ICAISET2026

Conference »Program

2026 International Conference on Artificial Intelligence, Systems, and Emerging Technologies (ICAISET) Program

Time (Cairo)	ROOM-A	ROOM-B	ROOM-C	ROOM-D	ROOM-E	COFFEE_ROOM	MAIN_ROOM	Elsewhere
Tuesday, April 21								
08:30-09:00								**REGISTRATI** *REGISTRATION* REGISTRATION
09:00-10:00							**OPENING CEREMONY**: *OFFICIAL OPENING CEREMONY* OPENING CEREMONY	
10:00-10:30							**KEYNOTE-1**: *DISTINGUISHED KEYNOTE SPEAKER: PROFESSOR DR. MOHAMAD SAWSAN - Westlake University, China*	

Time (Cairo)	ROOM-A	ROOM-B	ROOM-C	ROOM-D	ROOM-E	COFFEE_ROOM	MAIN_ROOM	Elsewhere
10:30-11:00						**NETWORKING**: *COFFEE BREAK & POSTER SESSION* STUDENTS POSTER SESSION (COMPETITION)		
11:00-11:30							**KEYNOTE-2**: *PROFESSOR DR. Wael WAEL BADAWY, Ph.D., P.Eng.* KEYNOTE-2	
11:30-12:00	**SESSION-1A**: *RESEARCH: UAV-Assisted Networking, Deployment Optimization & Aerial Applications* 6 PAPERS	**SESSION-1B**: *RESEARCH: IoT-Based Smart Systems & Environmental Monitoring* 6 PAPERS	**SESSION-1C**: *RESEARCH: Edge AI, Federated Learning & Distributed Intelligence Architectures* 6 PAPERS	**SESSION-1D**: *RESEARCH: : Cybersecurity: Threat Detection, Malware & Intrusion Prevention* 6 PAPERS	**SESSION-1E**: *RESEARCH: Security Frameworks, Critical Infrastructure Protection & Post-Quantum* 6 PAPERS			
12:00-13:00								
13:00-14:00						**LUNCH**: *LUNCH (ON SITE)* LUNCH		
14:00-15:30	**SESSION-2A**: *RESEARCH: Advancements in Intelligent Systems for Image Analysis,*	**SESSION-2B**: *RESEARCH: Innovative Approaches in AI-Driven Systems for*	**SESSION-2C**: *RESEARCH: Deep Reinforcement Learning &*	**SESSION-2D**: *RESEARCH:Geospatial Intelligence, Remote Sensing &*	**SESSION-2E**: *RESEARCH: Accessibility, Assistive Technology &*			

Time (Cairo)	ROOM-A	ROOM-B	ROOM-C	ROOM-D	ROOM-E	COFFEE_ROOM	MAIN_ROOM	Elsewhere
	Cybersecurity, and Educational Technologies 6 PAPERS	*Cybersecurity, Cloud Optimization, and Healthcare Surveillance* 6 PAPERS	*Evolutionary Optimization* 6 PAPERS	*Environmental Data Analytics* 6 PAPERS	*Sign Language AI* 6 PAPERS			
15:30-16:00						*COFEE BREAK & NETWORKING* COFEE BREAK ROOM - NETWORKING		
16:00-18:00	**SESSION-3A**: *RESEARCH: Frameworks and Innovations in Governance, AI Maturity, and Decision-Making in Complex Systems* 8 PAPERS	**SESSION-3B**: *RESEARCH: Innovations in AI Governance, Security, and Testing Frameworks for Cloud and IoT Systems*	**SESSION-3C**: *RESEARCH: Advancements in AI for Financial Risk, Healthcare Detection, and Cloud Data Governance* 8 PAPERS	**SESSION-3D**: *RESEARCH: Integrating AI for Fairness, Automation, and Security in Cloud and DevOps Environments* 8 PAPERS	**SESSION-3E**: *RESEARCH: Innovative Approaches in AI for Quality Assurance, Cybersecurity, and Predictive Analytics* 8 PAPERS			

Wednesday, April 22

Time (Cairo)	ROOM-A	ROOM-B	ROOM-C	ROOM-D	ROOM-E	COFFEE_ROOM	MAIN_ROOM	Elsewhere
08:30-10:00								
10:00-10:30							**KEYNOTE-3**: *KEYNOTE-3: ENG. SANDEEP SHIVAM, USA* KEYNOTE-3	
10:30-11:00						**NETWORKING**: *COFFEE-BREAK &POSTER SESSION* STUDENTS POSTER SESSION (COMPETITION)		**REGISTRATI** *REGISTRATION* REGISTRATION
11:00-11:30							**KEYNOTE-4**: *KEYNOTE-4: ENG. TEJAS PATEL - AMAZON, USA* KEYNOTE-4	

Time (Cairo)	ROOM-A	ROOM-B	ROOM-C	ROOM-D	ROOM-E	COFFEE_ROOM	MAIN_ROOM	Elsewhere
11:30-12:00 12:00-13:00	**SESSION-4A**: *RESEARCH:Renewable Energy, Smart Grids & Environmental Sustainability* 6 PAPERS	**SESSION-4B**: *RESEARCH: Robotics, Autonomous Systems & Adaptive Control* 6 PAPERS	**SESSION-4C**: *RESEARCH: Machine Learning for Disease Detection & Biomedical Classification* (6 PAGES)	**SESSION-4D**: *RESEARCH: Explainable AI (XAI): Methods, Interpretability & Trustworthiness* (6 PAGES)	**RSESSION-4E**: *RESEARCH: Wireless Communication: NOMA, RSMA & Multi-Access Schemes* 6 PAPERS			
13:00-14:00						**LUNCH**: *LUNCH (ON_SITE)* LUNCH		
14:00-15:30	**SESSION-5A**: *RESEARCH: Computer Vision: Object Detection, Recognition & Surveillance* 6 PAPERS	**SESSION-5B**: *RESEARCH:Generative AI, Digital Education & Societal Impact* 6 PAPERS	**SESSION-5C**: *RESEARCH: Large Language Models, RAG & Agentic AI Systems* (6 PAPERS)	**SESSION-5D**: *RESEARCH: Deepfake, Fake News, Spam & Anti-Phishing Detection* (6 PAPERS)	**SESSION-5E**: *RESEARCH: AI for Clinical Decision Support & Healthcare Risk Assessment* (6 PAPERS)			
15:30-15:45						*COFEE BREAK & NETWORKING* COFEE BREAK ROOM - NETWORKING		

Time (Cairo)	ROOM-A	ROOM-B	ROOM-C	ROOM-D	ROOM-E	COFFEE_ROOM	MAIN_ROOM	Elsewhere
15:45-18:00	**SESSION-6A**: *RESEARCH: Innovations in AI for Secure Identity Verification, Cost Optimization, and Enhanced Analytics* 9 PAPERS	**SESSION-6B**: *RESEARCH: Innovative Strategies in AI for Predictive Modeling, Compliance, and System Reliability* 9 PAPERS	**SESSION-6C**: *RESEARCH:Advanced Frameworks for Runtime Verification, AI Workloads, and Intelligent System Orchestration* 9 PAPERS	**SESSION-6D**: *RESEARCH: AI-Enhanced Governance, Data Reliability, and Intelligent Systems for Modern Applications* 9 PAPERS	**SESSION-6E**: *RESEARCH: AI Innovations in Digital Governance, Data Management, and Decision Support Systems* 9 PAPERS			

Time (Cairo)	ROOM-A	ROOM-B	ROOM-C	ROOM-D	ROOM-E	COFFEE_ROOM	MAIN_ROOM	Elsewhere
Thursday, April 23								
09:00-11:15	**SESSION-7A**: *RESEARCH:Agricultural, Food Science & Biological Applications of ML* 9 PAPERS	**SESSION-7B**: *RESEARCH:NLP, Sentiment Analysis & Domain-Specific AI Applications* 9 PAPERS	**SESSION-7C**: *RESEARCH:Workforce Analytics, Social Impact & Organizational Decision Support* 9 PAPERS	**SESSION-7D**: *RESEARCH:AI in Education: Learning Analytics, Prediction & Student Success* 9 PAPERS	**SESSION-7E**: *RESEARCH: Machine Learning for Business Intelligence, Finance & Fraud Detection* 10 PAPERS			
11:15-11:30								

Tuesday, April 21

Tuesday, April 21 8:30 - 12:00 (Africa/Cairo)

REGISTRATION: REGISTRATION

REGISTRATION

Tuesday, April 21 9:00 - 10:00 (Africa/Cairo)

OPENING CEREMONY: OFFICIAL OPENING CEREMONY

OPENING CEREMONY

Room: MAIN_ROOM

Tuesday, April 21 10:00 - 10:30 (Africa/Cairo)

KEYNOTE-1: DISTINGUISHED KEYNOTE SPEAKER: PROFESSOR DR. MOHAMAD SAWSAN - Westlake University, China

Room: MAIN_ROOM

Tuesday, April 21 10:30 - 11:00 (Africa/Cairo)

NETWORKING: COFFEE BREAK & POSTER SESSION

STUDENTS POSTER SESSION (COMPETITION)

Room: COFFEE_ROOM

Tuesday, April 21 11:00 - 11:30 (Africa/Cairo)

KEYNOTE-2: PROFESSOR DR. Wael WAEL BADAWY, Ph.D., P.Eng.

KEYNOTE-2

Room: MAIN_ROOM

Tuesday, April 21 11:30 - 13:00 (Africa/Cairo)

SESSION-1A: RESEARCH: UAV-Assisted Networking, Deployment Optimization & Aerial Applications

6 PAPERS

Room: ROOM-A

Chair: Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

11:30 ***A Preliminary Trajectory-Driven Framework for Queue Length and Control Delay Estimation at Multilane Roundabouts Using Drone Video Analytics***

Elawady Mohammed, Khaled Hamad and Lubna Obaid (University of Sharjah, United Arab Emirates)

This paper presents a preliminary trajectory driven framework for estimating queue length and control delay at multilane roundabouts using unmanned aerial vehicle based video data. The proposed approach is motivated by the limitations of conventional field data collection methods, which often provide incomplete spatial coverage and do not capture the full spatiotemporal evolution of queues. The framework integrates drone video acquisition, vehicle detection and tracking, geometric calibration, lane-based trajectory organization, and equation-based procedures for deriving operational performance measures. Queue length is estimated through back-of-queue identification along a defined lane study segment, while control delay is computed directly from vehicle speed profiles using a threshold-based formulation. To examine feasibility, the framework is applied in a preliminary manner to structured trajectory data obtained for Sheikh Abdul Kareem Al Bakri Square in Sharjah, United Arab Emirates. Initial implementation confirms that the workflow can generate lane-level queue and delay outputs from UAV derived trajectories. However, the current outputs remain preliminary and indicate the need for further calibration of queue membership logic, lane-based ordering, and delay thresholds. The study demonstrates the potential of drone video analytics as a scalable tool for roundabout operational assessment and provides a methodological basis for future validation and refinement.

11:45 ***Robust Energy-Causal Design for an FSO-Powered UAV SWIPT Relay in Disaster Recovery***

Aya Teguia (De Vinci Higher Education, France); Wafaa Mohammed Ridha Shakir (Al-Furat Al-Awsat Technical University, Iraq); Zeinab Mhanna (Ecole Supérieure d'Ingénieurs Léonard de Vinci, France)

Post-disaster communication outages are often accompanied by severe energy shortages at ground terminals, making conventional unmanned aerial vehicle (UAV) relays short-lived and operationally fragile. This paper investigates an FSO-powered UAV relay that harvests optical power from a ground optical base station (OBS) and provides radio frequency (RF) downlink service with simultaneous wireless information and power transfer (SWIPT) to multiple users. The OBS-UAV feeder link is modeled with unit-consistent atmospheric attenuation, Gamma-Gamma (GG) scintillation, and divergence-dependent pointing loss. To prevent energy starvation caused by deep irradiance fades, we replace the stochastic harvested inflow with a

deterministic risk-aware certainty-equivalent derived from fractional moments of the GG irradiance, resulting in a conservative energy-causality constraint under explicit battery dynamics that account for propulsion and circuit power. We formulate a max-min rate-energy fairness problem over a time-slotted 3D trajectory, per-slot divergence angle, RF transmit power, and receive-side splitting. The resulting nonconvex program is solved by alternating optimization: (i) per-slot divergence maximization of the expected pointing coupling, (ii) convex SWIPT resource allocation under fixed geometry, and (iii) safeguarded successive convex approximation (SCA) for trajectory refinement. Monte Carlo evaluation under scintillation shows substantial gains in 5th-percentile worst-user rate and harvested energy over mean-inflow and fixed-divergence baselines, while maintaining mission-feasible battery operation.

12:00 ***Digital-Twin and Ray-Casting-Supervised Surrogate Learning for Accelerated Radio-Map Generation and UAV-BS Placement Optimization in mmWave Campus Networks***

Karar Hamza Hussein (Al-Furat Al-Awsat Technical University, Iraq); Wasan Kadhim Saad (Engineering Technical College-Najaf, Iraq)

Accurate radio-map generation and UAV basestation (UAV-BS) placement planning in mmWave bands are strongly affected by site-specific blockage, yet high-fidelity 3D visibility evaluation can be computationally prohibitive at campus scale. This paper presents a digital-twin, ray-casting- supervised surrogate-learning framework to accelerate radiomap generation and placement optimization for UAV-assisted mmWave networks in a real campus environment. A 3D mesh of the Najaf Technical Institute campus is used to build a visibility-aware ground-truth oracle that determines LoS/NLoS via mesh-based ray casting and computes path loss using a link-budget model (FSPL with an NLoS penalty). Using oraclelabeled samples, a ray-free regression surrogate based on Extremely Randomized Trees is trained under group-split evaluation to generalize to unseen UAV placements, achieving MAE = 3.751 dB and RMSE = 5.841 dB on held-out placement groups. The trained surrogate enables fast screening of candidate UAV locations and supports coverage-driven and cell-edge-driven objectives, followed by Top-K full-grid oracle verification for reliability. At h = 200 m and γ = 5 dB, the celledge- optimized placement increases the 5th-percentile rate from 19.29 Mbps (center baseline) to 34.97 Mbps while maintaining near-saturated coverage (≈ 89.4%). Hierarchical surrogate planning reduces end-to-end search time from 600.36 s (wide surrogate grid-search) to 170.76 s (planning + oracle verification), and cuts expensive full-grid oracle evaluations by 391× via Top-K verification. These results demonstrate that digital-twin supervision coupled with surrogate acceleration provides a practical pathway for campus-scale mmWave UAV-BS planning with oracle-verified performance guarantees.

12:15 ***Saving Lives at Sea: Deep Learning-Based Aerial Detection for Search and Rescue Missions***

Mostafa Rizk (Lebanese University, Lebanon); Abbas Rammal (Phoenicia University, Lebanon); Amer Baghdadi (IMT Atlantique, France)

Maritime search and rescue (SAR) operations are vital in high-risk regions such as the Mediterranean Sea, where migrant incidents and maritime accidents occur frequently. Rapid and accurate detection of humans and vessels is essential for timely intervention and enhanced situational awareness. This study evaluates the performance of YOLOv7 and YOLOv8 models for detecting humans and boats in maritime environments. Experimental results indicate that YOLOv8 variants achieve superior generalization, with YOLOv8x attaining the highest mAP@50-95 and YOLOv8m offering a favorable balance between detection accuracy and computational cost. Meanwhile, YOLOv7 demonstrates the highest precision and recall, confirming its robustness in precision-critical scenarios. Lightweight models, including YOLOv8n and YOLOv7-tiny, exhibit strong potential for real-

time deployment on resource-constrained platforms. Overall, the findings highlight the effectiveness of vision-based intelligent systems in enhancing SAR efficiency and provide valuable insights for their integration with aerial and satellite surveillance frameworks.

12:30 ***UAV-Based SDR Signal Warfare: Experimental Constraints and Simulation-Based Evaluation***

Yousef AlMahmoud, Nabil Litayem, Mustafa Al Samara and Abdelhak Belhi (Joaan Bin Jassim Academy for Defence Studies, Qatar)

The increasing reliance on cellular communication infrastructures expands the cyber-electromagnetic attack surface, introducing new security risks in both civilian and defense contexts. This paper investigates the feasibility and effectiveness of UAV-enabled Software-Defined Radio (SDR) signal warfare targeting LTE-like networks. A two-phase methodology is adopted. First, controlled laboratory prototyping with low-cost SDR hardware identifies key operational constraints, including half-duplex limitations and power restrictions. Second, a MATLAB-based simulation framework evaluates a multi-stage attack chain-comprising jamming, interception, rogue base-station takeover, and credential harvesting-across 60 urban scenarios. Results show that direct jamming achieves limited disruption under realistic constraints, whereas higher-layer attacks, particularly rogue base-station association and credential harvesting, achieve high success rates, leading to large-scale user compromise even at moderate transmission power levels. These findings highlight the critical role of control-plane exploitation over pure radio-frequency interference. This work provides a reproducible evaluation framework and emphasizes the need for enhanced authentication mechanisms, RF-aware monitoring, and counter-UAV defenses to mitigate emerging aerial SDR threats.

12:45 ***Robust-APP: Adversarial-Robust Adaptive Path Planning for Noisy and Congested Curved Highway Environments***

Ehsan Wadood (Moscow Institute of Physics and Technology, Russia); Umer Mukhtar Andrabi (National Research University Higher School of Economics, Russia); Ilya Voronkov and Alexander Kharlamov (Moscow Institute of Physics and Technology, Russia); Sanjeev Kumar Ojha and Mahi Sharma (Manipal University Jaipur, India)

Autonomous vehicles have significant obstacles in curving highway settings characterized by sensor noise and elevated traffic congestion, with human error accounting for 94% of accidents. This research advances the Adaptive Path Planning framework to Robust-APP, a hybrid deep inverse reinforcement learning system that combines curvature-aware Markov Decision Processes with adversarial resilience. We present (i) adversarially robust state representation with certified ϵ-bounds through interval propagation, (ii) traffic-adaptive state compression with O(N log N) complexity utilizing KD-tree nearest-neighbor queries, (iii) robust MaxEnt inverse reinforcement learning employing projected gradient descent for adversarial training on reward and trajectory perturbations, and (iv) a distributed rollout architecture facilitating real-time simulation of over 120 vehicles.

Thorough SUMO-based assessments exhibit equivalent performance to the documented baselines while enhancing robustness: Robust-APP achieves an 89.7% overall success rate (24.2% higher than DQN's 72.3%), 92.1% success on curved segments, 18.4% lower path deviation, and 27% faster per-episode computation. Under sensor noise up to $\epsilon = 0.3$, it maintains 82% success (versus 39% for DQN). In dense traffic (2.5 veh/km/lane), success remains 89% (versus 25% for DQN). Computation scales to 0.5 s/episode at 120 vehicles while staying near real-time thresholds. These breakthroughs provide secure, scalable autonomous navigation on noisy and congested curved roadways.

Tuesday, April 21 11:30 - 13:00 (Africa/Cairo)

SESSION-1B: RESEARCH: IoT-Based Smart Systems & Environmental Monitoring

6 PAPERS

Room: ROOM-B

11:30 ***Edge-AI Driven Coastal Flood Early Warning System Using IoT-Satellite Data Fusion***

Samir Haddad (University of Balamand, Lebanon); Kassem Danach, Sr (Al Maaref University, Lebanon & Chairperson, Lebanon); Jinane Sayah (University of Balamand, Lebanon); Khouloud Eledlebi (Abu Dhabi University, United Arab Emirates); Joseph Merhej (Facultyof Sciences II, Lebanon); Chadi Kallab (Lebanese American University, Lebanon)

Island communities face escalating risks from coastal flooding due to climate change and sea-level rise, yet conventional warning systems fail to provide timely, accurate forecasts. To address this gap, we developed and deployed a novel AI-driven early warning system that fuses multi-source data from IoT sensors and satellite remote sensing within an edge-computing framework. By processing data in real-time, our system powers LSTM neural networks that generate probabilistic flood risk maps 6-72 hours in advance. A 12-month pilot study across three diverse coastal regions validated the system's performance, achieving 92.3% prediction accuracy and outperforming benchmarks by 15-20 percentage points. The system's high-fidelity forecasts enabled a 65% reduction in evacuation times and a 38% decrease in flood-related damages, with a 91% community satisfaction rate. Robust computational performance was evidenced by 4.7-second edge processing latency and 99.97% availability. This work provides a validated, scalable framework for climate adaptation, demonstrating how AI can significantly enhance disaster resilience for the world's most vulnerable coastal populations.

11:45 ***Design and Implementation of IoT-Based Flood Monitoring System with Web-Based Data Repository***

Jerry Incierto Teleron (Technological Institute of the Philippines - Quezon City, Philippines)

This study presents the design, development, and deployment of an Internet of Things (IoT)-based flood monitoring system for real-time river condition monitoring. The developed system integrates a NodeMCU microcontroller with environmental sensors, including the JSN-SR04T ultrasonic sensor for water level measurement and the DHT11 sensor for temperature and humidity monitoring. The collected environmental data are transmitted through a wireless network to an online database and visualized through a web-based monitoring dashboard that allows users to remotely monitor river conditions. The monitoring unit was installed beneath a bridge in a weather-protected enclosure to ensure reliable field operation. Calibration of the ultrasonic sensor was conducted using a reference water level gauge installed on the bridge footing, resulting in minimal measurement error. Functional testing confirmed that the sensors, microcontroller, database server, and web interface operated effectively, enabling continuous data acquisition, transmission, and visualization. User evaluation results indicated that the system is useful for flood preparedness and generally easy to use. The developed system demonstrates the practical application of IoT technology for real-time environmental monitoring and contributes to improving flood awareness and community-based disaster risk reduction.

