

IEEE-ICCA 2025

LIST OF BEST PAPERS SELECTED INDEPENDENTAL FROM THREE EXPERTS

#453 1571231513: LES-Quality Urban Wind and NFZ Mapping via Multi-Fidelity Learning.

Authors: Jinan Charafeddine (Engineering School Léonard de Vinci - ESILV, France); Aminallah Rabia (De Vinci Research Center, France); Hani Hamdan (Centrale Supélec, Université Paris-Saclay, France); Samir Yahiaoui (De Vinci Research Center, France)

Abstract: Reliable UAV operation in dense urban environments requires fast yet accurate estimation of building-induced turbulence. Conventional RANS/URANS simulations underpredict critical shear and TI/TKE peaks, whereas LES delivers high fidelity but remains computationally impractical for real-time use. This work introduces a data-driven, deep-learning multi-fidelity surrogate that upgrades low-cost CFD predictions into LES-quality flow fields and turbulence indicators relevant for NFZ assessment. The model learns cross-fidelity corrections from heterogeneous datasets spanning numerical simulations, wind-tunnel experiments, field measurements, and laboratory setups. Across all conditions, the proposed surrogate substantially improves turbulence estimation and hazard classification while operating in real time. These results demonstrate the potential of learning-based multi-fidelity approaches for reliable, turbulence-aware UAV guidance in complex urban environments.

Keywords: Urban turbulence; UAV safety; Multi-fidelity surrogates; Physics-guided learning; Turbulence intensity; No-fly- zone prediction

#46 1571167155: Enhancing Deep Learning-Based Fall Detection Under Non-Dominant Paths Using Wi-Fi CSI

Author: Israa Bayad (Université de Technologie de Troyes, France & Université Libanaise, Lebanon); Farah Mourad-Chehade and Sandy Mahfouz (Université de Technologie de Troyes, France); Khoulood Samrouth (Lebanese University, Lebanon); Hassan Amoud (Lebanese University & Faculty of Sciences, Lebanon)

Abstract: Falls are a leading cause of injury and hospitalization among older adults, especially in unsupervised settings. Wi-Fi Channel State Information (CSI) has emerged as a non-intrusive alternative to wearables and cameras, yet deep learning-based CSI models often struggle to generalize across different indoor environments due to spatial sensitivity and signal variability. This paper introduces a compact and efficient 2D-CNN model designed for robust fall detection across both familiar and unseen spaces, including challenging non-dominant path (NLoS) scenarios. The model is supported by a suite of targeted data augmentation techniques to simulate signal distortions, with asymmetric signal mixing playing a key role in enhancing intraclass variability and improving generalization. Evaluated on the ENetFall dataset across four environments, the proposed model outperforms larger EfficientNet baselines in difficult conditions, achieving up to 7% higher accuracy with 96% fewer parameters. These results demonstrate the feasibility of deploying CSI-based fall detection systems in real-world, multi-room settings.

Keywords: fall detection; WiFi sensing; lightweight deep learning; Channel State Information (CSI); adaptive systems

**#120 1571198030: Identifying Core Innovation Dimensions via Semantic Clustering
with a Novel Pairwise Validation Metric**

Authors Sarah Malaeb (Université Paris 8, France); Imad Saleh (University Paris 8, France); Maroun Jneid (Université Paris 8, France)

