating between heart failure and non-heart failure, as well as clear distension between different types of heart failure i.e. heart failure with reduced ejection fraction vs. heart failure with preserved ejection fraction. As part of STRATIFYHF project, AI methodology has been applied to processing and analysing recorded voice of patients as a powerful digital biomarker capable of predicting cardiac dysfunction. The most recently we have developed the project which will use AI to prevent cardiovascular dysfunction and disease in people living with Long COVID.

Chief Technology Officer at Supplyframe
Aleksandar Bradić
Siemens Digital Industries Software

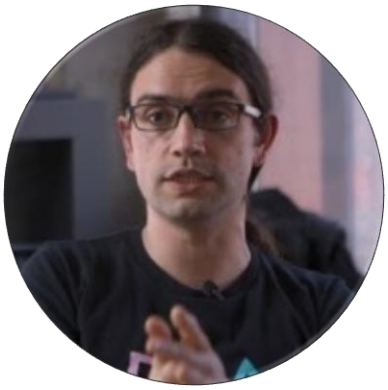

Time Series Analysis Meets Large Language Models: Challenges and Opportunities

Modern time-series forecasting, especially zero-shot forecasting with LLMs and foundation models, is often framed as a leaderboard problem: compare model families, identify the accuracy winner, and treat progress as the replacement of one dominant architecture by another. This framing is empirically useful but conceptually incomplete. Forecasting systems differ not only in average error, but also in their inductive biases, representations of time, failure modes, and computational costs. In this talk, we propose a framework for modern time-series analysis centered on bias-aware benchmarking, synthetic morphology generation, automated failure-mode discovery, and meta-level forecast selection under explicit accuracy and computational constraints. We hope such a framework can support a clearer understanding of the trade-offs involved in applying modern time-series methods in the foundation-model era.

Prof. David Naccache
Faculty of Engineering, University of Kragujevac, Serbia and DI ENS, École Normale Supérieure, PSL, CNRS,75005, Paris, France
(common work with Hadrien Barral, Gabriel Doriath Döhler and Aleksa Veličković)

LLMs as Partners in Conjecture Discovery

This paper describes an experiment during which we successfully derived a nontrivial mathematical relation with the help of an off-the-shelf Large Language Model (LLM). Our goal was to assess to what extent commercially available LLMs could be used for mathematical formula reconstruction. We hence constrained ourselves to:

- Use claude-sonnet-4-20250514 as is, i.e. avoid any ad hoc training.
- Use OEIS oeis.org), manually or through its API, to identify integer sequences encountered during the exploration.
- Target a nontrivial relation.

The intuition behind our experiment was that, despite their propensity for hallucination, LLMs demonstrate remarkable pattern recognition capabilities. The LLM's ability to repeat tedious operations and identify conceptual similarities was expected to detect patterns incompatible with human attention when massive amounts of individual experiments are to be conducted. The data processed during our experiments was assumed to follow some unknown logic but also present sporadic exceptions. We stress that our goal was not rigorous proofs but rather the exact formalization reverse-engineering of the underlying unknown logic. We end-up with both a formal resolution process by a human and an algorithm found, with the LLM's help, for deriving identical results. We did not formally prove the equivalence between the two results. This paper describes the discovery process.

Prof. Bharat Bhargava
Professor of CS at Purdue University

Situation Knowledge on Demand (SKOD)

The objective of this research is to fuse streaming data from multiple sources and identify rare events to alert the user and meet the mission requirements. The user can ask for specific information or the machine learning system will learn the needs/interest of user and forward new incoming data with a relevance score. Research questions such as trustworthiness of data, variation of data values from same source (such as sensor, video camera, user tweet, police incident report) are addressed due to uncertainty of data accuracy and noise. Application of the machine learning are to assist in security at military bases and the "missing person" problem. When a person is suspected missing police want to find him/her. The same problem arises in amber alerts, prison escapes, and missing children. When an incident report or 911 call arrives in police station, a physical description of the missing person (e.g., white male with medium built wearing a blue shirt, and black jeans) is available. Families may give additional details of missing child. Information such as specific medical conditions such as autism spectrum disorder or clinical depression may be available. This research is also improving the interaction of police with persons with mental issues.

Book of Abstracts

A HYBRID DEEP LEARNING AND GRADIENT BOOSTING MODEL FOR FULLY AUTOMATED NON-INVASIVE FRACTIONAL FLOW RESERVE ASSESSMENT

Aleksandar Milovanović[1*] [0000-0002-3251-8563], **Velibor Isailović** [2,3] [0000-0002-1417-9633], **Ognjen Matić**[2] [0009-0009-6697-029X] **and Igor Saveljić**[1,3] [0000-0002-0707-5174]

[1]Institute for Information Technologies, University of Kragujevac, Kragujevac, Serbia
[2]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
[3]Bioengineering Research and Development Center, Kragujevac, Serbia
e-mail: acakg85@hotmail.com
**corresponding author*

Abstract:

Fractional flow reserve (FFR) is the gold standard for assessing the functional significance of coronary artery stenosis. However, invasive FFR measurement carries risks and costs, while existing non-invasive analytical models suffer from limited accuracy due to imprecise geometric extraction from imaging and approximations of complex pressure loss integrals. Previous analytical approaches achieved only moderate correlation with clinical data ($r^2 = 0.726$) and still required manual adjustments. This study presents a novel hybrid AI-analytical model that combines deep learning for automatic geometry extraction with machine learning for coefficient optimization, delivering fast, fully automated, and highly accurate non-invasive FFR computation. Using 3D coronary artery models reconstructed from CT angiography, a convolutional neural network (CNN) was applied for automatic vessel segmentation and extraction of geometric parameters such as stenosis length, minimal cross-sectional area, proximal and distal reference areas, tapering angles, and entrance length. A gradient-boosting regressor trained on data from 120 patients was used to optimize coefficients for convective, diffusive, and expansion pressure-drop components and to approximate the numerical integrals describing diffuse and entrance-region losses. Final FFR values were calculated analytically assuming an aortic pressure of 100 mmHg and validated against invasively measured FFR (mFFR). The proposed hybrid model achieved significantly improved performance, with $r^2 = 0.892$, mean difference of –0.009, and standard deviation of 0.031. The misclassification rate for functionally significant stenosis (FFR ≤ 0.80) decreased from 17.5% to 3.3%. The model demonstrated particularly high precision in long stenosis where diffuse pressure losses predominate. Computation time per patient was reduced to less than 25 seconds on a standard CPU. By integrating CNN-based geometry extraction and gradient-boosting optimization into the analytical FFR framework, the hybrid model overcomes the main limitations of both purely analytical and computationally intensive CFD-based methods. Hybrid model offers a practical, clinically ready solution for rapid and accurate non-invasive assessment of coronary stenosis, with strong potential for real-time decision support in cardiology.

Keywords: coronary artery stenosis, fractional flow reserve, analytical method, artificial intelligence, machine learning, deep learning, non-invasive diagnostic.

COMPARATIVE ANALYSIS OF LSTM AND BILSTM MODELS FOR PATIENT BLOOD GLUCOSE PREDICTION

Ognjen G. Matić[1]* [0009-0009-6697-029X], **Tijana Geroski[1]** [0000-0003-1417-0521], **Igor Saveljić[2]** [0000-0002-0707-5174] **and Vesna Ranković[1]** [0000-0002-5445-9971]

[1]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjić Street, Kragujevac, Serbia
e-mail: ognjenmatic95@gmail.com, tijanas@kg.ac.rs, vesnar@kg.ac.rs
[2]Institute for Information Technologies, University of Kragujevac, 1A Liceja Kneževine Srbije Street, Kragujevac, Serbia
e-mail: isaveljic@kg.ac.rs
**corresponding author*

Abstract:

Accurate prediction of blood glucose levels is critical for timely management of hyperglycemia and hypoglycemia in patients with Type 1 diabetes mellitus. Diabetes represents a growing global health challenge, with ineffective glycemic control leading to severe long-term complications including cardiovascular disease, neuropathy and kidney failure. Numerous studies have demonstrated the effectiveness of deep learning approaches for blood glucose prediction, as they can capture complex nonlinear dynamics and long-term temporal dependencies in physiological signals. In this study, we present a comparative analysis of Long Short-Term Memory (LSTM) and Bidirectional Long Short-Term Memory (BiLSTM) architectures for blood glucose prediction 30 minutes ahead, using a personalized modeling approach with one model trained per patient.
Both models were evaluated on the OhioT1DM dataset (2018 and 2020 releases), comprising continuous glucose monitoring (CGM) data from 12 patients with Type 1 diabetes, recorded at 5-minute intervals. The dataset includes physiological and behavioral signals such as CGM readings, insulin doses, meal information, and physical activity data. Both models take a 2-hour historical window (24 time steps, 7 features) as input. The LSTM uses two layers with 64 hidden units each, while the BiLSTM employs two bidirectional layers with an effective hidden size of 128 units per layer. Both models were trained for 50 epochs using the Adam optimizer (lr = 0.001), Mean Squared Error (MSE) loss, dropout (0.2), and gradient clipping. Model performance was assessed using RMSE (Root Mean Squared Error), Mean Absolute Error (MAE), Mean Absolute Relative Difference (MARD), and Time in Good Range (TG). The LSTM achieved a mean RMSE of 31.44 ± 24.16 mg/dL, MAE of 18.73 ± 9.23 mg/dL, MARD of 11.00 ± 3.25%, and TG of 78.17 ± 7.24%. The BiLSTM achieved a mean RMSE of 29.83 ± 17.42 mg/dL, MAE of 18.78 ± 7.44 mg/dL, MARD of 11.30 ± 3.35%, and TG of 76.21 ± 8.63%. Results indicate that neither architecture consistently outperforms the other across all patients, suggesting that individual glucose dynamics play a significant role in model performance. These findings highlight the potential of personalized deep learning approaches in clinical decision support systems for diabetes management.

Keywords: blood glucose prediction, LSTM, BiLSTM, deep learning, continuous glucose monitoring, Type 1 diabetes mellitus, OhioT1DM, hyperglycemia, hypoglycemia.

Acknowledgement: Authors acknowledge the funding by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract numbers [451-03-34/2026-03/200107 (Faculty of Engineering, University of Kragujevac)].

AIOT INTEGRATION IN HEALTHCARE: WEARABLE CARDIAC MONITORING CASE STUDY

Radovan Stojanović[1*] [0000-0003-0906-2198] **and Jovan Djurković**[2] [0009-0007-3962-7328]

[1]Faculty of Electrical Engineering, University of Montenegro, Montenegro
e-mail: stox@ucg.ac.me
[2]Mediterranean Excellence in Computing and Ontology - MECOnet d.o.o. Podgorica, Montenegro
e-mail: jovan@meconet.me
**corresponding author*

Abstract:

This study explores the practical potential and limitations of integrating Artificial Intelligence (AI) with the Internet of Things (IoT) in modern healthcare systems through wearable devices. A real-world case study is presented using the CardiaFylax system, developed by MECOnet, designed for continuous monitoring of key cardiological signals including electrocardiogram (ECG), photoplethysmography (PPG), and accelerometric (ACC) data.
The device operates within the ambulatory/wearable diagnostic ECG class, featuring a sampling frequency of ≥250 Hz and A/D resolution of 10–12 bits. It supports real-time signal acquisition and on-device processing, including QRS detection, heart rhythm monitoring for arrhythmia detection, stress assessment via heart rate variability (HRV), and noise filtering, including motion artifacts. In PPG mode, additional parameters such as SpO_2 can be monitored.
CardioFylax enables multi-level system integration: as a standalone IoT node, as part of an edge system via an Android application, and within cloud-based environments including AI services and telemedicine platforms. Communication is achieved through BLE and USB at the edge level, and standard internet protocols for cloud connectivity.
AI functionalities are distributed across node, edge, and cloud layers, allowing scalable processing complexity. The integration with AI systems, including large language models, demonstrates enhanced diagnostic support and user interaction. This case study highlights how AI. –IoT integration can enable proactive, personalized, and remote healthcare solutions, while also addressing challenges related to data accuracy, interoperability, and system design.
System, integration and component levels will be discussed and demonstrated.
More about: https://meconet.me/smarthealth/

Keywords: AI, healthcare, wearables, IoT, cardiac

FROM INBOX TO KNOWLEDGE GRAPH: AN AGENTIC PIPELINE FOR EMAIL ENTITY AND RELATIONSHIP EXTRACTION

Nikola Prodanović[1][0000-0001-8698-9830], Ivan Vitezović[2][0009-0007-7651-4875], Bojana Bajić[3][0000-0001-7843-7091], Aleksandar Rikalović[4][0000-0003-2777-5706], Miloš Jovičić*[,5][0009-0006-2984-7047]

[1,2,5] Institute for Information Technologies, University of Kragujevac, Liceja Kneževine Srbije 1A, Kragujevac, Serbia
e-mail: iit@uni.kg.ac.rs
[3,4] The Institute for Artificial Intelligence Research and Development of Serbia, Novi Sad, Fruškogorska 1, Serbia
e-mail: nikola.prodanovic@gmail.com, vitezovicivan@gmail.com, bojana.bajic@ivi.ac.rs, aleksandar.rikalovic@ivi.ac.rs, milos.jovicic@uni.kg.ac.rs
**corresponding author*

Abstract:

The AI revolution is significantly affecting and will continue to affect all aspects of data processing and analytics. Corporate unstructured data presents an excellent example as a massive data segment that is very often underutilized by individuals and organizations and where AI tools can bring additional structure and value for the owners of such data. A special segment of such data lives in email inboxes representing individual professional activities and communications. We present an agentic workflow for transforming unstructured email communication into a structured knowledge graph for downstream analysis, retrieval, and decision support. The proposed approach orchestrates the following components: language-model-driven entity extraction, confidence-aware validation, relationship inference, and graph persistence into a unified pipeline for large-scale email processing. Raw emails are first ingested and parsed to extract structural metadata, including sender, recipients, subject, date. In the next run, email is being classified by a document classification component. With gathered metadata, the email body is then analyzed by an agentic extraction layer that identifies salient entities such as persons, organizations, projects, and other domain relevant concepts. To reduce extraction errors, the workflow employs a confidence-based routing mechanism in which uncertain outputs are passed to a validation stage before graph construction. The validated entities are subsequently used for semantic relationship extraction, enabling the system to capture not only isolated information items but also the contextual links among them. These outputs are normalized and persisted in a Neo4j knowledge graph. Our agentic workflow based system demonstrates how large collections of unstructured emails can be converted into a form that is significantly more accessible for intelligent analysis. The approach is particularly relevant for AI-driven knowledge management, enterprise information extraction, and graph-enhanced retrieval settings, where structured understanding of communication data is essential. The presented workflow demonstrates how modern agentic and graph-oriented methods can be combined to improve the usability and analytical value of large collections of unstructured enterprise communication.

Keywords: AI Agent, Entity Extraction, Email Processing Automation, Labeled Property Graph, Neo4j, GLiNER2.

APPLICATIONS OF ARTIFICIAL INTELLIGENCE IN BALANCE DISORDERS AND REHABILITATION

Filip Filipović[1*] [0009-0002-5321-5016], **Aleksandra Vulović**[1,2] [0000-0002-6726-3134], **Đorđe Ilić**[1] [0009-0009-8856-8226], **Nenad Filipović**[1,2] [0000-0001-9964-5615]

[1]Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
e-mail: filipfilipovic1999@gmail.com, djordje.ilic@kg.ac.rs
[2]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
e-mail: aleksandra.vulovic@kg.ac.rs, fica@kg.ac.rs
**corresponding author*

Abstract:

Artificial intelligence (AI) is increasingly transforming healthcare by enabling data-driven decision-making, personalized treatment planning, and continuous patient monitoring. In the context of balance disorders, AI offers tools to better understand complex sensorimotor impairments that affect posture, gait, and overall functional stability. Balance disorders significantly reduce quality of life, primarily due to an increased risk of falls and fall-related injuries, particularly in populations affected by neurological conditions such as stroke. The growing prevalence of such conditions has intensified the demand for effective and accessible rehabilitation strategies, while conventional clinic-based approaches remain limited in scalability and accessibility. Recent advances in this field focus on the integration of machine learning algorithms, computer vision, and wearable sensor technologies to objectively quantify balance performance and detect subtle deficits that may not be easily identified through standard clinical assessments. There is growing interest in adaptive rehabilitation systems that use AI to analyze patient progress over time and dynamically adjust therapeutic interventions. Despite these developments, challenges remain regarding data variability, clinical validation, and the interpretability of AI-driven recommendations in real-world rehabilitation settings. This review is aligned with an ongoing tele-rehabilitation project that employs avatar-guided exercise therapy, where patients perform structured movements demonstrated by a virtual avatar in a home-based or supervised environment. By integrating AI into tele-rehabilitation, the goal is to support clinicians in making more informed decisions, enhance personalization of therapy, and improve long-term rehabilitation outcomes for patients with balance disorders.

Keywords: artificial intelligence, balance disorders, tele-rehabilitation, stroke rehabilitation.

Acknowledgement: This research is co-funded by the European Union, under the Horizon Europe programme, GA number 101057747 (TeleRehaB project).

GRAPH ALGORITHMS AS FIRST-CLASS REASONING PRIMITIVES FOR SMALL LANGUAGE MODELS: A GRAPH-OF-THOUGHTS WORKFLOW WITH EMBEDED GRAPH ALGORITHMS FOR FINANCIAL FRAUD DETECTION

Miloš Jovičić[*,1][0009-0006-2984-7047] **,Nikola Prodanović**[2][0000-0001-8698-9830], **Ivan Vitezović**[3][0009-0007-7651-4875], **Bojana Bajić**[4][0000-0001-7843-7091], **Aleksandar Rikalović**[5][0000-0003-2777-5706]

[1,2,3] Institute for Information Technologies, University of Kragujevac, Liceja Kneževine Srbije 1A, Kragujevac, Serbia
e-mail: iit@uni.kg.ac.rs
[4,5] The Institute for Artificial Intelligence Research and Development of Serbia, Novi Sad, Fruškogorska 1, Serbia
e-mail: milos.jovicic@uni.kg.ac.rs, nikola.prodanovic@gmail.com, vitezovicivan@gmail.com, bojana.bajic@ivi.ac.rs, aleksandar.rikalovic@ivi.ac.rs
**corresponding author*

Abstract:

Large language model reasoning techniques such as Chain-of-Thought (CoT), Tree-of-Thoughts (ToT), Graph-of-Thoughts (GoT), and similar approaches have shown drastic improvements in the reasoning capabilities of LLMs and their ability to solve complex tasks. We improve on current research by embedding graph algorithms inside our hybrid GoT framework. GoT allows a smaller LLM to externalize reasoning into a graph-like structure, where intermediate thoughts can be generated, connected, aggregated, refined, and selectively revisited across multiple inference steps. Additionally, in our approach graph algorithms such as PageRank, Weakly Connected Components (WCC) and Louvain algorithm, are happening in parallel inside the GoT framework. We then seal the design with the verification layer as a leaf component of the GoT. Our research argues that established graph powered algorithms can be effectively used as reasoning primitives with smaller language models which very often lack deeper reasoning capabilities, and that inference level frameworks such as ToT, GoT or CoT are not enough for deep reasoning operations. Importantly, combining these approaches gives smaller language models the ability to perform operations such as fraud detection that are usually reserved for machine learning algorithms or state-of-the-art Graph Neural Networks (GNNs). We instantiate this position in a transactional fraud detection system built on two components: Temporal workflows act as a durable control plane orchestrating multiple graph algorithms running independently in parallel, and a small 4B-parameter model (Qwen 3.5-4B) fills six narrow agent roles, each backed by a deterministic rule-based fallback. Graph-of-Thoughts branch fan-out is implemented as parallel Temporal child workflows. In an initial evaluation, the system attains state-of-the-art performance in detecting fraudulent rings in a labeled property graph and outperforms a Mixture-of-Experts (MoE) reasoning model 70× its size from the DeepSeek-V4 series.

Keywords: Graph Data Science, Large Language Model, Small Language Model, AI Agent, Qwen.

SPECIALIZED AGENTS OVER PRODUCTION KNOWLEDGE GRAPHS: A VERIFIED GRAPH-RAG ARCHITECTURE FOR MANUFACTURING OPERATIONS

Miloš Jovičić*,1[0009-0006-2984-7047] ,Nikola Prodanović2[0000-0001-8698-9830], Ivan Vitezović3[0009-0007-7651-4875], Bojana Bajić4[0000-0001-7843-7091], Aleksandar Rikalović5[0000-0003-2777-5706]

[1,2,3] Institute for Information Technologies, University of Kragujevac, Liceja Kneževine Srbije 1A, Kragujevac, Serbia
e-mail: iit@uni.kg.ac.rs
[4,5] The Institute for Artificial Intelligence Research and Development of Serbia, Novi Sad, Fruškogorska 1, Serbia
e-mail: milos.jovicic@uni.kg.ac.rs, nikola.prodanovic@gmail.com, vitezovicivan@gmail.com, bojana.bajic@ivi.ac.rs, aleksandar.rikalovic@ivi.ac.rs
**corresponding author*

Abstract:

Large language models offer compelling natural language interfaces to enterprise operational data, but direct LLM-to-database interaction in manufacturing settings remains brittle: free-form query generation produces hallucinated entities, monolithic prompts conflate unrelated reasoning tasks, and outputs lack traceable grounding to source records, and all of the mentioned error can happen even with frontier LLMs. We present a bill of materials multi-agent system (BOM360) that connects natural-language queries from plant supervisors and operations managers to a live Neo4j production knowledge graph encoding bills of materials, production lines, supplier networks, and work instructions. Rather than a single generalist chatbot, BOM360 employs a team of domain-specialized analyst agents, each scoped to a single operational subdomain, coordinated by a LangGraph workflow that performs intent classification, deterministic graph retrieval via parameterized Cypher templates, agent execution with PydanticAI-typed structured outputs, and a verification step that cross checks every claim against the retrieved subgraph before returning a response. This architecture decomposes an inherently heterogeneous query surface (status reporting, risk assessment, procedural generation) into independently testable units while preserving end-to-end grounded output through enforced retrieval-then-verify discipline. We demonstrate the system on a manufacturing graph comprising 470 nodes and 1,970 relationships across product, process, and supply-chain entities, evaluating answer accuracy, hallucination rate, and latency across the three representative query classes above. We discuss design trade-offs between agent specialization granularity and orchestration overhead, and outline an evaluation protocol for grounded multi-agent GraphRAG systems in industrial settings

Keywords: Graph Database, Large Language Model, Retrieval Augmented Generation, AI Agent, GraphRAG, Industry 4.0.

NEXT LEVEL OF AUDITING, RISK MANAGEMENT AND COMPLIANCE IN THE TELECOMMUNICATIONS INDUSTRY - AI AS A FACILITATOR

Ognen Firfov

Makedonski Telekom, Kej 13 Noemvri, 6, Skopje, North Macedonia
e-mail: ognen.firfov@telekom.mk

Abstract:

Telecommunications operators operate in highly regulated, data-intensive environments characterised by complex infrastructure, large transaction volumes and cross-border regulatory obligations. Traditional risk management, compliance and audit models, largely based on periodic sample-driven reviews, often fail to provide timely assurance. This article analyses how artificial intelligence tools are applied in auditing, compliance and risk management across major global telecommunications operators. The study compares adoption models, implementation approaches and observed outcomes while discussing governance challenges associated with AI-enabled assurance.