12:00 ***Image Processing Model for Pechay and Chinese Pechay Health Detection Using Mobile IoT***

Jerry Incierto Teleron (Technological Institute of the Philippines - Quezon City, Philippines)

This study presents the design and development of a mobile Internet of Things (IoT)-based system integrated with an image processing model for detecting the health condition of pechay (Brassica rapa) and Chinese pechay leaves. The system utilizes a mobile application developed using MIT App Inventor to capture leaf images, which are transmitted to a cloud-based platform using Firebase for real-time processing. The image processing model, implemented in Python, consists of image acquisition, pre-processing, segmentation, feature extraction, and classification using a Support Vector Machine (SVM) classifier to determine whether a leaf is healthy or diseased. A dataset of 1000 labeled leaf images were used for model training and evaluation, with 200 images allocated for testing. The system achieved an overall classification accuracy of 92%, with a precision of 90.38%, recall of 94%, and F1-score of 92.14%. The average response time was 2.2 seconds, demonstrating efficient real-time performance. Usability evaluation results indicate that the application is user-friendly and suitable for practical agricultural use. The integration of IoT and image processing provides a reliable and scalable solution for early detection of plant diseases, contributing to improved crop management and smart agriculture practices.

12:15 ***IoT-Based Automated Mushroom House with Early-Stage Disease Detection***

Shivanshi Sharma, Tanvi Yadav, Kavneet Kaur and Vikas Upadhyaya (NIIT University, India)

This paper presents a fully automated IoT system for growing mushrooms and identifying diseases in their early stages by implementing real-time monitoring and cloud-based computer vision. All the essential elements such as temperature, humidity, moisture in soil, and CO_2 are continuously monitored via sensors which are regulated by ESP32, enabling closed-loop control. The images acquired by Raspberry Pi camera are classified and detected by the cloud-based machine learning algorithms ResNet50 and YOLOv8, respectively, which have provided high levels of precision and efficiency.

12:30 ***Toward Smart IoT-Enabled Marine Observatories: An Enterprise Architecture Modeling Approach***

Noura Mansour (Universidad del País Vasco, Spain); Fadi Dornaika (University of the Basque Country, Spain); Charbel Geryes Aoun (American University of the Middle East, Kuwait)

The design of Marine Observatories (MO) that rely on underwater sensor networks presents significant challenges due to the involvement of diverse expertise areas and the distribution of physical and logical components. Modeling mistakes at the design stage can result in expensive operational failures during deployment. In this paper, we propose an enhancement to the ArchiMO modeling tool by introducing Hardware-to-Software Assignment Constraints at the Business and Application layers of the ArchiMate Enterprise Architecture framework. The proposed constraints ensure compliance with the Distributed Fusion Architecture (DFA) principle by controlling the allocation of software functions to appropriate hardware nodes such as Smart Sensors and Fusion Servers. Through real-time connectivity with a Structured Relational Database, the tool automatically checks each assignment operation and prevents incorrect allocations by alerting the designer. The proposed approach is demonstrated through the MeDON (Marine e-Data Observatory Network) project for Marine Mammal localization, following a Model-Driven Engineering methodology. The constrained design model is compiled and verified through NS-3 Simulator, confirming correct runtime behavior.

12:45 ***IoT-Enabled Smart Medicine Container for Expiration Safety and Refill Alerts Using Weight and RFID/NFC Sensing***

Karar Hamza Hussein and Mohammed Jalil Mohammed Ali Alhasan (Al-Furat Al-Awsat Technical University, Iraq)

Medication nonadherence and unsafe home medication practices remain persistent challenges, particularly for patients managing multiple medicines over long periods. In parallel, household surveys show that expired or unnecessary medicines are commonly present in homes, increasing the likelihood of unintentional use when expiry dates are overlooked. This paper proposes a low-cost IoT-enabled medicine container that provides (i) proactive expiration warnings and (ii) remaining-quantity and refill notifications, and (iii) scheduled dose-time reminders with in-app patient acknowledgement based on weight sensing, with optional RFID/NFC or barcode-based registration to reduce user burden. The design adopts an edge-cloud architecture: an embedded node acquires and filters sensor readings, estimates remaining doses, detects consumption events, and synchronizes concise state updates to a cloud service that enforces alert policies and delivers notifications to the patient and (optionally) a caregiver. A practical methodology for dose estimation, depletion prediction, and alert triggering is presented, together with an implementation plan for a proof-of-concept prototype and evaluation protocol. The proposed system aims to unify safety (expiry awareness) and availability (refill readiness) in a single home-use workflow with minimal manual interaction.

Tuesday, April 21 11:30 - 13:00 (Africa/Cairo)

SESSION-1C: RESEARCH: Edge AI, Federated Learning & Distributed Intelligence Architectures

6 PAPERS

Room: ROOM-C

Chair: Golnaz Shahtahmassebi (Nottingham Trent University, United Kingdom (Great Britain))

11:30 ***Digital Sovereignty Beyond the Cloud: An Edge AI Framework for Small Island Developing States***

Mohamed Shareef (Nexia Maldives, Maldives)

Small Island Developing States face an architectural choice in artificial intelligence deployment. Most AI systems route computation through distant data centers, creating dependencies that fail under common island conditions. When submarine cables break, cloud-dependent systems stop working at the point when disaster response and public decision-making are most critical. This paper presents the SIDS Sovereign Edge Framework (SSEF), which combines resilient edge devices, compact language models, and indigenous knowledge systems. The analysis draws on an integrative review and comparative examination of deployments in Indonesia, India, and Pacific island nations. Across these cases, edge AI systems operate on hardware costing under $10 and continue functioning without internet connectivity. Autonomous geographic systems maintained disaster response during extended cable failures. Solar-powered diagnostic tools supported continuous healthcare operation. Hybrid forecasting that integrated traditional environmental indicators improved precipitation accuracy by 15 to 30 percent. These findings indicate that digital sovereignty through edge AI is

technically feasible for SIDS, but only where architectural choices prioritize local operation and task specificity rather than reliance on remote cloud platforms.

11:45 ***Toward Scalable AllReduce in AI Data Centers: A Cross-Layer Taxonomy and Survey***

Mariam Abdullah and Imtiaz Ahmad (Kuwait University, Kuwait); Sa'ed Abed (Kuwait University & College of Engineering and Petroleum, Kuwait)

Massive deep learning models rely on data-parallel training across thousands of accelerators, requiring frequent parameter synchronization via AllReduce. As model sizes scale to billions and trillions of parameters, communication has become the dominant bottleneck, limiting training efficiency and scalability. Existing optimization strategies often focus on isolated layers such as-algorithms or network fabrics-resulting in unpredictable performance in large, heterogeneous data centers. This survey presents a unified cross-layer perspective, introducing a taxonomy of AllReduce methods, analyzing cost models, algorithmic designs, and architectural factors, and comparing approaches using metrics such as scalability, latency, and throughput. It highlights systemic constraints and proposes an integrated path combining dynamic cost modeling, topology aware algorithms, in-network aggregation, and telemetry-driven scheduling, arguing that coordinated improvements across models, algorithms, and architectures are essential to overcome communication bottlenecks at the trillion-parameter scale.

12:00 ***KineticID: A Privacy-Preserving IoT Access Control System Utilizing Markerless Gait Biometrics and Edge Computing***

Afeera Firdoose (Middlesex University, United Kingdom (Great Britain)); Mervat Madi (American University of Dubai, United Arab Emirates); Judhi Prasetyo (Middlesex University Dubai, United Arab Emirates)

The push for passive biometrics to mitigate hygiene risks associated with touch-based sensors has particularly increased after the COVID-19 outbreak. While touch based sensors pose hygiene risks, alternatives like facial recognition introduce severe "honeypot" privacy risks and spoofing vulnerabilities. To solve this, KineticID, a low-cost Cyber-Physical System (CPS) that uses the way a person walks (their gait) as a secure, zero-contact key was developed. This paper proposes a system that uses the MediaPipe BlazePose framework, a system that extracts an abstract 33-landmark skeletal topology in real-time, functioning on standard CPU hardware. To ensure environmental robustness, a geometric algorithm calculates a scale-invariant biometric ratio (R=Stride/Height), allowing for consistent identification regardless of the subject's distance from the camera. KineticID processes all video data in volatile RAM, ensuring no personally identifiable images are stored, making it a "Privacy-by-Design" architecture. Integration between the Python-based vision engine and an ESP32 microcontroller is achieved via the lightweight MQTT protocol, enabling physical lock actuation with sub-100ms latency. Experimental results validate the system's efficacy as a hygienic, GDPR compliant, and cost-effective alternative to traditional biometrics.

12:15 ***Distributed Power Analytics: Leveraging Federated Harris Hawks Optimization for Confidential Energy Prediction***

Nader Bakir (Beirut Arab University, Lebanon); Ahmed M Gawish (Arab Open University, Saudi Arabia); Khawla El Hassan (Arab Open University Kuwait, Kuwait); Vadim Pak (Peter the Great Saint Petersburg Polytechnic University, Russia); Marcelle Merhy (Lebanese University, Lebanon); Abdallah El Chakik (Beirut Arab University, Lebanon)

Smart buildings consume 20% of global energy, necessitating accurate consumption prediction for operational efficiency. However, centralized machine learning requires aggregating sensitive occupancy and usage data, raising significant privacy concerns. Existing federated learning solutions preserve privacy but employ static optimization parameters that cannot adapt to the dynamic, heterogeneous nature of building energy patterns, resulting in suboptimal predictions. This paper presents FL-HHO, a novel federated learning framework that integrates Harris Hawks Optimization (HHO) for privacy-preserving energy prediction. By leveraging HHO's dynamic escaping energy mechanism, the framework adaptively balances exploration and exploitation to prevent convergence to local optima characteristic of multi-modal consumption patterns. Evaluation on the CU-BEMS dataset, featuring 18 months of real building data with 25.9% missing values, demonstrates that FL-HHO achieves a Mean Squared Error of 1.18. This outperforms existing federated methods while maintaining complete data locality. This advancement enables commercial buildings to deploy intelligent energy management without compromising occupant privacy, addressing a critical barrier to smart building adoption.

12:30 ***Micro-Solutions to Mega-Problems: A Design-Thinking Framework for Affordable Offline AI-IoT Access in Low-Resource Contexts***

Constantine Andoniou (Abu Dhabi University, United Arab Emirates)

The essential services of communication, education, healthcare, electricity and food security remain inaccessible to billions of people across the world. The implementation of large-scale technological systems, such as telecommunications towers, centralized cloud AI, grid-dependent devices and high-cost IoT deployments, proves impossible in low-resource areas because of financial, geographic and technical barriers. The current study presents a new approach through micro-solutions - small affordable AI-IoT devices - that function independently from power sources and need no specialized knowledge to operate in various cultural settings. The research develops micro-solutions through design thinking to create solutions for Edge AI, frugal fabrication and circular design which benefit underprivileged communities. The paper combines recent developments in resource-efficient AI systems with affordable IoT networks and sustainable product development to create a ten-step design-thinking framework which enables the development of micro-solutions for various industries. The design principles lead to develop multiple scalable low-complexity interventions through several examples. The research defines three fundamental requirements for offline AI systems and demonstrates methods to develop solutions through community participation, local manufacturing and policy development for maintaining micro-solution systems. Furthermore, it develops a technology solution framework which maintains accessibility and operational stability in limited resource settings through human-centered design instead of system-dependent infrastructure needs.

12:45 ***Resilient Edge Intelligence: Integrating Swarm Logic with Lightweight Agents and Localised SLMs***

Antony Lees (BPP University, United Kingdom (Great Britain))

The proliferation of Internet of Things (IoT) devices and the rapid growth of mobile edge computing have fostered a complex ecosystem where computational tasks are executed on devices with unpredictable availability. The intermittent connectivity prevalent in edge locations presents a significant challenge to conventional task scheduling and resource allocation models. This paper introduces a novel paradigm that integrates the Adaptive Swarm for Intermittent Edge Tasks (ASIET) - a bio-inspired swarm intelligence algorithm engineered to adapt dynamically within highly unreliable environments using ultra-lightweight artificial intelligence (AI) agent orchestrators (PicoClaw) and transported Small Language Models (SLMs).

Current distributed computing paradigms often struggle with the inherent unpredictability of edge environments. By encapsulating cognitive capabilities within self-transporting, single-binary agents, this architecture helps redefine systemic fault tolerance. The integration of ASIET's explicit, mathematically defined waiting periods allows the system to autonomously manage transient disconnections, whilst feedback-driven collective learning intelligently informs future task scheduling. This comprehensive study details the theoretical principles of the integrated tripartite system, outlines an empirical simulation leveraging a 1.5B parameter SLM on an ARM-architected edge gateway, and explores practical applications in remote, disaster, and mobile ad hoc environments, demonstrating a fundamentally resilient pathway for autonomous kinetic edge computing.

Tuesday, April 21 11:30 - 13:00 (Africa/Cairo)

SESSION-1D: RESEARCH: : Cybersecurity: Threat Detection, Malware & Intrusion Prevention

6 PAPERS

Room: ROOM-D

11:30 ***Extracting and Interpreting Emerging Cyber Threats from Dark Web Forums Using Transformer-Based NLP Models***

Moses Mupeta and Lukumba Phiri (University of Zambia, Zambia); Simon Tembo (University of Lusaka, Zambia)

The expansion of the dark web has proven to be a digital marketplace for the distribution of stolen data, malicious software, and hacker tools. For long-standing methods to detect and mitigate digital threats, the unorganized and vast communication structures of the dark web present unprecedented challenges. Time-sensitive, malicious, and coded communications, often embedded with slang and technical diversion, complicate matters even further.

This study designed a set of cyber threat extraction and cyber threat interpretation frameworks that leverage NLP and transformer models to peer into covert cyber threat discussions within dark web forums and dark web marketplaces. For the purpose of this research, a forums dataset of more than four million posts was created and analyzed. Over four million posts were analyzed in a computational study to build cyber threat models that identify and extract relevant discussions of data leak, cyber exploits, fraud, malware, and discussions relating to the sale of drugs, weapons, and other cyber threats. The BERT-Base transformer model achieved 73 percent accurate results for forum and threat-type classification, demonstrating the capacity of these models to classify the forum discussions with a high classification rate of 98 percent for malicious activity.

The designed models were used to construct the first working automated cyber intelligence system to analyze dark web data in real time. These results demonstrate the efficiency of transformer models in automation of cyber threat detection, cyber threat intelligence gathering, and integration into proactive cyber defense systems.

11:45 ***Feature-Based Machine Learning for Real-Time Detection of Malicious QR and NFC Content***

Eisa Mohammed Al-mannai, Nabil Litayem, Abdelhak Belhi and Mustafa Al Samara (Joaan Bin Jassim Academy for Defence Studies, Qatar)

Quick Response (QR) codes and Near Field Communication (NFC) tags are widely used in payments, transportation, retail, authentication, and public services. However, their convenience introduces a critical security challenge: users often interact with embedded links or actions without visibility into the underlying destination or content. This paper presents a practical web-based security scanning platform that performs pre-interaction analysis of QR and NFC content and classifies associated URLs as safe, suspicious, or malicious. The system integrates a browser-based frontend, a backend analysis service, feature-based URL processing, and Machine Learning (ML) classification, while external threat-intelligence services (e.g., VirusTotal) are used as a secondary verification layer. A dataset of 20,000 URLs is transformed into structured feature vectors using a standardized tri-value scoring scheme and used to evaluate nine ML models under a stratified 80/20 train-test split. Results show that XGBoost achieves the best performance with 90.15% accuracy, outperforming Random Forest (89.98%) and Multi- Layer Perceptron (89.53%), while also achieving the highest precision, recall, and F1-score values. The prototype is further validated in realistic QR/NFC usage scenarios using Androidbased browser testing, with iOS/device-policy and hardware compatibility limitations documented. The results demonstrate that proactive, browser-accessible screening can effectively enhance security in everyday QR/NFC interactions.

12:00 ***A lightweight CNN Model for Detecting Industrial Internet of Things (IIoT) Cyber-attacks***

Wahiba Abakker Mohammed Ismaiel and Dhabia Turki (Taif University, Saudi Arabia); Hind Muidh, Salha Mohammed and Muneera Fayhan (University College of Ranyah, Saudi Arabia); Omer Dawood (Prince Sattam Bin Abdulaziz University, Saudi Arabia)

Nowadays, the Industrial Internet of Things (IIoT) is a rapidly evolving technology with diverse applications. Its infrastructure comprises numerous sensors, controllers, industrial machines, and other devices connected to 5G networks and integrated with cloud computing for data storage. This technology faces challenges from adversaries attempting to infiltrate its networks to access its data and devices. Consequently, this study investigates a methodology that integrates Principal Component Analysis (PCA) for feature dimensionality reduction with a streamlined Convolutional Neural Network (CNN) architecture. The objective is to detect cyber-attacks within the Industrial Internet of Things (IIoT) environment, utilizing the WUSTL-IIOT-2021 dataset to classify reconnaissance, command injection, denial of service (DoS), and backdoor attacks. Furthermore, the proposed method utilized evaluation metrics such as accuracy, precision, F1-Score, and Recall with a confusion matrix to measure the model's performance in detecting cyber-attacks. Through multi-class evaluation experiments, CNN achieved 100% accuracy for command injection, DoS attacks, and normal traffic, while achieving 92.2% for command injection and 58.1% for backdoor attacks. It also achieved binary classification based on the distinction between legitimate and attack traffic. The prediction accuracy reached 99.96% for testing with an error rate of 0.0149%, while the training achieved an accuracy of 99.88% with an error rate of 0.018%. Through the prediction results, the model confirms its ability to detect cyberattacks with a high percentage in datasets that have balanced distributions of their categories

12:15 ***A Lightweight Approach to Spam Link Identification Using Automated Domain Reputation Scoring***

Petar Alilović, Davor Cafuta, Danijela Pongrac, Brigitta Cafuta and Ivica Dodig (Zagreb University of Applied Sciences, Croatia)

This paper examines the extraction and reputation analysis of domains identified within an email dataset containing both spam and ham messages. Using the Enron corpus, data processing was conducted in a virtualised Ubuntu environment (VMware) to ensure security. A custom Python application extracted domains from the message bodies, with the structured output stored in .parquet format for subsequent analysis on the host system. The reputation of each identified domain was assessed using the Open PageRank API. The methodology includes comprehensive statistical analysis, such as calculating average scores and examining the distribution of domains across a categorised ranking scale (1-10). Experimental results reveal a significant disparity between the two categories, with domains associated with spam exhibiting a notably lower average reputation. The study outlines the technical implementation, the technology stack, and provides an evaluation of the benefits and limitations of this approach, highlighting its potential application in improving automated spam filtering systems.

12:30 ***Temporal Generalization of Static Malware Detectors: A Large-Scale Empirical Study***

Sahil Mann, Shahroz Abbas and Ajmery Sultana (Algoma University, Canada)

Most static machine learning (ML) malware detectors are evaluated with randomly shuffled train-test splits, which can overestimate real-world performance by ignoring temporal distribution shift as malware evolves. To address this, in this paper, we conduct a large-scale temporal generalization study on the EMBER2024 Win32 dataset using strict chronological evaluation. Using 586 signed-log transformed Portable Executable (PE) features, we build three time-based splits with cutoff dates 2023-12-01, 2024-01-01, and 2024-03-01, training only on pre-cutoff samples and testing on future samples. We compare SGD Logistic Regression, SGD Linear SVM, LightGBM, XGBoost, and an MLP. LightGBM and XGBoost are consistently the strongest and most stable models across temporal cutoffs, both reaching peak AUC values of approximately 0.9922. In contrast, the linear baselines exhibit stronger recall degradation under temporal drift, with SGD Logistic Regression recall dropping from 0.8984 at the 2024-01-01 cutoff to 0.7255 at the 2024-03-01 cutoff. A random-split baseline substantially inflates performance, with LightGBM reaching 0.9996 AUC, XGBoost reaching 0.9969 AUC, and the linear baselines exceeding 0.992 AUC, confirming that shuffled evaluation can misrepresent deploy-time effectiveness. Permutation importance indicates that the boosted tree models rely on a largely stable subset of features across time, with top-10 feature overlap of 8 for LightGBM, 9 for XGBoost, and 4 for the MLP between the 2024-01-01 and 2024-03-01 cutoffs. Overall, chronological validation materially changes conclusions and should be standard for benchmarking static malware detectors.

12:45 ***Honeytoken and Honeypot Systems for IoT Networks: Proof of Concept and Implementation using ESP32-S3***

Nabil Litayem (Joaan Bin Jassim Academy for Defence Studies, Qatar)

In this work, we propose a resource-constrained deception framework for IoT environments that combines a micro-honeypot with embedded honeytokens on a low-cost ESP32-S3 platform. The design emulates commonly probed services and distributes unique decoy artifacts to generate high-signal telemetry under tight CPU, memory, and power budgets. We detail the system architecture, threat model, and implementation, and we report evaluation results using representative reconnaissance and exploitation workflows. Results indicate reliable detection of port scanning, unauthorized service interaction, and credential misuse events while maintaining low overhead suitable for continuous operation on the SoC. The proposed approach leverages the ESP32-S3's hardware cryptographic accelerators to ensure secure telemetry transmission without burdening the main application cores. Furthermore, it aligns with current practice emphasizing edge-resident detection, clean telemetry for cloud correlation, and mapping to MITRE ATT&CK for ICS/IoT. The framework complements baseline controls, including secure boot, SBOM, segmentation, MUD, and

privacy-by-design data minimization, and reduces the operational cost of security visibility across heterogeneous IoT fleets and Cyber-Physical Systems (CPS). Ultimately, this proof-of-concept validates that localized, edge-native deception provides a viable and scalable early-warning mechanism to significantly enhance the cyber-defense posture of critical infrastructure networks, particularly where passive traffic analysis falls short.