Abstract: Innovation capability frameworks remain fragmented, with inconsistent taxonomies that constrain maturity assessment and comparability across academic, industrial, and policy contexts. Addressing this challenge requires the systematic identification of core innovation dimensions from heterogeneous textual sources. This task is difficult, as it relies on clustering short, sparse, and semantically ambiguous texts, where ground-truth labels are absent and existing validation metrics fall short. Traditional internal and external indices rely on geometric properties or labeled data, while semantic measures improve interpretability but fail to assess cross-method stability. To overcome this gap, we propose the Pairwise Clustering Match Ratio (PCMR), a novel validation metric that quantifies cross-method agreement. PCMR serves a dual role: as a validation tool assessing stability and interpretability of clustering results, and as a model selection tool identifying the most robust algorithm. Applied within a hybrid methodology combining unsupervised clustering (K-means, Agglomerative Hierarchical Clustering) with a semi-supervised semantic classification using Sentence-BERT embeddings and cosine similarity, PCMR achieved a global mean score of 0.79 and identified K-means as the best-performing method. Through iterative validation and pruning of weak categories, we consolidated a stable set of 14 core innovation dimensions. These findings provide a semantically coherent, empirically validated dimensional structure to support unified innovation taxonomies, organizational maturity assessment, and the evaluation of innovation-related educational programs.

Keywords: innovation capabilities; semantic clustering; clustering validation; innovation maturity assessment; semi-supervised machine learning; natural language processing

#53 1571172891: Enhancing Low-Frequency Trading with Rule-Guided Reinforcement Learning

Authors Mohammed Ibrahim El-hajj (Arab Open University - Lebanon (AOU), Lebanon)

Abstract: The increasing integration of Machine Learning (ML) into financial trading has predominantly focused on high-frequency strategies reliant on raw data processing, often at the cost of computational efficiency and interpretability. This work explores an alternative paradigm for low-frequency trading by combining rule-based structural features—such as pivot points and support/resistance levels—with adaptive Reinforcement Learning (RL). While many RL-driven trading systems operate as opaque "black boxes," struggling with generalization and regulatory scrutiny, this approach embeds domain-specific knowledge directly into the agent's state representation and action space, enhancing transparency without sacrificing adaptability. A Proximal Policy Optimization (PPO) agent was trained on EUR/USD forex data (2015-2022) and evaluated on out-of-sample data (2023-2024), leveraging constrained action mechanics and dynamic reward shaping aligned with market structure. The agent achieved a Sharpe ratio of 4.41, a Calmar ratio of 6.69, and a maximum drawdown of 1.53%, demonstrating robust risk-adjusted performance. Key innovations include pivot-driven state encoding, limit-order execution logic, and asymmetric payoff profiles that balance frequent small losses from stop-loss triggers with infrequent but substantial gains. By reducing computational costs by 70-80% compared to high-frequency systems, this framework offers a practical solution for real-world deployment, bridging the divide between rigid rule-based strategies and purely data-driven models. Results validate that structural priors improve generalization in volatile markets while maintaining profitability, suggesting a path toward more resilient and auditable trading systems for low-frequency regimes.

Keywords: Reinforcement Learning (RL); Low-Frequency Trading (LFT); Rule-Based Systems

Multi-Dimensional Tail Classification in Procurement Spend Analysis: A Data-Driven Approach to Strategic Sourcing Optimization

Authors Magzhan Kairanbay (Array Innovation, Bahrain); Najma Ghuloom Taqi (Arab Open University Bahrain, Bahrain); Basaveswar reddy Mora Reddy (Arab Open University, Bahrain & Aluminium Bahrain, Bahrain); Ahmed Bahar (Aluminium Bahrain, Bahrain)

Abstract: Effective procurement tail spend management represents a significant opportunity for organizational cost optimization, yet identifying tail spend categories remains challenging due to the multi-dimensional nature of procurement data. This study presents a comprehensive analytical framework for classifying procurement tail spend using three distinct methodologies applied to a large industrial dataset of 24,983 transactions. Our analysis reveals that 88.5% of transactions exhibit tail characteristics under at least one classification method, with significant overlap across multiple dimensions. By employing F1-score evaluation, we identify 21 critical material groups representing 38.59% of transactions but only 9.56% of total spend value, achieving an F1-score of 0.65 and IoU of 0.48 against domain expert classifications. These findings demonstrate the importance of multi-dimensional analysis in procurement optimization and provide actionable insights for strategic sourcing initiatives in the aluminum manufacturing industry.