Keywords: Artificial intelligence, AI governance, auditing, risk management, telecommunications, compliance, continuous auditing, process mining.

ENERGY MANAGEMENT CHALLENGES AND SOLUTIONS FOR HYPERSCALE AI DATA CENTER

Aleksandar Selakov [0000-0001-6384-5674]

Faculty of Technical Sciences, University of Novi Sad, 6 Trg Dositeja Obradovica, Novi Sad, Serbia
e-mail: aselakov@uns.ac.rs

Abstract:

The compute infrastructure powering modern AI is rapidly becoming one of the largest and most demanding loads on the global power grid. A single hyperscale AI facility can consume as much electricity as a small country, with rapid fluctuations driven by the alternating compute and communication phases of large-scale training jobs. These load patterns are fundamentally different from anything the power industry has encountered, and existing Energy Management Systems are not designed to handle them. This paper presents the operational and architectural challenges that emerge when AI data centers connect directly to the high-voltage transmission network. It introduces a four-layer control framework, organized by the time scale on which decisions must be made, ranging from millisecond-level stability protection to long-term operational analytics. The framework is grounded in direct engagement with one of the first deployments of its kind, and is intended to bridge the conversation between the AI community that builds these facilities and the power systems community that must operate them.

Keywords: data center, EMS, energy management, grid-to-token, AI infrastructure, microgrid.

IMPACT OF IMAGE COMPRESSION LEVELS ON WILDFIRE DETECTION ABILITY USING YOLO MODELS

Luka Jovanović[1*] [0000-0001-9402-7391], **Petar Spalević**[*1] [0000-0002-6867-7259], **Branimir Jakšić**[1] [0000-0001-6683-0021], **Lazar Karanović**[2] [0009-0002-1937-9498], **Bogdan Ignjatović**[3] [0009-0004-4425-1378] **and Dejan Milić**[4] [0000-0002-0229-5960]

[1]Department of Electrical Engineering and Computer Science, Faculty of Technical Sciences in Kosovska Mitrovica, University of Priština in Kosovska Mitrovica, 7 Knjaz Miloš Street, Kosovska Mitrovica, Serbia
e-mail: jovanovic.luka@pr.ac.rs, petar.spalevic@pr.ac.rs*, branimir.jaksic@pr.ac.rs
[2]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjić Street, Kragujevac, Serbia
e-mail: lazarkaranovic68@gmail.com
[3]Faculty of Business and Law, MB University, Belgrade, Serbia
e-mail: bokiexe@gmail.com
[4]National Security Academy, Belgrade, Serbia
e-mail: dejanmilic00@gmail.com
**corresponding author*

Abstract:

This study examines the impact of progressive image compression on object detection performance. Using an existing YOLO-based dataset, images are subjected to ten levels of compression, and the performance of a pre-trained model is systematically evaluated. This approach avoids repeated training while enabling controlled analysis of input degradation effects. The motivation stems from large-scale monitoring systems, where thousands of distributed cameras generate high data volumes, creating bandwidth and latency constraints. For near real-time operation, reducing data size is essential, with video stream compression being a practical solution. However, its effect on object detection performance remains insufficiently explored. This paper addresses that gap by providing empirical insights into how varying compression levels influence detection accuracy, informing the design of efficient and scalable real-world surveillance systems.

Keywords: computer vision, compression, wildfire, agriculture, forestry.

AIOT FOR SMALL FARMS AND HOUSEHOLDS: A FEASIBLE APPROACH

Radovan Stojanović[1*] [0000-0003-0906-2198], **Vesna Maraš**[2] [0000-0002-4694-7377] **and Dejan Zejak**[3,4] [0000-0002-1184-8846] **and Nedeljko Latinović**[2] [0000-0002-5355-7764]

[1] Faculty of Electrical Engineering, University of Montenegro, Montenegro
e-mail: stox@ucg.ac.me
[2] University of Montenegro, Faculty of Biotechnology, Podgorica, Montenegro
e-mail: vesnamaras6@gmail.com
[3] Biotechnical Center, Rakonje, XV/3, Bijelo Polje, Montenegro
[4] Institute for interdisciplinary and multidisciplinary studies, University of Montenegro, Montenegro
e-mail: zejakd@gmail.com
**corresponding author*

Abstract:

Precision agriculture integrates advanced technologies and data-driven methods to optimize crop management, increase yields, and minimize environmental impact. At its core, IoT-based sensing enables real-time, location-specific monitoring of soil and environmental conditions, supporting targeted interventions such as irrigation, fertilization, and pest control. However, small farms and household-scale agricultural systems face constraints including limited budgets, small field sizes, and the need for simple, energy-efficient solutions.

This paper proposes a feasible and cost-effective AIoT architecture tailored to such settings. The system is based on off-the-shelf sensors, low-power microcontrollers, and lightweight communication protocols, designed for deployments where the maximum distance between sensor nodes and gateway does not exceed 200 m. Emphasis is placed on low energy consumption through optimized hardware design, embedding AI in the general purpose microprocessors and communication strategies that reduce power usage and minimize transmission collisions.

The proposed system supports long-term operation using battery power or small-scale solar energy sources, ensuring practicality in resource-constrained environments. The AI support is simple and optimized. Preliminary experimental results demonstrate reliable and energy-efficient communication, confirming the system's suitability for small-scale precision agriculture applications.

A case study focusing on a small vineyard illustrates the practical implementation and benefits of the approach. The results highlight the high feasibility of deploying low-cost AIoT sensor networks to enable data-driven decision-making in compact agricultural environments, contributing to more sustainable and efficient farming practices.

Keywords: precision agriculture, AIoT, optimized design, small vineyard.

A DATA-DRIVEN MACHINE LEARNING FRAMEWORK FOR DEFINING ECOLOGICAL PATTERNS AND DETECTING CLIMATE-DRIVEN RESPONSES IN FRESHWATER ECOSYSTEMS

Marija M. Jakovljević[1*] [0000-0002-2168-7548], **Marijana D. Nikolić**[1] [0000-0003-1079-9453], **Vladimir Đurđević**[2] [0000-0001-9882-1189], **Irida Lazić**[2] [0000-0001-7355-2410] **and Vladica M. Simić**[1] [0000-0001-6787-9783]

[1]Institute of Biology and Ecology, Faculty of Science, University of Kragujevac, 12 Radoja Domanovića Street, Kragujevac, Serbia
e-mail: marija.jakovljevic@pmf.kg.ac.rs, marijana.nikolic@pmf.kg.ac.rs, vladica.simic@pmf.kg.ac.rs
[2]Institute of Meteorology, Faculty of Physics, University of Belgrade, 16 Dobračinska Street, Belgrade, Serbia
e-mail: vdj@ff.bg.ac.rs, irida.lazic@ff.bg.ac.rs
**corresponding author*

Abstract:

Climate change is reshaping freshwater ecosystems through complex interactions with persistent anthropogenic stressors, yet ecological datasets often lack predefined output classes, constraining supervised machine learning. To address this limitation, we developed a data-driven analytical framework based on two previously published studies on river ecosystems, using *Alburnoides bipunctatus* as a model species.
In the first study, preliminary relationships were examined using traditional statistical approaches, after which models were developed using seven biological and seven environmental variables. Uniform Manifold Approximation and Projection (UMAP) was introduced as a novel analytical tool to detect latent ecological structure and generate data-driven output classes from unlabeled ecological data. In the second study, this concept was expanded through ecological modelling based on seven population-response variables derived from two decades of species monitoring, combined with anthropogenic stressors and climate predictors. UMAP-derived classes were integrated with Decision Tree, Random Forest, Linear Discriminant Analysis, and k-Nearest Neighbours models. The performance was evaluated using accuracy, precision, recall, and F1-score, with Linear Discriminant Analysis achieving the highest predictive performance (accuracy = 0.91; precision = 0.89; recall = 1.00; F1-score = 0.94). Two dominant ecological response groups were identified, indicating a dual ecological strategy. Mean annual water temperature was consistently associated with class separation, while climate change acted as a background amplifier of existing stressors. The proposed framework converts unlabeled ecological data into interpretable prediction targets, enables robust supervised modelling, and provides a transferable tool for freshwater biomonitoring, supports early-warning assessment and reveals indicator potential rather than assuming indicator status *per se*.

Keywords: data-driven framework, machine learning, UMAP, climate change, freshwater biomonitoring.

SCALABLE SKYSCAPE ANALYSIS USING THE BARNES-HUT ADAPTIVE CLUSTERING ALGORITHM

Helga Hochbauer* [0009-0004-7125-6174] **and Marc Frincu** [0000-0003-1034-8409]

Faculty of Computer Science, West University of Timisoara, Romania
e-mail: helga.hochbauer@e-uvt.ro, marc.frincu@e-uvt.ro
**corresponding author*

Abstract:

Advancements in high-resolution satellite imagery have enabled the large-scale remote analysis of ancient heritage sites. In skyscape archaeology, calculating the astronomical declination of these structures is an essential step for understanding their cosmological significance. However, computing the required horizon profiles from digital elevation models is computationally expensive and scales linearly with the number of monuments. To address this bottleneck, this paper leverages spatial clustering techniques to intelligently group geographically close objects, allowing a single horizon profile to be computed for each cluster's geometric centroid. We propose a novel adaptation of the Barnes-Hut (BH) algorithm, re-engineered with a Topographic Homogeneity Check to group static archaeological sites while respecting physical terrain barriers. The adapted BH method is evaluated against baseline calculations and standard clustering algorithms (K-Means, Agglomerative, DBSCAN, and CLIQUE) using a dataset of 400 prehistoric pendant tombs in north-west Arabia. The results demonstrate that the BH algorithm provides a highly scalable solution, yielding significant computational speedups while acting as the most topographically stable and reliable method across varying dataset sizes.

Keywords: skyscape archaeology, spatial clustering, geospatial data processing, unsupervised learning, Barnes-Hut algorithm, computational geometry.

Acknowledgments. This work has been supported by a grant from the Ministry of Research, Innovation and Digitization, CNCS/CCCDI - UEFISCDI, project number ERANET-CHISTERA-IV-AI4MultiGIS, within PNCDI IV.

AI-ASSISTED AUTHORING FOR GAMIFIED MICROLEARNING: EMBEDDING GENERATION IN THE CONTENT CREATOR'S WORKFLOW

Nikola Prodanović[1][0000-0001-8698-9830], Ivan Vitezović[2][0009-0007-7651-4875], Bojana Bajić[3][0000-0001-7843-7091], Aleksandar Rikalović[4][0000-0003-2777-5706], Miloš Jovičić*[,5][0009-0006-2984-7047]

[1,2,5] Institute for Information Technologies, University of Kragujevac, Liceja Kneževine Srbije 1A, Kragujevac, Serbia
e-mail: iit@uni.kg.ac.rs
[3,4] The Institute for Artificial Intelligence Research and Development of Serbia, Novi Sad, Fruškogorska 1, Serbia
e-mail: nikola.prodanovic@gmail.com, vitezovicivan@gmail.com, bojana.bajic@ivi.ac.rs, aleksandar.rikalovic@ivi.ac.rs, milos.jovicic@uni.kg.ac.rs
**corresponding author*

Abstract:

Microlearning short, focused bursts of instructional content paired with quick formative checks has emerged as one of the most effective formats for workplace and continuing education, with retention and engagement gains amplified further when the experience is gamified through points, streaks, and confidence-based feedback. Yet the bottleneck has shifted from delivery to authoring: producing a single well-calibrated topic typically requires a subject-matter expert to write a tight 1–2 page lesson and then hand-craft roughly ten heterogeneous questions (single-choice, multiple-choice, true/false, short-answer) that are non-redundant, anchored in the lesson, and pedagogically sound. At organizational scale, across hundreds of topics and continual content refresh cycles, this manual workload is the primary reason gamified microlearning programs stall before they reach critical mass. Large language models can absorb the bulk of this effort, but only when they are integrated into the authoring surface itself rather than offered as an external tool the creator must context-switch to. We present Quizkko, a microlearning platform whose content-admin UI embeds AI generation directly next to the artifacts being authored: a "draft theory" action seeded by a short brief produces a publish ready lesson body, and a "generate questions" action fills a topic to its quota while honoring the existing theory, the per-topic media pool, and already-authored questions as deduplication context. Generation is slot-allocated (one media-grounded question per uncovered asset, theory-only questions for the remainder), structurally validated with a single retry, and every call is logged for telemetry. The result is a workflow in which the human stays in the role they are best at reviewing, correcting, and approving while the cost of producing a new gamified microlearning topic drops from hours to minutes, making large, continuously-refreshed content libraries operationally feasible.

Keywords: microlearning; gamification; quiz-based learning; AI-assisted content authoring; large language models; human-in-the-loop; instructional design; educational technology; authoring tools; learning engagement

MONITORING THE IMPACT OF URBANIZATION ON LAND COVER CHANGE IN SERBIAN CITIES USING SENTINEL-2 DATA AND MACHINE LEARNING

Karlo Bala [0000-0002-9267-236X], **Jelena Petković** [0009-0004-7431-8158], **and Filip Josifljević*** [0009-0000-3271-8138]

[1]The Institute for Artificial Intelligence Research and Development of Serbia, Naučno-tehnološki park Novi Sad, 1 Fruškogorska Street, Novi Sad, Serbia
e-mail: karlo.bala@ivi.ac.rs, jelena.petkovic@ivi.ac.rs, filipajosifa@gmail.com
**corresponding author*

Abstract:

Urbanization leads to significant changes in land cover structure, particularly through the reduction of vegetation and agricultural areas. Monitoring these changes is crucial for understanding the environmental impact of urban expansion and for supporting sustainable development planning. In this study, we analyze land cover changes in selected cities across Serbia from 2016 to 2025 using multispectral satellite imagery from the Sentinel-2 mission. The input data include all available spectral bands from Sentinel-2. Several machine learning models were applied for land cover classification, including logistic regression as a baseline, as well as Random Forest, LightGBM, and a multilayer perceptron. The models were trained using a combination of ground truth data and Sentinel-2 Level-2A imagery from Slovenia (2019), selected as a representative dataset from a geographically and climatically similar region. The trained models were then used to predict land cover types for satellite imagery covering Serbia from 2016 to 2025. This study provides a systematic analysis of land cover changes in urban areas of Serbia, based on multi-year Sentinel-2 data and a comparative evaluation of multiple machine learning models. The results indicate a significant increase in urban areas, accompanied by a decline in vegetation across the analyzed regions. These findings can support the development of environmental protection strategies and improve urban planning processes through the application of modern remote sensing and data analysis methods.

Keywords: remote sensing, multispectral analysis, change detection, spatial analysis

A MULTIMODAL AI-BASED FRAMEWORK FOR DETECTION OF MUSCLE FATIGUE AND KNEE BIOMECHANICS USING SEMG AND CNN–LSTM MODELS

Djordje Kosanic[1*] [0009-0005-5996-6031], **Elvis Mahmutovic**[1] [0000-0003-1849-868X] **and Nenad Filipovic**[2] [0000-0001-9964-5615]

[1]State University of Novi Pazar, nn Vuka Karadzica Street, Novi Pazar, Serbia
e-mail: djordjekosanickg@gmail.com
[2]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjic Street, Kragujevac, Serbia
e-mail: fica@kg.ac.rs
**corresponding author*

Abstract:

Muscle fatigue is a multifactorial phenomenon involving complex central and peripheral mechanisms that alter neuromuscular performance and increase injury risk in professional athletes. The aim of this study was to develop and validate a multimodal framework based on artificial intelligence (AI) for the detection of compensatory muscle fatigue and biomechanical characteristics using surface electromyography (sEMG), clinical palpation, subjective fatigue perception, and a dedicated hardware–software system for knee biomechanics and cartilage loading analysis.

Twenty healthy male professional athletes underwent a standardized isometric upper-body fatigue protocol. sEMG signals were recorded bilaterally from selected trunk, shoulder, and lower limb muscles. Time-domain, frequency-domain, and nonlinear features were extracted and analyzed using Support Vector Machines (SVM), Random Forest (RF), and a hybrid Convolutional Neural Network–Long Short-Term Memory (CNN–LSTM) model.

Fatigue resulted in a significant decrease in median frequency (-18.7%, $p < 0.001$) and mean frequency (-15.9%, $p < 0.001$), accompanied by an increase in RMS amplitude ($+22.3\%$, $p < 0.001$). Nonlinear determinism significantly increased ($p < 0.01$), indicating altered motor unit synchronization. The CNN–LSTM model achieved the highest classification accuracy (91.4%, AUC = 0.95) and demonstrated strong correlation with clinical palpation scores ($r = 0.78$, $p < 0.001$). Compensatory activation patterns were observed in synergistic muscles, supporting non-local fatigue mechanisms. Biomechanical analysis revealed force distributions at the knee joint as well as cartilage deformation and stress patterns. This study presents a comprehensive AI-based framework for fatigue detection and biomechanical assessment, with strong translational potential in sports science, rehabilitation, and performance optimization.

Keywords: muscle fatigue, biomechanics, artificial intelligence, surface electromyography (sEMG), CNN–LSTM.

DEEP LEARNING FOR ALZHEIMER'S DISEASE CLASSIFICATION: A COMPARATIVE ANALYSIS OF CONVOLUTIONAL NEURAL NETWORK ARCHITECTURES

Nikola Stanković* [0009-0001-8484-5478], **Maja Vukasović** [0000-0003-0647-1922], **and Dražen Drašković** [0000-0003-2564-4526]

School of Electrical Engineering, University of Belgrade, 73 Bulevar kralja Aleksandra, Belgrade, Serbia
e-mail: nikolas@etf.bg.ac.rs, majav@etf.bg.ac.rs, drazen.draskovic@etf.bg.ac.rs
**corresponding author*

Abstract:

Alzheimer's disease is a neurodegenerative disorder for which early and accurate diagnosis is critical. In recent years, convolutional neural networks (CNNs) have demonstrated strong performance in medical image classification tasks, particularly when applied to magnetic resonance imaging (MRI) data. This paper presents a systematic analysis of CNN architectures and hyperparameter configurations for Alzheimer's disease classification. Five widely used architectures: AlexNet, VGGNet, ResNet, GoogLeNet, and MobileNet were evaluated across a range of hyperparameter settings, including learning rate, optimizer, dropout probability, and L2 regularization. To address class imbalance, data augmentation techniques were applied to the training dataset. A total of 720 models were trained and evaluated using standard classification metrics, including accuracy, precision, recall, and F1-score. The results show that ResNet and GoogLeNet achieve the highest classification performance, while MobileNet demonstrates a trade-off between efficiency and accuracy. Variability analysis indicates that smaller learning rates and the SGD optimizer yield more stable and accurate results. Statistical evaluation using one-way ANOVA confirms that the architecture and learning rate have the greatest impact on model performance. These findings highlight the importance of robust architectural design and hyperparameter selection when applying deep learning models to complex medical imaging data and provide practical guidelines for Alzheimer's disease classification tasks.

Keywords: Alzheimer's disease, convolutional neural networks, deep learning, hyperparameter optimization, image classification, medical image analysis, MRI.

LLM-BASED MASTER THESIS REPORT GENERATION

Matija Dodović* [0009-0005-2795-5709], **Janko Tufegdžić** [0009-0001-7439-7034], **Luka Hrvačević** [0009-0008-6456-9390] **and Marko Mišić** [0000-0002-7369-4010]

University of Belgrade, School of Electrical Engineering, 73 Bulevar kralja Aleksandra, Belgrade, Serbia
e-mail: matija.dodovic@etf.bg.ac.rs, janko.tufegdzic@etf.bg.ac.rs, luka.hrvacevic@etf.bg.ac.rs, marko.misic@etf.bg.ac.rs

Abstract:

Typical business processes in higher education institutions involve a substantial amount of administrative work performed by teaching staff. These tasks include course and examination paperwork, as well as various administrative reports related to internships, student theses, and similar activities. One representative example arises in master-level studies, where a student's thesis must be formally evaluated in written form by a committee prior to the oral defense. The preparation of such formal reports is mandatory, but largely repetitive and time-consuming. These documents typically follow a well-defined structure and rely on information already available across multiple sources, including the thesis manuscript, the student's biography, and administrative records (e.g. thesis title, approval date, and committee decisions). Despite their structured nature, these reports are usually prepared manually, leading to inefficiencies and potential inconsistencies. This paper proposes a hybrid pipeline for the automated generation of master thesis reports, combining deterministic programmatic processing with large language model (LLM)-based text generation. The core idea is to leverage rule-based extraction and template filling for structured components, while utilizing LLMs for tasks requiring natural language understanding and generation, such as summarization and stylistic adaptation. The system integrates multiple input sources, including structured parameters (e.g. candidate name, multilingual thesis title, approval date), semi-structured documents (e.g. student biography and thesis application forms), and unstructured text (the thesis itself). Report sections are classified according to the required generation strategy. Biographical data are directly reused from existing sources, while standard sections (e.g. preamble and conclusion) are generated using parameterized templates. Quantitative and factual elements - such as the number of pages, figures, tables, references, and chapter organization - are extracted programmatically to ensure accuracy. In contrast, descriptive sections, including summaries on study and research work, chapter overviews, and analyses of key results and contributions, are generated through controlled LLM-based summarization. A central design principle of the proposed system is the strict separation between fact extraction and language generation. All verifiable and quantitative information is obtained through deterministic methods to guarantee correctness, whereas LLMs are employed exclusively for producing coherent, contextually appropriate text. Particular attention is given to linguistic and stylistic aspects, including maintaining a formal academic tone and handling grammatical variations (e.g. name inflection and neutral phrasing). Prompt engineering techniques are used to guide the LLM toward generating outputs that conform to institutional standards, including predefined structure and stylistic constraints.
The main contribution of this paper is a practical and extensible framework for automating the generation of formal academic reports from heterogeneous data sources. Although the case study focuses on master thesis evaluation reports, the proposed approach is readily generalizable to other types of structured academic and administrative documents, demonstrating the broader applicability of hybrid systems that integrate programmatic processing with large language models.

Keywords: large language models, prompt engineering, automated document generation, information extraction, text summarization.

A ROBUST POSTPROCESSING GROUNDING FRAMEWORK FOR MONOCULAR 3D MOTION RECOVERY: BENCHMARKS ON MLB BASEBALL FOOTAGE

Valerijan Matvejev* [0009-0000-7125-2078] **and Dražen Drašković** [0000-0003-2564-4526]

University of Belrade, School of Electrical Engineering, 73 Bulevar kralja Aleksandra,
Belgrade, Serbia
e-mail: mv245055p@student.etf.bg.ac.rs, drazen.draskovic@etf.bg.ac.rs
**corresponding author*

Abstract:

Modern 3D motion capture pipelines require hundreds of thousands of dollars of investments in high-quality equipment in order to have accurate detection, tracking and temporal segmentation. Recent attempts to produce a successful monocular-to-3d human motion estimation and recovery system from a single smartphone camera, have found innovative ways to overcome their fundamental limitation - depth ambiguity. A cutting-edge machine learning model for 3D image segmentation, such as Meta's latest SAM3D body model, is the key component of every monocular-to-3d pipeline. These models' main drawback is the lack of temporal consistency, hence they require adequate postprocessing in order to acquire physically consistent data. This research paper focuses on benchmarking the most recent data optimization techniques for human 3D data grounding, while proposing a novel independent and robust postprocessing grounding algorithm. The data used in the research was gathered from high quality baseball RGB videos of top MLB teams (New York Yankees, Los Angeles Dodgers, Texas Rangers etc.) and the results of the research suggest complete elimination of ground penetration, foot floating and a significant reduction of other important anti-foot-skating metrics. The paper proceeds with listing limitations of the algorithm-level optimization and concrete suggestions for future improvements. It concludes by highlighting the necessary steps towards developing more custom machine learning models for different human body representation formats.