Tuesday, April 21 11:30 - 13:00 (Africa/Cairo)

SESSION-1E: RESEARCH: Security Frameworks, Critical Infrastructure Protection & Post-Quantum

6 PAPERS

Room: ROOM-E

11:30 ***Post-Quantum Cryptographic Framework for Critical Infrastructure Protection***

Wael Badawy (Egyptain Russian University, Egypt)

Quantum computing is rapidly approaching a stage where classical cryptographic systems, particularly RSA, ECC, and Diffie-Hellman, will no longer provide adequate security for critical infrastructure. The emergence of Shor's and Grover's algorithms introduces unprecedented vulnerability to national cyber-defense ecosystems, underpinning power grids, transportation, military communications, and banking systems. This paper presents a unified post-quantum cryptographic (PQC) security framework integrating lattice-based primitives, hybrid key-exchange mechanisms, and secure orchestration for resource-constrained operational-technology environments. The proposed architecture leverages NIST-recommended schemes, including CRYSTALS-Kyber, CRYSTALS-Dilithium, and SPHINCS+, to ensure long-term confidentiality, integrity, and interoperability across legacy and modern control environments. A defense-in-depth trust-fabric is introduced to protect critical data flows across edge controllers, automation PLCs, substations, cloud-SCADA gateways, and defense communication systems. Performance evaluation demonstrates a 39.4% reduction in key-exchange latency, 46% faster authentication cycles, and zero-downtime cryptographic transition capability. Our findings position this framework as a practical migration blueprint supporting national PQC readiness and cyber-resilience mandates.

11:45 ***National Cyber Shield Platform: A Scalable Cyber Resilience Engineering Framework for Sustainable Digital Governance***

Huda J. Nael, Mohammed Almulla and Ali Fenjan (American International University, Kuwait)

The accelerating digital transformation of governments and critical infrastructure worldwide has intensified the need for robust, measurable, and sustainable national cyber-resilience frameworks. Despite the proliferation of national cybersecurity strategies, many countries face a persistent implementation gap between high-level policy formulation and measurable institutional execution. This paper proposes a National Cyber Shield Platform (NCSP), a scalable, governance-driven, and data-centric cyber resilience engineering framework designed for adoption by any nation seeking to operationalize its cybersecurity strategy. The proposed architecture integrates three core components: (1) Compliance Intelligence Engine, (2)

Adaptive Cyber Training Lab, and (3) National Cyber Resilience Dashboard. These components operate within a closed-loop continuous improvement cycle, linking regulatory compliance, human readiness, and recovery maturity into a unified operational ecosystem. Furthermore, the study introduces a National Cyber Resilience Index (NCRI), a composite metric designed to quantify national cyber resilience dynamically. To demonstrate applicability, Kuwait is presented as a case study illustrating how the framework can be aligned with existing national strategies and regulatory structures. The proposed model is designed to be modular, scalable, and adaptable across diverse regulatory, economic, and governance contexts. This research contributes a replicable blueprint for next-generation national cyber resilience engineering.

12:00 ***Decentralizing Reputation Scoring Using Blockchain Behavioral Data and Machine Learning***

Md Shakil Ahmed, Shohan Islam Joy and Nabarun Halder (Independent University, Bangladesh, Bangladesh); Ashraful Islam (Independent University Bangladesh & Center for Computational and Data Sciences, Bangladesh); Sanzar A Alam (Independent University, Bangladesh & Center for Computational & Data Sciences, IUB, Bangladesh)

Existing financial scoring systems heavily rely on centralized credit records and banking history. This dependency leads to the exclusion of many individuals as they lack formal financial histories. This paper offers a decentralized behavioral reputation framework that leverages blockchain transaction behavior and machine learning (ML) to estimate user reliability in decentralized finance (DeFi) ecosystems. In order to derive metrics of borrower behavior, we utilized Aave protocol transactional data on Ethereum (Versions 2 and 3). The research compares the performance of different models of ensemble learning (Random Forest, LightGBM, and AdaBoost) against linear and tree-based models. The experiments were performed on 1.10 million transactions to ensure that the model is resistant to entity-level data leakage, which demonstrates that Random Forest achieves the most robust predictive performance with an R^2 = 0.953 and the lowest Root Mean Square Error (RMSE). In addition to that, the deceitful potency due to the leakage in blockchain datasets was also quantified formally to demonstrate how post-execution variables (e.g., block indices) and temporal artifacts (e.g., confirmations) artificially inflate offline accuracy metrics, rendering models undeployable in a real-time mempool environment. To emphasize the generalization of the models to unseen wallets, extensive ablation and robustness tests were also performed. In conclusion, we outline a conceptual Zero-Knowledge Machine Learning (zkML) architecture as a future pathway to implement verifiable, reputation scoring natively on-chain that also preserves user privacy.

12:15 ***An Assessment of Cybersecurity Awareness, Competency Levels, Behavior and Practices among Technical Trainees- An Objective Study***

Sayfuldeen Alsuhaymi and Muhammad Asif Khan (Taibah University, Saudi Arabia)

The increasing digital attacks on business organizations and individuals have made cybersecurity an integral part of our connected world. With the rapid technology advancement, cybercriminals are becoming smarter and they are using innovative tactics to breach data security. A successful cyberattack causes loss of data, loss of money and loss of trust and reputation. Cyberthreats may exist at any level in organization, therefore, cybersecurity awareness among personnel is very crucial. This study explores cybersecurity awareness among trainees at Khaybar Technical College in Saudi Arabia. With Saudi Vision 2030 pushing cybersecurity forward, it is essential to prepare future professionals with the right knowledge and habits of cybersecurity. This study presents a quantitative research design to gather data on cybersecurity awareness among the trainees. This method helps to maintain objectivity when assessing awareness levels across a large and varied group of participants. The data was analyzed using statistical

techniques and tools such as SPSS. The results showed a clear gap of cybersecurity awareness among the participants. The study showed that cybersecurity education cannot be limited to technical fields. Every trainee needs high cybersecurity awareness to stay safe online. To close this gap, institutions should offer targeted awareness programs, simple training, and clear policies that help build a strong culture of digital safety.

12:30 ***Architectural Trust Mediation for Containing IT-OT Escalation in Legacy Critical Infrastructure***

Cornelius Chipasha (University of Zambia, Zambia); Simon Tembo (University of Lusaka, Zambia); Shimaponda Mulundumina (University of Zambia, Zambia)

The integration of enterprise information technology (IT) with operational technology (OT) environments has introduced new cybersecurity risks in legacy critical infrastructure systems. In many industrial architectures, implicit trust relationships allow attackers who compromise enterprise systems to interact with operational control networks using legitimate industrial protocols. This paper investigates whether architectural mediation of trust can constrain such escalation pathways without modifying legacy OT devices. A protocol-breaking mediation architecture is proposed that terminates communication context at the IT-OT boundary and reconstructs only explicitly permitted interactions. A packet-level simulation environment is used to evaluate the security and operational impact of this approach. Experimental results demonstrate that architectural mediation eliminates unauthorized control operations, reducing attack success from 100% to 0%, while maintaining legitimate monitoring functionality. Performance evaluation shows bounded latency overhead below 1.13 ms under stress conditions. These findings indicate that terminating protocol context at the IT-OT architectural boundary provides a practical strategy for improving cybersecurity resilience in legacy critical infrastructure environments.

12:45 ***Latency-Aware Architectural Security for Enterprise-Operational Network Integration in Critical Digital Infrastructure***

Cornelius Chipasha and Shimaponda Mulundumina (University of Zambia, Zambia); Simon Tembo (University of Lusaka, Zambia)

Critical Digital Infrastructure (CDI) increasingly integrates enterprise information technology (IT) environments with operational technology (OT) systems responsible for monitoring and controlling industrial processes. While such integration improves operational coordination and visibility, it also introduces cybersecurity risks associated with lateral movement across interconnected infrastructure components. Architectural mediation mechanisms have been proposed to constrain adversarial reachability between enterprise and operational networks. However, their deployment in industrial environments depends on their operational feasibility, particularly with respect to communication latency. This paper evaluates the performance impact of a mediation-based security architecture deployed between enterprise and operational infrastructure domains. Using a Mininet-based network emulation environment representing enterprise OT infrastructures, we measure round-trip latency, service response time, and operational communication overhead under baseline and mediated network configurations. Experimental results demonstrate that the mediation architecture introduces only modest latency overhead while preserving operational service availability and preventing unauthorised control interactions. Mediation-based architectural containment substantially reduces compromise propagation with minimal latency overhead, a deployable security approach for operational technology environments.

Tuesday, April 21 13:00 - 14:00 (Africa/Cairo)

LUNCH: LUNCH (ON SITE)

LUNCH

Room: COFFEE_ROOM

Tuesday, April 21 14:00 - 15:30 (Africa/Cairo)

SESSION-2A: RESEARCH: Advancements in Intelligent Systems for Image Analysis, Cybersecurity, and Educational Technologies

6 PAPERS

Room: ROOM-A

Chairs: Salah A. Aly (Fayoum University, Egypt), Shrey Tyagi (Salesforce Inc, USA)

14:00 ***A Systematic Review of Lightweight Multi-Scale Feature Extraction Models for Medical Image Segmentation***

Mina Magdy Kamel, Mai S. Mabrouk and Ahmed F. Elnokrashy (Nile University, Egypt)

Medical image segmentation functions as an essential tool for both clinical diagnosis and treatment preparation. Deep learning models achieve better segmentation accuracy but their high computational needs restrict their deployment in real-time or edge-device applications. Lightweight models with multi-scale feature extraction provide an efficient solution that maintains high-performance levels. The evaluation of their accuracy together with computational cost has not received sufficient systematic assessment. This review assesses lightweight segmentation models that use multi-scale features that were published between 2020 and 2024. The models are divided into three distinct categories: (1) pure CNNs, (2) CNNs with attention mechanisms, and (3) hybrid models that integrate Transformers or MLPs. The evaluation process includes performance benchmarking of three public datasets CVC-ClinicDB, Kvasir-SEG, and ISIC2018 through Dice score, mean IoU, parameter count, and FLOPs (G) metrics. The MFLUNet and MSLUNet CNN-based models deliver high Dice scores exceeding 93% on Kvasir-SEG while maintaining extremely low complexity at under 2.2M parameters and 0.05 FLOPs (G). The attention-based models PMFSNet and PIS-Net enhance global context understanding but transformer-based models like LM-Net deliver superior accuracy at increased computational expense. The review demonstrates that segmentation accuracy depends on resource availability. Pure CNNs provide the best solution for fast efficient deployment yet hybrid and transformer models deliver better precision when resources become available. Future research needs to develop better generalization capabilities for imaging modalities while optimizing the accuracy-efficiency tradeoff for practical clinical use

14:15 ***Hierarchical Explainable Risk Scoring for Security Alert Prioritization in Security Operations Centers***

Siva Prasad Nandi (Oracle, USA); Rajesh Purushothaman (Zscaler Inc, USA); Sai Manohar Nethi (Intuit, USA); Pavan Nutalapati (Oracle Corp, USA); Naga Satya Praveen Kumar Yadati (Meta, USA); Anil Kumar Thimmapuram (USA)

Security operations centers have a constant struggle to prioritize the large number of heterogeneous alerts, in which traditional detection systems usually give classifications without adequate decision support or interpretability. This work introduces a hierarchical explainable risk scoring framework for security alert prioritization that treats incident-grade prediction as an analyst-oriented triage problem instead of a traditional binary detection problem. Using the Microsoft GUIDE security incident dataset, the method integrates structured SOC alert attributes, temporal context and security semantics in a lightweight LightGBM-based multiclass model followed by risk-based ranking and explanation generation with feature importance and SHAP analysis. The proposed framework achieved an accuracy of 0.9030, macro precision of 0.9034, macro recall of 0.8906, macro F1-score of 0.8962, ROC-AUC of 0.9845 and PR- AUC of 0.9697 in addition to a good class-wise performance, including the F1-score of 0.91 for TruePositive incidents. The results further showed that organizational context, detector behavior, alert semantics, temporal cues, and attack-technique information were the dominant drivers of prioritization, highlighting the novelty of combining accurate multiclass SOC triage with transparent, explanation-aware cyber defense support.

14:30 ***Evaluating the Efficacy of Microsoft Teams as an Online Learning Platform: A Case Study in a Philippine State University***

Julius Cesar O. Mamaril, Jerry C. Diaz and Clark Kim C. Castro (Pangasinan State University, Philippines)

The COVID-19 pandemic profoundly disrupted educational systems worldwide, compelling higher educational institutions to transition rapidly to online learning modalities. At Pangasinan State University (PSU)-a Philippine state university-Microsoft Teams was adopted as the primary learning management system (LMS) to facilitate both synchronous and asynchronous instruction. This study evaluates the perceived usability and user acceptance of Microsoft Teams among N = 162 PSU stakeholders (118 students and 44 faculty members) using data collected during the first quarter of 2022 (Q1 2022), a period marking the transition phase of online learning in the Philippines. A dual-methodology approach was employed: the System Usability Scale (SUS) as a Human-Computer Interaction (HCI) instrument (Cronbach's $\alpha = 0.84$) and the Technology Acceptance Model (TAM) as an information systems adoption framework (Cronbach's $\alpha = 0.88$). TAM was applied in a post-usage context, theoretically grounded in Bhattacherjee's Expectation-Confirmation Model (ECM) of information systems continuance, which legitimizes post-adoption measurement of Perceived Usefulness (PU) and Perceived Ease of Use (PEOU) as continuance-intention predictors. The Likelihood to Recommend (LTR) and Overall Experience (OE) metrics were incorporated to capture post-usage satisfaction on a 0-10 scale. Findings reveal that both SUS (Mean = 66.83) and TAM (Mean = 67.15) yielded equivalent below-average scores (CGS grade: C). The consumption platform (laptop vs. smartphone) was the primary differentiating factor in user experience-a finding attributed to a 'Mobile-Desktop Mismatch' in feature availability and interface design. Significant interaction effects involving user role, gender, and age were identified, including a 'Reliability Anxiety' pattern among female students. Regression models explained 20.1%-34.7% of variance in LTR and OE ($R^2 = 0.201$-0.347), confirming that contextual factors beyond TAM constructs substantially shape user experience in resource-constrained Philippine HEI settings.

14:45 ***Multivariate Deep Temporal Forecasting of Enterprise Cyber Risk Using Vulnerability and Exploitation Intelligence***

Pavan Nutalapati (Oracle Corp, USA); Siva Prasad Nandi (Oracle, USA); Naga Satya Praveen Kumar Yadati (Meta, USA); Rajesh Purushothaman (Zscaler Inc, USA); Sai Manohar Nethi (Intuit, USA); Gayathri Balakumar (Capital One, USA)

Proactive cyber defense means that cybersecurity must anticipate short-term changes in enterprise risk instead of reacting after exploitation has happened, but most previous cybersecurity analytics focus on intrusion detection or static vulnerability scoring instead of time-series cyber risk prediction. This paper fills that void with a lightweight and operationally grounded approach that models weekly cyber risk via fused public threat intelligence, combining CISA Known Exploited Vulnerabilities (KEV) and EPSS temporal signals into a unified forecasting pipeline. The approach is new in treating cyber risk as a combination of a continuous, weekly prediction problem, as well as a next week, high-risk warning problem, in order to connect predictive analytics with actionable security decision support. For weekly risk regression, Random Forest had MAE 0.2580, RMSE 0.2842 and R2 0.3266, LSTM had MAE 0.2259, RMSE 0.2811 and R2 0.4700, and XGBoost showed the best risk forecasting performance with MAE 0.2440, RMSE 0.2960 and R2 0.5830. For high-risk next-week classification, the optimal operational performance was given by Logistic Regression with threshold 0.45, with Accuracy 0.8110, Precision 0.5356, Recall 0.7563, F1-score 0.6870 and ROC-AUC 0.8220. These results show the proposed formulation to be a practical, reproducible and computationally efficient way of achieving proactive cyber risk forecasting and early warning in enterprise environments.

15:00 ***Detection and Classification of Aedes Aegypti Mosquito Eggs Using Convolutional Neural Networks and Advanced Computer Vision***

Arnel C. Fajardo (Isabela State University, Philippines)

Dengue hemorrhagic fever (DHF) is a mosquito- borne viral disease primarily transmitted by Aedes aegypti, a species well-adapted to indoor human environments. Early detection of mosquito breeding sites is essential for effective prevention and control. This study proposes a computer vision-based approach for automated detection and counting of Aedes aegypti eggs using ovitrap images. The methodology combines patch-based sliding window techniques with OpenCV preprocessing methods, including noise reduction and image enhancement, to improve early-stage feature detection, while a black cup ovitrap design was utilized due to its cost-effectiveness and widespread use in mosquito surveillance. Three object detection models were evaluated: Single Shot Detector (SSD), an ensembled model of YOLOv11 variants (Gen3 and Gen6), and YOLOv11 Gen6. Results show that SSD achieved an accuracy of 78.7%, while the ensembled model reached 86.7%, and YOLOv11 Gen6 demonstrated the highest performance with 98.8% accuracy in detecting and counting Aedes aegypti eggs. These findings highlight the potential of deep learning-based computer vision systems to enhance mosquito surveillance and support early intervention strategies in dengue-prone areas.

15:15 ***Improved Energy Demand Forecasting in LGU (Local Government Unit) Using Ensemble Learning as Compared to Single Model Machine Learning***

Arnel C. Fajardo, Armald C. Marcos and Genesis Carlo Pangan (Isabela State University, Philippines)

Accurate energy demand forecasting is essential for efficient power system planning and management, particularly in rapidly developing regions. This study proposes a quad-hybrid ensemble framework that integrates Pi-Mamba, Informer, Frequency Enhanced Decomposition Transformer (FEDformer), and a Quantum-Informed Neural Network (QINN) to improve forecasting accuracy in Local Government Units (LGUs). Using electricity demand data from the National Grid Corporation of the Philippines (NGCP) and Isabela Electric Cooperative (ISELCO) spanning 2021 to 2026, the dataset was preprocessed through normalization, temporal encoding, and feature extraction techniques. Each model contributes distinct strengths: Pi-Mamba captures long-term dependencies, Informer enhances computational efficiency, FEDformer improves frequency-based decomposition, and QINN introduces quantum-enhanced feature representation. A dynamic gating mechanism was further employed to optimize model integration and prediction weighting. Experimental results demonstrate that the proposed ensemble outperforms individual models, achieving lower error rates and improved stability across multiple forecasting scenarios. The findings highlight the effectiveness of hybrid deep learning and quantum-inspired

techniques in addressing complex, nonlinear, and seasonal energy demand patterns, making the framework suitable for real-world grid forecasting applications.

Tuesday, April 21 14:00 - 15:30 (Africa/Cairo)

SESSION-2B: RESEARCH: Innovative Approaches in AI-Driven Systems for Cybersecurity, Cloud Optimization, and Healthcare Surveillance

6 PAPERS

Room: ROOM-B

Chair: Vivek Venkatesan (Vanguard, USA)

14:00 ***Deep Learning for Intelligent Healthcare Surveillance: Detection and Classification of Patients and Medical Personnel Using YOLO***

Mostafa Rizk (Lebanese University, Lebanon); Abbas Rammal (Phoenicia University, Lebanon); Batoul Cheikh (Lebanese University, Lebanon)

This paper presents a deep learning-based framework for intelligent healthcare surveillance, focusing on the detection and classification of patients and medical personnel, including doctors and nurses. The proposed system employs advanced YOLOv12 object detection models to achieve accurate and real-time recognition in complex hospital environments. A custom annotated dataset is developed to capture variations in appearance, lighting, and occlusions. Transfer learning and data augmentation techniques are applied to enhance model robustness. Experimental results demonstrate that the YOLOv12 variants provide high precision, recall, and mean Average Precision (mAP) while maintaining computational efficiency suitable for real-time deployment. The proposed approach supports smart hospital applications such as patient monitoring and workflow optimization.

14:15 ***Enhanced AI-Based Disposable Email Detection System***

Gasser Sayed (Arab Open University, Egypt); Eid Amery (Amery, Egypt); Hala Abbas (Helwan University & Arab Open University, Egypt)

Disposable and temporary email addresses represent a significant and growing threat vector in modern cybersecurity, enabling attackers to bypass identity verification, evade blacklists, and conduct phishing campaigns with minimal traceability. Internet Service Providers (ISPs) are particularly vulnerable due to the massive scale of email traffic they process daily. This paper proposes a novel Multi-Scale Convolutional Neural Network (Multi-Scale CNN) for detecting disposable and temporary email addresses at ISP scale. The proposed architecture employs three parallel convolutional branches with kernel sizes of 2, 3, and 5, enabling simultaneous capture of character-level patterns at multiple granularities-from local bigram transitions to longer morphological structures. The model operates within a fourstage layered detection pipeline integrating blacklist lookup, heuristic rule evaluation, domain reputation API querying, and the proposed deep learning classifier. Experiments were conducted on a rigorously audited and balanced dataset of 236,162 domains (118,081 disposable, 118,081 legitimate), evaluated using 5-fold cross-validation. The proposed Multi-Scale CNN achieves a mean accuracy of 79.04%, F1-score of 0.7691, and AUC of 0.8786 in cross-validation, and 79.35% accuracy with an F1-score of 0.7735 and AUC of 0.8789 on the held-out test set. Benchmarking against two CNN baseline architectures-a Baseline CNN and an Improved CNN with

deeper convolutional blocks and batch normalization-demonstrates that the proposed model achieves the highest recall (0.7058) and F1-score among all configurations, improving recall by 4.1% over the Baseline CNN without sacrificing precision. Statistical significance testing confirms that progressive architectural enhancements yield meaningful performance gains ($p = 0.034$). The proposed system offers ISP Security Operations Centers a scalable, adaptable, and modular solution for disposable email detection that requires no manual feature engineering and operates effectively at production email traffic volumes.