Keywords: Tail spend management; procurement analytics; multi-dimensional classification; F1-score optimization; supply chain optimization; aluminium industry

#254 1571222809: Cross-Agent Multimodal Provenance-Aware Framework for Robust Prompt Injection Defense in Large Language and Vision-Language Models

Authors: Toqeer Ali and Mishal Ateeq Almutairi (Islamic University of Madinah, Saudi Arabia); Mahmoud Abdelmajed Ahmed, Sr (Arab Open University Bahrain, Bahrain & Arab Open University, Bahrain)

Abstract: Language Models (LLMs) and Vision- Language Models (VLMs) have achieved remarkable progress in multimodal understanding, reasoning, and autonomous task execution. However, their openness to natural language and visual inputs exposes them to an emerging class of security threats known as multimodal prompt injection (PI) attacks. These attacks embed hidden or adversarial instructions not only in text but also within images-through captions, metadata, or steganographic patterns-to override model policies, exfiltrate sensitive information, or manipulate reasoning processes. This paper proposes a Cross-Agent Multimodal Provenance-Aware Framework that unifies text and image sanitization, multi-agent cross-validation, and provenance tracking to detect and mitigate multimodal prompt injection threats. The architecture comprises a Text Sanitizer Agent, a Visual Sanitizer Agent, a Main Multimodal Task Model, and an Output Validator Agent, all coordinated via a shared provenance ledger that records modality, source, and trust-level metadata. Experimental evaluations demonstrate significant improvements in injection detection accuracy and cross-modal trust consistency compared to single-model or text only defense approaches. By extending provenance tracking and validation across multiple modalities, the framework contributes to the advancement of explainable, accountable, and trustworthy multimodal AI security.

Keywords: Prompt Injection; Multi-Agent Systems; Provenance Tracking; LLM Security; Trust Validation; AI Safety

#261 1571223670: Model Context Protocol Business Analyst (MCP-BA): A Governed and Explainable Framework for Enterprise Analytics

Authors: Tejas Pravinbhai Patel (Independent Researcher, IEEE); Sandeep Shivam (Tavant Technologies, USA); Chaitanya Kulkarni (Oracle America Inc, USA); Chandrashekhar Medicherla (Salesforce Inc, USA); Amit Kumar Padhy (University of Illinois Urbana-Champaign, USA & Adobe Inc., USA); Bharadwaj Vulugundam (Oracle, USA)

Abstract: Large Language Models (LLMs) have shown strong potential in enterprise analytics, but their deployment remains limited by ungoverned data access, inconsistent context alignment, and missing audit trails. This paper introduces the Model Context Protocol Business Analyst (MCP-BA) - a governed reasoning framework that operationalizes the Model Context Protocol (MCP) to connect user intent, organizational data, and analytical tools under identity-aware, policy-controlled execution. The architecture combines a crawler-vector layer for continuous context refresh with an audit-enforced reasoning loop that transforms LLMs into compliant analytical micro-services. A synthetic enterprise evaluation demonstrated substantial gains over RAG and Atlas baselines - improving context precision by 17 points, governance accuracy by 42 %, and audit completeness to 100 %, while reducing reasoning latency by 29 %. The proposed system establishes a repeatable, explainable design pattern for high-trust enterprise AI, bridging intelligent automation with corporate accountability. Experimental validation shows that embedding MCP into reasoning workflows yields measurable benefits in precision, traceability, and compliance, confirming that explainability can act as an operational efficiency multiplier when engineered as a first-class design principle.