Keywords: machine learning, computer vision, 3D reconstruction, monocular 3D human pose estimation, data grounding.

GENERATING BDD SCENARIOS FROM SOURCE CODE USING LARGE LANGUAGE MODELS

Teodora Radaljac* [0009-0004-2970-6389], **Aleksa Srbljanović** [0000-0002-9937-0073], **Danko Miladinović*** [0009-0002-7365-4905] **and Dražen Drašković*** [0000-0003-2564-4526]

School of Electrical Engineering, University of Belgrade, Belgrade, Serbia
e-mail: teodora.radaljac@etf.bg.ac.rs, aleksa.srbljanovic@etf.bg.ac.rs, danko.miladinovic@etf.bg.ac.rs, drazen.draskovic@etf.bg.ac.rs
**corresponding author*

Abstract:

Behavior-Driven Development (BDD) is a software development methodology that enables improved collaboration between development teams, testers, and business analysts by defining system behavior in the form of readable scenarios written in the Given/When/Then format. These scenarios simultaneously serve as requirements specifications, system documentation, and a basis for automated tests, thereby establishing a link between technical implementation and business logic. Despite the advantages offered by the BDD approach, manually writing scenarios based on existing source code requires significant time and code familiarity, especially in large-scale projects. The emergence of Large Language Models (LLMs) opens up the possibility of automating this process, which has been confirmed in recent studies [1, 2, 3].

Within this research, an application was developed in the Python programming language that enables automatic forwarding of source code to multiple language models, requesting each model to generate corresponding BDD scenarios in Gherkin syntax based on code analysis. Parts of a web application developed in the Django framework were used as the source code. A set of representative code units of varying complexity was assembled. Five language models were employed, of which two are commercial general-purpose models accessed via public API services: GPT-4o and Claude 3.5 Sonnet, and three are open-source models specialized for working with code, run locally via the Ollama platform: Qwen2.5-Coder, DeepSeek-Coder-V2, and CodeLlama. This selection enabled a comparison of commercial general-purpose models with locally deployed code-specialized models, along with an analysis of the impact of model size and type on the quality of generated scenarios. For each code sample, generation was carried out in two separate passes. In the first pass, the entire prompt was provided in English, with the requirement that the Gherkin scenarios be generated in English. In the second pass, the same request was formulated entirely in Serbian, including localized Gherkin keywords, in order to examine the models' ability to generate scenarios in both languages.

The evaluation of generated scenarios was conducted across three dimensions. The first dimension concerns syntactic correctness, i.e., whether the generated scenarios are valid according to the Gherkin specification, which was verified through automatic parsing using the Gherkin-official library in Python. Scenarios that did not conform to Gherkin grammar were classified as syntactically incorrect, and for each model and language, the percentage of successfully parsed scenarios was calculated. The second dimension assesses whether the scenarios accurately describe the actual behavior of the program. The third dimension evaluates linguistic adequacy, i.e., the correctness of grammar, terminology, and naturalness of formulations, particularly for scenarios in Serbian where language models have more limited training resources. For scenarios in both languages, a consistency comparison was conducted, examining whether scenarios generated in Serbian and English by the same model for the same code describe identical behavior. Results are presented comparatively by model and by language, with quantitative analysis of differences.

This research provides a systematic insight into the potential of applying large language models for automatic generation of BDD scenarios from existing source code, with a particular contribution in the analysis of the multilingual context. The results indicate differences in the quality and completeness of generated scenarios depending on the model and language.

Keywords: behavior-driven development, large language models, test scenario generation, Gherkin, software testing.

COMPARATIVE ANALYSIS OF DIFFERENT LLM MODELS IN DETECTING SECURITY VULNERABILITIES IN SOURCE CODE

Aleksa Srbljanović* [0000-0002-9937-0073], **Danko Miladinović** [0009-0002-7365-4905], **Teodora Radaljac** [0009-0004-2970-6389] **and Dražen Drašković** [0000-0003-2564-4526]

School of Electrical Engineering, University of Belgrade, Belgrade, Serbia
e-mail: aleksa.srbljanovic@etf.bg.ac.rs, danko.miladinovic@etf.bg.ac.rs, teodora.radaljac@etf.bg.ac.rs, drazen.draskovic@etf.bg.ac.rs
**corresponding author*

Abstract:

Source code security represents one of the key challenges of modern software engineering. Vulnerabilities such as SQL injection, Cross-Site Scripting (XSS), OS command injection, path traversal, insecure deserialization, and hardcoded secrets are among the most common and most serious threats in software systems. The OWASP Top 10 classification encompasses a wide spectrum of vulnerability categories. Within this research, four categories from the OWASP Top 10 were selected: Broken Access Control (A01), Injection (A05), Authentication Failures (A07), and Software or Data Integrity Failures (A08), thereby covering six specific vulnerability types at the CWE classification level (CWE-89, CWE-79, CWE-78, CWE-22, CWE-502, CWE-798). Traditional approaches to detecting these vulnerabilities rely on Static Application Security Testing (SAST) tools that apply predefined rules and patterns for detecting known vulnerability categories. The emergence of Large Language Models (LLMs) opens up the possibility of automating this process, which has been confirmed in recent studies [1, 2, 3, 4].

Within this research, a controlled set of source code samples containing known security vulnerabilities was formed. As the basis for forming the dataset, the publicly available NIST SARD (Software Assurance Reference Dataset, https://samate.nist.gov/SARD/) was used, which contains over 450,000 test cases with documented weaknesses classified according to the CWE standard. From this database, samples for the Java and C programming languages were obtained and adapted, ensuring diversity in terms of syntactic and semantic code characteristics. For each sample, a ground truth classification was defined, encompassing the exact location of the vulnerability, its type according to the CWE and OWASP taxonomy, and a verified correct fix. Each sample was forwarded to multiple language models under identical conditions, with the same prompt and the same instructions. The prompt was constructed to require the model to identify vulnerabilities in the given code, classify the vulnerability type, and propose a concrete solution for its remediation.

The evaluation of model responses was conducted across three dimensions. The first dimension concerns the detection capability, i.e., whether the model identified the presence of a vulnerability in the code at all. The second dimension evaluates classification accuracy, i.e., whether the model correctly determined the vulnerability type according to OWASP categories. The third dimension analyzes whether the proposed solution actually eliminates the vulnerability without introducing new problems. The results were quantified through precision, recall, and F1 score metrics, grouped by vulnerability type and by model. As a baseline for comparison, the results of classical SAST tools applied to the same set of samples were used, enabling a direct comparison of language models with existing industry tools. The research provides a systematic and quantitative evaluation of Large Language Models as tools for security code review. The results enable the formulation of practical recommendations for integrating language models into existing code security analysis workflows, with a clear delineation of scenarios in which their application is justified.

Keywords: large language models, security vulnerabilities, static code analysis, OWASP, vulnerability detection, comparative analysis.

HIERARCHICAL FEATURE FUSION FOR IMPROVED OPTICAL CHARACTER RECOGNITION USING CNN-TRANSFORMER ARCHITECTURE

Veljko Petković [0009-0008-6082-3224]

Faculty of Sciences and Mathematics, University of Niš, 33 Višegradska Street, Niš, Serbia
e-mail: veljko.petko0022@gmail.com, veljko.petkovic@pmf.edu.rs

Abstract:

Scene text recognition remains a challenging problem in computer vision, particularly under conditions of visual degradation, irregular fonts, and complex backgrounds. In this paper, we propose a Hierarchical Feature Fusion CNN-Transformer (HFF-CNN-Transformer) architecture for optical character recognition (OCR) that improves upon a standard CNN-Transformer baseline by incorporating multi-scale convolutional feature representations. Unlike conventional approaches that pass only the final CNN feature map to the Transformer encoder, the proposed method fuses feature maps extracted from both shallow and deep convolutional layers together with the final CNN output, enabling the model to simultaneously capture fine-grained local patterns and high-level semantic information. Both models are trained using Connectionist Temporal Classification (CTC) loss on a synthetically generated dataset of 100,000 labeled text images, stratified into three augmentation levels: light (20%), moderate (60%), and heavy (20%), designed to simulate varying degrees of real-world visual difficulty. Experimental results demonstrate that the HFF-CNN-Transformer achieves a better validation Character Error Rate (CER) of 2.39% at epoch 13 and a test CER of 2.44%, outperforming the baseline CNN-Transformer, which reaches a validation CER of 2.60% at epoch 14 and a test CER of 2.62%. The HFF model further achieves a lower test loss of 0.0570 compared to 0.0622 for the baseline, confirming consistent improvement across all evaluation metrics. Total training times were 67.22 minutes and 62.22 minutes respectively, indicating a modest computational overhead introduced by the hierarchical fusion mechanism. These findings highlight the potential of hierarchical feature fusion as a lightweight yet effective enhancement to transformer-based OCR pipelines.

Keywords: optical character recognition, hierarchical feature fusion, convolutional neural networks, transformer encoder, character error rate, deep learning, computer vision.

ON THE ALGEBRAIC COMPLEXITY OF OPTIMAL POLYNOMIAL APPROXIMATION CONSTANTS

Filip Filipovic[1] [0009-0002-5321-5016], **Rémi Géraud-Stewart**[2] [0000-0001-8719-1724], **Aleksa Velickovic**[3] [0009-0001-1496-9977] **and David Naccache**[2,4*] [0000-0002-8651-6555]

[1]Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
[2]DI ENS, École Normale Supérieure, PSL, CNRS, Paris, France
[3]Ecole normale supérieure and Be-YS Research, Paris, France
[4]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjica Street, Kragujevac, Serbia
e-mail: filipfilipovic1999@gmail.com, david.naccache@gmail.com
**corresponding author*

Abstract:

We investigate the algebraic nature of constants arising from Chebyshev equiripple (minimax) polynomial approximation of Lp norms on [0, 1]. For the Euclidean case $\sqrt{1+t^2}$, we compute equiripple solutions from degree 1 through 8 and determine exact minimal polynomials and Galois groups for degrees 1 and 2 in both absolute and relative error formulations. A sharp phase transition occurs: the degree-1 constants are solvable by radicals (Galois groups C4, D4), while the degree-2 constants probably are not (Galois groups S12, $S_{10} \times C2$). We explain this by a structural dichotomy (decoupling versus coupling of critical points) and extend to L_3 norms, where the minimal polynomial degree jumps to 246. A general impossibility result follows from Hilbert's irreducibility theorem. These results appear to establish the first connection between Chebyshev approximation theory and the non-solvability of algebraic equations. This provides a short-cut formula which is extremely useful in graphic and AI calculations, easy to implement in hardware.

Keywords: algebra, complexity calculation, Chebyshev approximation, equiripple, Galois theory, Lp norms, polynomial approximation.

STABILIZING SERVICE OPERATIONS THROUGH APPLIED ARTIFICIAL INTELLIGENCE UNDER WORKFORCE PRESSURE: EVIDENCE FROM HOSPITALITY ORGANIZATIONS

Marko Petrović [0009-0004-6192-1919]

Faculty of Project Management, Satbayec University, Almaty, Kazakhstan
e-mail: djaniluter@yahoo.com

Abstract:

Service industries are increasingly integrating artificial intelligence (AI) technologies within labor-intensive environments characterized by operational uncertainty and workforce volatility. While AI is often framed as a workforce replacement mechanism, its acceptance within hospitality organizations depends largely on employees' perceptions.

Drawing on survey data from hospitality employees (N = 122) and qualitative insights from general managers, this study examines the determinants of employees' acceptance of AI systems. Regression analysis indicates that AI Reliability (B = 0.434, $p < .001$) and Perceived Safety (B = 0.364, $p < .001$) significantly enhance acceptance, while Fear and Anxiety negatively affect supportive attitudes (B = -0.261, $p = .003$). Interestingly, perceived Job Security Threat is positively associated with acceptance (B = 0.408, $p < .001$). The model explains 59.2% of the variance in acceptance ($R^2 = 0.592$).

The findings suggest that AI acceptance in hospitality contexts is driven less by substitution dynamics and more by perceptions of reliability, safety, and psychological reassurance. By reframing AI as a performance-supporting and stabilizing infrastructure, organizations may foster sustainable technology adoption in service environments undergoing digital transformation.

Keywords: artificial intelligence (AI), employee acceptance, AI reliability, perceived safety, psychological factors, hospitality management.

THE ROLE OF ARTIFICIAL INTELLIGENCE AND DIGITAL MARKETING IN SHAPING GREEN INVESTMENT TRENDS AND ORGANIZATIONAL COMPETITIVENESS: EVIDENCE FROM SERBIA AND THE EUROPEAN UNION

Bojan S. Jokić[1] [0009-0004-8324-981X], **Bojana S. Jokić**[1] [0009-0000-8493-5118], **Jelena R. Šaranović**[1] [0009-0004-3638-4226] **and Sonja D. Radenković**[2*] [0000-0001-6830-7533]

[1]Academy of Vocational Studies Kosovo and Metohija, nn Dositeja Obradovica, Leposavic, Serbia
e-mail: bojan.jokic@akademijakm.edu.rs, bojana.jokic@akademijakm.edu.rs, jelena.saranovic@akademijakm.edu.rs
[2]Faculty of Banking, Insurance and Finance, 12 Zmaj Jovina Street, Belgrade Banking Academy Belgrade, Serbia
e-mail: sonja.radenkovic@bba.edu.rs
**corresponding author*

Abstract:

Artificial intelligence and digital marketing are becoming increasingly important factors in modern business, mainly in shaping green investment trends and improving organizational competitiveness. This paper analyzes the role of artificial intelligence and digital marketing in encouraging sustainable investment activities and strengthening the competitive position of organizations, with special reference to the experiences of Serbia and the European Union. It is assumed that integration of intelligent digital technologies and modern marketing approaches can significantly contribute to more efficient directing of investments towards sustainable projects, better positioning of organizations on the market and more successful adaptation to the requirements of the green transition. The methodological approach is based on the analysis of relevant scientific and professional literature, as well as on a comparative overview of development tendencies in Serbia and the European Union. Special focus is put on the application of artificial intelligence in the process of data analysis, prediction of investment patterns, automation of business activities and decision support, as well as the role of digital marketing in the promotion of sustainable products, strengthening relations with consumers and improving the market visibility of organizations. The results indicate that the European Union has a more developed institutional and technological framework for connecting artificial intelligence, digital marketing and green investments, while Serbia shows positive progress, but still faces limitations in terms of financial support, technological equipment and organizational readiness. The conclusion is that the synergy of artificial intelligence and digital marketing can be an important driver of sustainable competitiveness and green transformation of organizations.

Keywords: artificial intelligence, digital marketing, green investments, organizational competitiveness, sustainable development, Serbia, European Union.

AI-SUPPORTED ANALYSIS OF ORGANIZATIONAL, INVESTMENT AND DIGITAL MARKETING DETERMINANTS OF SUSTAINABLE COMPETITIVENESS IN SERBIA AND THE EUROPEAN UNION

Bojana S. Jokić[1] [0009-0000-8493-5118], **Bojan S. Jokić**[1] [0009-0004-8324-981X], **Jelena R. Šaranović**[1] [0009-0004-3638-4226] **and Sonja D. Radenković**[2*] [0000-0001-6830-7533]

[1]Academy of Vocational Studies Kosovo and Metohija, nn Dositeja Obradovica, Leposavic, Serbia
e-mail: bojan.jokic@akademijakm.edu.rs, bojana.jokic@akademijakm.edu.rs, jelena.saranovic@akademijakm.edu.rs
[2]Faculty of Banking, Insurance and Finance, 12 Zmaj Jovina Street, Belgrade Banking Academy, Belgrade, Serbia
e-mail: sonja.radenkovic@bba.edu.rs
**corresponding author*

Abstract:

The sustainable competitiveness of organizations in modern business environment increasingly depends on the ability to connect organizational capacities, investment activities as well as digital marketing strategies enhanced by opportunities provided by artificial intelligence. This paper aims to analyze organizational, investment and digital-marketing determinants of sustainable competitiveness in Serbia and the European Union, applying an approach based on artificial intelligence. The goal is to examine how modern AI tools can contribute to identifying key factors that influence the strengthening of competitive organizations' positions in the conditions of the green and digital transition. The paper is based on the analysis of relevant scientific and professional literature, as well as on a comparative overview of development trends in both Serbia and the European Union. Special attention is paid to the role of artificial intelligence in processing large amounts of data, recognizing patterns, improving the decision-making process and evaluating the effects of investment and marketing activities on sustainable competitiveness. In addition, organizational adaptability, investments in sustainable and digital projects, as well as the application of digital marketing in building market recognition and long-term value for consumers are considered. The results indicate that the European Union has more a developed institutional, investment and technological framework for the application of artificial intelligence in the function of sustainable competitiveness, while Serbia shows positive progress, but still faces certain limitations. The conclusion is that AI-supported analysis can significantly contribute to a better understanding and the improvement of sustainable competitiveness factors in the national and European contexts.

Keywords: artificial intelligence, sustainable competitiveness, organizational factors, investments, digital marketing, Serbia, European Union.

SYNERGY BETWEEN CORPORATE ECONOMIC ANALYSIS AND BANKING MANAGEMENT IN THE CREDIT MANAGEMENT PROCESS

Kojić Nenad[1*] [0000-0001-5399-8132], **Mrvaljević Mirjana**[1] [0000-0002-0124-0423] **and Ujkanović Enis**[2] [0009-0005-4302-4402]

[1]Kosovo and Metohija Academy of Applies Studies, Peć Department in Leposavić, Serbia
e-mail: nenad.kojic@akademijakm.edu.rs, mirjana.mrvaljevic@akademijakm.edu.rs
[2]State University of Novi Pazar, Novi Pazar, Serbia
e-mail: e.ujkanovic@uninp.edu.rs
**corresponding author*

Abstract:

Modern banking management operates within an environment of significant uncertainty, where a bank's ability to identify high-quality borrowers serves as its primary competitive edge. This paper investigates the critical role of corporate economic analysis as the foundation upon which banking management builds its lending decisions. The core of the research lies in the transformation of accounting data from financial statements into strategic intelligence vital for banking operations. The first section focuses on quantitative analysis, utilizing liquidity, solvency, and profitability ratios to assess the "financial health" of a company. Special attention is given to Cash Flow analysis, which bank managers view as a more realistic indicator of debt-service capacity than net profit alone. The second part explores the decision-making process, examining how analytical findings are used to assign credit ratings, determine interest rates, and establish collateral requirements. The objective of this study is to demonstrate that integrating modern analytical methods - such as the Altman Z-score model and qualitative market position assessments - empowers banking management to act proactively and prevent the accumulation of non-performing loans (NPLs). In conclusion, the paper emphasizes that a robust link between economic analysis and bank management is a vital factor for financial system stability and broader economic development.

Keywords: banking management, economic analysis, financial statements, credit rating, risk management, balance sheets.

FROM PREDICTIVE TO PRESCRIPTIVE AI: VALUE OPTIMIZATION IN COMPLEX DECISION SYSTEMS

Ana M. Spasojević[1*] [0000-0003-1373-7370] **and Dragan M. Vukolić**[2] [0000-0002-6364-9849]

[1]Faculty of Economics, University of Kragujevac, 1A Liceja Kneževine Srbije Street, Kragujevac, Serbia
e-mail: ana.spasojevic1985@gmail.com
[2]Faculty of Tourism and Hotel Management, University of Business Studies, Jovana Dučića 23a, Banja Luka, Bosnia and Herzegovina
e-mail: vukolicd@yahoo.com
**corresponding author*

Abstract:

Artificial intelligence (AI) has become an important tool for data analysis and prediction across domains such as economics, healthcare, and complex organizational systems. However, most applications remain focused on forecasting what might happen, without directly supporting decision-making, which creates a gap between insight and action. This paper examines the transition from predictive to prescriptive AI, emphasizing its role in optimizing value within complex decision systems. Unlike traditional models, prescriptive AI combines prediction with optimization to recommend effective actions under uncertainty and multiple competing objectives, shifting the focus from analysis to decision support. From an economic perspective, this approach enhances value creation through better resource allocation, increased efficiency, and more informed strategic decisions. The findings indicate a statistically significant impact of prescriptive AI on improving decision quality and organizational performance, including cost reduction and overall efficiency gains. The study also highlights challenges such as model interpretability, data quality, computational complexity, and ethical concerns. Despite these limitations, the transition toward prescriptive AI represents a significant step forward, enabling a shift from information generation to value-driven decision-making in complex systems.

Keywords: prescriptive AI, predictive analytics, decision-making systems, value optimization, artificial intelligence.

AN INFORMER-BASED DEEP LEARNING MODEL FOR STOCK MARKET INDEX PREDICTION

Vesna M. Ranković[1*] [0000-0002-5445-9971], **Ljiljana N. Matić**[2] [0009-0002-0473-0214] **and Zoran S. Kalinić**[2] [0000-0001-8137-9005]

[1]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjić Street, Kragujevac, Serbia
e-mail: vesnar@kg.ac.rs
[2]Faculty of Economics, University of Kragujevac, 3 Liceja Kneževine Srbije Street, Kragujevac, Serbia
e-mail: ljiljana.matic@ef.kg.ac.rs, zkalinic@kg.ac.rs
**corresponding author*

Abstract:

Artificial intelligence has become an integral component of the financial sector, offering robust frameworks for data analysis and informed decision-making. Modeling financial time series, including stock prices and market indices, is essential for optimizing trading strategies and mitigating risks. However, these series are inherently characterized by high volatility, nonlinearity, and non-stationarity, which present significant challenges for accurate prediction. The stock market index remains a central concern in finance due to the combined effects of macroeconomic indicators, monetary policies, and global economic trends. In this study, the Informer model was employed for stock market index prediction due to its ability to capture long-range dependencies in long-sequence time series. For demonstration, S&P500 (Standard & Poor's 500) index was chosen as one of the most famous and widely followed stock market indices in the world. It tracks the performance of 500 leading publicly traded companies listed on the US stock exchanges. The model was trained and evaluated using daily S&P 500 data from 2015 to 2026. Compared to conventional Transformer architectures, the Informer reduces computational and memory complexity in the attention mechanism while efficiently modeling non-stationary and nonlinear dynamics. The predictive performance of the Informer model depends on the systematic tuning of its hyperparameters, which directly affect its capacity to learn long-term patterns and generalize across varying market conditions. Key hyperparameters include input sequence length, embedding dimension, the number of encoder and decoder layers, and the number of attention heads. In this paper, five hyperparameters were systematically optimized to maximize model performance. Evaluation was performed using standard metrics, including Mean Absolute Error (MAE), Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and Directional Accuracy (DA). The model achieved an MAE of 23.24 and 37.62, an RMSE of 36.54 and 55.36, and a MAPE of 0.76% and 0.67% on the training and test sets, respectively. Directional Accuracy (DA) reached 64.83% on the training set and 73.44% on the test set, confirming high predictive precision and robust generalization across periods of varying market volatility.