14:30 ***Uncertainty-Aware Intelligent Workload Scheduling for Cloud-Native Systems Under Bursty Demand***

Venkata Raja Satya Teja Nagavalli, Sapan Bharadwaj Bonala and Sharath Chandra Samineni (USA)

Cloud-native systems often have bursty and highly imbalanced workload dynamics that undermine the performance of traditional resource schedulers especially when peak demand is challenging to predict on a short-term basis. This paper introduces an uncertainty-aware, burst-risk-directed workload scheduling scheme in cloud systems that integrates window-based workload modeling, burst-aware temporal feature engineering, long short-term memory learning, and risk-sensitive CPU allocation. In contrast to normal heuristic or forecast-only schedulers, the one presented actually marks high-risk burst windows and becomes more conservative in allocation when the burst likelihood is high in order to align the scheduling decisions with the operational risk. Large-scale cloud workload trace experimental evaluation indicated good burst-risk detection performance with accuracy of 0.9707, precision of 0.9594, recall of 0.9829, F1-score of 0.9710, and ROC-AUC of 0.9868. The proposed method minimized the SLA violation rate to 0.5426, reduced mean under-provision to 0.0112, and boosted burst-window protection to 0.1441 in terms of minimizing the violation rate, minimizing mean under-provision, and maximizing burst-window protection, respectively, compared to heuristic and forecast-based baselines. These findings point to the originality and practical significance of incorporating burst-risk intelligence into dependable cloud-native resource management.

14:45 ***Beyond Agent Design: A Systematic Framework for MCP Server Runtime Orchestration***

Vijayakumar Venganti (Cisco Systems, Inc., USA); Deepak Kole (IEEE Senior Member, USA); Siva Prasad Nandi (Oracle, USA)

The Model Context Protocol (MCP) has rapidly emerged as a standard interface for connecting AI agents to external tools, databases, and services. While considerable research has addressed agent design, prompt engineering, and tool selection, the runtime infrastructure layer that hosts MCP servers remains unstudied. This paper argues that production failures in agentic systems arise primarily from runtime and orchestration mismatches rather than from deficiencies in agent logic. We present a systematic decision framework comprising a six-dimension workload characterization rubric, a four-tier runtime taxonomy (serverless/FaaS, container-based, VM/bare metal, and managed orchestration), a decision matrix mapping workload profiles to runtime tiers, and a comparative evaluation across four production-relevant metrics. Two illustrative case studies demonstrate that applying the framework reduces latency variability by up to 73% and operational incident rate by over 60% compared to ad hoc runtime selection. This work establishes runtime orchestration as a first-class design concern in production agentic systems and provides practitioners with actionable, cloud-agnostic guidance for MCP server deployment.

15:00 ***Neural-Heatmap: An Innovative Predictive Data Visualization Technique Using Deep Learning to Strengthen Cybersecurity***

Nidadavolu Venkat Durga Sai Siva Vara Prasad Raju (Birla Institute of Technology and Science (BITS), India); Nuruzzaman Faruqui (Daffodil International University, Bangladesh); Moutaz Alazab (Al-Balqa Applied University, Jordan); Olivia-Roxana Alecsoiu (Romania)

Enterprise application security is essential in today's dynamic cyber landscape. As cybersecurity techniques advance, intrusion patterns become more advanced and complicated. That is why monitoring the effectiveness of existing Artificial Intelligence (AI)- driven Intrusion Detection Systems (IDSs) plays a crucial role in defending against new and effective cyberattacks. However, most AI-based automatic IDSs are focused on detection performance, ignoring the importance of data visualization. This paper proposes an innovative neural network-based data visualization technique named Neural-Heatmap. It is a real-time intrusion monitoring system for enterprise applications that is responsible for concurrent cybersecurity at multiple divisions. It combines a Bidirectional Long Short-Term Memory (BiLSTM) network with a heatmap and results in a unique real-time predictive data visualization method. The Neural-Heatmap detects intrusion with an average accuracy of 98.82% and reduces the response delay by 89.82% after a probable security incident. It achieves an impressive precision, recall, and F1-score of 98.82%, 98.76%, and 98.79%, respectively. With an innovative concept and outstanding performance, the proposed Neural Heatmap demonstrates the potential to be an effective data visualization technique for strengthening cybersecurity.

15:15 ***An Agentic Event-Driven Architecture for Compliance Exception Triage in Auditable Regulatory Systems***

Krishna Kandi (Convoke, USA); Lavi Kumar (Discover Financial Services, USA); Gouri Sankar Dash (Tata Consultancy Services, USA)

Regulatory platforms generate a steady stream of operational events, validation failures, audit alerts, and policy exceptions that must be reviewed quickly and handled consistently. In many organizations, this process is still manual or dependent on rigid rule-based workflows, which can slow response times and make it harder to maintain clear audit trails. This paper presents a design for using specialized software agents within an event-driven system to support compliance exception triage in audit-sensitive environments. The proposed approach allows agents to monitor incoming events, classify exceptions, gather relevant context, prioritize cases based on risk and policy requirements, and recommend next actions while preserving a full record of how each decision was made. The architecture separates event ingestion, agent coordination, decision logic, and audit logging so that the system remains scalable, transparent, and resilient under high event volume. A prototype implementation is evaluated on a simulated compliance operations dataset using metrics such as triage latency, prioritization quality, escalation accuracy, and completeness of audit records. The results suggest that this approach can improve the speed and consistency of exception handling while preserving the traceability required in compliance-sensitive systems. The study highlights how agent-based orchestration can support faster operational decision-making without reducing accountability.

Tuesday, April 21 14:00 - 15:30 (Africa/Cairo)

SESSION-2C: RESEARCH: Deep Reinforcement Learning & Evolutionary Optimization

6 PAPERS

Room: ROOM-C

Chair: Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA)

14:00 ***Bayesian Change Point Detection and Its Applications in Artificial Intelligence***

Golnaz Shahtahmassebi (Nottingham Trent University, United Kingdom (Great Britain)); Jose Maria Sarabia (University of Cantabria, Spain)

Bayesian Change Point Detection (BCPD) provides a principled probabilistic framework for identifying structural shifts and regime changes in time-series and sequential data. Such shifts naturally arise in many Artificial Intelligence (AI) and Data Science applications due to nonstationarity, concept drift, and evolving system dynamics. This paper presents a structured overview of modern BCPD methodologies, including offline and online formulations, with particular emphasis on Bayesian Online Change Point Detection and hierarchical Bayesian models.

We discuss how BCPD supports robust learning and decision-making in contemporary AI domains such as reinforcement learning, deep learning systems operating on streaming data, natural language processing, and AI-driven healthcare monitoring. From a data science perspective, BCPD enables real-time analytics, predictive modelling under distributional change, anomaly detection in large-scale data streams, and interpretable segmentation for business intelligence and data governance.

Furthermore, we highlight recent advances in Bayesian modelling for change point analysis, which enhances sensitivity to latent regime transitions while retaining interpretability and uncertainty quantification. A concise mathematical formulation based on posterior inference and Bayesian online updates is provided, illustrating the suitability of BCPD for explainable, adaptive, and trustworthy AI systems in complex and evolving environments.

14:15 ***Entropy-Constrained Hyperheuristics for the Australian Post p-Hub Median Problem***

Kassem Danach, Sr (Al Maaref University, Lebanon & Chairperson, Lebanon); Samir Haddad and Jinane Sayah (University of Balamand, Lebanon); Khouloud Eledlebi (Abu Dhabi University, United Arab Emirates); Joseph Merhej (Facultyof Sciences II, Lebanon); Chadi Kallab (Lebanese American University, Lebanon)

The p-hub median problem (PHMP) is a fundamental NP-hard network design problem with wide applications in logistics, transportation, and telecommunication systems. Hyperheuristics have demonstrated competitive performance for large-scale hub location instances by dynamically selecting low-level heuristics during the search process. However, existing approaches lack explicit mechanisms to regulate heuristic diversity, often leading to premature convergence and structural stagnation. This paper proposes an Entropy-Constrained Hyperheuristic (ECHH) that integrates an information-theoretic diversity control mechanism into heuristic selection. By enforcing a lower bound on the Shannon entropy of the heuristic usage distribution, the proposed framework transforms diversification into a measurable and controllable property. The entropy threshold parameter governs the exploration-exploitation balance and induces a structural phase transition in search behavior. The method is evaluated on the Australian Post 200-node benchmark dataset under multiple hub densities and entropy levels. Experimental results show that moderate entropy thresholds significantly improve robustness, basin exploration, and average solution quality compared to unconstrained selection, while excessively high entropy reduces convergence speed. The findings demonstrate that structural diversity control acts as a regularization mechanism in hyperheuristic search and provides a principled approach to balancing exploration and exploitation in large-scale combinatorial optimization.

14:30 ***Contrastive Representation Learning for Cross-Project Software Vulnerability Discovery***

Rajesh Purushothaman (Zscaler Inc, USA); Naga Satya Praveen Kumar Yadati (Meta, USA); Suresh Dodda (USA); Sai Manohar Nethi (Intuit, USA); Pavan Nutalapati (Oracle Corp, USA); Anil Kumar Thimmapuram (USA)

Software vulnerability discovery is a challenging task in the secure software engineering field because of the severe class imbalance, large cross-project variability, and the diverse semantics of the weakness types in the real-world. This paper introduces a CodeBERT-based software vulnerability discovery framework for the function-level analysis of DiverseVul dataset, which covers wide project diversity and 150 CWE categories. The proposed approach focuses on vulnerability-aware representation learning instead of surface-level pattern matching so it is able to perform robust discrimination between vulnerable and non-vulnerable source code under realistic data conditions. Experimental evaluation shows good and balanced detection ability with the precision of 0.68, recall of 0.67, F1-score of 0.67, MCC of 0.41, ROC-AUC of 0.89 and PR-AUC of 0.70. These results illustrate the power of transformer-based code representations for practical vulnerability analysis while solving the problem of imbalanced and heterogeneous software security data. The study adds a realistic, scalable and security-oriented vulnerability discovery framework with an obvious relevance to the topic of AI-driven software assurance.

14:45 ***Semantic Transformer-Based Detection of Security Misconfigurations in Cloud Infrastructure-as-Code***

Naga Satya Praveen Kumar Yadati (Meta, USA); Sai Manohar Nethi (Intuit, USA); Siva Prasad Nandi (Oracle, USA); Rajesh Purushothaman (Zscaler Inc, USA); Pavan Nutalapati (Oracle Corp, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA)

Cloud infrastructure misconfigurations continue to be a leading cause of security breaches, but current rule- based detection methods are not flexible and do not generalize across diverse Infrastructure-as-Code (IaC) patterns. This work proposes a deep learning-based framework for automated detection of security misconfigurations of Terraform configurations, leveraging transformer based semantic understanding of code. A lightweight but effective pipeline is built by using DistilBERT for sequence classification, which allows for efficient learning from IaC text while leaving practical training constraints. The approach is tested on a balance dataset of 2000 samples, achieving good performance with an F1-score of 0.7598, recall of 0.8700, ROC-AUC of 0.8224 and precision-recall-AUC of 0.8344, which outperforms the traditional TF-IDF with Logistic Regression in recall and overall detection capability. The results indicate the capability of the model in capturing the security patterns in the given context beyond the syntactic rules, presenting a scalable and intelligent solution to proactively assure security in clouds.

15:00 ***Machine Learning Based Analysis of Global Fisheries Production Consumption and Sustainability Trends***

Jennie Fernandez (Pangasinan State University, Philippines)

This study aimed to predict global fisheries production and assess sustainability trends based on key factors such as aquaculture production, capture fisheries output, seafood consumption, and environmental indicators using machine learning algorithms. A dataset consisting of production, consumption, and sustainability variables was analyzed using several machine learning models, including Linear Regression, Random Forest, Support Vector Machines (SVM), XGBoost, and Neural Networks (DNN). The results showed that XGBoost and Random Forest achieved the highest performance in both regression and classification tasks, with accuracy values of 0.90 for XGBoost and 0.87 for Random Forest, while Linear Regression was particularly effective for predicting continuous outcomes such as fisheries production, achieving an R^2 value of 0.85. Feature importance analysis revealed that aquaculture production (0.20-0.22), capture fisheries output (0.18-0.20), and seafood consumption (0.15-0.17) were the most significant predictors of fisheries performance. Additionally, countries were classified into High Risk, Medium Risk, and Low Risk categories based on

sustainability levels, with High Risk countries exhibiting low fish stock sustainability (below 40%) and declining production trends. The findings suggest that machine learning can effectively model complex fisheries systems and provide insights for targeted, data-driven interventions. Overall, the study highlights the importance of balanced production systems and sustainable resource management in improving fisheries outcomes.

15:15 ***Comparative Analysis of Organic Matter Dynamics and Coral Reef Condition Using Machine Learning***

Jennie Fernandez, Lorena F. Aquino and Nerda C. De Vera (Pangasinan State University, Philippines); Jackie Millama (Pangasinan State University-Binmaley, Philippines); Rosalie Belano and Maria Angelica S. Swin (Pangasinan State University, Philippines)

This study aimed to analyze and predict sediment organic matter (SOM) dynamics and coral reef ecosystem condition based on key environmental and biological factors such as total organic carbon (TOC), chlorophyll concentration, sediment grain size, nutrient levels, and hydrodynamic conditions using machine learning algorithms. Data sourced from publicly available datasets on Kaggle were analyzed using Linear Regression, Random Forest, Support Vector Machines (SVM), XGBoost, and Neural Networks (DNN). The results showed that XGBoost and Random Forest achieved the highest performance in both regression and classification tasks, with Boost obtaining an accuracy of 0.90 and an R^2 value of 0.89, while Linear Regression demonstrated reliable performance in predicting continuous variables such as TOC (R^2 = 0.85). Feature importance analysis revealed that sediment grain size, chlorophyll concentration, and nutrient levels were the most significant predictors of organic matter distribution and reef condition. Reef sites were classified into Healthy, Moderate, and Degraded categories, with degraded areas characterized by high nutrient levels, increased turbidity, and environmental stress. The findings highlight the effectiveness of machine learning in modeling complex reef ecosystem processes and provide valuable insights for supporting data-driven coral reef management and conservation strategies.

Tuesday, April 21 14:00 - 15:30 (Africa/Cairo)

SESSION-2D: RESEARCH:Geospatial Intelligence, Remote Sensing & Environmental Data Analytics

6 PAPERS

Room: ROOM-D

14:00 ***Energy-Aware GPU Skinning for Real-Time Skeletal Animation Using glTF 2.0***

Anas Refaat Sabry and Alaa Adel Zaki (October University for Modern Sciences and Arts (MSA), Egypt); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

Skeletal animation is a fundamental technique in real-time computer graphics that enables efficient motion of articulated characters using a joint hierarchy. From an environmental computing perspective, real-time animation can become wasteful when CPU-side deformation, redundant per-frame computations, and excessive data movement increase power consumption and thermal load. This paper presents an energy-aware implementation of GPU-based skeletal animation and linear blend skinning using OpenGL 3.3 and the glTF 2.0 asset format. Animation sampling is

evaluated on the CPU with minimized per-frame overhead, while vertex deformation is executed on the GPU via vertex shaders to reduce CPU utilization and memory traffic. The proposed system supports textured meshes, hierarchical skeleton visualization, animation playback control, and real-time rendering. Experimental results show stable real-time performance while reducing CPU workload and maintaining visual accuracy, highlighting a practical pathway toward lower-cost, lower-waste character animation pipelines.

14:15 ***UAV-based Wildlife Detection using Deep Learning and Resource-Constrained Edge Devices***

Mohamad Kassem, Manar Anwer Khaleel Abusirdaneh and Manar Abu Talib (University of Sharjah, United Arab Emirates); Qassim MH Nasir (University Of Sharjah, United Arab Emirates); Fouad Lamghari (Fujairah Research Centre, United Arab Emirates); Mohammad AlShabi (University of Sharjah, United Arab Emirates)

Monitoring and conserving wildlife involves various hurdles to overcome, which include needing to observe the wildlife from a safe distance while remaining effective. Strides in artificial intelligence, namely in the field of computer vision, have made it possible to achieve that. In this work, multiple publicly available UAV-based wildlife datasets are reviewed and evaluated for their suitability in aerial wildlife monitoring, while also referencing them for future researchers to help with the scarcity of datasets. Based on this analysis, two datasets are selected: the Wildlife Aerial Images from Drones (WAID) dataset, composed of 6 animal classes and 14,375 images, and an Arabian Tahr dataset, composed of 4 classes and 1,959 images. These two datasets are used to create two models: YOLOv12n and YOLOv12m. The proposed models have achieved strong detection performance, with mAP scores of 93.6% and 96.3%, respectively. These models are experimented on a Raspberry PI 5 to simulate running them on edge devices, with software optimization techniques being implemented, which have shown modest improvements in performance. Lastly, to get an insight on the performance of the PI 5 compared to a moderately powered device, a comparison between the inference times of the models on a PI 5 and a MacBook Air M1 is done. This comparison showed a difference of up to 1.49 seconds between them. Future work involves the exploration of hardware-based accelerators, such as the edge TPU and Jetson Nano to improve performance.

14:30 ***ARIA: An Agentic Release Intelligence Architecture for Autonomous End-to-End Mobile DevOps - from Build Prediction to Incident Resolution***

Satyanarayana Gudimetla (Nike India Technology Private Ltd, India); Suresh Gangula (Nike, Inc, USA); Chandrakanth Challa (Jawaharlal Nehru Technological University Hyderabad, India); R Adinarayana (Andhra University, India); Sudhakar Vunnam (India)

Enterprise mobile delivery pipelines remain fragmented: build, CI, release, deployment, and incident response operate as disconnected stages, with post-deployment learning rarely propagating upstream. We present ARIA (Agentic Release Intelligence Architecture), a multi-agent architecture whose central contribution is the ARIA Feedback Protocol (AFP)-a confidence-gated, time-decayed mechanism providing two complementary benefit pathways: (i) per-incident observability acceleration that enriches diagnosis context for every incident, and (ii) multi-incident feedback accumulation that improves upstream test selection over time. ARIA coordinates five specialized agents through a shared provenance substrate and a graduated autonomy model whose governance components are mapped to NIST AI Risk Management Framework practices. A controlled simulation study parameterized from the public CI/CD histories of Signal Android, Firefox for Android (Fenix), and Wikipedia Android reports baseline, ablation, and sensitivity analyses over 30 stochastic replications. The per-incident mechanism yields a statistically significant MTTR reduction from 51.0 to 44.6 minutes (−12.6%, $p<0.001$, $d=0.86$) and governance improvement from 8.3 to 9.1 (+10.1%, $p<0.001$, $d=1.30$). The multi-incident mechanism (DDR +1.0%, RIR −5.8%) shows favorable trends; a post-hoc power analysis determines that ≈670 observations (≈22 applications) are needed for 80%

statistical power, providing a concrete sample-size target for future production studies. An offline replay against 30 held-out CI failure events yields 83.3% prediction accuracy (vs. 60.0% baseline), supporting simulation fidelity.

14:45 ***AI Based Air Quality Forecasting in the Gulf Region***

Mohammed Abdul Rahim, Hatem Tamimi and Jim O Otieno (Higher Colleges of Technology, United Arab Emirates)

Air pollution presents a formidable environmental and public health threat in arid regions, where natural dust events sharply exacerbate particulate matter fluctuations. In order to address these challenges, this study introduces a machine learning framework for dissecting and projecting air quality across the Gulf region by drawing on annual observations of PM2.5, PM10, and NO2. Using support Vector Regression models featuring radial basis function kernels which we derived from engineered temporal features, such as lagged concentrations and rolling averages, we evaluated their efficacy through conventional error metrics while probing the merits of Principal Component Analysis for dimensionality reduction. Results, notably, show that normalised SVR-RBF configurations outperformed PCA-augmented versions for every pollutant examined. NO2 proved most amenable to prediction, whereas particulate matter yielded elevated errors owing to dust-driven volatility. In addition, results showcased that country-specific breakdowns exposed entrenched spatial variations, alongside long-term NO2 outlooks that show stabilisation or gentle downturns by 2030. However, hurdles tied to yearly data sparsity and excluded meteorological factors temper these advances. Nonetheless, the results underscore SVR frameworks' versatility in dusty climes, laying a repeatable foundation for enduring air quality scrutiny and policy insights in the Gulf. Overall, these interconnected modelling endeavours and insights mutually reinforce SVR-RBF's edge over enhanced alternatives, while spotlighting dust's outsized role in hindering particulate forecasts amid arid unpredictability.