Keywords: Model Context Protocol; Governance and Explainable AI; Enterprise Analytics; Vector Databases; Auditability; Responsible AI

#363 1571229590: The Genome as a Quantum Error-Correcting Code

Authors: Irfan Alam Irfan (Arab Open University Bahrain, Bahrain & Computer Studies, Bahrain); Sultan Alam and Anish Pawar (VIT Bhopal University, India); Mohammad Abdussami (SRM University AP, India)

Abstract: DNA replication upholds genomic integrity with remarkable precision, rectifying inaccuracies that emerge from quantum level phenomena. This paper introduces a novel framework that models this biological process using the formalism of quantum error correction. We propose that genetic information encoded in DNA base pairs can be represented as the code space of a Biological Stabilizer Code, where cellular proofreading and mismatch repair mechanisms function as an active error correction protocol. We provide comprehensive proofs that illustrate the code's capacity to preferentially rectify the most likely types of mutations, thereby guaranteeing elevated logical fidelity. We establish a biological equivalent of the quantum threshold theorem, demonstrating that high-fidelity replication can be achieved if the inherent mutation rate falls below a critical threshold. The performance of the BSC is corroborated through simulations conducted in Qiskit, which reveal its superior efficacy compared to symmetric codes under biologically plausible noise conditions. This research establishes a quantitative relationship between the principles of fault-tolerant quantum computation and the mechanisms that uphold the stability of life's genetic code.

#471 1571231638: SPIFFE-Based Zero-Trust Authentication for AI Agent Ecosystems

Authors **Karthik Pappu (Dakota State University, USA); Badal Bhushan (I.E.T. M.J.P. Rohilkhand University, India); Akshay Mittal (University of the Cumberlands, USA)**

Abstract: AI agent ecosystems commonly rely on API key-based authentication, which introduces security risks related to credential exposure, manual rotation procedures, and weak guarantees of workload identity. This paper presents a zero-trust authentication framework that eliminates static API keys by using SPIFFE (Secure Production Identity Framework for Everyone) as a certificate-based identity layer for AI agents. We implement a multi-agent security pipeline coordinated through a central orchestrator, where each agent receives a short-lived X.509 SVID automatically issued and rotated by SPIRE. All inter-agent communication is authenticated through mutual TLS using verifiable SPIFFE identities, eliminating long-lived secrets from the system. Experimental results show that the architecture maintains continuous operation under frequent certificate rotation while preserving end-to-end authentication integrity. The findings demonstrate that SPIFFE-based workload identity provides a practical and secure alternative to API keys for AI agent frameworks, improving security posture, operational automation, and reliability.

Keywords: SPIFFE; zero-trust authentication; AI agents; workload identity

#119 1571198022: Kinect-Based Markerless Motion Capture for Telerehabilitation: A Comparative Review of Technologies, Clinical Applications, and Research Gaps

Authors: Hawraa Hotiet (University of Genova, Italy); Alaa Wehbe (Università degli Studi di Genova, Italy); Mohamad Abou Ali (University of the Basque Country, Spain & Lebanese International University, Lebanon); Abdallah Kassem (Notre Dame University, Lebanon); Camilla Baffigo (University of Genoa, Italy); Silvana G. Dellepiane (University of Genova, Italy)

Abstract: Telerehabilitation has become a pivotal modality for expanding access to motor therapy, particularly in remote and underserved environments. Markerless motion capture (MMC) technologies-most notably those leveraging Microsoft Kinect-offer a scalable, cost-effective alternative to traditional sensor-based systems, enabling real-time movement analysis without the encumbrance of wearable devices. This paper provides a comprehensive review of Kinect-based MMC systems across three hardware generations (Kinect v1, v2, and Azure Kinect), and benchmarks their capabilities against alternative platforms such as Intel RealSense and OpenPose. Clinical applications are systematically examined across diverse neurorehabilitation contexts, including stroke recovery and cerebral palsy, with a focus on insights from randomized controlled trials and validation studies. The review also highlights persistent challenges, including sensor precision limitations, lack of interoperability standards, and barriers to generalizable machine learning. In response, we propose forward-looking strategies encompassing federated learning, standardized datasets, and the integration of clinically interpretable AI. The paper concludes with actionable recommendations to accelerate the clinical integration and translational impact of MMC systems in next-generation telerehabilitation.

Keywords: Telerehabilitation; Markerless Motion Capture; Kinect; Virtual Reality; Stroke Rehabilitation; Cerebral Palsy