Keywords: deep learning, financial time series, stock market index, Informer model.

Acknowledgement: Authors acknowledge the funding by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract numbers[451-03-34/2026-03/200107 (Faculty of Engineering, University of Kragujevac)]and [451-03-33/2026-03/200099(Faculty of Economics, University of Kragujevac)].

RECOVERING DEMOGRAPHICS FROM SURVEY RESPONSES: AN LLM-BASED APPROACH

Aleksandar Kartelj[1*] [0000-0001-9839-6039], **Miljana Mladenović**[2] [0000-0002-7973-452X] **and Staša Vujičić Stanković**[1] [0000-0002-7200-3724]

[1]Faculty of Mathematics, University of Belgrade, Belgrade, Serbia
e-mail: aleksandar.kartelj@gmail.com, stasa.vujicic.stankovic@matf.bg.ac.rs
[2]Faculty of Education, University of Nis, Vranje, Serbia
e-mail: ml.miljana@gmail.com
**corresponding author*

Abstract:

We present the first empirical evaluation of large language model (LLM) persona simulation on a Serbian survey, addressing a gap in the non-English and cross-cultural evaluation of such methods. Using a media-habits survey of 1,000 Serbian respondents, we recast the standard forward task (generating opinions given a demographic profile) as an inverse recovery problem: given an LLM-written summary of a respondent's non-demographic answers, how well can models recover that respondent's seven categorical demographic attributes? We benchmark four OpenAI models (GPT-4.1-mini, GPT-4o-mini, GPT-5-mini, and GPT-5.4-mini), in both zero-shot and retrieval-augmented (top k-NN) configurations, against three classical baselines (majority-class, k-NN, and XGBoost). Our results show that (i) XGBoost remains the strongest method overall (0.645 mean accuracy, 0.499 macro-F1); (ii) the best k-NN-augmented LLM is statistically indistinguishable from XGBoost on six of seven targets under a McNemar test; (iii) zero-shot LLM personas underperform a plain k-NN baseline, indicating that the predictive lift of retrieval-augmented prompting comes from retrieval itself rather than from the LLM; (iv) additional reasoning effort and chain-of-thought prompting have no measurable effect; and (v) raw mean accuracy systematically overstates performance on skewed targets, so macro-F1 and a per-target uniformity-deviation analysis are essential. Our findings support the critical line of literature that cautions against treating LLM personas as sources of independent inference in survey simulation, and extend this argument from English to Serbian language contexts.

Keywords: LLM persona simulation, silicon samples, retrieval-augmented prompting, k-nearest neighbors, XGBoost, Serbian survey data, demographic recovery, non-English evaluation.

CONVERSATIONAL AI AGENT FOR INTELLIGENT SEARCH OVER OPEN ACCESS REPOSITORIES OF RESEARCH OUTPUTS

Boško Laković[1*] [0009-0003-0056-8181], **Radovan Drašković**[2] [0009-0007-4607-3239], **Miloš Ivanović**[1,2] [0000-0002-8974-2267] **and Boban Stojanović**[1,2] [0000-0003-2901-0061]

[1]Vodéna, 19b/5 Kralja Milana IV Street, 34000 Kragujevac, Serbia
e-mail: bosko.lakovic@vodena.rs
[2]Faculty of Science, University of Kragujevac, 12 Radoja Domanovića Street, 34000 Kragujevac, Serbia
e-mail: radovan.draskovic@pmf.kg.ac.rs, milos.ivanovic@pmf.kg.ac.rs, boban.stojanovic@pmf.kg.ac.rs
**corresponding author*

Abstract:

Accessing structured academic data typically requires manual database queries or navigation through multiple systems, which is time-consuming, error-prone, and often inaccessible to non-technical users. This paper presents SCIDAR Assistant, an intelligent conversational agent designed for natural language querying of scientific publications and researcher profiles on any Dublin-Core metadata format compatible repository. The system enables users to pose complex queries, which may involve various filters and constraints, without needing technical expertise, even if these filters and constraints are not strictly defined. To support reliable question answering over structured academic data, we propose an agent-based architecture implemented in the N8N platform. A Large Language Model (LLM) orchestrates a set of specialized tools to retrieve data from the institution's structured research database. The system obtains the information required for answer generation by leveraging fuzzy matching in combination with deterministic unique identifier lookups, ensuring precise linking between researchers, publications and other related entities. Experimental evaluation of SCIDAR Assistant demonstrates that the system provides consistent and accurate results for a variety of user queries. These findings highlight the importance of a chosen LLM, explicit identifier management, and well-defined tool interfaces for enabling robust multi-step reasoning in LLM-based systems for academic information access.

Keywords: conversational agent, research publications, large language models (LLMs), agent-Based architecture, N8N.

Acknowledgement: This paper is supported through the EIT's Higher Education Initiative A-SIDE projects, coordinated by EIT RawMaterials, funded by the European Union and supported by the projects No. 451-03-34/2026-03/200107, and 451-03-34/2026-03/200122, Ministry of Science, Technological Development and Innovation, Republic of Serbia.

A LOCAL RETRIEVAL AUGMENTED GENERATION APPROACH FOR INTERNAL MEDICINE QUESTION ANSWERING

Andjela Blagojević[1,2*] [0000-0002-8652-3827], **Tijana Geroski**[1,2] [0000-0003-1417-0521], **Nenad Filipović**[1,2] [0000-0001-9964-5615] **and Vesna Ranković**[1] [0000-0002-5445-9971]

[1]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
e-mail: andjela.blagojevic@kg.ac.rs, tijanas@kg.ac.rs, fica@kg.ac.rs, vesnar@kg.ac.rs
[2] Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
e-mail: andjela.blagojevic@kg.ac.rs, tijanas@kg.ac.rs, fica@kg.ac.rs
**corresponding author*

Abstract:

Large language models (LLMs) have shown impressive capabilities in natural language understanding and generation, but their application in specialized medical domains remains challenging due to limited domain knowledge, risk of hallucinations, and the scarcity of high-quality resources in non-English languages.
This paper presents a methodology for developing a local, retrieval-augmented generation medical chatbot using Mistral 7B for question answering in the field of internal medicine. The system uses internal medicine textbook in Croatian language as its knowledge source. The cleaned text is split into overlapping chunks, which are then embedded using the multilingual E5 model from Sentence Transformers, capturing semantic similarity across languages. The Mistral 7B Instruct model generates responses conditioned on the retrieved knowledge chunks, improving contextual relevance.
The system was evaluated on a manually created set of medical QA pairs. Results show high semantic similarity (average cosine similarity 0.9566, BERT score 0.8679), while BLEU and ROUGE-L scores are lower due to sensitivity to exact word matches. The model also mitigates hallucinations by providing a truthful fallback message when relevant information is unavailable.
Overall, the proposed Retrieval Augmented Generation (RAG) system demonstrates effective retrieval and generation capabilities in the Croatian language, providing reliable and semantically accurate responses in the domain of internal medicine.

Keywords: natural language processing, retrieval augmented generation, large language model, internal medicine, medical chatbot.

Acknowledgement: This research is supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract number [451-03-33/2026-03/200107 (Faculty of Engineering, University of Kragujevac)].

HYBRID CODING ASSISTANCE: EVALUATING LOCAL-PLUS-PAID LLM WORKFLOWS AGAINST PAID-ONLY CODING

Danko Miladinović* [0009-0002-7365-4905], **Teodora Radaljac** [0009-0004-2970-6389], **Aleksa Srbljanović** [0000-0002-9937-0073] and **Jovan Đukić** [0009-0000-8047-1379]

The Department of Computer Science and Information Tehnology, University of Belgrade, 73 Bulevar kralja Aleksandra, Belgrade, Serbia

e-mail: danko@etf.rs, teodora@etf.bg.ac.rs, aleksas@etf.bg.ac.rs, dj@etf.bg.ac.rs

**corresponding author*

Abstract:

The rapid adoption of large language models (LLMs) for software development has created an important practical choice for both individual developers and organizations, and that is whether to rely entirely on a paid frontier model for coding tasks or to combine a local model with a paid model in a hybrid workflow. While paid models are often associated with stronger reasoning, broader contextual understanding, and better performance on difficult software engineering tasks, local models offer advantages in privacy, controllability, offline availability, and potentially lower long-term operational cost. This paper examines whether integrating a local coding model with a paid LLM provides measurable benefits over using only a paid model for coding-related work. Rather than treating model choice as a purely technical comparison of raw benchmark performance, this study frames the problem as a workflow design issue centered on the allocation of coding subtasks between local and paid models.

The central hypothesis of this paper is that the main benefit of a hybrid local-plus-paid setup is not necessarily superior peak intelligence, but more efficient task allocation. In such a workflow, repetitive, lower-complexity, or privacy-sensitive tasks can be handled locally, while higher-complexity tasks that require deeper reasoning, broader repository understanding, or more reliable synthesis can be delegated to the paid model. This separation may allow developers to preserve much of the quality advantage of frontier paid systems while reducing external token usage, lowering the amount of proprietary code sent to remote services, and improving responsiveness for frequent everyday interactions. By contrast, a paid-only workflow offers conceptual simplicity and may achieve stronger results on difficult end-to-end tasks, but it may also introduce higher recurring cost and greater dependence on remote infrastructure.

To examine this tradeoff, the paper compares two coding-assistance workflows: a paid-only workflow, in which all coding requests are processed by a paid LLM, and a hybrid workflow, in which a local model performs first-pass assistance for selected subtasks and escalates more challenging requests to the paid model. The comparison is designed around realistic software development activities rather than isolated code generation prompts. These activities include code completion, test generation, small bug fixing, refactoring, repository summarization, and multi-file feature implementation. The evaluation considers several dimensions of performance, including task success, code correctness, latency, and cost. This broader framing is important because developer experience is shaped not only by final code quality, but also by iteration speed, reliability, and the practical constraints of handling sensitive codebases.

Overall, this paper argues that the most meaningful comparison is not local versus paid models in isolation, but hybrid versus single-model workflows in real coding practice. By analyzing the tradeoffs among quality, cost, and usability, the study aims to provide a more practical framework for understanding how developers should integrate LLMs into software engineering. The findings are intended to inform both academic discussion and real-world deployment decisions, particularly for developers and teams choosing between convenience, capability, and control in AI-assisted programming.

Keywords: large language models, hybrid AI systems, AI for software engineering, local deployment, cloud-based LLMs, code completion, software development workflows, cost efficiency, privacy, model orchestration.

NEURAL ALGORITHMIC REASONING FOR ZERO-SUM GAME THEORY: MASTERING MINIMAX ADVERSARIAL SEARCH

Marko Njegomir* [0000-0003-2939-902X]

Faculty of Technical Sciences, University of Novi Sad, 6 Trg Dositeja Obradovića, Novi Sad, Serbia
e-mail: marko.njegomir@uns.ac.rs
**corresponding author*

Abstract:

Neural Algorithmic Reasoning (NAR) has emerged as a powerful paradigm for training neural networks to execute classical algorithms, aiming to achieve robust out-of-distribution generalization. While existing research predominantly focuses on polynomial-time problems, extending NAR into computational game theory presents profound structural challenges. Specifically, adversarial search algorithms like Minimax, which is the foundational decision rule for two-player zero-sum games, require strict recursive reasoning across deep graph structures. This introduces an inherent depth-parity effect, where the algorithm's objective dynamically alternates between maximization and minimization at each tree depth.
In this paper, we systematically evaluate the capacity of Graph Neural Networks (GNNs) to execute the Minimax algorithm within the CLRS-30 algorithmic reasoning framework. Recognizing that standard algorithmic processors struggle to adapt to shifting depth parities out-of-distribution, we introduce a Parity-Aware processor architecture designed to structurally align with the bipartite nature of turn-based games. Furthermore, we explore the integration of explicit pointer hints and step-by-step Depth-First Search (DFS) execution trajectories into our models to investigate their potential for guiding the neural network's latent representations. Our framework establishes a baseline for embedding discrete adversarial game theory algorithms into continuous latent spaces, offering new insights into structural generalization in neural reasoning.

Keywords: neural algorithmic reasoning, game theory, minimax, graph neural networks, zero-sum games.

LLM-BASED OPTIMIZATION OF X-RAY SPECTRAL FILTERING FOR IMPROVED CT DOSE EFFICIENCY

Jasna Radulović[1] [0000-0001-8777-8031], **Nikola Mijailović**[2] [0000-0002-2665-9933], **Marijana Gavrilović Božović**[1] [0000-0002-0144-0767], **Milan Čabarkapa**[1] [0000-0002-2094-9649], **Ivan Krstić**[1] [0000-0001-7583-3152], **and Mina Vasković Jovanović**[1] [0000-0003-3644-8723]

[1]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjic Street, Kragujevac, Serbia
e-mail: jasna@kg.ac.rs, marijana.gavrilovic@kg.ac.rs, mcabarkapa@kg.ac.rs, ivan.krstic@kg.ac.rs, mina.vaskovic.jovanovic@kg.ac.rs
[2]HTEC group, Belgrade, Serbia
e-mail: nmijailovic3@gmail.com
**corresponding author*

Abstract:

Designing CT imaging systems involves achieving an optimal balance between high-quality image reconstruction and minimizing patient exposure to ionizing radiation. These two objectives are often at odds, as lowering the radiation dose can negatively impact image clarity and diagnostic reliability. However, the energy-dependent interaction of X-ray photons with different biological tissues provides an opportunity to tailor scanning strategies in a way that preserves image quality while reducing dose.
In this study, we propose an enhanced framework that incorporates large language models (LLMs) as a decision-support tool for optimizing filter parameters in CT imaging. By adjusting filter characteristics, such as material composition and thickness, it is possible to reshape the emitted X-ray spectrum and improve contrast between specific tissue types, particularly soft tissue and cortical bone. X-ray spectra are generated through simulations performed using SpekCalc.
The novelty of this approach lies in the integration of LLMs to guide the selection of optimal filter configurations based on multiple criteria, including dose efficiency and image contrast. By analyzing simulation outputs and leveraging learned patterns, the LLM can suggest parameter adjustments that align with desired imaging outcomes.
The analysis focuses on two representative tissue categories: soft tissue and cortical bone. The findings indicate that combining physics-based simulations with LLM-assisted optimization can streamline the parameter tuning process and contribute to the development of adaptive, tissue-specific scanning protocols. This approach has the potential to reduce overall radiation exposure while maintaining clinically relevant image quality.

Keywords: CT tomography, X-ray dose radiation, LLM optimization.

THE DEEP-MICRO-CORE PROJECT: MACHINE LEARNING ALGORITHMS FOR THE IDENTIFICATION OF THE CORE MICROBIOME FROM INTEGRATED DATA ON MULTIPLE SPECIES

Mladen Seničar [0009-0007-6321-2501]

AI Institute of Serbia, Fruškogorska 1, Novi Sad, Serbia

e-mail: senicar93@gmail.com

Abstract:

The DeepMicroCore project aims to leverage machine learning (ML), in particular explainable AI (Artificial Intelligence), to identify the core microbiome across substrates and animal species. On a large integrated dataset, ML/AI models are developed to accurately predict substrates and species of origin of any given microbiome: variable importance techniques then reveal the relevant microbial species for prediction, which are supposed to constitute the core microbiome. These results are benchmarked against current approaches for the identification of core microbiomes.
Our integrated dataset currently include ~1000 samples from the milk, rumen and gut microbiomes of cows, pigs, goats. Data were obtained from ten 16S rRNA-gene sequencing research projects and were integrated using a custom bioinformatic pipeline based on a modified Nextflow nf-core/ampliseq analysis workflow. Initial predictive models have been prototyped on a reduced dataset of cow milk, gut and rumen microbiomes, using Lasso-penalised logistic regression (LLR), Random Forest (RF), Extreme Gradient Boosting (XGB) and convolutional neural networks (1d-CNN). Two different data representations were used: i) filtered and normalised count (ASV/OTU) tables, for the LLR, RF and XGB models; ii) direct sequence data, either independent or joined R1/R2 reads, for 1d-CNN. A validation split approach was used to evaluate the performance of predictive models: 80% of the data were used for training, 20% for testing. Within the training set, a 5-fold cross-validation scheme was used to tune the hyperparameters: the degree of penalization in LLR; number and depth of trees, number of resampled variables, and size of terminal nodes for RF and XGB -plus shrinkage parameter for XGB only; network architecture, learning rate, batch size and number of epochs for 1d-CNN. Results showed high accuracy of predictions from all multi-class classification models, with very few misclassifications (average accuracy = 0.984; average Cohen's kappa = 0.97; average Matthew's correlation coefficient = 0.973). Such high accuracy justifies the use of important variables (microbial taxa) to identify the core microbiome of the given substrate and/or animal species. Next steps include the application of predictive models to the larger integrated datasets and the comparison of results with current approaches to the definition of core microbiomes.

Keywords: machine learning, microbiome, explainable AI, random forest, extreme gradient boosting, convolutional neural networks, multi-class classification.

THREE-DOMAIN EEG FEATURE SELECTION FOR DISTINGUISHING NEUTRAL AND INCONGRUENT STROOP STATES

Tamara Parojčić[1,2] [0009-0009-6628-7976], **Milica M. Janković**[2] [0000-0002-7506-4995] **and Marija Novičić**[2] [0000-0001-9713-810X]

[1]Vlatacom Institute, Belgrade, Serbia
[2]School of Electrical Engineering, University of Belgrade, Belgrade, Serbia
e-mail: pt255049p@student.etf.bg.ac.rs, piperski@etf.rs, no[vicic@etf.bg.ac.rs

Abstract:

The Stroop task is a widely used method for assessing cognitive state. It involves showing participants color words printed in either matching (neutral state) or conflicting (incongruent state) ink colors, requiring them to name the ink color while suppressing the automatic response of reading the word. In the incongruent state, higher reaction times and error rates, known as the Stroop effect, indicate the increased cognitive demand needed to process and resolve the task conflict. Various brain signals recorded during the Stroop task have been utilized to assess cognitive states and diagnose neuropsychological disorders. This study investigates three-domain feature selection for the electroencephalogram (EEG) signals obtained from a publicly available database (21 participants), aiming to identify the most informative features for distinguishing between neutral and incongruent states. Before feature extraction, EEG signals were bandpass filtered (0.5-45 Hz), and electrooculogram (EOG) artifacts were identified and removed using Independent Component Analysis (ICA). The cleaned signals were segmented into blocks corresponding to neutral and incongruent states. From these segments, different features were extracted for each subject across three domains: time, frequency, and wavelet, resulting in a total of 406 features per subject. The Support Vector Machine-Based Recursive Feature Elimination (SVM-RFE) method was used to identify the most important EEG features for distinguishing between neutral and incongruent states. The most informative features in the time domain were amplitude variance of FC3, CPZ, C4, F3 channels, and mean amplitude of FP1. In the frequency domain, the five most informative features were the maximum of the Fourier coefficients of PZ, FP2, FZ, CPY, and T7 channels. The five most informative features in the wavelet domain were the wavelet coefficient ratios of CPZ, FCZ, F3, FC3, and F4 channels. This study identifies a heterogeneous set of informative EEG features across the time, frequency, and wavelet domains, highlighting the prominent role of the frontal, central and parietal regions in distinguishing cognitive states during the Stroop task. The selected features and regions demonstrate significant potential for EEG-based classification, suggesting their utility as robust biomarkers for distinguishing cognitive states.

Keywords: electroencephalogram, Stroop task, features, time domain, frequency domain, wavelet domain, feature selection, support vector machine, recursive feature elimination.

Acknowledgment: This work was financially supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia under contract number: 451-03-34/2026-03/200103. The research was partially conducted in the premises of the Palace of Science, Miodrag Kostić Endowment.

PHYSICS-INFORMED NEURAL NETWORKS FOR VANCOMYCIN PHARMACOKINETIC PARAMETER ESTIMATION FROM SPARSE CONCENTRATION DATA

Nikolija Cuckić [0009-0002-2771-5631]

Faculty of Sciences and Mathematics, University of Nis, 33 Višegradska Street, Niš, Serbia
e-mail: nikolija.cuckic@pmf.edu.rs

Abstract:

Accurate estimation of pharmacokinetic (PK) parameters from sparse clinical measurements is a central challenge in Therapeutic Drug Monitoring (TDM). Vancomycin, a glycopeptide antibiotic widely used against methicillin-resistant Staphylococcus aureus (MRSA), exhibits a narrow therapeutic window and substantial inter-individual variability, making individualized dosing critical for both efficacy and safety. In clinical practice, only 3–8 blood samples per patient are typically available, which renders classical nonlinear least-squares fitting unreliable for complex pharmacokinetic models.
We propose a Physics-Informed Neural Network (PINN) approach to the inverse problem of estimating individual PK parameters from sparse, noisy concentration–time profiles. By embedding the governing ordinary differential equations directly into the neural network loss function, the physics of drug elimination serves as a regularizer that constrains the solution space even when data are scarce. We evaluate the method against a classical least-squares benchmark across a systematic grid of measurement counts ($N \in \{3, 5, 8, 10, 12\}$) and noise levels ($\sigma \in \{0\%, 5\%, 10\%, 20\%\}$) using synthetic data generated from population PK parameters reported in the literature. For the two-compartment model, physics-informed regularization yields substantially more reliable parameter recovery under clinically realistic sparse and noisy conditions (3–5 samples, 5–10% noise), whereas for the one-compartment model PINN and classical fitting perform similarly. We additionally assess robustness to initial parameter values through a sensitivity experiment with deliberately misspecified initialization, demonstrating that the advantage of physics-informed regularization at minimal sampling density is preserved even when PINN parameters are initialized away from the population mean.

Keywords: physics-informed neural networks, pharmacokinetics, vancomycin, parameter estimation, therapeutic drug monitoring, inverse problem, sparse data.