15:00 ***Comparative Analysis of LSTM, GRU, Hybrid LSTM-GRU, and XGBoost Models for IoT-Based Flood Forecasting Using Multi-Sensor Environmental Data***

Jerry Incierto Teleron (Technological Institute of the Philippines - Quezon City, Philippines)

This study presents a comparative machine learning approach for flood prediction as an extension of a previously developed IoT-based flood monitoring system. The earlier system enabled real-time acquisition, transmission, and visualization of environmental parameters, including water level, temperature, and humidity. Building upon this foundation, the current study utilizes the collected dataset from 2024 to develop predictive models for future flood forecasting. Four models-Long Short-Term Memory (LSTM), Gated Recurrent Unit (GRU), hybrid LSTM-GRU, and Extreme Gradient Boosting (XGBoost)-were implemented and evaluated using standard performance metrics such as accuracy, F1 score, and Receiver Operating Characteristic (ROC) curve. Experimental results show that XGBoost outperformed the deep learning models, achieving the highest accuracy (73.76%) and F1 score (0.73), while also demonstrating significantly faster training time. Furthermore, the selected model was used to generate flood forecasts for the period of May 2026 to March 2027, highlighting its capability to predict future flood events based on historical IoT sensor data. The findings demonstrate that integrating machine learning with IoT-based monitoring systems enhances predictive capabilities and supports proactive disaster risk management and early warning systems.

15:15 ***Multi-Layer Moisture Content Prediction for Oyster Mushroom Substrate Using Bidirectional LSTM***

Aileen A Sieras (Technological Institute of the Philippines, Philippines)

Accurate moisture content prediction during substrate preparation is critical for optimizing yields and minimizing contamination in mushroom spawn cultivation. To address this challenge, this study presents a direct-wired, multi-layer embedded sensing system coupled with a Bidirectional Long Short-Term Memory (Bi-LSTM) neural network, enabling the prediction of moisture content at three distinct substrate depths (bottom, middle, top) within rice-grain medium. The hardware architecture utilizes a Raspberry Pi Pico for high-frequency (3.4 s) USB-based data acquisition. Building on this, the Bi-LSTM leverages two stacked layers (64 and 32 units) to process sequences bidirectionally, thereby capturing complex, non-monotonic temporal dependencies. Importantly, empirical data from six independent trials comprising 2,044 valid records revealed a consistent vertical moisture gradient driven by capillary action, with the bottom layer exhibiting maximum retention. To operationalize these readings for decision-making, a correlation-optimized Substrate Readiness Index (SRI) (SRI= 0.40 x MCB + 0.35 x MCM + 0.25 x MCT) was formulated, effectively differentiating acceptable batches by a margin of 5.81 percentage points (Wilcoxon $p < 0.001$, $\sigma = 1.17\%$). Model generalization was rigorously evaluated using Leave-One-Session-Out Cross-Validation (LOSO-CV) across eight algorithms. Notably, while traditional models failed to generalize (yielding negative R2), the proposed Bi-LSTM significantly outperformed all baselines, achieving an RMSE = 0.480±0.507%, R^2 = 0.972±0.036, and Within-±2% Accuracy (W2A) = 98.1±2.8%. Furthermore, statistical superiority over traditional algorithms and Transformer-Encoders was confirmed via Friedman's (χ^2 = 42.67, $p < 0.001$) and Holm-Bonferroni-corrected pairwise tests. These findings demonstrate predictive accuracy equivalent to commercial sensors, but with enhanced multi-layer spatial resolution.

Tuesday, April 21 14:00 - 15:30 (Africa/Cairo)

SESSION-2E: RESEARCH: Accessibility, Assistive Technology & Sign Language AI

6 PAPERS

Room: ROOM-E

Chair: Moses Mupeta (University of Zambia, Zambia)

14:00 ***A Vision-Based Deep Learning Framework for Egyptian Sign Language Recognition Using CNN and VGG-16***

Fatma G. El-megharbel and Mennatalla T. Elgamal (MSA University, Egypt); Mazen A. Ebrahim (MSA University, Egypt & University of Greenwich, United Kingdom (Great Britain)); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

Sign language is considered the main linguistic bridge for the Deaf and Hard-of-Hearing community; however, a continuous communication gap still exists because only a limited number of hearing individuals have achieved fluency in sign languages. This study presents an efficient vision-based deep learning system for recognizing Egyptian Sign Language (ESL) using convolutional neural networks (CNNs) and transfer learning techniques. The proposed models were developed using a dataset that contains 20 signs that are frequently used in everyday life in Egypt, expressed in two different forms according to the Egyptian dialect and Standard Arabic. The Model's accuracy was evaluated through the integration of spatial feature extraction using CNNs with temporal modeling to capture sign motion, specifically employing the VGG16 architecture. To measure real world hand and body

landmarks, Media Pipe Holistic was utilized to enhance recognition performance during testing. The experimental results indicate that the VGG16-based model achieved an overall accuracy of 96.12%, while the CNN-based model achieved an overall accuracy of 96%. To further develop and validate the proposed algorithm, a real-time graphical user interface application was created for sign prediction. These results establish the proposed framework as a scalable and effective solution for automating sign language recognition and improving the overall quality of life for individuals in the Deaf and Hard-of-Hearing communities. In this paper, Arabic Sign Language is referred to as ArSL to avoid confusion with American Sign Language (ASL). Index Terms-Keywords Egyptian Sign Language(ESL), hearing impairments, Arabic Sign Language (ARSL), deaf community, Sign Language Recognition(SLR), Convolution al Neural Network (CNN), VGG16, Deep Learning (DL)

14:15 ***Real-Time Egyptian Currency Recognition for Visually Impaired Users***

Yehia Samir and Eyad Mohmed (October University for Modern Sciences and Arts (MSA), Egypt); Abdelrahman Ezzeldin Nagib (MSA, Egypt); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

This paper presents a deep learning-based currency recognition system developed to assist people with vision impairments to identify Egyptian currency in real-time. The proposed system used the YOLOv8 architecture to detect and classify the different denominations of Egyptian banknotes that were captured by a camera under real-world conditions. The dataset used for training the model consists of images from twelve different denominations, with varying backgrounds, different light conditions and a variety of orientations of the banknotes. A number of data augmentation techniques (e.g. rotation, scaling and translation, flipping and perspective-distortion) were utilized in order to enhance generalization and robustness to orientation. Significant detection accuracy and stability was demonstrated by the trained model regardless of the banknote being positioned horizontally, vertically or inclined. The results of the experimental evaluation conducted on the trained model provided evidence that the proposed system provides an average of 98%, with real-time performance on commercial-grade GPU hardware. Users are provided with audio feedback announcing which denomination of currency is being detected by the system. This allows for practical, hands-free use of the system. The results of the study indicate that the proposed system represents an effective assistive solution for identifying currency, and would be appropriate for use in conjunction with wearable assistive technologies for persons who are visually impaired.

14:30 ***Building a Large-Scale Silver Corpus for Modern Standard Arabic Word Segmentation from Kalimat***

Alaaeddine Ramadan (American University of Bahrain, Bahrain); Hussein Ali Awdeh and Mai Khalil (University of Paris 8, France); Omar Saadiyeh (University of Paris 8, Lebanon); Gilles Bernard (University of Paris 8, France); Mohammad Hajjar (Lebanese University, Lebanon)

Reliable large-scale resources for Modern Standard Arabic (MSA) grammatical word segmentation remain scarce. This study presents a silver-standard resource built on Kalimat, a multi-purpose Arabic corpus enriched with automated annotations such as part-of-speech tagging and morphological analysis. The original annotations were found to be insufficiently reliable to be used directly for machine learning. Higher quality segmentations and POS tags were produced by combining Farasa, Stanford POS tagger, Al-Khalil morphological analyzer, with a dedicated category-driven segmentation, on one hand, adjusting the combinations with the help of a small gold annotated subset of Kalimat, and dedicated reliability assessment procedures on the other hand. The resulting resource provides grammatical segmentation for approximately 10 million words. Its qualification as a silver-standard corpus was supported through the evaluation of multiple segmenters and segmenter combinations, together with an indirect methodology for

assessing output reliability. The resource is made freely available for training and evaluating MSA segmentation models, along with a manually controlled subset of context-dependent segmentation for helping fine-tuning and error analysis.

14:45 ***3D Alphanumeric Recognition in Arabic Sign Language Using WiLoR-Based 3D Hand Pose Features***

Yosr Mansour (Elshorouk Academy, Egypt); Asmaa Elsaid (El Shorouk Aacademy, Egypt); Mariam Alaa (El Shorouk Academy, Egypt)

This research proposes a robust Arabic Sign Language (ArSL) recognition approach utilizing 3D hand pose features extracted via the WiLoR framework and MANO parametric model. Departing from traditional RGB or 2D landmark methods, this study leverages a 151-dimensional feature space-including global orientation, articulated pose, and geometric distances-to achieve viewpoint-invariant hand characterization. The discriminative power of these 3D representations was validated through PCA and feature importance analyses, then evaluated using Random Forest, XGBoost, and MLP classifiers across alphabetical and numerical ArSL datasets. Notably, the optimized MLP achieved 95.72% accuracy on letters, while the Random Forest reached 99.75% on numbers. Furthermore, the system demonstrates exceptional computational efficiency, maintaining inference speeds between 837 and 1,126 FPS, proving its viability for real-time deployment on resource-constrained hardware. Visualizations of the latent feature manifolds further confirm that 3D structural data provides a highly effective foundation for high-speed, high-accuracy sign language recognition.

15:00 ***Attention Modulation of P300 and MMN: A Pilot Study Using a 2-Electrode Oddball Paradigm***

Ibrahim Rida Kassem (Lebanese University, Lebanon); Tarek Houri (Beirut Arab University, Lebanon); Katia Sawaya (Beirut Arab University (BAU), Lebanon)

The P300 component is a well-established electro physiological marker of attention and memory updating. However, individual variability in how conscious attention modulates P300 amplitude remains poorly understood. This pilot study examined P300 and Mismatch Negativity (MMN) during automatic (passive) versus conscious (active) auditory oddball tasks using a simplified 2-electrode setup (Fz, Cz). EEG was recorded from six healthy individuals during passive and active auditory oddball paradigms. Automatic artifact rejection removed epochs exceeding 150 µV, flat signals, or inverted polarity. P300 was identified as the maximal positive peak between 300-600 ms. MMN was calculated as the negative difference wave (Deviant minus Standard) between 100-250 ms. Active P300 (n=5) showed mean amplitude of 18.85 ± 10.18 µV at 440 ± 5 ms. Passive P300 (n=3) showed mean amplitude of 6.22 ± 2.16 µV at 445 ± 10 ms. MMN was reliably detected in both conditions (Passive: -3.63 ± 2.15 µV; Active: -2.72 ± 1.19 µV). Notably, all participants with usable paired data (n=2) showed P300 enhancement during conscious attention (mean effect: +4.11 µV). Detection rates exceeded 97% across usable conditions. This pilot study demonstrates the feasibility of a simplified 2-electrode oddball paradigm for assessing attention modulation of P300 and MMN. Preliminary findings suggest consistent P300 enhancement with conscious attention. Results support further investigation with larger samples and provide practical insights for simplified EEG protocols.

15:15 ***Implementation of a Smartphone-Based Multi-sensor Counterfeit Currency Detection***

Anupam Raj and Vikas Upadhyaya (NIIT University, India)

Currency counterfeiting has become a serious problem for financial systems today. Detecting fake notes is difficult for both banks and common people. Earlier methods mostly used machine learning image processing or special hardware devices. However, these solutions are often not easy for

normal users to access or they check only limited features. This paper proposes an improved mobile application that can detect fake currency in real time. The System builds on existing multi sensor methods such as LiDAR, ultraviolet (UV), infrared (IR), magnetic (MG) sensor and Time-of-Flight (TOF) are used to improve accuracy under different conditions. This application is designed to be lightweight and easy to use even people without technical knowledge can use it. The proposed system combines multiple sensor data with other techniques like ORB feature matching, color analysis, anomaly detection and size and dimension checking. This makes the detection process stronger and more reliable.

Tuesday, April 21 15:30 - 16:00 (Africa/Cairo)

COFEE BREAK & NETWORKING

COFEE BREAK ROOM - NETWORKING

Room: COFFEE_ROOM

Tuesday, April 21 16:00 - 18:00 (Africa/Cairo)

SESSION-3A: RESEARCH: Frameworks and Innovations in Governance, AI Maturity, and Decision-Making in Complex Systems

8 PAPERS

Room: ROOM-A

Chair: Tejas Pravinbhai Patel (Amazon, USA)

16:00 ***Who Is Next in a Crowded ED? Design of a Medical Triage Stochastic Decision-Making Game for Exploring Policy Trade-offs***

Adedolapo Kehinde Adebayo (Norfolk State University, USA); Isaac O Osunmakinde (Norfolk State University, USA & None, USA)

Emergency departments must allocate limited beds and staff under uncertainty while balancing patient survival, workforce strain, and institutional reputation. Existing triage games focus on case-level decisions, and most simulation studies optimize narrow operational metrics, offering limited support for exploring multi-objective trade-offs in stochastic environments. This paper presents a Medical Triage Stochastic Decision-Making Game that integrates a simplified emergency department model with a configurable triage policy module and a web-based interactive interface. The system models noisy severity cues, probabilistic deterioration and death, and resource constraints across discrete rounds, with four available actions per patient: Treat Now, Monitor, Defer, and Transfer. Performance is summarized using three composite metrics-Survival Score, Staff Stress, and Reputation. We report results from 2,000 stochastic replications per condition comparing three decision strategies across Basic and Stochastic difficulty regimes. Treat-heavy and monitor-heavy strategies fail to meet win thresholds due to stress or reputation collapse, while a capacity-aware Treat+Transfer strategy achieves the highest success rates (57.4% in Basic; 44.8% in Stochastic). These findings illustrate how explicit triage policies

interact with congestion and uncertainty, and demonstrate the value of simulation-based serious games for teaching policy-level trade-offs in emergency care.

16:15 ***Governance-Aware Agentic Retrieval-Augmented Generation with Evidence Abstention for Reliable AI Policy Question Answering***

Ashish Hastimal Jain (University of Central Florida, USA); Suhas Reddy Yarrabothu, Nanda Kishore Kande and Dharmateja Reddy Kethireddy (USA)

Answering reliable questions over AI governance documents is still challenging since the traditional retrieval-augmented generation systems tend to provide responses that cannot be supported, do not differentiate between answerable questions and overly broad policy queries, and do not utilize governance indicators, including legal status, jurisdiction, and source authority. The paper introduces a governance-conscious agentic retrieval-augmented generation model with evidence abstinence to policy-sensitive question answering that answers only when the evidence retrieved is well-grounded and refers the cases of uncertainty to ABSTAIN / HUMAN REVIEW. It is based on the AGORA corpus of AI governance documents, retrieval of segments, a hybrid retriever that uses semantic similarity, lexical overlap and governance sensitive reranking, and a query risk layer of broad or globally framed queries. To preserve meaning of ground, the final answer module employs an extractive evidence-based composer rather than unconstrained generation. The proposed framework attained an almost perfect decision accuracy, 0.8583 mean retrieval keyword coverage on answerable queries, and 0.6833 mean answer keyword coverage with the selected extractive composer (0.5083) on a curated benchmark. These findings reveal a new and useful direction of the creation of reliable and policy-appropriate GenAI systems of the governance-seeking question answering.

16:30 ***Architecting Multi-Model Agentic AI Systems: A Taxonomy and Reference Architecture for Specialized LLMs in Production Agents***

Sandeep Shivam (Tavant, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA); Venkat Nutalapati (Walmart Global Tech, USA); Vinay R Soni (USA); Vijayakumar Venganti (Cisco Systems, Inc., USA)

Modern AI agents deployed in enterprise and cloud environments are rarely powered by a single monolithic large language model. Production systems require heterogeneous compositions of specialized models, each optimized for distinct computational roles: reasoning, perception, planning, memory, and action. This paper makes three primary contributions. First, we introduce the Multi-Model Agentic Systems Taxonomy (MMAS-T), formally characterizing six model classes General-Purpose LLMs, Mixture-of-Experts (MoE), Vision-Language Models (VLMs), Large Reasoning Models (LRMs), Small Language Models (SLMs), and Large Action Models (LAMs) along four dimensions: functional role, computational profile, modality, and deployment context. This taxonomy extends our prior Enterprise AI Model Selection Taxonomy (EMST) [A] into the multi-agent orchestration domain. Second, we propose a four-layer Reference Architecture (MMAS-RA) decomposing agentic systems into Perception, Cognition, Memory, and Action layers, with explicit model-to-layer assignments and four reusable design patterns. Third, we present a comparative empirical evaluation of three agent configurations on a realistic enterprise workflow-automation scenario. Results show that heterogeneous multi-model architectures achieve 38.8-41.8% end-to-end latency reduction and 32.0-41.0% operational cost reduction versus single-model baselines, with task success rates improving from 91.2% to 96.8%. Statistical significance is confirmed via paired t-tests ($p < 0.01$) on 200 task executions.

16:45 ***A Responsible GenAI Maturity Model for Delivery Governance: A Five-Level Framework for PMOs, Agile Coaches, and Engineering Leaders***

Ritesh H Ruparel (CSG International, USA); Sai Mohan Dasari (University of the Cumberlands, USA); Vinod Bottu (Microsoft, USA); Vigneshwarr Venkatesan (USA); Suraj Vangala (Charter Communications, USA)

Project teams use generative AI to write status updates, summarize meetings, draft risks, and prepare stakeholder messages. These uses can save time, but they also introduce risks. Teams may share sensitive data, publish statements that are not supported by project evidence, or rely too much on AI output. Many organizations do not have a practical way to assess and improve responsible AI use in daily delivery workflows. This paper presents a five level maturity model for responsible generative AI use in delivery governance, from ad hoc use to a managed and continuously improved program. We ground the model using desk evidence and alignment with established risk, security, privacy, and project governance guidance. We provide design principles, a threat model, standards mapping, assessment dimensions, and actionable artifacts including a checklist and a worked example. We also report a small pilot practitioner review (N=3) that provides early evidence of clarity and operational fit.

17:00 ***Evidence-Based Generative AI for Incident Command Updates: A Practical Method to Prevent Overclaim***

Ritesh H Ruparel (CSG International, USA); Sai Mohan Dasari (University of the Cumberlands, USA); Vigneshwarr Venkatesan (USA); Vinod Bottu (Microsoft, USA); Suraj Vangala (Charter Communications, USA)

Engineering teams increasingly use generative AI to draft incident updates for internal and external stakeholders. These drafts can save time during high pressure events, but they also introduce risk. AI generated updates may overclaim resolution, imply a root cause without evidence, commit to timelines that are not justified, or include sensitive details. This paper presents an evidence-based method for safe AI assisted incident communication. The method combines (i) an overclaim taxonomy, (ii) a risk tiering and approval matrix, (iii) an incident update template that enforces evidence links and explicit unknowns, and (iv) a safe-to-send scoring rubric. We ground the method through desk evidence and alignment with established risk, security, and incident management guidance. We provide a worked example, a scenario pack, and a small pilot practitioner review that provides early evidence of clarity and operational fit without using proprietary incident data.

17:15 ***Dependency-Aware Post-Quantum Cryptography Migration Planning: Graph-Based Optimization for Enterprise Infrastructure***

Saurabh Sharma (Cisco Systems, USA); Karthik Pappu (Dakota State University, USA); Ankit Parashar and Navin Suvarna (Cisco Systems, USA)

The transition to post-quantum cryptographic (PQC) algorithms requires enterprises to migrate thousands of cryptographic assets from vulnerable algorithms (RSA, ECDSA, DH, DSA) to NIST-standardized replacements (ML-KEM, ML-DSA), but uncoordinated migration introduces interoperability failures when migrated assets cannot communicate with unmigrated dependents. This paper presents an automated migration planning platform that builds a cryptographic dependency graph, computes quantum risk scores with graph centrality, and generates an optimized migration plan using reversed topological ordering with risk-weighted prioritization. A factorial experiment (3 scales × 3 strategies × 10 seeds = 90 runs) on simulated enterprises with 100 to 10,000 assets shows that the dependency-aware strategy reduces cumulative quantum risk exposure 26% faster than random ordering ($p < 0.001$), eliminates 88.6% of interoperability failures ($p = 0.028$), and scales from 14% AUC improvement at 100 nodes to 41% at 10,000

nodes. The core mechanism is processing dependency graph leaves first, ensuring that when a server transitions to PQC, its dependent applications already speak the new protocol.

17:30 ***Autonomous AI-Powered Resource Management for Apache Flink on Amazon EKS***

Hari Krishna Pokala (CIGNA, USA); Venkata Pavan Kumar Gummadi (Broadridge, USA)

Apache Flink supports millions of real-time streaming applications daily, powering applications such as fraud detection, sensor streaming, and live advertising bidding engines. Such streaming applications need to scale dynamically as workloads change drastically within seconds while also managing cloud costs. Existing solutions are often too slow to react. Scaling occurs only after incurred delays or idle resource containers sit around wasting money. In this paper, we introduce a practical AI-powered solution to forecast workload changes, detect anomalies early both in data streams and application execution, and intelligently act on anomalies to reduce costs while meeting strong performance requirements. We combine forecasting methods, anomaly pattern detection, and learning-based cost optimization to demonstrate improvements across three realistic streaming workloads. Compared to popular reactive solutions, our system achieves higher cluster resource utilization, lower application latency, faster anomaly detection rates, and up to 40% reduction in cloud costs. Implementation details are discussed with practicality in mind. Explanations are provided on how to build using existing AWS managed services along with open-source solutions, and how it can be deployed in just a few steps on any standard Amazon EKS cluster.

17:45 ***Uncertainty-Aware Deep Neural Networks for Reliable and Selective Decision-Making in Medical Image Classification***

Krishna Kishor Tirupati (USA); Naresh Kumar Methuku (Fidelity Investments, USA); Pradeep Kumar Chilukury and Sai Reddy Busi Reddy (USA)

Artificial intelligence systems used in safety critical areas tend to make overconfident predictions without knowing their own uncertainty, making them unreliable for decision making in real life. This work proposes an uncertainty-aware deep learning framework for robust medical image classification based on the DermaMNIST dataset, and employs the Monte Carlo Dropout technique within a convolutional neural network to model the uncertainty of the prediction. Besides traditional classification, the approach allows calibrated confidence estimation and selective prediction. Experimental results obtain accuracy of 0.7362 with macro F1-score of 0.4311 and good calibration of Expected Calibration Error of 0.0363. Uncertainty analysis shows just how much incorrect predictions have uncertainty which almost double (0.0020109) the uncertainty of correct predictions (0.0010636). Furthermore, selective prediction increases reliability significantly in accuracy of prediction with 0.9302 at 40% coverage and 1.0000 at 10% coverage. These findings signal the benefits of uncertainty-aware models in trustworthy, risk-sensitive, AI decision-making.

Tuesday, April 21 16:00 - 18:00 (Africa/Cairo)

SESSION-3B: RESEARCH: Innovations in AI Governance, Security, and Testing Frameworks for Cloud and IoT Systems

Room: ROOM-B

Chairs: Chaitanya Kulkarni (Oracle America Inc, USA), Jay Bharat Mehta (Cleveland State University, Alumni, USA)

16:00 ***Future-Proofing Identity Security for Agentic AI Systems: Design, Implementation, and Evaluation of an Identity Fabric***

Karthik Pappu (Dakota State University, USA); Badal Bhushan (I.E.T. M.J.P. Rohilkhand University, India); Nilesh Jaiswal (Electronic Arts, USA)

Agentic AI systems composed of dynamically instantiated, tool-enabled agents operating across trust boundaries introduce identity failures that exceed the assumptions of existing identity and access management frameworks. This paper identifies five identity gaps in agentic AI systems and introduces six composable security primitives that address ephemeral identity, delegation integrity, capability-level control, cross-domain trust, behavioral assurance, and content provenance: Ephemeral Attested Agent Identity (EAAI), Intent-Bound Delegation Tokens (IBDT), Identity-Constrained Capability Sandboxes, Decentralized Agent Identity Registries (DAIR), Continuous Behavioral Attestation, and Content Provenance Binding. We implement the architecture on Kubernetes using SPIFFE/SPIRE workload identity and mutual TLS for all inter-service communication. Across 33 mTLS test cases (n = 100 per operation), the system achieved 100% delegation fidelity (20/20 narrowing tests), sub-25 ms mean latency for credential issuance and delegation, and 6/6 abuse-resistance rejections. In an end-to-end pipeline with Claude Sonnet, total identity overhead was 203 ms, representing 2.6% of inference latency.

16:15 ***Confidence-Aware Adaptive Neural Computation for Energy-Efficient Deep Learning***

Dharmateja Reddy Kethireddy and Nanda Kishore Kande (USA); Ashish Hastimal Jain (University of Central Florida, USA); Suhas Reddy Yarrrabothu (USA)

Deep neural networks provide good predictive performance but have fixed-depth inference pipelines which introduce unnecessary computational burden to easy samples and shall limit energy-efficient deployment in resource constrained intelligent systems. This paper introduces a confidence-aware adaptive neural computation framework for energy-efficient image classification. where the multi-exit deep architecture includes dynamically changing inference depth which depends on the difficulty of the input without executing the entire model for all instances. Using Tiny ImageNet as a challenging benchmark, the proposed approach relies on the combination of early-exit classifiers and threshold-based adaptive routing to ensure a balance between the predictive performance and computational efficiency. Experimental results demonstrate an exit-wise improvement and the test accuracy of Exit 1 is 17.82%, the test accuracy of Exit 2 is 30.13% and the test accuracy of the final exit is 61.49%. The best value of the threshold setting [0.5, 0.6] in adaptive inference achieves an accuracy of 61.33% with an average exit depth of 2.70, compared with an accuracy of 60.18% of inference full with depth of second exit, 3.00, while the routing is 8% and 14% of samples through first and second exits, respectively. These results show the novelty and practical importance of confidence guided adaptive inference as a promising direction for sustainable, computation-aware deep learning.

16:30 ***Governance of Generative AI in the Enterprise: Prompt Lifecycle Management and Secure Middleware Design***

Subhash Tatavarthi (Kasmo, USA); Saikrishna Tarakampet (California's Correctional Healthcare Services, USA); Venkatasatyaravikiran Bikkavolu, Raghunath Reddy Koilakonda and Manoj Gudala (USA)

Organizations are rapidly adopting Generative AI into enterprise decision-making processes; however, implementation has exposed governance gaps not addressed by traditional software control approaches. Large Language Model behavior is determined not only by model weights but by prompts, contextual instructions, retrieval configurations, and user intent. These elements evolve continuously without structured tracking, versioning, or auditable review mechanisms. Integration of Generative AI with enterprise platforms such as Snowflake, Databricks, Oracle, and ERP systems introduces risks including policy drift, unintentional data exposure, unpredictable reasoning behavior, and erosion of trust in AIsupported decisions. This paper proposes a governance framework combining Prompt Lifecycle Management (PLM) and Secure Middleware Design to establish stable, compliant, and scalable enterprise GenAI systems. Treating prompts as governed enterprise assets and middleware as the enforcement layer enables predictable, auditable, and secure AI operations across multiple domains and models.

16:45 ***A TinyML-Based Out-of-Distribution Security Monitor for Energy-Constrained Industrial and Medical IoT Systems***

Viswanathan Ranganathan (Netflix, USA); Arun Kumar Elengovan (Okta Inc. and IEEE Senior Member, USA); Venkat Nutalapati (Walmart Global Tech, USA); Deepak Kole (IEEE Senior Member, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA); Nandagopal Seshagiri (Okta Inc., USA)

The sudden increase in industrial and medical Internet of Things (IoT) systems deployment has created serious cybersecurity and reliability problems under tight energy and resource constraints, where unknown attacks and device failures are to be detected at the edge itself. In this work, a TinyML-based out-of-distribution (OOD) security monitoring framework is introduced for energy-constrained IoT devices, where normal operational behavior is learned in an unsupervised way and deviation is detected without the need for previous attack signatures. A lightweight autoencoder in conjunction with a confidence-weighted OOD scoring mechanism and persistence-based decision logic is used to achieve robust detection without continuous high-rate sensing. The proposed approach achieves a good performance in the detection, with a ROC-AUC of 0.9257 and a PR-AUC of 0.8913 for OOD detection, with no false escalation in normal operation. Deployment-time evaluation shows timely detection with an average latency of around 41.5 windows and 61.65% system-level energy reduction as compared to always-on monitoring. The whole framework is implemented with an INT8-quantized TinyML model with only 2.98 KB memory footprint, that can be used for practical on-device deployment. By jointly integrating TinyML, confidence-aware OOD detection and adaptive energy-aware security actions, this work addresses one critical gap in sustainable and reliable protection of industrial and medical cyber-physical IoT systems.

17:00 ***Extending TMMi for FinOps: A Test Maturity Framework for Cloud Cost Governance***

Lakshmi Vidya Peri (TMMI, USA); Deepak Pai (USA); Yogesh S Thanvi (Akamai Technologies, USA)

The Test Maturity Model integration (TMMi) defines five maturity levels and 16 process areas for software test process improvement, but does not address cloud cost governance testing. This paper extends TMMi by proposing 16 FinOps Test Process Areas (FTPAs) distributed across Levels 2 through 5, mirroring TMMi's 5+5+3+3 structure. The FTPAs are grounded in the FinOps Foundation Framework and the FinOps Open Cost and Usage Specification (FOCUS). The extension includes four formal testable properties (completeness of cost attribution, budget compliance, idle resource

constraint, and detection of cost drift), eight Level 4 metrics, including Mean Time to Savings (MTTS), and TAMAR-compatible assessment criteria, allowing organizations to evaluate and improve their maturity in cloud cost governance testing within the established TMMi framework.

17:15 ***NVMe Validation Maturity: A TMMi-Aligned Framework for Testing Process Improvement***

Lakshmi Vidya Peri (TMMI, USA)

Non-Volatile Memory Express (NVMe) storage devices have become foundational components in modern computing infrastructure, yet the validation processes ensuring their quality remain fragmented and lack systematic maturity assessment frameworks. While the Test Maturity Model integration (TMMi) has successfully guided software testing process improvement across industries, its application to hardware-firmware validation contexts remains unexplored. This paper addresses this gap by proposing a TMMi-aligned maturity framework specifically adapted for NVMe validation. Through the synthesis of NVMe validation practices and TMMi adoption patterns in the literature, we define five maturity levels with hardware-specific process areas that cover firmware testing, performance characterization, compliance verification, and fault injection techniques. The framework provides implementation guidance, tooling recommendations, and projected benefits based on analogous domain adoption patterns. This work presents the first TMMi-aligned maturity framework for NVMe validation.

17:30 ***Agentic Dashboard Generation from Natural Language on Lakehouse Platforms***

Hari Krishna Pokala (CIGNA, USA)

The emergence of agentic AI systems is reshaping business intelligence (BI) workflows by enabling natural- language-driven generation of analytical artifacts. Recent advances in large language models (LLMs) and retrieval-augmented generation (RAG) have enabled systems that combine reasoning with grounded data access [1], [2]. This paper presents an empirical evaluation of an agentic dashboard generation approach on a lakehouse data platform, where natural-language queries are translated into multi-visualization dashboards under governance constraints. We propose a reference architecture that integrates metadata aware retrieval, retrieval-augmented generation, and governed data access mechanisms to support automated dashboard construction. The system leverages catalog-level metadata, semantic retrieval, and iterative agent planning to generate SQL queries, select visualization types, and assemble dashboards. An experimental evaluation is conducted across representative enterprise workloads, including financial analytics, sales reporting, and operational monitoring. Metrics include task completion rate, output correctness, latency, and error characteristics. Results indicate that agentic generation can reduce dashboard authoring time relative to manual workflows while preserving access controls enforced by the underlying governance layer. We further analyze failure modes such as schema ambiguity and visualization mismatches and discuss trade-offs between automation and analytical accuracy. The findings highlight both the potential and current limitations of agentic BI systems, particularly in relation to metadata quality and system interpretability.

17:45 ***A Secure CI/CD Pipeline using GitHub Actions and Open Policy Agent (OPA) for Kubernetes Applications***

Hari Krishna Pokala (CIGNA, USA)

As Kubernetes deployments become pervasive, risk of developer/configuration error causing severe misconfigurations continues to grow. Common mistakes leading to privilege escalation, resource abuse, supply-chain attacks, and non-compliance are unfortunately easy to introduce given YAML

manifests and Helm templates. Current tooling shifts security left via scanning in CI/CD pipelines, but even this occurs after faulty configurations are checked in. This intermediate-level paper/demo shares an actionable secure CI/CD pipeline that uses GitHub Actions and integrates OPA as Code using Contests. Automatically validate Kubernetes manifests and Helm templates at pull-request time using expressive, declarative Rego policies. Enforce Policies-as-Code around core risks such as latest image tag, resource requests/limits, privileged containers, non-approved registries, etc. This Cloud Native Ready solution is tested with Amazon Elastic Kubernetes Service (Amazon EKS), leverages images hosted in Amazon Elastic Container Registry (Amazon ECR), and utilizes IAM roles for service accounts (IRSA). Optionally scan container image vulnerabilities with Trivy. Evaluation of violation detection rate (100% for 10 seeded violations) and runtime overhead demonstrates clear feedback and warning messages back to developers with minimal impact on workflow (fourteen seconds average complete runtime). Ideal balance between static linting tools and runtime admission controllers/Gatekeeper for GitHub/Kubernetes intermediate teams adopting DevSecOps principles.

Tuesday, April 21 16:00 - 18:00 (Africa/Cairo)

SESSION-3C: RESEARCH: Advancements in AI for Financial Risk, Healthcare Detection, and Cloud Data Governance

8 PAPERS

Room: ROOM-C

16:00 ***Zero-Hop Semantic Retrieval on AWS Lambda Using In-Memory Vector Indexing***

Mythili Annamalai Sekar (Amazon Inc, USA); Siva Kumar Chintham (LTM, USA); Mayilsamy Palanigounder (NTT Data, USA); Murali Shankar Dulam, Narender Reddy Bitla and Naga Surya Kamala Dattatreya Pasupuleti (JPMorgan Chase, USA); Abhirup Mazumder (Amazon, USA); Anurag Kumar (Cisco Systems, USA)

Serverless computing provides elastic scalability for AI inference workloads but introduces performance instability due to cold starts, execution environment replication, and repeated dependency initialization. E-commerce product search systems deployed on AWS Lambda typically query external vector databases during each invocation, significantly increasing tail latency under burst concurrency and multi-tenant workloads. This paper presents an execution-environment-aware architecture using AWS Lambda Managed Instances that initializes vector indices and a lightweight embedding model in-memory during cold start, eliminating all runtime external service dependencies. Warm containers via Capacity Provider enable zero-hop similarity search using local embedding inference and in-memory numpy dot product across product catalogs. Experimental results on a 100K product catalog demonstrate 68-94% reduction in p99 latency depending on cache state, 100% elimination of runtime external calls, and 61% lower total infrastructure cost compared to OpenSearch-based baselines, while maintaining stable throughput under a workload of 100 concurrent requests at 1,000 queries per minute. The proposed method provides a practical, cost-efficient design for serverless vector retrieval and RAG pipelines.

16:15 ***Detection of COPD-like Breathing Patterns from Non-Invasive Pressure and Flow Data Using Machine Learning and Deep Learning***

Sarvina Mutiya Seidu (Georgia Southern University, USA); Hayden Wimmer (Georgia Southern University, USA & Institute for Health Logistics and Analytics, USA); Jie Du (Grand Valley State University, USA)

Chronic obstructive pulmonary disease (COPD) remains a major public health burden, with early detection being critical to improving outcomes. Traditional diagnostic tools often rely on spirometry, which can be invasive, effort-dependent, and unsuitable for large-scale screening. This study explores the use of non-invasive respiratory pressure and flow signals to classify COPD-like breathing patterns using machine learning (ML) and deep learning (DL) approaches. A dataset of simulated respiratory waveforms was segmented into fixed-length windows, from which statistical and physiological features were extracted. Four traditional ML classifiers, including Random Forest, Logistic Regression, Support Vector Machine, and k-Nearest Neighbors, were evaluated alongside DL architectures such as Convolutional Neural Networks and Long Short-Term Memory networks. Model interpretability was enhanced using SHapley Additive exPlanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME), which highlighted pressure mean, volume, and pressure variability as key predictive features. These findings demonstrate that integrating signal-derived features with advanced ML and DL models enables accurate and interpretable detection of COPD-like respiratory dynamics. The proposed framework provides a step toward non-invasive, automated screening methods that could complement clinical assessments and support early intervention strategies.

16:30 ***The Yellow Brick Road to AGI: A Detailed Path to Achieving Autonomous General Intelligence with Rapid Testing and Results Management***

Gregory David Spehar (GiDanc AI LLC, USA); Akshay Mittal (University of the Cumberlands, USA)

The race toward Artificial General Intelligence (AGI) has produced extraordinary advances in cognitive capabilities while leaving a critical gap: the governance infrastructure necessary for safe AGI deployment remains virtually nonexistent. This paper presents the first comprehensive framework linking AGI requirements analysis with runtime behavioral governance, arguing that governance is not an optional safety layer applied after AGI is achieved, but a constitutive requirement for AGI itself. We synthesize five competing schools of AGI definition---the Capabilities School (DeepMind Levels), the Cognitive Science School (Hendrycks CHC-based AGI Score), the Economic Value School (OpenAI charter), the Process-Oriented School (Marcus et al.), and the Embodiment School---into a unified taxonomy of 25 requirements across seven domains. Through systematic assessment, we demonstrate that the global AI industry achieves near-zero capability in the Safety, Ethics, and Governance domain (Domain~7), with the Future of Life Institute's AI Safety Index confirming that no frontier company scored above D in existential safety planning. We then present the AI Assess Tech Governance Framework---comprising the LCSH (Lying, Cheating, Stealing, Harm) multi-dimensional behavioral assessment, constitutional separation of powers across specialized AI agents, economic mortality alignment, temporal ethical drift detection, and cryptographic verification infrastructure---as a concrete, patent-protected, production-deployed implementation achieving 100% coverage of Domain~7 requirements. We propose the ``Yellow Brick Road'' model: a phased path to AGI that treats governance infrastructure as foundational waypoints rather than afterthoughts, analogous to how aviation safety protocols are constitutive of the aviation system rather than add-ons. The framework is deployed at \texttt{aiassesstech.com}, verified on Ethereum mainnet, and protected under eight provisional patent applications.

16:45 ***Governed Agentic Cloud Data Pipelines with MCP Gateways***

Aswathnarayan Muthukrishnan Kirubakaran (Meta Inc, USA); Vinoth Punniyamoorthy (JPMorgan Chase, USA); Akshay Deshpande (Apple, USA); Adithya Parthasarathy (Newyork University, USA); Nachiappan Chockalingam (Meta, USA); Mythili Annamalai Sekar (Amazon Inc, USA); Naga Surya Kamala Dattatreya Pasupuleti (JPMorgan Chase, USA); Shiva Kumar Reddy Carimireddy (Fidelity Investments, USA)

Autonomous agents are increasingly used to orchestrate cloud data and machine learning pipelines to improve adaptability and operational efficiency. However, most agent-tool integrations rely on loosely specified interfaces and prompt-driven invocation, providing limited guarantees on correctness, security, and governance. This restricts adoption in production and regulated environments where auditability, policy compliance, and operational reliability are mandatory. We present a governed orchestration architecture in which all agent actions are mediated by a Model Context Protocol (MCP) gateway that enforces bidirectional JSON Schema validation, semantic compatibility, and policy constraints while recording append-only telemetry for auditing and root cause analysis. The gateway establishes a protocol-level trust boundary between autonomous reasoning and production tools. We implement a prototype and evaluate it on open datasets and workloads across batch and micro-batch pipelines with controlled failure and violation injection. Results show that MCP mediation improves pipeline success rate by up to 7.6%, reduces mean time to recovery by 58.3%, blocks over 96% of injected schema, policy, and security violations, and introduces less than 6.2% p95 end-to-end latency overhead. These findings indicate that MCP-based gateways enable reliable and auditable deployment of agentic orchestration in production cloud pipelines.

17:00 ***Transformer-Based Predictive Maintenance for Risk-Aware Instrument Calibration***

Adithya Parthasarathy (Newyork University, USA); Aswathnarayan Muthukrishnan Kirubakaran (Meta Inc, USA); Akshay Deshpande (Apple, USA); Ram Sekhar Bodala (Amtrak, USA); Suhas Malempati (The Cato Corporation, USA); Nachiappan Chockalingam (Meta, USA); Vinoth Punniyamoorthy (JPMorgan Chase, USA); Seema Gangaiah Aarella (Austin College, USA & University of North Texas, USA)

Accurate calibration is essential for instruments whose measurements must remain traceable, reliable, and compliant over long operating periods. Fixed-interval programs are easy to administer, but they ignore that instruments drift at different rates under different conditions. This paper studies calibration scheduling as a predictive maintenance problem: given recent sensor histories, estimate time-to-drift (TTD) and intervene before a violation occurs. We adapt the NASA C-MAPSS benchmark into a calibration setting by selecting drift-sensitive sensors, defining virtual calibration thresholds, and inserting synthetic reset events that emulate repeated recalibration. We then compare classical regressors, recurrent and convolutional sequence models, and a compact Transformer for TTD prediction. The Transformer provides the strongest point forecasts on the primary FD001 split and remains competitive on the harder FD002-FD004 splits, while a quantile-based uncertainty model supports conservative scheduling when drift behavior is noisier. Under a violation-aware cost model, predictive scheduling lowers cost relative to reactive and fixed policies, and uncertainty-aware triggers sharply reduce violations when point forecasts are less reliable. The results show that condition-based calibration can be framed as a joint forecasting and decision problem, and that combining sequence models with risk-aware policies is a practical route toward smarter calibration planning.

17:15 ***OpsAgent: A Governance-Aware Agentic AI Framework for Autonomous Operational Reliability***

Sabitha Muppuri and Chiranjeevisantosh Madugundi (Palo Alto Networks, USA); Ramkinker Singh (Palo Alto Networks, Inc, USA)

This paper presents OpsAgent, a four-layer agentic AI framework for autonomous operational reliability that integrates Observation, Reasoning, Execution, and Governance into a unified architecture. OpsAgent formulates operational reliability as a constrained multi-objective optimization balancing system availability, mean time to recover (MTTR), and policy compliance rate (PCR). A formally defined Hierarchical Policy Graph (HPG),

implemented as an Open Policy Agent (OPA) Rego bundle, encodes approval levels, change-management windows, and regulatory requirements. Evaluation across 1,074 fault-injection scenarios in three trace-grounded simulated environments over 30 days shows that OpsAgent reduces MTTR by 67.3% ± 4.1%, improves availability by 0.38-0.43 percentage points, and sustains PCR ≥ 98.5% relative to a rule-based baseline. Ablation studies confirm that the LLM Orchestrator contributes the largest MTTR reduction (−3.5 min), while the HPG Governance Layer provides the largest PCR gain (+11.1 pp) at minimal latency cost. Sensitivity analysis, cost characterization, and escalationpath validation further demonstrate the framework's practical viability.