HEPATOTOXICITY CLASSIFICATION UPON DRUG CHEMICAL COMPOSITION AND MOLECULAR STRUCTURE – DEEP LEARNING APPROACH

Milica Nikolić[1,2,3*] [0000-0002-8537-7944], **Miroslava Jordović Pavlović**[2] [0000-0003-4038-6537], **TijanaGeroski**[1,4,5] [0000-0003-1417-0521], **Thomas Lautenschlager**[6] [0009-0002-3214-9363], **Ralf Mikut**[6] [0000-0001-9100-5496], **Mohammad Ghasemzadeh Hasankolaei**[7] [0000-0003-2053-5566], **AmirGhaemmaghami**[7] [0000-0003-3160-8759] **and Nenad Filipović**[1,4,5] [0000-0001-9964-5615]

[1]Steinbeis Advanced Risk Technologies Institute doo Kragujevac, Kragujevac, Serbia
e-mail: mnikolic@risk-technologies.com, tsustersic@risk-technologies.com, nf@risk-technologies.com
[2]Faculty of Mechanical and Civil Engineering in Kraljevo, University of Kragujevac, Kraljevo, Serbia
e-mail: mnikolic@risk-technologies.com, jordovicpavlovic.m@mfkv.kg.ac.rs
[3]Institute for Information Technologies Kragujevac, University of Kragujevac,Kragujevac, Serbia
e-mail: mnikolic@risk-technologies.com
[4]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
e-mail: tsustersic@risk-technologies.com, nf@risk-technologies.com
[5]Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
e-mail: tsustersic@risk-technologies.com, nf@risk-technologies.com
[6]Institute for Automation and Applied Informatics (IAI), Karlsruhe Institute of Technology, Karlsruhe, Germany
e-mail: thomas.lautenschlager@kit.edu, ralf.mikut@kit.edu
[7]Immunology & Immuno-bioengineering Group School of Life Sciences, University of Nottingham, Nottingham, United Kingdom
e-mail: Mohammad.Ghasemzadeh Hasankolaei@nottingham.ac.uk, Amir.Ghaemmaghami@nottingham.ac.uk

**corresponding author*

Abstract: The development of *in vitro* testing systems, such as liver spheroids and other tissue models, has enabled researchers to test drugs more effectively in accordance with the 3R principles (Replace, Reduce, Refine) in the ethical guidelines in animal research, leading to substantial progress in creating new drugs that are non-toxic or less toxic. Further optimisation is supported by development of *in silico* computational models. In this work, we examined how drug chemical composition and structure links to cellular responses in human liver spheroids. The dataset, used as an input to the model, and taken from public database, contains 20,000 brightfield microscopy images of human liver spheroids exposed to 108 drug compounds for assessing cell viability. Our work relies on work performed by Dubinsky et al., who trained a deep learning model to predict ATP values for a non-invasive estimation of cell viability over a period of up to 7 days of drug exposure. We enriched the model input by incorporating chemical composition and molecular structure of drug compounds (i.e. molecular descriptors, fingerprints and physicochemical properties), rather than relying solely on the concentration of the active substance, leading to improved latent representations within the model. Descriptors and fingerprints are computed using the AllChem module in RDKit. We created a training set that learns relationships between compound characteristics and liver spheroid toxicity, allowing us to predict the toxicity for similar drugs that have not yet been tested experimentally. The model assigns a compound, represented by its chemical and structural composition features, to one of four predicted toxicity classes – very toxic, moderately toxic, slightly toxic, and non-toxic. As part of future investigations, we will create *in silico* evaluations for additional 20 drugs not used in training, to identify less toxic alternatives within specific drug groups.

Keywords: liver spheroid, hepatotoxicity, deep learning, drug chemical composition and structure, drug safety.

Acknowledgement: This research is funded by the project that has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101138387 (TOXBOX). This article reflects only the authors' views. The European Commission is not responsible for any use that may be made for the information it contains. This research was also supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract numbers 451-03-34/2026-03/200108 (Faculty of Mechanical and Civil Engineering in Kraljevo, University of Kragujevac), 45103-33/2026-03/200378 (Institute for Information Technologies Kragujevac, University of Kragujevac) and 451-03-33/202603/200107 (Faculty of Engineering, University of Kragujevac). This research is in line with the strategic plan of the Republic of Serbia 2030, particularly in the section 2.2.2. Development of human resources, goal 3: Good health. The Helmholtz Association funds this project under the program "Natural, Artificial and Cognitive Information Processing (NACIP)" and the "Helmholtz Information & Data Science School for Health (HIDSS4Health)".

DEEP LEARNING FOR ZEBRAFISH EMBRYO PHENOTYPE CLASSIFICATION USING THE EMBRYONET DATASET

Dajana Jovanović[1,2*] [0009-0008-0580-9008], **Tijana Geroski[1,2,3]** [0000-0003-1417-0521], **Anđela Stojadinović[1,2]** [0009-0003-3444-1445], **Milica Nikolić[3,4,5]** [0000-0002-8537-7944], **Thomas Lautenschlager[6]** [0009-0002-3214-9363], **Ralf Mikut[6]** [0000-0001-9100-5496] **and Nenad Filipović[1,2,3]** [0000-0001-9964-5615]

[1]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
e-mail: dajana.jovanovic@uni.kg.ac.rs, tsustersic@risk-technologies.com, andjela.stojadinovic@uni.kg.ac.rs, nf@risk-technologies.com
[2]Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
e-mail: dajana.jovanovic@uni.kg.ac.rs, tsustersic@risk-technologies.com, andjela.stojadinovic@uni.kg.ac.rs, nf@risk-technologies.com
[3]Steinbeis Advanced Risk Technologies Institute Kragujevac, Kragujevac, Serbia
e-mail: tsustersic@risk-technologies.com, mnikolic@risk-technologies.com, nf@risk-technologies.com
[4]Faculty of Mechanical and Civil Engineering in Kraljevo, University of Kragujevac, Kraljevo, Serbia
e-mail: mnikolic@risk-technologies.com
[5]Institute for Information Technologies Kragujevac, University of Kragujevac, Kragujevac, Serbia
e-mail: mnikolic@risk-technologies.com
[6]Institute for Automation and Applied Informatics (IAI), Karlsruhe Institute of Technology, Karlsruhe, Germany
e-mail: thomas.lautenschlager@kit.edu, ralf.mikut@kit.edu

**corresponding author*

Abstract: In developmental biology, identifying phenotypic abnormalities in zebrafish embryos is key to understanding signalling pathways. Zebrafish are widely used due to their transparent embryos and suitability for large-scale imaging, enabling systematic phenotype analysis across different experimental conditions. However, analysis of time-lapse microscopy data remains mostly manual and difficult to scale. Deep learning offers a promising solution, although robustness across imaging conditions remains a challenge. In this work, we investigate binary classification of zebrafish embryo phenotypes from timelapse images, focusing on normal development and Nodal-related abnormalities. The publicly available EmbryoNet dataset consists of image sequences acquired over 2-26 hours post fertilization, with multiple embryos per frame annotated by bounding boxes and phenotype labels. The data are collected using different microscopy systems (Acquifer and Keyence), introducing domain shift. We evaluate four structured experiments: (i) within-domain training and testing, (ii) multi-source training, and (iii) cross-domain evaluation. A pretrained ResNet-18 model is used, and in experiment (iv), temporal information is added as an extra channel and as a scalar feature. The model demonstrates reliable performance in (i) (F1 0.853, ROC-AUC 0.941 with TTA). Multi-source training improves results (F1 0.918, ROC-AUC 0.971). In (iii), performance remains high (F1 0.899, ROC-AUC 0.970), indicating that results are influenced not only by domain shift but also by dataset characteristics, such as bit depth, exposure time, and file format. In experiment (iv), temporal information leads to a slight decrease (F1 0.831, ROC-AUC 0.926), indicating that the current temporal encoding strategy does not provide consistent improvements and requires further refinement. These results highlight the impact of domain variability and the benefit of multi-source training. Further, this illustrates the challenges of deploying pretrained models across diverse imaging systems in real laboratory environments. Future work will extend the approach to multi-class classification and improve modelling of temporal consistency and domain invariance, including transition-based and attention-based approaches.

Keywords: zebrafish embryos, phenotype classification, deep learning, convolutional neural networks, domain shift, temporal information, explainability.

Acknowledgement: This research is funded by the project that has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No. 101138387 (TOXBOX). This article reflects only the authors' views. The European Commission is not responsible for any use that may be made for the information it contains. This research was also supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract numbers 451-03-33/2026-03/200107 (Faculty of Engineering, University of Kragujevac), 451-03-34/2026-03/200108 (Faculty of Mechanical and Civil Engineering in Kraljevo, University of Kragujevac) and 451-03-33/2026-03/200378 (Institute for Information Technologies Kragujevac, University of Kragujevac). This research is in line with the strategic plan of the Republic of Serbia 2030, particularly in the section 2.2.2. Development of human resources, goal 3: Good health. The Helmholtz Association funds this project under the program "Natural, Artificial and Cognitive Information Processing (NACIP)" and the "Helmholtz Information & Data Science School for Health (HIDSS4Health)".

PHYSICS-INFORMED NEURAL NETWORKS FOR MODELING THE HODGKIN-HUXLEY ELECTROPHYSIOLOGICAL SYSTEM

Bogdan Milićević[1,2*] [0000-0002-0315-8263], **Miljan Milošević[1,2,3]** [0000-0003-3789-2404], **Edin Dolićanin[4]** [0000-0002-9896-8575], **Lazar Velicki[5]** [000-0002-2907-819X], **Andrej Preveden[5]** [0000-0003-3572-2604], **Danica Milošević[6]** [N/A], **Svetlana Apostolović[6]** [0000-0001-9015-297X], **Marija Polovina[7]** [0000-0001-9589-0023], **Stefan Simović[8]** [0000-0003-0911-2028], **Goran Davidović[8]** [0000-0001-5851-5857], **Dragana Ignjatović Ristić[8]** [0000-0002-2814-3105] **and Nenad Filipović[2,9]** [0000-0001-9964-5615]

[1]Institute for Information Technologies, University of Kragujevac, Kragujevac, Serbia
e-mail: bogdan.milicevic@uni.kg.ac.rs
[2]Research and Development Center for Bioengineering, Kragujevac, Serbia
[3]Belgrade Metropolitan University, Street, Belgrade, Serbia
e-mail: miljan.m@kg.ac.rs
[4]State University of Novi Pazar, Novi Pazar, Serbia
e-mail: edin@np.ac.rs
[5]Faculty of Medicine, University of Novi Sad, Novi Sad, Serbia
e-mail: lazar.velicki@mf.uns.ac.rs, andrej.preveden@mf.uns.ac.rs
[6]Faculty of Medicine, University of Niš, Niš, Serbia
e-mail: drapostolovic@gmail.com, danicamilosevic333@gmail.com
[7]Faculty of Medicine, University of Belgrade, Belgrade, Serbia
e-mail: maki.marijapolovina@gmail.com
[8]Faculty of Medical Sciences, University of Kragujevac, Kragujevac, Serbia
e-mail: simovicst@gmail.com, medicusbg@yahoo.com, draganaristic4@gmail.com
[9]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
e-mail: fica@kg.ac.rs
**corresponding author*

Abstract:

Physics-Informed Neural Networks represent a bridge between deep learning and biophysical modeling by embedding governing differential equations directly into the learning process. This study applies these networks to the Hodgkin-Huxley model to evaluate their capacity for solving electrophysiological equations without extensive manual intervention. A significant challenge in modeling excitable media is the spectral bias of deep learning, where neural networks tend to learn smooth curves before resolving the sharp, non-linear spikes essential to action potentials. To accurately capture the Hodgkin-Huxley system's interplay between rapid depolarization and slower recovery processes, we implement a sequential causal training approach. By dividing the time domain into connected segments, the network is forced to solve each segment accurately, respecting the physical conditions inherited from the previous step, before advancing. Furthermore, a PDE Point Resampler is utilized to dynamically shift training points toward areas of high residual error. This research demonstrates that PINNs can effectively replace or augment traditional numerical solvers in scientific computing, providing predictions at specific timestamps without the need to solve for all previous time steps.

Keywords: physics-informed neural networks, cardiac modeling, electrophysiological models, Huxley-Hodgkin model.

Acknowledgement: This research was supported by INTELHEART project that has received funding from the Science Fund of the Republic of Serbia, under Grant Agreement No 7580. The INTELHEART project team is solely responsible for the content of the paper and the content does not reflect the views of the Science Fund of the Republic of Serbia. This research was also supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract number [451-03-33/2026-03/200378 (Institute of Information Technologies, University of Kragujevac) and 451-03-34/2026-03/200107 (Faculty of Engineering, University of Kragujevac)].

DIGITALIZATION AND DIGITAL TRANSFORMATION OF THE ECONOMY: TWO SIDES OF THE SAME COIN

Marko M. Mihić [0000-0003-3994-2268]

Faculty of Organizational Sciences, University of Belgrade, 154 Jove Ilića Street, Belgrade, Serbia
e-mail: marko.mihic@fon.bg.ac.rs

Abstract:

The paper analyses the evolution of the digital economy, distinguishing between the mere digitalization of processes and the comprehensive digital transformation of socio-economic systems. Focusing on trends over the last three decades, the paper considers the intensity of change through the following indicators: from the increase in the digital economy's share of GDP in advanced economies from 3–4% in the mid-1990s to 8–12% today, to the dramatic leap in the global online population from less than 1% in 1993 to 67% in 2023. Attention is given to the expansion of e-commerce, which has grown from a marginal 0.5% share of retail sales in 2000 to over 20% in 2023, followed by a significant rise in average annual online consumer spending from $1,060 to over $2,500. From the perspective of the European Union (EU) Digital Decade 2030 vision, the paper examines the feasibility of the set goals through the prism of 'two sides of the same coin'. One dimension includes rapid technological development, increased efficiency, and new economic opportunities in emerging digital sectors. The other side highlights critical risks such as job displacement, social inequality, ethical challenges in AI governance, and the expanding skills gap due to delayed workforce reskilling. The main part of the research analyses future of the labor market, underlying predictions for the Serbian labor market. It is estimated that by 2030, approximately 70,000 jobs in Serbia will be automated due to AI, while at least 130,000 new opportunities are expected to emerge. In the public sector, where artificial intelligence is estimated to be capable of performing up to 82% of tasks, its impact is likely to be most profound in administrative services, particularly among record-keeping clerks, secretaries, counter staff or similar. Furthermore, around 27% of engineering and technical occupations will be affected. While up to 75% of current jobs are expected to be preserved, they will undergo significant shifts in job descriptions and required skill sets. In that sense, the paper analyzes the dualism between job positions highly vulnerable to automation and those likely to persist or grow through human judgment, creativity, or emotional intelligence. The main conclusions indicate that building a successful digital society by 2030 is not a matter of technological implementation, but of a society's capacity to manage integration between global digital brands and human capital, focusing on workforce reskilling to capitalize on the projected net job growth.

Keywords: digital transformation, digital economy, EU Digital Decade 2030, AI governance, Serbian labor market.

TRANSFORMING MACHINE LEARNING MODELS TO DECISION SUPPORT SYSTEMS

Boris Delibašić*[0000-0002-6153-5119], **Sandro Radovanović** [0000-0002-9975-8844] **and Milija Suknović** [0009-0003-0449-5258]

Faculty of Organizational Sciences, University of Belgrade, 154 Jove Ilića Street, Belgrade, Serbia
e-mail: boris.delibasic@fon.bg.ac.rs
**corresponding author*

Abstract:

We present a systematic approach for transforming machine learning (ML) models into decision support systems (DSS) suitable for real-world deployment. While ML models often achieve strong predictive performance, prediction alone is frequently insufficient in operational settings where decisions must comply with business rules, policy constraints, and requirements for stakeholder interpretability and justification. Moreover, many practical problems require actionable recommendations that go beyond patterns observed in historical data, introducing a decision layer not inherently captured by standard ML models. To address this gap, we propose an efficient and straightforward transformation process that augments ML models with decision-support capabilities. The approach integrates domain constraints, policy alignment, and explicit decision logic while preserving the predictive strengths of the original models. Empirical evidence suggests that this transformation can be achieved with minimal loss in predictive performance. We argue that such DSS-enhanced models provide a more appropriate framework for decision-making in complex organizational environments, where correctness, transparency, and actionability are as critical as accuracy. Our work positions DSS not as an alternative to ML, but as a necessary extension that bridges the gap between prediction and decision, particularly in business contexts where accountable and policy-aligned decisions are essential.

Keywords: Machine learning, predictive analytics, decision support systems, decision expert modeling.

PERCEPTION OF GENERATIVE AI AMONG SERBIAN STUDENTS IN A GLOBAL CONTEXT

Gordana Savic [0000-0001-6233-1815]

Faculty of Organizational Sciences, University of Belgrade, Belgrade, Serbia
e-mail: gordana.savic@fon.bg.ac.rs

Abstract:

The rapid development of generative artificial intelligence (GenAI), particularly large language models such as ChatGPT, is transforming higher education and redefining the competencies required in modern business environments. As early adopters, students play a critical role in shaping how these technologies are integrated into learning and professional practice. This study explores the perceptions and usage of generative AI among Serbian students and situates these findings within broader global trends.
The research adopts a comparative perspective, drawing on recent large-scale international evidence, including a global survey of over 23,000 students from 109 countries (Ravšelj, D., Keržič, D., Tomaževič, N., Umek, L., Brezovar, N., Iahad, N. A., & Aristovnik, A. (2025). Higher education students' perceptions of ChatGPT: A global study of early reactions. *PLoS One*, *20*(2), e0315011.). Global results indicate that students primarily use GenAI tools for tasks such as brainstorming, summarizing content, and supporting academic writing. They generally perceive these tools as useful for improving efficiency and access to knowledge, while also expressing concerns about reliability, academic integrity, and potential overreliance.
Against this global backdrop, the study explores whether Serbian students exhibit similar patterns of perception and use. The analysis focuses on key dimensions, including usage behavior, perceived benefits and limitations, ethical considerations, and expectations regarding the role of AI in future careers. The findings suggest a high level of alignment between Serbian students and global trends. Serbian students, like their international peers, predominantly use generative AI to support learning processes, particularly for idea generation, content organization, and exam preparation. Positive attitudes toward GenAI are evident, especially regarding its ability to enhance productivity and simplify complex information. At the same time, students express concerns similar to those reported globally, including risks of plagiarism, misuse in academic settings, and the need for clearer institutional guidelines.
From a management and business perspective, the results confirm that generative AI is widely perceived as a key factor influencing future employability. However, the findings also highlight a gap between frequent usage and the development of higher-order competencies, such as critical thinking and decision-making. This underscores the need for a structured and responsible integration of GenAI into higher education curricula.

Keywords: generative AI, ChatGPT, student perception, higher education, Serbia, global comparison, AI in business.

ENTREPRENEURSHIP IN TWELVE STEPS: AN LLM-ORIENTED CASE STUDY

Veljko Milutinović[1] [0000-0002-9380-5232], **Janko Tufegdžić**[2*] [0009-0001-7439-7034], **Matija Dodović**[2] [0009-0005-2795-5709] **and Dražen Drašković**[2] [0000-0003-2564-4526]

[1]Indiana University, 107 S. Indiana Avenue, Bloomington, IN 47405-7000, US
e-mail: veljko.milutinovic.1951@gmail.com
[2]School of Electrical Engineering, University of Belgrade, 73 Bulevar Kralja Aleksandra, Belgrade, Serbia
e-mail: janko@etf.bg.ac.rs, mdodovic@etf.bg.ac.rs, drazen.draskovic@etf.bg.ac.rs
**corresponding author*

Abstract:

This research starts from an existing course on Entrepreneurship, with its twelve technology-supported dimensions, and discusses their application in the context of selected LLM-oriented research efforts. The course has been taught so far in the USA at Purdue, IU, in reduced versions also at MIT and Harvard, at some EU and Middle East universities, but also in Kragujevac and Belgrade. It covers Proposing, Pinpointing, CMMI, AgiPlanning, SBA, PTO, ePresence, MindGenomics, Survey and Research publishing, DataMining, and eBranding. The stress is on pitfalls to avoid if the research focus is on LLMs and if the activity leans on AI.

Keywords: Entrepreneurship, LLM, AI.

TWELVE ENTREPRENEURSHIP-ORIENTED PARADIGMS SUPPORTED WITH AI

David Aksović* [0009-0003-5554-3491], **Marija Božić** [0009-0007-5913-2592], **Luka Bura** [0009-0001-2308-0782], **Mila Jakovljević** [0009-0001-4635-6024], **Aleksa Jovanović** [0009-0003-3393-3587], **Mina Kovandžić** [0009-0000-9633-1143], **Jelena Komazec** [0009-0007-4876-3161], **Aleksandra Labović** *[0009-0006-1917-1717], **Lazar Lazović** [0009-0007-9328-3423], **Nikola Milošević** [0009-0004-1414-3608], **NemanjaRšumović** [0009-0002-0716-9354], **Iva Čitlučanin** [0009-0004-0880-5863] and **Mihajlo Živković** [0009-0006-8563-6340]

Faculty of Mathematics, University of Belgrade, Belgrade, Serbia
**corresponding author*

Abstract:

This presentation covers twelve paradigmatic issues of importance for professional development after graduation, as well as for the management of large and complex projects. The twelve issues are: a) Research proposal preparation, b) Solving of GMATH problems, c) CMMI, d) Agile Methods, e) SBA, f) PTO, g) Writing of survey articles, h) Writing of research articles, i) Generation of web presentations, j) Mind Genomics, k) Data Mining, and l) Heritage Branding. The activity on each one of the issues is supported with AI. The topics covered by the group of authors from the School of Mathematics, University of Belgrade, include: making the educational process more effective, the creation of mission-specific software, public health, city intelligence, agriculture, more efficient communications, and improving the understanding of the environment. Special attention in this presentation is dedicated to methods that enhance creativity in research.

Keywords: Entrepreneurship-Oriented Paradigms, Artificial Intelligence, Agile Methods, Data Mining, Creativity in Research, Educational Innovation.

MAPPING A SPECIFIC SCIENCE-ORIENTED SOFTWARE ONTO 12 PARADIGMS OF INTEREST FOR ENTREPRENEURSHIP

Mihajlo Živkovic [0009-0006-8563-6340]

Faculty of Mathematics, University of Belgrade, Studentski trg 16, Belgrade, Serbia
e-mail: mihajlozivkovic97@gmail.com

Abstract:

This presentation covers the creation of a software system for the treatment of a scientific issue of importance. The selected problem has been treated through the following 12 dimensions: (a) Proposal writing, (b) Concise presentation of the problem/solution essence, (c) CMMI, (d) Agile planning, (e) SBA, (f) PTO, (g) Writing of a survey article, (h) Writing of a research article, (i) Generation of web presence, (j) Mind genomics, (k) Data mining, and (l) Heritage branding. It will be explained how AI could be used in each one of the 12 dimensions covered.

Keywords: agile software engineering, AI-assisted development, software lifecycle, software entrepreneurship, innovation paradigms.

CONSTITUTION AND ARTIFICIAL INTELLIGENCE

Srđan S. Đorđević [0000-0002-7005-1810]

Faculty of Law, University of Kragujevac, 1 Liceja Kneževine Srbije Street, Kragujevac, Serbia
e-mail: srdjordjevic@jura.kg.ac.rs

Abstract:

The modern development of artificial intelligence is characterized by the expansion of its influence over the entire legal order. In this regard, the constitution is no exception and raises a series of issues that require investigation. The theoretical nature of this paper aims to establish a methodological framework for future positive-legal research. The relationship between artificial intelligence and the constitution is not one-dimensional, as it can be observed within the context of the legal regulation of modern technologies, but also as a source of data for processing by artificial intelligence. The final and most sensitive problem facing constitutional law relates to the potential of artificial intelligence to master the function of interpreting constitutional provisions. Researching these designated topics poses a challenge for constitutional science, which is expected to provide answers regarding the position of the constitution in the digital era of the 21st century. The advancement of sophisticated generations of artificial intelligence necessitates a critical reassessment of established theoretical results in understanding the constitution within the modern state. The findings of the research presented in this paper lead to the conclusion that human creation must remain a key factor in constitutional processes and the process of interpreting constitutional provisions. Assistance from artificial intelligence is acceptable, but only under the management, supervision, and control of standard constitutional actors, without surrendering constitutional power to an artificial entity.

Keywords: constitution, constitutional text, artificial intelligence, law, interpretation.

AI MEMORIZATION REPRESENTS COPYRIGHT INFRINGEMENT: REGIONAL COURT IN MUNICH JUDGMENT *GEMA VS. OPEN AI*

Nikola Milosavljević [0000-0003-0208-2807]

Faculty of Law, University of Kragujevac, 1 Liceja Kneževine Srbije Street, 34000 Kragujevac, Serbia
e-mail: nmilosavljevic@jura.kg.ac.rs

Abstract:

The legal battle between copyright owners and Artificial intelligence (AI) providers rages on. With each passing day, there are more and more judgments and settlements between these two stakeholders, without a final solution in sight. In this paper, the author aims to analyze the judgment of the Regional Court in Munich (Ger. *Landsgericht München*) in *GEMA vs. Open AI*, which could position itself as one of the key cases for solving the problem of protecting copyrighted works used in AI development. In judgment, the problem of memorization is assessed by the court, taking into account the so-called text and data mining (TDM) exceptions and the possibility of their application. After the presentation of the most important facts of the case, the author then scrutinizes the impact of the court's reasoning and the possibilities it offers for further interpretation. Using the legal dogmatic and method of literature review, the author aims to answer the question of whether the reasoning of the court is correct and can be sustainable, having in mind the academic debate concerning these questions, as well as the plethora of social factors that influence this field.

Keywords: Artificial intelligence, memorization, GEMA vs. OpenAI, copyright protection, copyright infringement.

ARTIFICIAL INTELLIGENCE: LAWMAKING AND THE APPLICATION OF LAW

Dejan S. Matić [0000-0001-8190-0745]

Faculty of Law, University of Kragujevac, Kragujevac, Serbia
e-mail: dmatic@jura.kg.ac.rs

Abstract:

This paper explores the transformative impact of artificial intelligence on the legal landscape, focusing on two critical dimensions: the creation of legal norms and their practical application. As artificial intelligence technologies transition from speculative tools to active participants in decision-making processes, they challenge traditional legal doctrines and procedural frameworks. The research analyzes the potential for artificial intelligence to streamline legislative drafting (lawmaking) while scrutinizing the ethical and constitutional implications of algorithmic judicial reasoning (law application). By examining the intersection of computational logic and legal interpretation, the author highlights the necessity of human-centric oversight to ensure that the principles of justice, transparency, and the rule of law remain uncompromised in the digital age. To provide a comprehensive analysis of artificial intelligence's role in lawmaking and law application, the research employs a multifaceted methodological approach. By utilizing the legal (normative) method, the study scrutinizes existing regulatory frameworks and their adaptability to autonomous systems. The sociological method is applied to evaluate the broader social implications of algorithmic decision-making and its impact on legal subjects. Furthermore, the author incorporates an axiological approach to address the fundamental values of justice, fairness, and human dignity that must be preserved within automated processes. Through the synthesis of these methods, the paper highlights the tension between technological efficiency and legal certainty, ultimately arguing for a robust regulatory model that ensures artificial intelligence remains a tool for enhancing, rather than replacing, the human-centric nature of the rule of law.

Keywords: artificial intelligence, lawmaking, law application, legal theory, digital jurisprudence.

THE COMPLIANCE FUNCTION IN AN AI ENVIRONMENT: BETWEEN AI GOVERNANCE, DATA PROTECTION AND CYBERSECURITY

Zoran M. Đorović [0009-0005-3083-4078]

Data Cloud Technology d.o.o. Kragujevac, Kragujevac, Serbia
e-mail: zoran.djorovic@gmail.com

Abstract:

The development and deployment of artificial intelligence (AI) systems, particularly general-purpose AI (GPAI) models and large language models (LLMs), have opened a complex field of legal, organisational and security-related issues. In such an environment, an incident can no longer be viewed solely as a technical event; rather, it may at the same time constitute a regulatory issue and involve matters of data protection, contractual liability and risk management. The subject of this paper is the position and function of the compliance mechanism within such a normative and organisational framework. The central thesis of the paper is that the compliance function in an AI environment is not an auxiliary legal function, but an operational centre for the demonstrability of compliance. The paper does not merely present the obligations arising from the Artificial Intelligence Act (AI Act), the General Data Protection Regulation (GDPR) and Directive (EU) 2022/2555, commonly known as the NIS2 Directive, but derives an integrated model of AI compliance governance from them. In this model, the compliance function represents the central coordination mechanism for demonstrating compliance, the Chief Information Security Officer (CISO) represents the technical pillar of control and response, the Data Protection Officer (DPO)/GDPR function ensures the lawfulness of processing and data protection accountability, while governing bodies represent the level of approval, oversight and risk acceptance. The paper specifically analyses the problem of model opacity, the compliance lifecycle, the importance of documentation and audit trails, the allocation of responsibility between upstream and downstream actors, as well as recommendations for organisational action before, during and after an AI incident.

Keywords: AI governance, compliance function, AI regulation, GDPR, cybersecurity.

TRUST UNDER PRESSURE: EXPLAINABLE AI FOR CRISIS COMMUNICATIONS -- A LITERATURE REVIEW

Marko Selaković[1*] [0000-0002-6568-6627], **Tanja Kerševan**[2] [0000-0001-8318-4551], **Marijana Krkic**[3][0009-0006-8839-229X] **and Nenad Filipović**[4] [0000-0001-9964-5615]

[1]S P Jain School of Global Management, Block 5, Dubai International Academic City, Dubai, UAE
e-mail: marko.selakovic@spjain.org
[2]Faculty of Social Sciences, University of Ljubljana, Kardeljeva ploščad 5, Ljubljana, Slovenia
e-mail: tanja.kersevan@fdv.uni-lj.si
[3]Faculty of Political Sciences, University of Belgrade, Belgrade, Serbia
[4]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjic Street, Kragujevac, Serbia
e-mail: fica@kg.ac.rs
**corresponding author*

Abstract:

This paper examines explainability as a critical ethical and operational requirement for the use of artificial intelligence in crisis communications planning. The purpose of the research is to identify if and how AI explainability supports trustworthy, accountable, and effective decision-making when AI is used to support communications function in real crisis settings. Analyzing from interdisciplinary perspectives of strategic communications, security studies, and crisis management, the paper argues that explainability is one of the core conditions for legitimate AI-assisted crisis communication responses.
Research methodology includes a systematic review of scientific literature in the fields of communications, security, crisis management, and AI governance. The review outlines how explainability is defined, identifies the relationship of explainability with ethical AI principles such as transparency, accountability, and human oversight, and synthesizes evidence on its relevance for crisis preparedness and response.
The findings of the review demonstrate consistent agreement across the disciplines that explainability is one of the key AI ethics principles and a foundational enabler of responsible AI-assisted crisis communication. The literature shows that AI explainability increases stakeholders' trust, supports faster and justified decisions, enables scrutiny of data quality, while preserving meaningful human oversight. At the same time, the review highlights conceptual inconsistency across fields and very limited applied research on explainability in real crisis communication settings. The paper recognizes the need to develop clearer interdisciplinary frameworks and offers instruction for crisis communications practitioners to embed explainability into AI procurement, governance, training, and crisis protocols.

Keywords: explainable AI, crisis communication, AI ethics, strategic communication, AI governance.

LEGAL ASPECTS OF LIABILITY OF AI SYSTEMS WITH SPECIAL REFERENCE TO THE EU AI ACT

Andrija Popović [0000-0002-7567-5593]

Faculty of Law, University of Kragujevac, Kragujevac, Serbia
e-mail: apopovic@jura.kg.ac.rs

Abstract:

The development and application of artificial intelligence systems bring a range of significant changes and drive a comprehensive transformation of all major legal, economic, political, and other important social structures. Due to their increasing influence on everyday life, the question arises not only of how artificial intelligence systems function, but also of the dilemmas concerning the decisions made through their operation, as well as the consequences of those decisions, and subsequently the broader social consequences that emerge in contemporary societal processes.
The focus of this scientific article is the analysis of the legal aspects of liability of AI systems, with a special emphasis on the regulatory framework of the European Union, primarily the AI Act. The essence is to determine the complex system of liability that may be established between manufacturers, operators, and users of AI systems, as well as the need to analyze these highly complex relationships, considering the wide range of subjects that may come into contact with AI systems.
The classification of AI systems will be highlighted, which is of crucial importance for addressing the topic, as well as the emphasis on the necessity of efficient and transparent data management that is subject to analysis. It will also be pointed out that the EU AI Act represents a starting point towards harmonization in this field, but with an emphasized need for the development and analysis of certain hybrid legal models that will address existing legal gaps as well as unpredictable legal consequences that may arise from the use of AI systems.

Keywords: EU AI Act, digital law, algorithm, artificial intelligence, AI system, legal liability.

A CYBERSECURITY TRUST CHECK MODULE FOR RISK TRIAGE OF SUSPICIOUS MESSAGES

Višnja Mladenović [0009-0006-0627-7942], **Saša Cvetković*** **and Vladimir Mladenović** [0000-0001-8530-2312]

Faculty of Technical Sciences Čačak, University of Kragujevac, 65 Svetog Save Street, Čačak, Serbia
e-mail: visnja.mladenovic@student.kg.ac.rs sasa.cvetkovic@ftn.kg.ac.rs,
vladimir.mladenovic@ftn.kg.ac.rs
**corresponding author*

Abstract:

This paper presents an early-stage feasibility study of a cybersecurity Trust Check module for risk triage of suspicious messages directed at non-expert users. The module formulates message assessment as a four-class text classification problem covering legitimate messages, credential phishing, impersonation or urgent-action fraud, and malicious links or attachments. Beyond returning a class label, the module supplements each prediction with a risk level, a plain-language explanation, and a recommended next step, functioning as a lightweight decision-support workflow rather than a fully autonomous detection system. To evaluate the approach, we compare four modeling strategies: TF-IDF + Linear SVM as a lexical baseline, RoBERTa-base as a general-purpose transformer baseline, SecureBERT as a cybersecurity-domain transformer, and a lightweight hybrid combining SecureBERT with rule-based indicators. Evaluation is performed on a purpose-built 1000-message scenario-based benchmark with 250 messages per class, designed to be harder than prior datasets through broader stylistic diversity, multiple communication channels and organizational contexts, stronger lexical overlap between legitimate and malicious messages, and deliberate suppression of obvious giveaway cues. Results show that SecureBERT achieves the strongest performance with 0.935 accuracy and 0.9339 macro-F1, followed by SecureBERT + hybrid rules (0.9250 / 0.9244), TF-IDF + Linear SVM (0.9200 / 0.9190), and RoBERTa-base (0.9100 / 0.9098). Error analysis confirms that residual errors occur in semantically plausible borderline cases, validating the benchmark's difficulty. The findings indicate that cybersecurity-domain pretraining becomes more beneficial as task difficulty increases, and that suspicious-message triage is a practical and feasible cybersecurity NLP problem.

Keywords: cybersecurity NLP, suspicious message triage, phishing detection, social engineering, text classification, domain-specific language models, user-facing security.

HOLISTIC SECURE PLATFORM FOR INTELLIGENT THREAT AWARENESS AND BUSINESS-CONTINUITY IN TRUSTED ADAPTIVE HEALTHCARE ENVIRONMENTS

Ioannis Chouchoulis[1*] [0009-0007-6887-5435], **Ioannis Makropodis**[1] [0009-0006-3715-9690], **Aristeidis Farao** [1,2] [0000-0001-6954-0791], **Vaios Bolgouras** [2] [0000-0002-4934-7625], **Apostolis Zarras** [2] [0000 0002-8480-6989] and **Christos Xenakis**[1,2] [0000-0001-6718-122X]

[1]InQbit Innovations SRL, Romania
e-mail: giannis.chouchoulis@inqbit.io, giannis.makropodis@inqbit.io, aris.farao@inqbit.io, chris@inqbit.io
[2]University of Piraeus, Greece
e-mail: arisfarao@unipi.gr, vbolgouras@unipi.gr, zarras@unipi.gr, xenakis@unipi.gr
**corresponding author*

Abstract

The EU healthcare sector is undergoing rapid digital transformation, improving care and efficiency but also expanding its cyberattack surface through EHRs, telemedicine, IoT devices, and cloud platforms. Ransomware, third-party risks, complex regulation, legacy systems, and AI-driven vulnerabilities make it necessary to develop a next-generation, scalable, and interoperable cybersecurity solution for healthcare. To address these challenges, we propose HOSPITABLE, a modular, interoperable, and scalable cybersecurity framework designed to strengthen the resilience of healthcare environments through the integrated use of artificial intelligence, automation, and compliance-aware security services. The proposed framework combines an AI-powered Security Operations Center with SIEM and SOAR capabilities to support real-time threat detection, automated incident response, and forensic logging. These core functions are complemented by a Privacy and Compliance Center for regulatory readiness, a zero-trust identity and access management layer, a Training and Awareness Center offering adaptive and gamified exercises for different staff roles, and risk and cost assessment tools that support cybersecurity planning and investment decisions. The architecture has been designed to operate across heterogeneous hospital infrastructures, including large hospital networks and smaller regional facilities, while preserving interoperability and deployment flexibility. This paper presents the architecture, methodology, and components of HOSPITABLE, with the aim of delivering an interoperable and holistic cybersecurity solution that addresses technical safeguards, staff preparedness, and compliance with the evolving EU regulatory landscape.

Keywords: healthcare cybersecurity, compliance, AI-powered SOC, SIEM/SOAR, ransomware resilience, zero-trust IAM, compliance, privacy, gamified training, risk & cost tools, awareness.

Acknowledgement: This research has been funded by the ANTIDOTE project funded under Grant Agreement No. 101183162 supported by Horizon Europe, by the NITRO and CyberAId with Grant Agreement No. 101145872 and 101249596 correspondingly, funded under the Digital Europe Programme.

METAHEURISTIC OPTIMIZATION OF LGBM FOR SOFTWARE DEFECT PREDICTION

Strahinja Markovic[1*] [0009-0004-7812-8781], **Tamara Zivkovic[2]** [0000-0002-8563-8518], **Igor Dugonjic[3]** [0009-0003-5210-1372], **Aleksandar Petrovic[2]** [0000-0003-3324-3909], **Petar Spalevic[1]** [0000-0002-6867-7259] **and Miodrag Zivkovic[2]** [0000-0002-4351-068X]

[1]Faculty of Technical Sciences, University of Pristina in Kosovska Mitrovica, Serbia
e-mail: stralemarkovic@gmail.com, petar.spalevic@pr.ac.rs
[2]Faculty of Informatics and Computing, Singidunum University, Belgrade, Serbia
e-mail: tzivkovic@singidunum.ac.rs, aleksandar.petrovic@singidunum.ac.rs, mzivkovic@singidunum.ac.rs
[3]Faculty of Information Technologies, Pan-European University "APEIRON", Banja Luka, Republic of Srpska, Bosnia and Herzegovina
e-mail: igor.r.dugonjic@apeiron-edu.eu
**corresponding author*

Abstract:

The importance of software defect prediction (SDP) is enhanced by the recent advancements in the field of artificial intelligence (AI). The amount of code created has dramatically increased, while the percentage of defects in it, due to the use of generative agents, has increased. This problem is not new, and regardless of the generated code, a robust SDP method is required to tackle this problem. This paper explores one possible solution to it with the use of AI. Light gradient boosting machine (LGBM) is optimized with 9 metaheuristic algorithms and their results are validated statistically with the Friedman rank test. The experiments are performed on the well-known appraisal-based, estimation, and evaluation of effort models (AEEEM) dataset and compared against standard classification metrics along Mathew's correlation coefficient (MCC). This is research is of smaller scale and represents the first iteration of experimentation with these algorithms and is to be improved in future work.

Keywords: LGBM, metaheuristics, optimization, software defect prediction.

VISUAL ANALYSIS OF USER INTERACTIONS OVER TIME FOR THE DETECTION OF MALICIOUS ACTIVITIES USING MACHINE LEARNING

Ramona Glötter, Maja Temerinac-Ott[*][0009-0001-8588-7151] **and Richard Zahoransky** [0000-0001-7935-5258]

Department of Computer Science and Applications, Furtwangen University,
Furtwangen, Germany
e-mail: ramonagloetter@gmail.com, maja.temerinac-ott@hs-furtwangen.de, richard.zahoransky@hs-furtwangen.de
**corresponding author*

Abstract

The present paper investigates the visualisation of user behaviour and the detection of malicious activity through machine learning-based methods. The video game *Among Us* functions as a controlled environment for the simulation of insider threat scenarios. The detection of insider threats is defined as the identification of malicious activities conducted by authorised users within a system. This process is considered to be a critical yet challenging objective, primarily due to the subtlety of behavioural deviations and the overlap with legitimate activity. In the course of this study, keyboard and mouse dynamics are analysed as participants engage with the social deduction game *Among Us*. The random assignment of roles in each round naturally fosters the presence of malicious intent. The interaction data is automatically captured and labelled during gameplay, thereby providing a comprehensive dataset of both benign and malevolent behaviours. A range of visualisation techniques are employed to facilitate the analysis of interaction patterns, thereby emphasising the distinctive characteristics that differentiate between user roles. By leveraging behavioural metrics, the study explores how deviations in user input dynamics can serve as indicators of malicious intent, contributing to the detection of insider threats in security-critical systems. For the purpose of classification, a number of machine learning models are trained and evaluated on the collected data.

Keywords: visual analytics, user behaviour analysis, insider threat detection, keystroke dynamics, mouse dynamics.

AI RISK MANAGEMENT: A STRUCTURED APPROACH TO IDENTIFYING, ASSESSING, AND MITIGATING AI-DRIVEN THREATS

Zoran Keković[1] [0000-0003-0971-5371], **Nikola Vojtek**[2*] [0000-0001-9380-7046], **Vladan Pantović**[3] [0009-0001-6319-8822], **Kristijan Lazić**[4] [0009-0007-6243-3310] **and Ema Kraktus**[5] [0009-0001-1478-4315]

[1]Faculty of Security Studies, 50 Gospodara Vučića Street, Belgrade, Serbia
e-mail: zorankekovic@yahoo.com
[2]Smart Consulting Agency, 8 Miloja Zakića Street, Belgrade, Serbia
e-mail: nikola@smartconsulting-agency.com
[3]Digital Business Consulting, 65 Jove Ilića Street, Belgrade, Serbia
e-mail: vladan@pantovic.rs
[4]Argo-IT, 21 Nikole Marakovića Street, Belgrade, Serbia
e-mail: kristijan.lazic@gmail.com
[5]Ministry of Education, 22-26 Nemanjina Street, Belgrade, Serbia
e-mail: ema.kraktus@gmail.com
**corresponding author*

Abstract:

The rapid adoption of artificial intelligence (AI) across industries and domains has introduced significant opportunities for innovation, efficiency, and improvements in the decision-making area. However, alongside these benefits, AI systems pose a wide range of risks related to security, ethics, reliability, and regulatory compliance. These risks are related to the data, AI models and usage of AI tools, and are often complex, dynamic, and difficult to detect using the traditional risk management approaches.
This paper explores the concept of AI risk management through a structured framework that integrates principles from established risk management practices with emerging AI governance standards (such as ISO/IEC 42001). The study focuses on identifying key categories of AI-related risks, including data-related risks, model risks, operational risks, reputation and societal impacts. Furthermore, it proposes a systematic approach for risk assessment and mitigation, emphasizing the importance of transparency, traceability, and continuous monitoring throughout the AI model lifecycle.
Special attention is given to the role of processes within the company, cross-functional collaboration between teams, and alignment with international standard ISO/IEC 42001. The paper highlights how integrating AI risk management into existing governance frameworks can support companies in building trustworthy and resilient AI systems.
The findings suggest that a structured approach to AI risk management is essential for ensuring safe, ethical, and sustainable deployment of AI technologies in modern companies.

Keywords: artificial intelligence, risk management, AI governance, ISO/IEC 42001, AI risk, trustworthy AI, AI lifecycle, compliance.

DESIGN OF STRATIFYHF: AI-DRIVEN AND COMPUTATIONAL MODELLING CLOUD-BASED DECISION SUPPORT SYSTEM FOR HEART FAILURE

Miljan Milošević[1,2*] [0000-0003-3789-2404], **Bogdan Milićević**[1,2] [0000-0002-0315-8263], **Miloš Ivanović**[2,3] [0000-0002-8974-2267], **Aleksandar Atanasijević**[2,4] [N/A] **and Nenad Filipović**[2,5] [0000-0001-9964-5615]

[1]Bioengineering Research and Development Center (BioIRC), 26 Andre Marinkovica Street, Kragujevac, Serbia
e-mail: miljan.m@kg.ac.rs, bogdan.milicevic@uni.kg.ac.rs
[2]Institute for Information Technologies, University of Kragujevac, 1A Liceja Knezevine Srbije Street, Kragujevac, Serbia
e-mail: miljan.m@kg.ac.rs, bogdan.milicevic@uni.kg.ac.rs, mivanovic@kg.ac.rs, aatanasijevic@uni.kg.ac.rs, fica@kg.ac.rs
[3]Faculty of Science, University of Kragujevac, 12 Radoja Domonovića Street, Kragujevac, Serbia
e-mail: mivanovic@kg.ac.rs
[4]University of Kragujevac Computing Center, 6 Sestre Janjic Street, Kragujevac, Serbia.
e-mail: aatanasijevic@uni.kg.ac.rs
[5]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjic Street, Kragujevac, Serbia
e-mail: fica@kg.ac.rs
**corresponding author*

Abstract:

In this paper we present the architecture of the STRATIFYHF cloud-based decision support system (DSS) and integration of the data and AI analytics. Moreover, we present the multilayer structure of the platform, the integrated modules and technologies, the incorporated engines and services and the information flows between them. The STRATIFYHF architecture integrates different functional modules, their corresponding engines and the advanced computational and simulation project developed tools. Basic modules are: User Access Management Module, Data Management Module, Workflow Manager Module, and Visual Analytics Module, while the basic tools are: Risk Stratification tool, Early Diagnosis tool, Disease progression and prognosis tool, 3D computer modeling tool, and Voice processing tool. Moreover, the specific service engines are: Data quality control engine, Workflow engine, Docker engine, Visual analytics engine, and REST API manager. The STRATIFYHF framework is realized in the form of a hierarchical multilayer schema which consists of five layers: the hardware layer where the cloud resources stand along with any additional resources, the security layer, the workflow manager with its core engines, the back-end layer, and the front-end layer which includes the user interface and the visual analytics.

Keywords: cloud-based platform, AI-driven decision support system, computational modeling.

Acknowledgement: This study has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101080905 (STRATIFYHF). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or granting authority – European Health and digital executive agency (HADEA). Neither the European Union nor the granting authority can be held responsible for them. This research was also supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract number [451-03-33/2026-03/200378 (Institute of Information Technologies, University of Kragujevac) and 451-03-34/2026-03/200107 (Faculty of Engineering, University of Kragujevac)].