17:30 ***Artificial Intelligence-Based Default Loan Prediction for Financial Risk Assessment in Digital Lending***

Sandeep Shivam (Tavant, USA)

This paper discusses the increasing number of loan applications in the banking industry and the obstacles encountered by financial organizations in making informed loan decisions. It presents a machine learning (ML) approach that uses historical loan data to predict loan approval using various classification models. This paper suggests an AI-based model of default loan prediction with the Lending Club dataset to improve the process of assessing financial risks in the digital lending market. To improve the model's performance, the approach uses a variety of data preprocessing techniques, including data cleaning, one-hot encoding, feature selection, data normalization, and data balancing. Two ensemble ML models include XGBoost and AdaBoost, and they are used to classify loan status and predict the possibility of default. The models are evaluated using a variety of performance measures, including ROC-AUC, F1-score, recall, accuracy, and precision. The experimental findings demonstrate that the XGBoost model is superior with an accuracy (acc) of 99.60, precision (prec) of 99.93, recall (rec) rate of 99.85, F1-score (F1) of 99.88, and AUC of 0.9986. However, the AdaBoost model also achieves decent results, with an accuracy of 98.65% and an AUC of 0.9957. The effectiveness of the suggested approach is additionally supported by comparison with the existing models. The results demonstrate that the suggested AI-based framework can contribute to enhancing loan default forecasting and enabling credible financial risk evaluation within digital lending mechanisms to a considerable degree.

17:45 ***A Comparative Performance Analysis: Vector Search vs Graph Databases for RAG Applications***

Siva Prasad Nandi (Oracle, USA); Vijayakumar Venganti (Cisco Systems, Inc., USA); Vinay R Soni (USA); Venkat Nutalapati (Walmart Global Tech, USA)

Retrieval-Augmented Generation (RAG) systems are increasingly critical in enterprise AI deployments, requiring efficient and semantically accurate retrieval mechanisms to supply Large Language Models (LLMs) with domain-specific context. While vector databases have emerged as the dominant paradigm for embedding-based retrieval, graph databases offer compelling advantages in explicit semantic relationship modeling and multi-hop reasoning across interconnected knowledge corpora. This paper presents a rigorous comparative performance analysis contrasting vector search systems specifically FAISS with HNSW indexing and Milvus against graph database retrieval using Neo4j, across three heterogeneous real-world datasets: Wikipedia QA (8.8M passages), PubMed biomedical abstracts (2.3M documents), and a financial regulatory corpus (500K documents). We evaluate across five performance dimensions: query latency (p50/p95/p99), throughput (QPS), Recall@10, nDCG@10, and multi-hop reasoning accuracy. Our experimental results demonstrate that vector databases achieve 8.4ms median latency with 5,200 QPS throughput and a Recall@10 of 0.92 under high-volume semantic search workloads. In contrast, graph databases excel in relational reasoning, achieving 89% multi-hop accuracy versus 63% for vector approaches, at the cost of higher latency (31.2-52.3ms) and lower scalability (1,350 QPS at 1M documents, degrading to 920 QPS at 10M). We also examine Oracle Autonomous Database's unified AI Vector Search and Property Graph capabilities as a representative hybrid

platform, discuss indexing strategies, embedding model selection, and graph construction methodologies, and derive actionable guidelines for RAG architects navigating technology selection.

Tuesday, April 21 16:00 - 18:00 (Africa/Cairo)

SESSION-3D: RESEARCH: Integrating AI for Fairness, Automation, and Security in Cloud and DevOps Environments

8 PAPERS

Room: ROOM-D

Chairs: Pavan Nutalapati (Oracle Corp, USA), Sandeep Shivam (Tavant Technologies, USA)

16:00 ***Behavioral vs. Demographic Predictors of Customer Purchase Frequency: A Multi-Model Analytical Approach***

Nabeel M Sheikh (Oracle America, USA); Taranpreet Kaur (Amazon, USA)

Customer purchase frequency is a critical metric for retention strategies, loyalty programs, and marketing opti- mization. This study analyzes a Customer Purchasing Behaviors dataset containing demographic variables (age, annual income, region) and behavioral attributes (purchase amount, loyalty score) to identify the most influential factors impacting purchase frequency. Using R Studio, we conducted a comprehensive analyt- ical workflow consisting of exploratory data analysis, correlation testing, ANOVA for regional differences, and three predictive modeling techniques: Linear Regression, Poisson Regression, and Random Forest Regression. Models were evaluated using RMSE, MAE, deviance, and feature importance scores. Results show that behavioral factors-specifically loyalty score and purchase amount-have the greatest impact on purchase frequency, while demographic variables (age, income) demonstrate comparatively less predictive power. Regional differences show minimal signif- icance based on ANOVA results. Random Forest feature impor- tance analysis confirms behavioral variables as primary drivers of purchasing behavior. These findings suggest that customer behaviors and spending preferences are more reliable indicators of buying frequency than demographic traits, enabling brands to design more effective targeting and customer segmentation strategies.

16:15 ***SATHOS: Self-Adaptive Trust-Hierarchical Orchestration for Zero-Trust DevSecOps Pipelines***

Nilesh Jaiswal (Electronic Arts, USA); Karthik Pappu (Dakota State University, USA); Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA)

Modern DevSecOps pipelines execute workloads across heterogeneous environments, including cloud runners, on-premises agents, and ephemeral containers, yet orchestration frameworks assume pre-trusted execution agents and rely on static policy gates. This assumption fails under agent compromise, configuration drift, and trust asymmetry between cloud and on-premises zones. We present SATHOS (Self-Adaptive Trust-Hierarchical Orchestration System), a zero-trust orchestration framework that models CI/CD pipelines as trust-governed directed acyclic graphs (DAGs) where each node's execution is conditioned on dynamically evolving four-dimensional trust vectors covering identity, platform, behavioral, and contextual evidence. SATHOS introduces a distributed trust negotiation protocol over mutual TLS with replay protection and cryptographic transcript verification,

combined with a self-adaptive trust evolution mechanism that adjusts trust scores based on execution outcomes without requiring policy redeployment. Evaluated on a Kubernetes-based testbed with three pipeline topologies (3-5 nodes), three experimental conditions, and three random seeds (27 runs), SATHOS blocks 100% of compromised agent execution requests with zero false positives ($p < 0.001$, Fisher's exact test). The five-message trust negotiation protocol adds a median per-node latency of 4.73 ms (p50), and the Wilcoxon signed-rank test confirms bounded overhead ($p = 0.065$, one-sided). The system achieves zero false allows and zero false denies across all trial configurations.

16:30 ***GitOps Driven Automation for Federated Learning Infrastructure***

Vinoth Punniyamoorthy (JPMorgan Chase, USA); Srivenkateswara Reddy Sankiti (Cleveland State University, USA); Nachiappan Chockalingam (Meta, USA); Ram Sekhar Bodala (Amtrak, USA); Aswathnarayan Muthukrishnan Kirubakaran (Meta Inc, USA); Balakrishna Pothineni (JPMorgan Chase, USA); Nitin Saksena (Albertsons Companies, USA); Abhirup Mazumder (Amazon, USA)

Federated learning (FL) enables collaborative model training across distributed data sources while preserving privacy, but deploying and operating FL systems at scale remains highly complex. Most existing frameworks emphasize algorithms rather than the practical challenges of provisioning cloud resources, configuring communication layers, and managing heterogeneous client environments, resulting in configuration drift, operational inefficiencies, and limited reproducibility. This paper presents a GitOps driven control plane that expresses the entire FL topology including servers, aggregators, clients, networking, and storage as declarative infrastructure. The system automates multi cloud provisioning, maintains continuous configuration consistency, and applies training aware self healing to remediate failures without manual intervention. Evaluation across multiple cloud environments demonstrates significant gains in provisioning speed, operational reliability, and reproducible deployment. The proposed approach offers a unified and automated foundation for managing large scale FL infrastructures and supports more dependable operation of distributed learning systems.

16:45 ***Carbon-Aware SLO Enforcement for Cloud-Native Microservices Using Adaptive Reliability Policies***

Jagadish Almaipeta, Subhashchandra Babu Madineni and Usha Rani Sheri (USA)

Cloud-native microservices architectures have transformed how modern enterprises build and operate distributed systems. However, the rapid expansion of elastic cloud workloads has significantly increased energy consumption and associated carbon emissions. Traditional Service Level Objective (SLO) enforcement mechanisms focus primarily on availability and latency guarantees without considering environmental impact. This paper introduces a carbon-aware SLO enforcement framework that dynamically adjusts reliability policies based on real-time carbon intensity signals and workload criticality. By integrating carbon telemetry with adaptive reliability controllers, the proposed system balances performance, resilience, and sustainability objectives. The framework enables enterprises to reduce carbon footprint while preserving business-critical reliability guarantees. Experimental validation in a Kubernetes-based microservices environment demonstrates measurable reductions in energy consumption without violating SLO constraints. The results highlight the feasibility of embedding sustainability directly into reliability engineering practices for next-generation cloud-native platforms.

17:00 ***AgentQE-Bench: A Reproducible Evaluation Framework for Agentic AI Reliability, Safety, and ROI in CI/CD Pipelines***

Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA); Lakshmi Vidya Peri (TMMI, USA)

Agentic AI systems are increasingly being deployed within CI/CD pipelines to diagnose failures, recommend remediations, and accelerate release cycles. However, no standardized evaluation framework exists to assess the reliability, safety, and return on investment (ROI) of such agents under controlled and reproducible conditions. We introduce AgentQE-Bench, an open source evaluation framework comprising a deterministic pipeline simulator, a five-rule safety policy engine, a composite scoring metric weighting reliability (0.45), safety (0.35), and ROI (0.20), and a pluggable agent interface that supports any LLM backend. The framework provides 20 episodes across three domains (hardware NVMe validation, retail microservice pipelines, and safety-critical constraints) with ground-truth remediations and deterministic seeding for reproducibility. We demonstrate the framework with four baseline agents: a rule-based oracle, an LLM agent (Claude Sonnet), a safety-wrapped LLM variant, and a lower bound of random-action. Across three seeds, the framework produces stable classifications (cross-seed composite std < 0.01) and reveals that reliability, not safety, is the primary discriminating dimension for LLM-based CI/CD agents. Index Terms-Agentic AI, valuation framework, CI/CD pipelines, reliability, safety, ROI, LLM agents, policy engine, DevOps, software quality, TMMI .

17:15 ***Policy-Driven Shift-Left Security for Hybrid OpenShift CI/CD Pipelines***

Nilesh Jaiswal (Electronic Arts, USA)

The adoption of hybrid cloud platforms by enterprises, such as Red Hat OpenShift, has increased the complexity of securing containerized workloads across heterogeneous infrastructure. Traditional CI/CD pipelines defer vulnerability detection and policy validation to late stages, allowing insecure artifacts to propagate toward production. This paper presents a policy-driven shift-left security framework that integrates three controls into the pre-deployment stage of hybrid OpenShift CI/CD pipelines: container image vulnerability scanning using Trivy, Kubernetes deployment policy enforcement using Open Policy Agent (OPA) with eight Rego-based denial rules, and governance tag validation against a four-label compliance schema. We evaluate the framework through a controlled experiment comparing a baseline pipeline (no security checks) against the shift-left pipeline across 50 builds deployed to a Kubernetes namespace, comprising 35 compliant, 8 vulnerable, 5 policy violating, and 2 governance-noncompliant configurations. The baseline pipeline deploys all 50 builds without checks, while the shift-left pipeline blocks all 15 insecure builds (100% detection rate) with zero false positives. The shift-left pipeline introduces a mean overhead of 28.3 seconds per build, with container image scanning accounting for 82.2% of the added time. The results demonstrate that early-stage integration of layered security controls eliminates insecure deployments in hybrid OpenShift environments without disrupting compliant workloads.

17:30 ***AI-Assisted Bias Detection and Fairness Auditing in Engineering Performance Evaluation Systems: A Responsible AI Governance Framework***

Bhaskar Manchuri (USA)

Performance evaluation systems play a critical role in the retention of engineering talent and career progression. Despite formal organizational policies promoting equal opportunity, evaluation outcomes often remain susceptible to subjective managerial bias through narrative performance reviews, forced ranking mechanisms, and subjective promotion decisions. This paper proposes an AI-assisted auditing framework that combines machine learning prediction models, natural language processing (NLP) of evaluation narratives, and anomaly detection algorithms to detect bias in engineering performance evaluation systems. Using a synthetic dataset of 1,000 engineer records representing performance metrics, peer feedback,

and managerial review narratives, the framework identifies inconsistencies between objective performance indicators and final ratings through a hybrid AI approach integrating gradient-boosted performance prediction, transformer-based sentiment analysis, and ensemble anomaly detection using Isolation Forest and Local Outlier Factor. The experimental evaluation demonstrates that the proposed hybrid AI approach can identify biased evaluations with up to 87% detection accuracy and a 7% false positive rate across the simulated scenarios. Unlike prior work focused on ensuring fairness in algorithmic decision-making, this framework applies AI to audit human evaluation decisions, addressing a critical gap in responsible AI governance. This work contributes to the emerging fields of algorithmic fairness and people analytics by providing a practical, scalable approach for organizations seeking to enhance trust, transparency, and accountability in technical performance management systems

17:45 ***An Empirical Study of Defect Detectability and Operational Impact of Component vs. End-to-End Testing in Regulated Banking CI/CD Pipelines***

Bhaskar Manchuri (USA)

End-to-end (E2E) tests are commonly assumed to provide the highest assurance in regulated banking CI/CD pipelines, often leading to test strategies dominated by slow and brittle full-system validation. However, empirical evidence comparing the effectiveness of E2E and component testing in production-validated banking CI/CD environments remains limited. This paper presents a 14-month empirical study across 54 microservices in a regulated financial platform, analyzing 2,341 defect records and 847,293 CI/CD pipeline executions, and evaluating a controlled transition from E2E-primary to component-primary testing in six services. Component tests demonstrated significantly higher assessed detectability for business logic defects than E2E tests (88.1% vs. 59.4%, $p < 0.001$) and detected defects earlier in the pipeline. Following the transition, services experienced substantial operational improvements, including a 67% reduction in flaky test failures and a 41% increase in deployment frequency without increasing defect escape rates (escape rate decreased from 12.3% to 8.7%). These results suggest that component centric testing strategies can deliver faster feedback and equal or better defect detection than E2E dominated approaches in regulated banking CI/CD environments.

Tuesday, April 21 16:00 - 18:00 (Africa/Cairo)

SESSION-3E: RESEARCH: Innovative Approaches in AI for Quality Assurance, Cybersecurity, and Predictive Analytics

8 PAPERS

Room: ROOM-E

16:00 ***Shaping QA Coverage with Production Telemetry: Risk-Based Test Selection Under CI Budget Constraints***

Yogesh S Thanvi (Akamai Technologies, USA); Lakshmi Vidya Peri (TMMI, USA); Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA)

Continuous integration (CI) pipelines for cloud-native microservices face a resource allocation problem in which test suites grow beyond time budgets, forcing subset selection. Static strategies (fixed ordering, historical pass rates) ignore runtime system behavior, producing test distributions

misaligned with production traffic. We propose a telemetry-driven test selection strategy that queries OpenTelemetry metrics to compute per-service risk scores and greedily selects tests by risk-adjusted priority under a fixed CI time budget. The telemetry-driven pipeline reduces the divergence of coverage from production traffic by 64.8% ($p < 0.001$), improves the fault detection rate from 67% to 94% ($p = 0.0253$), and increases the regression detection throughput by 60% ($p = 0.0013$), without increasing the CI runtime. The primary mechanism is improved test selection, that is, shaping the test execution distribution to match production risk, rather than faster execution ordering.

16:15 ***Self-Optimizing Cloud Data Analytics Platforms via Closed-Loop AI Control***

Nabeel M Sheikh (Oracle America, USA); Mohan Krishna Mannava (Independent Researcher, USA); Srinivas Gadam (SyntageTech, USA); Taranpreet Kaur (Amazon, USA); Vinil Pasupuleti (IBM, USA); Balaji Shesharao Ingole (IEEE Senior, USA)

Modern cloud data analytics platforms face persistent tension between performance, cost efficiency, and reliability, especially as workloads shift unpredictably. Traditional static provisioning and rule-based tuning often fail to keep pace with these dynamics, resulting in resource waste and unstable service quality. This paper introduces a self-optimizing cloud analytics architecture powered by closed-loop AI control. The system continually monitors operational telemetry, identifies emerging performance trends, and dynamically fine-tunes resource levels, task allocation, and configuration settings as conditions change in real time. By treating the data infrastructure as a responsive, feedback-driven system rather than a manually tuned environment, the approach replaces fixed assumptions with adaptive, evidence-based decision-making. Experimental evaluations demonstrate that closed-loop AI control significantly improves resource utilization, reduces operational costs, and maintains stable latency under variable demand. The results highlight a practical pathway toward fully autonomous, resilient, and cost-aware cloud analytics platforms.

16:30 ***Azure Windows Defender for Enterprise Cybersecurity: Architecture, Automation, and Measured Operational Impact***

Jagadish Almaipeta (USA)

Modern enterprises operate in increasingly complex hybrid and cloud-native environments, where traditional perimeter-based security controls are no longer sufficient. The rapid adoption of cloud platforms, remote work models, and distributed applications has expanded the attack surface and introduced new cybersecurity challenges. This paper presents a comprehensive analysis of leveraging Azure-based security services, with a particular focus on Windows Defender technologies, to build a scalable, automated, and resilient enterprise cybersecurity posture. The study explores system architecture, operational deployment strategies, automation capabilities, and measurable security benefits observed in real-world enterprise environments. The findings demonstrate that integrating Windows Defender with Azure security services enables proactive threat detection, reduced response times, and improved operational efficiency, while supporting compliance and governance requirements across large organizations. The evaluation is based on enterprise operational metrics collected over a three-month post-deployment observation period.

16:45 ***Predicting Heart Failure Outcomes Using Machine Learning Models on Clinical and Laboratory Data***

Nabeel M Sheikh (Oracle America, USA); Vansh Agarwal (Amazon, USA); Rucha Deshmukh (Academia, USA); Amit Kumar (Nanit, USA)

Heart failure is a leading cause of mortality world- wide, requiring timely and accurate prediction to guide clinical decision-making [1]. This study explores the use of machine learning models to predict mortality outcomes in patients with heart failure using clinical and laboratory data from 296 individ- uals. The dataset was obtained from the UCI Machine Learning Repository and includes variables such as age, comorbidities, vital signs, and

laboratory test results [3]. Four models were evaluated using five-fold cross-validation: logistic regression, decision tree, k-nearest neighbors (KNN), and extreme gradient boosting (XGBoost) [2]. Among these, XGBoost achieved the best overall performance in terms of precision and F1 score, highlighting its ability to capture complex, non-linear relationships in clinical data. Model interpretability was supported using SHAP values to explain the influence of individual features on predictions [4]. These findings reinforce the value of data-driven decision support in cardiology and align with recent efforts to integrate artificial intelligence into clinical practice [5], [13]. Future work should expand on this foundation by incorporat- ing larger, multicenter datasets and longitudinal health records to improve generalizability and early risk detection. Overall, this study underscores how machine learning, when carefully validated and interpreted, can provide meaningful insights into patient outcomes in heart failure care.

17:00 ***Big Data-Driven Sales Forecasting for Consumer IoT Products: A Machine Learning Approach***

Nabeel M Sheikh (Oracle America, USA); Amit Kumar (Nanit, USA); Vansh Agarwal (Amazon, USA)

Accurate sales forecasting is a critical competitive advantage in the rapidly expanding Consumer Internet-of-Things (IoT) market. This paper presents a comprehensive big data analytics framework for monthly revenue forecasting in the baby monitor segment, leveraging 36 months of retail point-of-sale data (January 2023-December 2025) encompassing eight major brands and over $540 million in cumulative market revenue. We implement and systematically compare four fore- casting models: Seasonal ARIMA, Facebook Prophet, XGBoost with lag-engineered features, and Random Forest. Evaluation on a six-month held-out test set demonstrates that XGBoost achieves the best predictive accuracy with a Root Mean Squared Error (RMSE) of $2,879,316, a Mean Absolute Error (MAE) of $1,925,601, and a Mean Absolute Percentage Error (MAPE) of 9.69%-outperforming ARIMA (MAPE 23.9%), Prophet (MAPE 20.5%), and Random Forest (MAPE 15.4%). Feature importance analysis reveals that temporal trend and lag features are the strongest predictors, confirming the value of supervised feature engineering for non-stationary retail time series. These findings provide actionable guidance for data-driven inventory planning and go-to-market strategy in consumer IoT markets.