ECHOCARDIOGRAPHY VIEW CLASSIFICATION AS AN IMPORTANT STEP TO HEART FAILURE DIAGNOSIS: A CASE STUDY OF THE ECHOJEPA FOUNDATION MODEL

Ognjen Pavić[1,2*] [0000-0003-2533-1079], **Lazar Dašić[1,2]** [0000-0002-8055-100X], **Tijana Geroski[2,3]** [0000-0003-1417-0521], **Anđela Bladojević[2,3]** [0000-0002-8652-3827], **Andrej Preveden[4]** [0009-0008-5567-3014], **Aleksandra Milovančev[4]** [0000-0002-4398-8093], **Lazar Velicki[4]** [0000-0002-2907-819X], **Nduka Okwose[5]** [0000-0003-0145-2538], **Anne Nelissen[6]** [0009-0008-5358-2573], **Đorđe Jakovljević[5]** [0000-0003-2686-6542] **and Nenad Filipović[2,3]** [0000-0001-9964-5615]

[1]Institute for Information Technologies, University of Kragujevac, Kragujevac, Serbia
e-mail: opavic@kg.ac.rs, lazar.dasic@kg.ac.rs
[2]Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
e-mail: opavic@kg.ac.rs, lazar.dasic@kg.ac.rs, tijanas@kg.ac.rs, andjela.blagojevic@kg.ac.rs, fica@kg.ac.rs
[3]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
e-mail: tijanas@kg.ac.rs, andjela.blagojevic@kg.ac.rs, fica@kg.ac.rs
[4]Faculty of Medicine, University of Novi Sad, Street, Novi Sad, Serbia
e-mail: a.preveden@gmail.com, msaska05@gmail.com, lazar.velicki@mf.uns.ac.rs
[5]Institute for Cardio Metabolic Medicine, University Hospitals Coventry and Warwickshire, the United Kingdom
e-mail: ad6707@coventry.ac.uk, ad5287@coventry.ac.uk
[6]University Medical Center Utrecht, Street, Netherlands
e-mail: A.P.Nelissen@umcutrecht.nl
**corresponding author*

Abstract:

Echocardiographic images provide important information on cardiac structure and function, which is crucial for heart failure (HF) diagnosis assessment. Left ventricular ejection fraction (LVEF) is needed as one of the most important parameters in determining HF class and subsequent risk of cardiovascular adverse events. This parameter can be extracted from echocardiography videos/images. However, there are multiple special echocardiographic views from which ultrasound images of the heart are recorded that provide views from different angles of the heart and they are obtained from different viewpoints. Apical 2 chamber (A2CH) and apical 4 chamber (A4CH) views in tandem with parasternal long-axis (PLAX) views provide a comprehensive picture of different parts of the myocardial walls and their movement and thus enable inclusive EF assessment. Manual view classification is time-consuming, for these reasons we utilize artificial intelligence (AI) to classify these views. In order to automate this process, we utilize the EchoJEPA foundation model for echocardiographic images. Since EchoJEPA is a foundation model trained on a large dataset which encompassed more than 18 million ultrasound recordings, it provides a good baseline for this assessment with smaller databases. In this paper we assess the EchoJEPA architecture with modified output layers for automatic classification of cardiac echocardiography views. This step is an integral part of a larger pipeline for automatic assessment of echocardiography videos/images for HF diagnosis. By automatic detection of the view, we also set a baseline for image filtering and uncovering hidden knowledge such as subtle motion patterns and myocardial textures in echocardiographic images that are not measured in standard clinical practice.

Keywords: view classification, echocardiography, imaging, computer vision, foundation model, heart failure.

Acknowledgement: This paper has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101080905. (STRATIFYHF project). The research is supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract numbers [451-03-33/2026-03/200107 (Faculty of Engineering, University of Kragujevac) and 451-03-33/2026-03/200378 (Institute for Information Technologies, University of Kragujevac)].

DIFFERENTIATING SUSPECTED AND CONFIRMED HEART FAILURE USING MACHINE LEARNING AND REFINED VOCAL FEATURES

Lazar Dašić[1,2*] [0000-0002-8055-100X], **Ognjen Pavić**[1,2] [0000-0003-2533-1079], **Tijana Geroski**[2,3] [0000-0003-1417-0521], **Anđela Blagojević**[2,3] [0000-0002-8652-3827], **Andrej Preveden**[4] [0009-0008-5567-3014], **Aleksandra Milovančev**[4] [0000-0002-4398-8093], **Lazar Velicki**[4] [0000-0002-2907-819X], **Nduka Okwose**[5,6] [0000-0003-0145-2538], **Anne Nelissen**[7] [0009-0008-5358-2573], **Renae J. Stefanetti**[8] [0000-0001-6402-6435], **Sarah Charman**[8] [0000-0002-3029-7672], **Alessandra Fornarno**[9] [N/A], **Marta Jimenez-Blanco Bravo**[10] [0000-0001-9825-8285], **Nebojša Zdravković**[11] [0000-0003-0877-3684], **Snežana Lukić**[12] [0000-0002-1150-2925], **Đorđe Jakovljević**[5,6] [0000-0003-2686-6542] **and Nenad Filipović**[2,3] [0000-0001-9964-5615]

[1]Institute for Information Technologies, University of Kragujevac, Kragujevac, Serbia
[2]Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
[3]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
[4]Faculty of Medicine, University of Novi Sad, Novi Sad, Serbia
[5]Research Centre for Discoveries in Life Sciences, Coventry University, United Kingdom
[6]Institute for Cardio Metabolic Medicine, University Hospitals Coventry and Warwickshire, United Kingdom
[7]University Medical Center Utrecht, Utrecht, Netherlands
[8]Translational and Clinical Research Institute, Faculty of Medical Sciences, Newcastle University, Newcastle upon Tyne, UK
[9]Cardiomyopathy Unit, Careggi University Hospital, Florence, Italy
[10]Hospital Universitario Ramón y Cajal, Madrid, Spain
[11]Department of Medical Statistics and Informatics, Faculty of Medical Sciences, University of Kragujevac, Kragujevac, Serbia
[12]Department of Radiology, Faculty of Medical Sciences, University of Kragujevac, Kragujevac, Serbia
e-mail: lazar.dasic@kg.ac.rs, opavic@kg.ac.rs, tijanas@kg.ac.rs, andjela.blagojevic@kg.ac.rs, andrej.preveden@mf.uns.ac.rs, msaska05@gmail.com, lazar.velicki@mf.uns.ac.rs, ad6707@coventry.ac.uk, A.P.Nelissen@umcutrecht.nl, Renae.Stefanetti@newcastle.ac.uk, fornaroa@aou-careggi.toscana.it, martajbb@googlemail.com, nzdravkovic@fmn.kg.ac.rs, snezanamlukic@gmail.com, ad5287@coventry.ac.uk, fica@kg.ac.rs
**corresponding author*

Abstract: Heart failure (HF) is characterized by the heart's inability to pump enough blood to meet the body's metabolic demands. One of the most intriguing and innovative approaches to diagnosis of HF is the use of voice characteristics, as it is already established that HF can significantly alter a patient's voice. This study aimed to develop a robust machine learning pipeline capable of differentiating between patients with suspected HF referred to secondary care and those with a confirmed diagnosis, solely from voice biomarkers. The study included data from 240 patients (50 suspected HF and 190 confirmed HF cases), collected from six different medical centers across Europe. The patients followed multi-test voice recording protocol: reading the text out loud, free speech, number counting up, number counting down and maximum phonation test. Out of these tests, 490 voice characteristics were extracted. To mitigate "the curse of dimensionality", we implemented a staged feature selection process involving collinearity filtering ($|r| > 0.80$) and LASSO regularization, which reduced the feature space to 22 key biomarkers. To address class imbalance, the Synthetic Minority Over-sampling Technique (SMOTE) was integrated into the training of multiple classifiers, including Random Forest, XGBoost, and Extra Trees. Models were evaluated using a stratified 80/20 train-test split and optimized for accuracy and macro-F1 scores. The Extra Trees classifier demonstrated the highest diagnostic performance, achieving an accuracy of 78.4% and a macro-F1 score of 0.76. The model exhibited a high sensitivity for confirmed HF (89.5%) and successfully differentiated 70.0% of suspected HF cases. These findings suggest that a refined subset of 22 vocal biomarkers can serve as a reliable, non-invasive tool for HF stratification and screening, providing a balanced diagnostic performance.

Keywords: classification, heart failure, machine learning, voice biomarkers.

Acknowledgement: This paper has received funding from the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101080905. (STRATIFYHF project). The research is supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract numbers [451-03-33/2026-03/200107 (Faculty of Engineering, University of Kragujevac) and 451-03-33/2026-03/200378 (Institute for Information Technologies, University of Kragujevac)].

COMPARATIVE ANALYSIS OF U-NET-BASED ARCHITECTURES FOR CORONARY ARTERY SEGMENTATION USING X-RAY ANGIOGRAPHY IMAGES

Anđela Stojadinović[1,2*] [0009-0003-3444-1445], **Tijana Geroski**[1,2] [0000-0003-1417-0521], **Dajana Jovanović**[1,2] [0009-0008-0580-9008] **and Nenad Filipović**[1,2] [0000-0001-9964-5615]

[1]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjic Street, Kragujevac, Serbia
[2]Bioengineering Research and Development Center (BioIRC), 6 Prvoslava Stojanovica Street, Kragujevac, Serbia
e-mails: andjela.stojadinovic@uni.kg.ac.rs, tijanas@kg.ac.rs, dajana.jovanovic@uni.kg.ac.rs, fica@kg.ac.rs
**corresponding author*

Abstract:

Cardiovascular diseases are one of the leading causes of death worldwide, with coronary artery disease being among the most common forms. Accurate representation of the coronary arteries can help identify abnormalities such as narrowing or blockages, which are critical for early intervention. However, this task remains challenging due to low image contrast, complex vessel morphology, and the presence of extremely thin and branching structures. In this study, we perform a comparative analysis of three widely used segmentation architectures: U-Net, U-Net++, and U-Net 3+. All models were trained on 1,000 images from the ARCADE dataset, validated on 200, and tested on a held-out set of 300 images at 512×512 resolution. The models were trained with hyperparameters optimized for each architecture to ensure a fair comparison, using a unified preprocessing pipeline with mask generation and data augmentation. The initial results show that all three architectures achieve comparable performance, indicating their ability to capture the overall vessel structure despite the dataset complexity. U-Net achieves a Dice score of 0.749 ± 0.118 and IoU of 0.612 ± 0.14, while U-Net++ achieves a Dice score of 0.748 ± 0.114 and IoU of 0.61 ± 0.134. U-Net 3+ achieves the highest performance, with a Dice score of 0.752 ± 0.101 and IoU of 0.614 ± 0.129, suggesting improved multi-scale feature representation. In addition, U-Net++ achieves higher sensitivity (0.781 ± 0.11), indicating better detection of smaller vessel regions, while U-Net shows higher precision (0.789 ± 0.134), reflecting fewer false positive predictions. These differences reflect the balance between accurate detection of fine vessel structures and minimizing false positive predictions. Overall, the results demonstrate that U-Net-based architectures show strong potential for accurate coronary artery segmentation despite the challenges posed by angiographic data.

Keywords: coronary artery segmentation, angiographic images, deep learning, U-Net, U-Net++, U-Net 3+, convolutional neural networks, medical image analysis.

Acknowledgement: This research is funded by the project that has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101080905 (STRATIFYHF). This article reflects only the authors' views. The European Commission is not responsible for any use that may be made for the information it contains. This research was also supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract number 451-03-33/2026-03/200107 (Faculty of Engineering, University of Kragujevac).

SOC-BASED IMPLEMENTATION OF CNN MODEL FOR END-DIASTOLIC VOLUME CLASSIFICATION FROM ECHOCARDIOGRAM VIA HLS4ML

Nemanja Marković[1,2*] [0009-0003-8991-780X], **Tijana Geroski**[1,2] [0000-0003-1417-0521], **Emil Jovanov**[3] [0000-0001-6754-3518] **and Nenad Filipović**[1,2] [0000-0001-9964-5615]

[1]Bioengineering Research and Development Center (BioIRC), Kragujevac, Serbia
e-mail: nemanja@kg.ac.rs., tijanas@kg.ac.rs, fica@kg.ac.rs
[2]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
[3]Department of Electrical and Computer Engineering, University of Alabama in Huntsville, Huntsville, AL, USA
e-mail: emil.jovanov@uah.edu
**corresponding author*

Abstract:

Heart failure is a life-threatening, progressive cardiovascular condition manifested in the heart's inability to pump enough blood to meet the body's metabolic demands, leading to severe morbidity and high mortality rates. One of the primary drivers of this condition is dilated cardiomyopathy (DCM), characterized by the abnormal enlargement and weakening of the left ventricle, severely compromising systolic function. The medical gold standard for evaluating DCM non-invasively is through echocardiogram, which allows cardiologists to assess the End-Diastolic Volume (EDV), presenting the maximum physical capacity of the ventricle just before contraction which is the quantifiable measurement of pathological ventricular dilation, establishing it as a foundational parameter for identifying DCM severity. In this project, we present a low-latency edge computing solution by deploying a highly optimized, 8-bit quantized and 50% pruned CNN directly onto a Xilinx Artix-7 FPGA using the hls4ml interface. Our CNN model architecture is specifically tuned to maximize recall, prioritizing the capture of all pathological cases and minimizing the risk of diagnostic oversight in acute patients. Hardware-accelerated model autonomously processes echocardiogram frames to classify EDV condition and detect ventricular enlargement. This embedded FPGA solution guarantees deterministic, real-time processing without relying on external cloud infrastructure, thereby ensuring strict patient data privacy. By integrating this low-power, area-efficient model directly into portable ultrasound machines, we enable instantaneous, point-of-care heart failure screening. This equips emergency departments and under-resourced remote clinics with automated diagnostic support, empowering clinicians to make rapid, life-saving interventions directly at the patient's bedside.

Keywords: heart failure, dilated cardiomyopathy, echocardiogram, convolutional neural network, hls4ml, FPGA.

Acknowledgement: This research is funded by the project that has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101080905 (STRATIFYHF). This article reflects only the authors' views. The European Commission is not responsible for any use that may be made for the information it contains. This research was also supported by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract number 451-03-33/2026-03/200107 (Faculty of Engineering, University of Kragujevac).

EVALUATION OF LATEST CHATGPT MODELS FOR DISC-BASED PERSONALITY-TYPE RECOGNITION FROM INTERVIEW RESPONSES

Jelena Mišić[1][0009-0004-2546-3041], **Aleksandar Miljković**[2], **Mina Vasković Jovanović**[3] [0000-0003-3644-8723] **and Milan Čabarkapa**[3] [0000-0002-2094-9649]

[1]Faculty of Diplomacy and Security, University Union - Nikola Tesla, Belgrade
e-mail: jelena.misic@fdb.edu.rs
[2]Faculty of Technical Sciences, University of Novi Sad, Novi Sad, Serbia
e-mail: aleksandar.miljkovic@mup.gov.rs
[3]Faculty of Engineering, University of Kragujevac, Kragujevac, Serbia
e-mail: mcabarkapa@kg.ac.rs, mina.vaskovic.jovanovic@uni.kg.ac.rs

Abstract:

At present, large language models (LLMs) are widely used in recruitment processes through AI-assisted interviews, automated screening, and preliminary candidate evaluation. Since many job applications on professional employment platforms now include some form of AI-based interaction, it is important to examine how reliable these models are when used as supporting tools in human resource assessment. Recent studies have demonstrated that LLMs can identify certain personality-related patterns in textual responses. However, their accuracy, consistency, and practical validity are still not fully clear, particularly for the recruitment-oriented tasks. To this end, this study examines the use of several recent ChatGPT models with a different thinking depth for DISC-based personality-type recognition from open-ended interview-style answers. The analysis is based on the four commonly used DISC behavioral categories: Dominance, Influence, Steadiness, and Conscientiousness. The purpose of the study is not to replace human judgment, but to evaluate whether LLM-based classifications can provide accurate additional information during the early stages of candidate assessment. Different ChatGPT model settings, including standard and reasoning-based modes, are evaluated by comparing the consistency of their classifications, the clarity of their explanations, and their agreement with questionnaire-derived reference labels. The study also discusses key methodological and ethical issues related to using LLMs as supportive tools for personality-oriented assessment in recruitment.

Keywords: large language models, ChatGPT, interview analysis, candidate evaluation, human-AI interaction.

ACCELERATION WITH TWELVE PARADIGMS: AN LLM-ORIENTED CASE STUDY

Veljko Milutinović[1] [0000-0002-9380-5232], **Janko Tufegdžić**[2*] [0009-0001-7439-7034], **Matija Dodović**[2] [0009-0005-2795-5709] **and Dražen Drašković**[2] [0000-0003-2564-4526]

[1]Indiana University, 107 S. Indiana Aventua, Bloomington IN 47405-7000, US
e-mail: veljko.milutinovic.1951@gmail.com
[2]School of Electrical Engineering, University of Belgrade, 73 Bulevar Kralja Aleksandra, Belgrade, Serbia
e-mail: janko@etf.bg.ac.rs, mdodovic@etf.bg.ac.rs, drazen.draskovic@etf.bg.ac.rs
**corresponding author*

Abstract:

This research starts from an existing course on Acceleration, with twelve technology-enabled paradigms, and discusses their application in the context of selected LLM-oriented research efforts.
The course has been taught so far in the USA at Purdue, IU, in reduced versions also at MIT and Harvard, at some EU and Middle East universities, but also in Kragujevac and Belgrade. It covers ControlFlow in its four variations, DataFlow in its four variations, plus Quantum, Optical, Bio, and Chemo computing.
The stress is on pitfalls to avoid when the research focus is on LLMs and the activity leans on AI.

Keywords: large language models, paradigm, AI acceleration, computing.

ANALYSIS OF ATTENTION MECHANISMS IN THE CONTEXT OF COMPUTATIONAL PARADIGMS

Nikola Bogdanović [0009-0008-9124-9253], **Ivan Bošković, Stefan Bukorović, Mladen Čolaković, David Milinković, Nikola Nikolić, Miloš Paunović, Lazar Radosavljević, Nikola Ristić, Pavle Šarenac and Uroš Savurdić**[1] [0009-0006-1357-4080]

School of Electrical Engineering, University of Belgrade, 73 Bulevar kralja Aleksandra, Belgrade, Serbia
e-mail: savurdic.uros@gmail.com

Abstract:

This work investigates the performance characteristics of attention mechanisms in modern deep learning models through an operation-level decomposition and execution analysis across different computational paradigms. The attention mechanism is broken down into matrix multiplication (QK^T), softmax normalization, and weighted value aggregation, and each component is analyzed in terms of computational cost and memory behavior. We evaluate how these operations map onto different execution models, including multi-core and many-core control-flow systems, as well as graph-oriented and systolics-oriented dataflow architectures, also, those based on GaAs (Risk-V), CdTe (Turing), opto, quantum, bio, and chemo technologies.
The study focuses on identifying performance bottlenecks such as memory bandwidth limitations and compute utilization imbalance under varying input sizes. By comparing execution behavior across paradigms, we aim to characterize how architectural design choices influence the efficiency of attention-based workloads. The results provide insight into which system properties are most critical for accelerating attention mechanisms in practical AI applications.

Keywords: attention, matmul, softmax, computation, accelerators, AI.

COMPARISON OF ACCELERATION PARADIGMS FOR SELECTED AI ALGORITHMS

Nikola Bogdanović [0009-0008-9124-9253], **Ivan Bošković, Stefan Bukorović, Mladen Čolaković, David Milinković, Nikola Nikolić, Miloš Paunović, Lazar Radosavljević, Nikola Ristić, Pavle Šarenac and Uroš Savurdić**[1] [0009-0006-1357-4080]

School of Electrical Engineering, University of Belgrade, 73 Bulevar kralja Aleksandra, Belgrade, Serbia

Abstract:

This paper presents a comparative study of computational paradigms for the acceleration of key algorithms in modern artificial intelligence, including deep neural networks, matrix multiplication, attention mechanisms, diffusion models, and a representative image understanding task. We focus on understanding how different execution models impact performance and scalability across these workloads. The study considers multiple paradigms, including multi-core and many-core control-flow architectures, distributed systems based on shared memory and message passing, dataflow approaches (e.g. Groq-type and Google-type architectures), edge-oriented systems such as IoT and WSN, as well as emerging paradigms including GaAs, CdTe, optical, quantum, biological, and chemical computing. We aim to explore their characteristics and potential for efficient AI computation.

Keywords: neural networks, matmul, attention, computation, accelerators, AI, vision models.

A SEMANTIC JSON SCHEMA FOR KNOWLEDGE GRAPHS OF ELECTROSPUN PVDF PROCESSING, PROPERTIES, AND APPLICATIONS

Dusan Milojevic[1] [0009-0005-0113-3918], **Ana Kaplarevic Malisic**[2*] [0000-0002-1106-0274], **Strahinja Milenkovic**[1,3] [0000-0002-5822-4846], **Fatima Zivic**[1*] [0000-0003-2509-187X] **and Nenad Grujovic**[1] [0000-0002-8765-2196]

[1]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjica Street, Kragujevac, Serbia
e-mail: dusanmilojevic@uni.kg.ac.rs, strahinja.milenkovic@kg.ac.rs, zivic@kg.ac.rs, gruja@kg.ac.rs
[2]Faculty of Science, University of Kragujevac, 12 Radoja Domanovica Street, Kragujevac, Serbia
e-mail: ana.kaplarevic@pmf.kg.ac.rs
[3]Institute for Information Technologies, University of Kragujevac, 1A Liceja Kneževine Srbije Street, Kragujevac, Serbia
e-mail: strahinja.milenkovic@kg.ac.rs
**corresponding author*

Abstract:

The development of piezoelectric polymer systems, particularly electrospun poly (vinylidene fluoride) (PVDF), is driven by complex interactions between material properties, processing conditions, and environmental parameters. However, experimental data describing these relationships remain fragmented and inconsistently reported, limiting their reuse in data-driven materials design. In this work, we propose a structured JSON-based schema for representing electrospun PVDF systems as a foundation for knowledge graph construction. The schema captures hierarchical relationships across three primary domains: material properties, electrospinning conditions, and application context (biosensors and energy harvesting). Material descriptors include elastic stiffness tensors, dielectric permittivity, and piezoelectric coefficients, while process parameters encompass solution composition, applied electric field, flow rate, collector configuration, and environmental conditions such as temperature and humidity. The proposed schema incorporates semantic features such as standardized symbols and aliases to improve interoperability and machine readability. A case study demonstrates the instantiation of the schema using representative literature data, highlighting its ability to systematically link processing variables to functional properties such as the piezoelectric coefficient d33. The results show that the JSON structure provides a flexible and extensible framework for organizing heterogeneous experimental data and supports integration with graph-based machine learning approaches. This work contributes toward the development of materials knowledge graphs for piezoelectric polymers and provides a step toward FAIR-compliant data infrastructures and accelerated discovery of high-performance electrospun materials.

Keywords: electrospun PVDF, knowledge graph, materials informatics, process–structure–property relationships, data standardization, energy harvesting, biosensors.

Acknowledgement: This paper is supported through the EIT's Higher Education Initiative A-SIDE projects, coordinated by EIT RawMaterials, funded by the European Union and supported by projects No. 451-03-34/2026-03/200107, No. 451-03-33/2026-03/ 200122 and No. 451-03-33/2026-03/200378 Ministry of Science, Technological Development and Innovation, Republic of Serbia.

A TWO-LEVEL CLOUD–EDGE FRAMEWORK FOR PERSONALIZED GLYCEMIC PREDICTION

Aleksandra I. Stojnev Ilić[1*] [0000-0003-1071-6144], **Dragan H. Stojanović[1]** [0000-0003-4893-3814], **Natalija M. Stojanović[1]** [0000-0002-3594-1826] **and Milos Ilić[2]** [0000-0002-9696-2874]

[1]Faculty of Electronic Engineering, University of Niš, 4 Aleksandra Medvedeva Street, Niš, Serbia
e-mail: aleksandra.stojnev@elfak.ni.ac.rs, dragan.stojanovic@elfak.ni.ac.rs, natalija.stojanovic@elfak.ni.ac.rs
[2]Toplica Academy of Applied Studies, 1 Cirila i Metodija Street, Prokuplje, Serbia
e-mail: milos.ilic@vpskp.edu.rs
**corresponding author*

Abstract: Accurate blood glucose prediction is essential for effective diabetes management and is nowadays mostly based on machine/deep learning (ML/DL) model performed over time series CGM, lifestyle and activities data. Yet ML/DL model design must balance two competing requirements. Personalized models can capture individual metabolic responses and lifestyle patterns, but they require large amounts of patient-specific data that are not available at the time of deployment. In contrast, population-level models provide immediate usability but often fail to capture individual metabolic variability. This limitation is particularly important in diabetes, where individuals exhibit distinct glycemic dynamics influenced by factors such as insulin sensitivity, daily routines, and physical activity.

A single global model is therefore insufficient to represent the diversity of patient profiles. Instead, patients can be grouped into subcategories with similar glucose dynamics, enabling the construction of multiple generalized models that better capture shared patterns within each group. However, even group-level models cannot fully account for individual-specific deviations. This creates the need for a structured approach that combines generalized knowledge with patient-level refinement and personalization.

This work proposes a cloud–edge system that addresses this problem through a two-level prediction strategy. At the population level, multiple generalized ML/DL models are constructed, each corresponding to a subgroup of patients with similar glycemic behavior. At the individual level, each patient relies on a personalized and adaptive ML model that refines predictions based on their own data. This design enables immediate deployment using group-level knowledge, followed by rapid personalization as patient-specific data become available.

The system is implemented in three stages. First, a cloud component identifies patient groups using clustering over glycemic profiles derived from a large diabetes and continuous glucose monitoring (CGM) dataset. For the purpose of this research, we use the HUPA-UCM dataset, which includes CGM data and wearable sensor signals for 25 patients with Type 1 diabetes. K-means clustering is applied, with the number of clusters selected using silhouette score. For each cluster, a Long Short-Term Memory (LSTM) model is trained to learn temporal glucose dynamics using CGM signals and contextual features such as heart rate, physical activity, and circadian patterns. These models serve as generalized predictors for their respective patient groups.

Second, an edge component enables patient-specific adaptation. Each patient is assigned to a group model and receives its predictions as a baseline. A lightweight online linear model is then applied to adjust these predictions using a set of contextual features, including recent prediction errors, heart rate, step count, exercise indicators, calories intake and time-of-day signals. The model is updated incrementally as new glucose measurements become available, with delayed feedback reflecting real-world sensing conditions. This approach ensures stable performance at deployment while allowing continuous improvement without retraining the underlying models.

Third, an evaluation stage assesses both statistical accuracy and clinical relevance. Experiments are conducted on 14-day test windows with prediction horizons of 30 and 60 minutes. The personalized adaptation improves performance over group models for 12 of 18 patients with available group assignments. Cold-start analysis shows that patient-specific adaptation reaches stable performance within 2–3 days. The system is fully containerized using Docker Compose, with separate cloud, edge, and evaluation services, ensuring reproducibility and modular deployment.

The main contribution is a practical framework that combines multiple generalized models with efficient patient-level adaptation. By structuring the transition from population knowledge to personalized prediction, the system achieves strong accuracy, rapid convergence, and low computational cost, which makes it suitable for deployment on resource-constrained wearable and mobile platforms.

Keywords: blood glucose prediction, personalized and group-based modeling, cloud–edge computing.

DISCOVERING PHYSICAL LAWS FROM EXPERIMENTAL DATA USING REGRESSION ANALYSIS

Dalibor Rajkovic [0009-0002-4947-0401]

Gymnasium Cuprija, 57 Karadjordjeva Street, Cuprija, Serbia
e-mail: d.rajkovicrs@gmail.com

Abstract:

Discovering physical laws from experimental data represents one of the fundamental challenges in the analysis of physical systems. The application of data analysis methods, regression models, and modern computational tools enables a transition from the classical verification of known formulas toward a data-driven approach in which physical laws are derived directly from measurements, experiments, and numerical simulations.
This paper presents a research model based on correlation analysis, regression analysis, and advanced regression techniques for identifying functional relationships between physical quantities. The main objective is not the confirmation of previously known equations, but the reconstruction of physical regularities using experimental data. By analyzing measured values, the form of dependence between observed variables is examined, where classical regression models provide an initial estimation of the relationship type, while advanced regression approaches enable the determination of the analytical form of the law without prior knowledge of its exact formula.
The particular value of this approach lies in the integration of physics, mathematics, programming, and artificial intelligence, demonstrating how data-driven methods can contribute to understanding natural phenomena and improving scientific reasoning. This model provides a modern framework for physical system analysis in which experimentation, numerical modeling, and algorithmic reasoning become the basis for discovering physical laws.

Keywords: regression analysis, data analysis, physical laws, symbolic regression, experimental data, artificial intelligence.

WHAT DOES UNSUPERVISED ML SAY? CLUSTERING VIRTUAL REALITY BEHAVIOR FOR TYPICAL AND PHYSICALLY DISABLED PERSONS

Kriszta Soos and Marc Frincu* [0000-0003-1034-8409]

Faculty of Computer Science, West University of Timisoara, Bd. Vasile Pârvan 4, Timișoara, Timiș, Romania
e-mail: kriszta.soos04@e-uvt.ro, marc.frincu@e-uvt.ro
**corresponding author*

Abstract:

Virtual Reality is increasingly utilized in fields such as education and medicine. Studies have evidenced that it has the ability to induce neurological effects comparable to those generated by real-world experiences. As a result, it can serve as an important support tool for individuals with various types of disabilities, e.g., motor impairments, by providing an effective and non-intrusive rehabilitation approach. Through artificial intelligence, behavioral patterns in virtual environments can be determined and analyzed, providing further insight into patients' well-being.
Therefore, the aim of this study is to investigate behavioral patterns based on virtual reality data from 19 subjects and correlate them with the subjects' medical motor disability. Three distinct scenarios covering various locomotion systems (i.e., controller-based, eye movement-based, and static–no movement) are investigated with real-time data collected from the virtual environment. Dimensionality reduction via Principal Component Analysis is employed, followed by unsupervised learning using the K-means algorithm.
Results show that the proposed methodology successfully detected all participants presenting motor disabilities, accurately capturing their specific navigational limitations through quantifiable kinematic signatures. The behavioral variation scoring system produced a structured three-tier classification, with the two clinically confirmed atypical participants receiving scores of 1.000 and 0.875, clearly separated from the remainder of the cohort. However, highly energetic or spatially exploratory participants were occasionally misclassified as atypical, indicating that the current model cannot yet reliably distinguish between restriction-driven and engagement-driven behavioral variation. These findings demonstrate the potential of VR-based unsupervised behavioral analysis as a non-intrusive screening tool for motor impairments, while highlighting the need for a secondary interpretative layer to improve classification specificity.

Keywords: virtual reality, motor disability, unsupervised learning.

Acknowledgments: The work of the second author has been supported by the "Romanian Hub for Artificial Intelligence - HRIA", Smart Growth, Digitization and Financial Instruments Program, 2021-2027, MySMIS no. 334906/351416.

MAPPING AI APPLICATIONS IN EDUCATION – A CONCEPTUAL SYNTHESIS OF THE LITERATURE

Daliborka Purić* [0000-0002-8621-9171] and **Mirjana Čutović** [0000-0002-1998-1094]

Faculty of Education in Uzice, University of Kragujevac, Uzice, Serbia
e-mail: daliborka.puric@gmail.com, mirjana.cutovic3@gmail.com
**corresponding author*

Abstract:

The rapid development of artificial intelligence highlights the need for its comprehensive study, particularly given the scope and scale of its current and potential impact across various aspects of modern society. The purpose of this study was to reveal and analyse the key aspects of AI applications in education. The study aimed to: identify major thematic clusters and conceptual patterns of AI applications in education; analyze their dimensions, implementation approaches, and influencing factors through a thematic synthesis of the literature; highlight underexplored areas and conceptual gaps, providing recommendations for future research. This study is based on a semi-systematic literature review aimed at synthesizing prior research on AI applications within educational settings. The research approach involved a structured search of relevant academic publications in scholarly databases using predefined keywords related to the possibilities, benefits and risks of applying AI in education, as well as to the roles and competencies of teachers in applying AI in teaching and learning. The methodology focused on the identification, analysis, and classification of principal issues and conceptual foundations of AI applications in education in order to contribute to a deeper understanding of this field within the educational system. The literature analysis indicates that significant attention has been devoted to the pedagogical and technological aspects of AI applications in education, particularly in the areas of personalized learning, the development of intelligent tutoring systems, and the improvement of automated and objective assessment, all of which enhance the quality of the educational process. On the other hand, the literature review identifies psychological, ethical, and socio-economic aspects as insufficiently explored. Under-researched issues include privacy and ethical concerns related to the use of student data, fairness in AI implementation, the impact on socio-emotional development, long-term effects on learning, and the (re)definition of teachers' roles. Given that a keyword-based research strategy may be biased toward frequently indexed constructs that do not necessarily represent specific educational topics, and that it cannot fully explain qualitatively similar or culturally nuanced constructs, the research findings should be interpreted as a plausible representation of conceptual patterns rather than as empirically precise conclusions.

Keywords: AI applications in education, pedagogical-psychological aspect, technological aspect, socio-economic aspect, ethical aspect, semi-systematic literature review.

AUTOMATED LICENSE PLATE RECOGNITION IN CONSTRAINED ENVIRONMENTS: A CASE STUDY ON AI-ASSISTED DEVELOPMENT, SECURITY IMPLICATIONS AND SUSTAINABILITY

Nenad Korolija[1*][0000-0003-2224-8157] **and Veljko Milutinovic**[2][0000-0002-9380-5232]

[1]University of Criminal Investigation and Police Studies, Belgrade, Serbia
e-mail: nenadko@gmail.com
[2]Indiana University, Bloomington, Indiana, USA,
e-mail: veljko.milutinovic.1951@gmail.com
**corresponding author*

Abstract:

AI-assisted software development can accelerate software development, but introduces novel security risks. Simultaneously, the computational cost of AI training and inference raises sustainability concerns in high-performance computing environments. This paper presents a dual contribution: (1) a constrained automated license plate recognition (ALPR) system using TensorFlow, and (2) a security and sustainability analysis of the AI-assisted development process itself. The ALPR system was developed using a hybrid software engineer and AI development approach. The process includes rectangle detection, black-white color verification, and red-separator localization under controlled conditions (no skew, no digit zero ambiguity). Security vulnerabilities were analyzed through the lens of the recent trend in changing the programming model from conventional to AI-assisted. The carbon efficiency was also analyzed. The system achieved functional recognition given certain constraints, with a total development time of around 4 person-hours. The security analysis revealed that debugging and automated testing (e.g. pre-execution hooks) are necessary, but not sufficient. We argue that the software engineering trend toward "less programming, less understanding" is ecologically and security-unsustainable. We advocate for mandatory pre-execution validation layers. Untrusted AI-generated code requires architectural zero-trust validation.

Keywords: AI-assisted software development, license plate recognition, ALPR, software supply chain security, HPC sustainability, ghost dependencies, pre-execution validation, TensorFlow, OpenCV, Tesseract OCR, Serbian license plates.

TOWARDS INTELLIGENT AUTOMATED ESSAY GRADING USING LARGE LANGUAGE MODELS

Vasileios C. Pezoulas[1] [0000-0002-1872-693X], **Fani Biskanaki[2]** [0000-0002-3537-7276], **Athanasios K. Sypsas[2]** [0000-0001-8301-4763], **Venediktos Hadjisavva[3], Andreas Paraskeva[3], Maria Evripidou[3], Jelena Pavić[4,5]** [0000-0002-5875-6906], **Ana Orlović Kovačević[6], Zlatan Car[7]** [0000-0003-2817-9252], **Neda Vidanović Miletić[8]** [0000-0003-4784-4473], **Costas Assimakopoulos[9]** [0000-0002-4326-1695] **and Themis P. Exarchos[1*]** [0000-0001-9420-9419]

[1]Dept. of Informatics, Ionian University, Tsirigoti Square 7, Corfu, Greece
e-mail: bpezoulas@gmail.com, exarchos@ionio.gr
[2]Directorate of Secondary Education, Amynta 9, Preveza, Greece
e-mail: fani.biskan@gmail.com, sipsas@gmail.com
[3]C' Technical and Vocational School of Education and Training Limassol, Limassol, Cyprus
e-mail: venedik@cytanet.com.cy, paraskeva_cy@hotmail.com, mariaevripidou16@gmail.com
[4]City of Kragujevac, 3 Trg slobode, Kragujevac, Serbia
e-mail: jelena.pavic@uni.kg.ac.rs
[5]The Institute for Information Technologies Kragujevac, 1A Liceja Kneževine Srbije Street, Kragujevac, Serbia
e-mail: jelena.pavic@uni.kg.ac.rs
[6]Croatian Academic and Research Network - CARNET, 5 Josipa Marohnića Street, Zagreb, Croatia
e-mail: ana.orlovic@carnet.hr
[7]Faculty of Engineering, Catholic University of Croatia, Ilica 244, Zagreb, Croatia
e-mail: zlatan.car@unicath.hr
[8]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjić Street, Kragujevac, Serbia
e-mail: neda@kg.ac.rs
[9]Organisations Management, Marketing and Tourism, International Hellenic University, International Hellenic University - Thessaloniki campus PO Box 141, Thessaloniki, Greece
e-mail: casim@ihu.gr
**corresponding author*

Abstract:

Automated essay grading (AEG) remains a challenging problem in educational technology. This paper presents a benchmarking study of Retrieval-Augmented Generation (RAG) combined with state-of-the-art Large Language Models (LLMs) for automated assessment of student essays across three pedagogically meaningful dimensions: Content, Organization, and Language. The proposed framework employs multilabel stratified sampling, strict student-level data partitioning, and dimension-specific FAISS-based knowledge bases using the MiniLM embedding model. Four LLMs - GPT-4o-mini, GPT-4.1-mini, DeepSeek-Chat, and DeepSeek-Coder - were evaluated on a corpus of 2,617 student essays from the multinational AISE Erasmus+ educational initiative. GPT-4o-mini consistently outperforms all competing models, achieving accuracies of 0.820, 0.790, and 0.740 on the Content, Organization, and Language dimensions, respectively, with macro F1-scores of 0.813, 0.774, and 0.724. Confusion matrix analysis reveals that GPT-4o-mini maintains well-calibrated decision boundaries, while competing models exhibit systematic directional biases and grade compression. Runtime measurements confirm GPT-4o-mini also offers superior computational efficiency, with mean inference latencies approximately three times lower than the DeepSeek models. The grading engine is deployed as a FastAPI service integrated into the AISE educational platform, demonstrating the feasibility of scalable, interpretable, and pedagogically aligned automated essay assessment.

Keywords: automated essay grading, RAG, large language models, GPT-4o-mini, educational AI, FAISS, FastAPI, AISE platform.

Acknowledgement: This work received funding from the ERASMUS+ programme 2023-1-EL01-KA220-SCH-000157157: "An artificial intelligent platform to support students by assessing their performance skills through predictive models created from their writing skills."

RGB-D BASED 4DOF GRASP PREDICTION VIA CNN SEGMENTATION FOR ASSISTIVE ROBOTICS

Alexander Mattes and Maja Temerinac-Ott* [0009-0001-8588-7151]

Department of Computer Science and Applications, Furtwangen University, Furtwangen, Germany
e-mail: maja.temerinac-ott@hs-furtwangen.de , a.mattes@hs-furtwangen.de
**corresponding author*

Abstract:

Recent advances in robotic arm safety have enabled their deployment in assisted living environments to support elderly and physically impaired individuals with everyday tasks. A key remaining challenge is reliable grasping of unfamiliar objects in unstructured real-world settings. This work addresses automated grasp prediction using RGB-D data and formulates the problem as a pixel-wise segmentation task with a Convolutional Neural Network. Each pixel is classified as graspable or non-graspable, while a rotation map predicts grasp orientation. The most graspable pixel, combined with depth information, determines the 3D grasp position and gripper rotation, resulting in a 4DoF grasp representation. A detailed literature review was conducted to evaluate existing grasping methods and guide architectural design choices. Based on this analysis, a novel RGB-D segmentation architecture was developed, introducing a new feature extraction module called Resplacement, which combines local and contextual information. Experimental results demonstrate improved grasp prediction performance compared to baseline approaches. The study also explored effective RGB-D fusion strategies and developed a real-time capable network emphasizing both spatial detail and contextual understanding. The approach was later extended to 6DoF grasp prediction using surface normals, though with slightly lower accuracy. Evaluated in a realistic robotic setup, the system achieved an 82.2% grasp success rate on unknown objects.

Keywords: robotic grasping, semantic segmentation, convolutional neural networks, rgb-d perception, 4DoF prediction.

ADVANCES IN ARTIFICIAL INTELLIGENCE AND DEEP LEARNING FOR NEUROSONOGRAPHIC AND CHEST X-RAY ANALYSIS OF PREMATURE INFANTS

Julija Bošković[1] [0009-0006-3031-2234], **Suzana Živojinović**[2,3] [0000-0002-6844-2150], **Jelena Ceković Đorđević**[2,3] [0000-0003-3312-6307], **Nela Papović**[4], **Nikola Prodanović**[5,6] [0000-0003-0760-3931], **Goran Devedžić**[1] [0000-0002-6589-5883], **Suzana Petrović Savić**[1*] [0000-0002-5629-6221] **and Tijana Prodanović**[2,3] [0000-0002-7399-2480]

[1]Faculty of Engineering, University of Kragujevac, 6 Sestre Janjic Street, Kragujevac, Serbia
e-mail: julijabkg@gmail.com, devedzic@kg.ac.rs, petrovic.savic@kg.ac.rs
[2] Faculty of Medical Sciences, University of Kragujevac, Department of Pediatrics, 69 Svetozara Markovica Street, Kragujevac, Serbia
[3]University Clinical Center Kragujevac, Pediatric Clinic, Center for Neonatology, 30 Zmaj Jovina Street, Kragujevac, Serbia
e-mail:zivojinovicsuzana@yahoo.com, j.cekovic86@gmail.com, tijanaprodanovic86@gmail.com
[4]University Clinical Center Kragujevac, Pediatric Clinic, 30 Zmaj Jovina Street, Kragujevac, Serbia
e-mail:nelakg83@gmail.com
[5]Faculty of Medical Sciences University of Kragujevac, Department of Surgery, 69 Svetozara Markovica Street, Kragujevac, Serbia
[6]University Clinical Center Kragujevac, Clinic for Orthopaedic and Trauma Surgery, 30 Zmaj Jovina Street, Kragujevac, Serbia
e-mail: nikolaprodanovickg@gmail.com
**corresponding author*

Abstract:

Premature infants represent one of the most vulnerable patient populations in neonatology, frequently affected by severe neurological and respiratory complications that require rapid, reliable, and accurate diagnosis. Conditions such as hypoxic-ischemic encephalopathy (HIE) and respiratory distress syndrome (RDS) are among the leading causes of morbidity and mortality in premature newborns, often associated with long-term developmental and health consequences. Early detection and continuous monitoring of these conditions are essential for improving therapeutic outcomes and reducing the risk of severe complications. In recent years, advances in artificial intelligence (AI), deep learning, and computer-assisted medical imaging have opened new possibilities for improving diagnostic precision, reducing subjectivity in image interpretation, and supporting clinical decision-making in neonatal intensive care units (NICUs).

This paper presents an overview of contemporary AI-based approaches in neonatal imaging, with a particular focus on neurosonographic analysis of HIE and chest X-ray analysis of RDS in premature infants. The presented methodologies include convolutional neural network (CNN)-based classification of neurosonographic findings using echogenicity analysis and Delta E CIE76 quantification, as well as computer-assisted lung segmentation and radiographic analysis combined with blood gas parameters' evaluation. The neurosonographic classification model demonstrated high performance in differentiating normal, moderate, and severe pathological findings, while lung segmentation algorithms achieved promising accuracy and robustness in monitoring respiratory recovery and evaluating disease progression. In addition, the integration of quantitative image analysis with clinical parameters enables more objective assessment of neonatal conditions and supports the development of intelligent diagnostic support systems.

By combining medical image processing, machine learning, and clinical data analysis, these approaches demonstrate the considerable potential of AI technologies for early diagnosis, treatment planning, and continuous monitoring of premature newborns. Furthermore, the study highlights the growing importance of intelligent diagnostic systems in modern neonatology and emphasizes the future potential of explainable and multimodal AI models for improving neonatal healthcare outcomes, optimizing therapeutic strategies, and advancing personalized neonatal care.

Keywords: neonatal imaging, premature infants, neurosonography, chest X-ray analysis, hypoxic-ischemic encephalopathy, respiratory distress syndrome, convolutional neural networks, medical image processing.

MAPPING AI IN OSTEOCHONDRAL TISSUE ENGINEERING: A REPRODUCIBLE PYTHON FRAMEWORK

Marko N. Živanović [0000-0002-8833-8035]

University of Kragujevac, Institute for Information Technologies Kragujevac, Department of Science, Liceja Kneževine Srbije 1A, Kragujevac, Serbia
e-mail: marko.zivanovic@uni.kg.ac.rs

Abstract:

Osteochondral tissue engineering (OCTE) faces persistent challenges in scaffold design, yet the systematic integration of Artificial Intelligence (AI) into the field remains largely unexplored. To map this gap reproducibly, an open-source Python framework, OCTE-AI-Mining, is presented. The framework automates literature collection, citation enrichment and visualization across four target domains (AI, scaffolds, computational modeling, OCTE) and their combinations for the period 2010–2025. It comprises three modules: literature_octe_ai.py performs year-stratified NCBI Entrez queries with Title/Abstract filters, PMID deduplication and HTTP 429/502 retry logic; top_cited.py enriches DOI-bearing records through batched OpenAlex calls and exports per-category top-cited lists; growth_from_matrix.py renders annual growth curves, heatmaps and pie charts at 300 dpi from a Year × Category matrix. Deterministic execution is ensured via fixed package versions and seeded operations. All outputs (CSV corpora, audit-trail matrix, top-cited lists, figures) are released publicly through Zenodo, providing a transparent baseline that other groups can re-run, extend, and benchmark against in AI–biomaterials mapping studies.

Keywords: literature mining, osteochondral tissue engineering, artificial intelligence, reproducible workflow, Python.

Acknowledgments: This study is supported by the European Union's HORIZON TMA MSCA Staff Exchanges program under grant agreement No. 101183041 — EngVIPO. Funded by the European Union. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. This research was funded by the Ministry of Science, Technological Development and Innovation of the Republic of Serbia, contract number 451-03-33/2026-03/200107 (Faculty of Engineering, University of Kragujevac).

Notes

The Fifth Serbian International Conference on Applied Artificial Intelligence (SICAAI)