17:15 ***Agentic SDLC Automation for ERP Release Readiness: Design Patterns, Governance Guardrails, and Measurable Operational Outcomes***

Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA)

Enterprise Resource Planning (ERP) release readiness remains one of the most coordination-intensive stages of the software development lifecycle. Although deterministic DevSecOps controls can detect configuration drift and access violations at deploy time, interpretive readiness activities such as change summarization, impact hypothesis generation, testplan drafting, and evidence compilation still depend on manual expert effort. This paper presents a two-layer framework that combines a deterministic release-gating layer with a governed agentic augmentation layer comprising five task-specific agents: Change Summary, Impact Hypothesis, Test Planning, Evidence Assembly, and Release Narrative. Each agent operates under governance guardrails that enforce source grounding, confidence scoring, traceability, and mandatory human approval at critical decision points. We evaluate the framework on a Kubernetes testbed with three experimental conditions: (A) baseline with no controls, (B) deterministic pipeline only, and (C) agentic layer followed by deterministic gating. In a controlled experiment with 63 runs (7 scenarios × 3 seeds × 3 pipelines), the agenticgoverned pipeline achieves a 100% detection rate (18/18 faults) compared to 94.4% for the deterministic-only pipeline and 0% for the baseline (Fisher's exact test, $p < 0.001$). Impact-based test selection reduces test effort by 58.3% without increasing leakage, and the deterministic gate maintains precision of

1.00 and recall of 1.00 when preceded by agentic reasoning. The agentic layer adds 190.8s mean overhead per run, acceptable for release-readiness decisions that occur once per release cycle.

17:30 ***Adversarial Attacks on AI Systems and Multi-Layer Defense Mechanisms for User Data Protection***

Naga Satya Praveen Kumar Yadati (Meta, USA); Gayathri Balakumar (Capital One, USA); Ramadas Pulapaka (Constellation Energy, USA); Vinod Bottu (Microsoft, USA); Vidisha Vijay (Birla Institute of Technology and Science, India); Anand Patel (Facebook, USA)

The proliferation of artificial intelligence (AI) systems in high-stakes domains has introduced a critical surface for vulnerabilities: adversarial attacks that systematically subvert model behavior to compromise user data confidentiality, integrity, and availability. This paper presents a comprehensive analysis of the adversarial threat landscape targeting production AI systems and proposes a multi-layer defense framework for industry practitioners. We evaluate four principal attack categories-evasion, data poisoning, model inversion, and membership inference-against six contemporary defense mechanisms on standardized benchmark datasets. Our experiments demonstrate that no single defense is universally sufficient; a combined strategy integrating adversarial training, differential privacy (ε = 1.0), and certified randomized smoothing reduces the aggregate attack success rate from 97.1% to 24.9% with only a 9.9 percentage-point reduction in clean accuracy. We further provide actionable deployment guidelines, a compliance mapping to GDPR/CCPA requirements, and an open evaluation methodology aligned with the NIST AI Risk Management Framework. Our findings offer security engineers and AI architects concrete, evidence-based strategies for hardening AI deployments against adversarial threats.

17:45 ***Safety-Gated Continual Learning for Drift-Aware Anomaly Detection in Real-Time Cyber-Physical Systems***

Pradeep Kumar Chilukury, Krishna Kishor Tirupati and Sai Reddy Busi Reddy (USA); Naresh Kumar Methuku (Fidelity Investments, USA)

Cyber-physical systems function under changing process conditions where static anomaly detectors quickly lose reliability, especially under conditions where normal-abnormal boundaries change over time due to distribution drift. This work is addressing that limitation by a safety-gated self-adaptive anomaly detection framework for dynamic CPS environments using the HAI 22.04 industrial benchmark. The approach uses a combination of GRU-based sequence autoencoding, temporal feature augmentation, persistent drift detection, and controlled online adaptation approach to update the detector only in a high-confidence adaptation region, improving the robustness of the detector while reducing unsafe model update. The proposed method is characterized by its drift-aware and safety-gated adaptation strategy as an alternative to the conventional fixed offline modeling. Experimentally, the baseline GRU autoencoder obtained F1-score of 0.2201 at window level, F1-score of 0.2560 was obtained by augmented pre-adaptation, and after the adaptation, the proposed method achieved precision to 0.6023, recall to 0.3925, F1-score to 0.4860, ROC-AUC to 0.6907, and PR-AUC to 0.3604. At the event level, it achieved 0.7111 precision, 0.7793 recall, 0.7418 F1-score, a 20.05 window detection delay, and 35.7448 false alarms per 10k windows, carrying out effective, operationally relevant CPS conditions anomaly monitoring.

Wednesday, April 22

Wednesday, April 22 8:30 - 12:00 (Africa/Cairo)

REGISTRATION: REGISTRATION

REGISTRATION

Wednesday, April 22 10:00 - 10:30 (Africa/Cairo)

KEYNOTE-3: KEYNOTE-3: ENG. SANDEEP SHIVAM, USA

KEYNOTE-3

Room: MAIN_ROOM

Wednesday, April 22 10:30 - 11:00 (Africa/Cairo)

NETWORKING: COFFEE-BREAK &POSTER SESSION

STUDENTS POSTER SESSION (COMPETITION)

Room: COFFEE_ROOM

Wednesday, April 22 11:00 - 11:30 (Africa/Cairo)

KEYNOTE-4: KEYNOTE-4: ENG. TEJAS PATEL - AMAZON, USA

KEYNOTE-4

Room: MAIN_ROOM

Wednesday, April 22 11:30 - 13:00 (Africa/Cairo)

RSESSION-4E: RESEARCH: Wireless Communication: NOMA, RSMA & Multi-Access Schemes

6 PAPERS

Room: ROOM-E

Chair: Salah A. Aly (Fayoum University, Egypt)

11:30 ***Short-Packet Communication in CDRT Networks with Imperfect CSI Over Nakagami-m Fading Channels: RSMA or NOMA?***

Kai Wang, Xiaochen Shen, Zhihua Zhang, Liang Shi, Chao He and Xiang Zhang (Defense Innovation Institute)

In this paper, we investigate a downlink rate splitting multiple access (RSMA) based coordinated direct and relay transmission (CDRT) system, considering short-packet communication over Nakagami-m channels. In the proposed system, a base station directly transmits signals to the nearby user while serving the distant user through a half-duplex and decode-and-forward relay. Closed-form expressions for the average block error rates (BLERs) of both users are theoretically derived in the presence of imperfect channel state information. Numerical results validate the theoretical findings and reveal the following: 1) The RSMA-CDRT system outperforms its non-cooperative counterpart in terms of BLER performance. 2) The BLER of the nearby user exhibits significant improvement compared to non-orthogonal multiple access-CDRT and space division multiple access-CDRT schemes at low-to-moderate transmit signal-to-noise ratios (e.g., 0-17 dB), while the BLER performance of the distant user is degraded.

11:45 ***Intelligent and Sample-Efficient Handover in High-Speed Rail Networks Using Deep Reinforcement Learning***

Mariam Abdul-Zahra and Asaad S Daghal (Al-Furat Al-Awsat Technical University, Iraq)

The realization of reliable and effective HO plays an essential role in HSR communication systems since they involve fast movement and dynamic channels. However, while DRL shows good performance in optimizing HO strategies, most algorithms depend on massive training samples and always consider fixed scenario settings. Hence, a novel HSR HO decision algorithm based on Double Deep Q-Network coupled with Quantum Neural Network (DDQN-QNN) is presented. The proposed technique makes full use of the stability of the DDQN strategy and the powerful representation ability of QNN to obtain effective HO policies. Then, the model is trained by using various percentages of the dataset (20%, 50%, and 100%) and evaluated through the same test scenario. Simulation results show that the DDQN-QNN model exhibits excellent effectiveness in predicting HO decisions, with the highest accuracy obtained even when only small numbers of samples are available. Besides, the proposed algorithm exhibits satisfactory stability with minor performance gains achieved by increasing the amount of training samples.

12:00 ***Slice-Aware AoI Control for UAV-Assisted 5G and Beyond-5G Waste Hotspot Mapping Using Kernel Fused Density Updates***

Mohammed Jalil Mohammed Ali Alhasan and Karar Hamza Hussein (Al-Furat Al-Awsat Technical University, Iraq)

Long-lived waste (e.g., plastics and slowdegrading residues) often accumulates into spatial hotspots, yet periodic IoT reporting may deliver outdated hotspot information when municipal actions are taken. This paper proposes a slice-aware Age-of-Information (AoI) reporting framework for UAV-assisted waste hotspot mapping that explicitly targets actionable freshness at the server. A UAV samples a waste-density proxy over a gridded area and triggers transmissions using an AoI/change-aware rule to avoid redundant updates and prioritize informative measurements. When an update is generated, a hotspot-aware policy maps urgent reports to a URLLC-oriented slice while routing routine traffic through an mMTC/eMBB-oriented slice, leveraging 5G/Beyond-5G service differentiation. At the server, sparse measurements are fused using kernel footprint aggregation to reconstruct a continuous waste-density field and support density- and freshness-driven municipal alerts. Simulation results indicate that the proposed method improves end-of-run coverage from 0.6269 to 0.9443 and reduces RMSE from 0.1692 to 0.1328 over the full grid (0.3814 to 0.2826 over hotspots), while decreasing mean AoI from 547.27 s to 405.62 s and hotspot AoI P95 from 1010.0 s to 756.5 s, at the cost of higher communication energy.

These results indicate that sliceaware AoI control can substantially reduce hotspot staleness and improve map fidelity, enabling more timely and reliable municipal response.

12:15 ***Shielded Primal-Dual Multi-Objective PPO for Energy-Efficient and Fair OFDMA Scheduling in UAV-Enabled 5G/6G Networks***

Mohammed Jalil Mohammed Ali Alhasan, Wafaa Mohammed Ridha Shakir and Asaad S Daghal (Al-Furat Al-Awsat Technical University, Iraq)

UAV-enabled wireless networks offer rapid, flexible coverage for beyond-5G and early-6G scenarios, but their limited onboard energy and the stringent quality-of-service (QoS) requirements of cell-edge users make downlink scheduling challenging. This paper proposes Shielded Primal-Dual Multi-Objective Proximal Policy Optimization (SPDMO-PPO), a safe deep reinforcement learning framework for UAV downlink OFDMA scheduling and power allocation under an explicit cell-edge QoS constraint. The cell-edge constraint is defined using the 5th-percentile user throughput computed from episode-average rates. A primal-dual mechanism updates a Lagrange multiplier to regulate long-term feasibility, while a lightweight action shield overrides unsafe allocations when the estimated cell-edge margin becomes small, improving practical safety during learning and execution. Simulation results indicate that SPDMO-PPO improves cell-edge throughput (R5) and Jain's fairness compared with Max-Rate and proportional-fair baselines, with competitive energy efficiency and modest aggregate throughput reduction. A QoS-target sweep further illustrates the feasibility frontier and tail-QoS robustness of shielding combined with primal-dual updates.

12:30 ***Dynamic SIC Ordering in Downlink NOMA With User-Dependent CSI Uncertainty***

Zainab Ulhassan, Asaad S Daghal and Wafaa Mohammed Ridha Shakir (Al-Furat Al-Awsat Technical University, Iraq)

This paper investigates how uncertainty in channel state information (CSI) affects the performance of a three-user downlink non-orthogonal multiple access (NOMA) system. Conventional NOMA designs rely on fixed SIC ordering based on ideal or perfectly known channel gains. This leads to severe performance degradation when CSI is imperfect. To address this, the paper proposes an estimation-aware, robust framework in which the SIC decoding sequence is dynamically determined according to estimated channel gains rather than assumed ideal ordering. This framework incorporates structured CSI uncertainty modeling to reduce ordering mismatch and mitigate residual interference accumulation along the SIC chain. Comprehensive Monte Carlo-based performance evaluations demonstrate that the proposed approach significantly enhances decoding stability, improves resilience against CSI distortion, and exhibits strong robustness in spectral efficiency and reliability metrics compared to conventional fixed-order systems. These results confirm that adapting the decoding hierarchy dynamically to estimated channel conditions substantially improves system robustness. Thus, the proposed design is well-suited for high-reliability, low-latency communication scenarios in 5G and beyond.

12:45 ***Robust SIC Design in Downlink NOMA Under User-Dependent CSI Uncertainty***

Zainab Ulhassan, Asaad S Daghal and Wafaa Mohammed Ridha Shakir (Al-Furat Al-Awsat Technical University, Iraq)

Successive interference cancellation (SIC) in downlink non-orthogonal multiple access (NOMA) is highly sensitive to channel state information (CSI) mismatch. However, many existing studies still assume identical uncertainty across users, which can obscure the hierarchical nature of SIC error propagation. This paper investigates a user-dependent CSI uncertainty model for downlink NOMA and combines it with inverse-gain-based power allocation and adaptive SIC ordering. In the considered three-user Rayleigh-fading downlink scenario, weaker ordered users are assigned larger

uncertainty levels to reflect more severe estimation distortion, while stronger users retain smaller residual interference factors. Numerical results show that, compared with a baseline dynamic NOMA scheme based on uniform CSI uncertainty, the proposed framework achieves higher robustness under mismatch. At a transmit power of 30 dBm and ρ_max= 0.05, the sum rate increases from 4.70 to 9.78 bps/Hz, while the outage probability decreases from 0.1419 to 0.0203. These results indicate that heterogeneous CSI uncertainty modeling can improve SIC reliability and spectral efficiency under imperfect CSI, although the gains depend on the specific uncertainty profile and power allocation design.

Wednesday, April 22 11:30 - 13:00 (Africa/Cairo)

SESSION-4A: RESEARCH:Renewable Energy, Smart Grids & Environmental Sustainability

6 PAPERS

Room: ROOM-A

Chair: Moses Mupeta (University of Zambia, Zambia)

11:30 ***Random Forest Weight Auditing for Gulf Sustainability Indices: PM2.5 Integration and Structural Heterogeneity in a Six-Country Panel***

Nahla Nabil Skaik, Hasan Ali Razzaqi and Mohammad Riyaz Belgaum (Arab Open University - Bahrain, Bahrain); Alyaa Al Aradi (Arab Open University, Bahrain)

The Organisation for Economic Co-operation and Development (OECD) calibrated benchmarks were not designed for Gulf conditions. The consideration of PM2.5 is very rare though it has a significant environmental burden in the region. Despite this, the six Gulf Cooperation Council (GCC) states all rank close to the bottom of global sustainability tables. In this study, 186 country-year observations (1990 to 2020) are used to create a five-indicator District Sustainability Index (DSI) for each GCC state. All data is taken directly from World Bank World Development Indicators, with no values imputed. As an audit rather than a design tool, random forest feature importance is applied after the index is constructed rather than during it. The main query is whether the stated weights accurately represent the data that each indicator contains. The within-panel fit achieves $R^2 = 0.991$. In LOCO results, the range is −7.63 for Qatar to +0.75 for Saudi Arabia, with four states showing negative R^2, confirming that the declared weighting scheme fails to reflect the structural distances separating these economies. The discrepancy during the audit revealed that the electricity CO_2 pair demands 55% of declared weight, it accounts for only 36.3% of signal, while PM2.5 explains 41.1% of predictive variation yet carries only a 25% declared weight. At ±40% weight perturbation, fourteen of fifteen pairwise rankings exceed 99.9% stability, with twelve pairs holding at exactly 100%.

11:45 ***A Multi-Head Neural Network for Simultaneous Prediction of Perovskite Solar Cell Parameters***

Fatmah AlAwadhi, Fatemah Lari, Amani Albuloushi, Mariam Hussain, Zainab Sadeq, Khaled Chahine and Mohamad Arnaout (American University of the Middle East, Kuwait); Marc AlAtem (American University of Middle East, Kuwait)

Machine learning approaches for perovskite solar cell (PSC) optimization typically model photovoltaic parameters-open-circuit voltage (Voc), short-circuit current density (Jsc), fill factor (FF), and power conversion efficiency (PCE)-as independent targets, ignoring the physical coupling defined by PCE = Voc × Jsc × FF. We propose a multi-head neural network (MH-NN) with a shared backbone that learns a unified device-state representation before branching into task-specific prediction heads whose depths are calibrated to each target's physical complexity. Six physics-informed features-including compositional ratio proxies for lattice strain, bandgap tuning, and volume recombination-augment the raw input space. Evaluated on 7,176 SCAPS-1D simulated configurations via 5-fold cross-validation, the MH-NN achieves R^2 = 0.996 ± 0.002 (Voc), 0.998 ± 0.001 (Jsc), 0.995 ± 0.001 (FF), and 0.997 ± 0.001 (PCE). Compared to architecturally matched single-task baselines, multi-task learning reduces RMSE by 74-85% across all targets and reduces the physical inconsistency between predicted PCE and the Voc×Jsc×FF product. SHAP analysis reveals physically coherent feature-target relationships, including the expected halide-mediated Voc-Jsc trade-off. These results establish multi-task learning as a superior paradigm for photovoltaic parameter prediction.

12:00 ***Uncertainty-Aware Multi-Horizon Wave Energy Assessment Using Probabilistic Deep Learning***

Taha Houda (Prince Mohammad Bin Fahd University, Saudi Arabia); Wafaa Mohammed Ridha Shakir (Al-Furat Al-Awsat Technical University, Iraq); Abdullah Baradaaji (Lebanese International University, Lebanon); Jihad Jaam (Liverpool J. Mores University, United Kingdom (Great Britain))

Ocean wave energy is a promising renewable resource due to its high energy density and temporal complementarity with wind and solar power. Accurate short-term wave forecasting is essential for reliable wave energy assessment; however, most existing approaches rely on deterministic predictions and neglect forecast uncertainty, which can lead to overconfident energy estimates and increased operational risk. This paper proposes an uncertainty-aware multi-horizon wave energy assessment framework based on probabilistic deep learning. A quantile-based Long Short-Term Memory (LSTM) model is developed to generate probabilistic forecasts of significant wave height and peak wave period over horizons ranging from 1 to 24 hours. Predictive uncertainty is explicitly quantified through prediction intervals and propagated into wave energy estimation, enabling the computation of conservative, expected, and optimistic energy scenarios. Experiments conducted on publicly available offshore buoy data demonstrate that the proposed probabilistic model achieves deterministic accuracy comparable to standard LSTM predictors while providing well-calibrated uncertainty estimates. The results show that energy uncertainty increases significantly with forecast horizon, highlighting the limitations of deterministic wave energy assessments. The proposed framework provides a more realistic and risk-informed evaluation of wave energy potential, supporting improved operational planning and decision-making.

12:15 ***Decentralized Droop Control for Stabilizing DC Bus Voltage in Renewable DC Microgrids***

Ahmed Ibrahim (The Knowledge Hub Universities, Egypt); Ehab Bayoumi (Mechatronics and Robotics Section The British Univer, Egypt); Nathalie Nazih (The British University in Egypt, Egypt)

The increased penetration of distributed renewable energy sources and DC-based loads has led to an increased adoption of DC microgrids, which can be an efficient solution in comparison to traditional AC-based microgrids. Nevertheless, a significant technical challenge in DC microgrids pertains to

the maintenance of a reliable DC bus voltage dealing with intermittent energy sources and variable DC loads. This study suggests an integrated regulatory strategy for an independent DC microgrid, which encompasses a 20-kW photo voltaic (PV) system, a 30-kW wind energy system, and a 30-kWh bidirectional battery energy storage system (BESS) that caters to a variable DC load within the 0 kW to 60 kW range. The voltage of the direct current (DC) bus is regulated at a steady level by implementing a droop control technique in the BESS. This approach allows autonomous power equilibrium and allows for bidirectional power transfer based on the current operational circumstances. The DC microgrid that is being investigated is simulated and analyzed using the MATLAB software. The results obtained from the simulation are analyzed based on their dynamic performance characteristics when subjected to intermittent renewable energy sources and DC loads. The results obtained have confirmed that the proposed droop control strategy, when applied to the Battery Energy Storage System, is able to maintain the DC bus voltage within acceptable limits and improve the stability of the DC microgrid.

12:30 ***Assessing Green, Resilient Supply Chains Through Sustainable Engineering Innovations***

Mony Trad (NDU, Lebanon); Atef Harb and Abdallah Kassem (Notre Dame University, Lebanon)

Sustainable engineering innovations offer a promising path forward. By leveraging new materials, technologies, and processes that minimize environmental impact while maintaining performance, companies can build more circular, transparent, and adaptable supply networks. Technologies like bio-based materials, 3D printing, energy-efficient systems, and AI optimization are revolutionizing how supply chains operate. Additionally, incorporating lifecycle thinking and circular economy principles helps companies extend product longevity, capture value from waste, and reduce their overall ecological footprint. This paper explores practical ways to integrate sustainable engineering advances into modern supply chain frameworks. Through real-world case studies from manufacturing, logistics, and energy sectors, we identify key enablers and barriers to implementation. Our research contributes to the ongoing dialogue about technology's role in creating more resilient and environmentally conscious industrial systems.

12:45 ***High-Performance GPU Particle System Editor for VR Environments: Achieving Low-Latency VFX Through Optimized Rendering***

Yousef Wael (October University for Modern Sciences and Arts (MSA), Egypt); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

This research aims to meet the need to maintain high frame rates (above 90 FPS) and minimize motion-tophoton latency in Virtual Reality (VR) applications to ensure the sustainability of the experience. Spikes and instabilities in frame time are known to cause visual discomfort and negatively affect the overall quality of the experience. Particle-based effects are known to be one of the major contributors to frame rate variability due to their dynamic nature, high rate of updates, and the use of blending with transparency. The authors have proposed a particle system editor designed specifically with the needs of a VR application and built using the OpenXR and OpenGL 4.1 libraries. The proposed system uses controllerbased methods to create particle effects such as paint, fire, smoke, and sparks in three-dimensional space. A lightweight simulation model and an efficient rendering approach are used to focus on the performance and stability of the application while satisfying the low latency requirements. The performance of the proposed system is evaluated using runtime logging and frame time analysis, along with image comparisons.

Wednesday, April 22 11:30 - 13:00 (Africa/Cairo)

SESSION-4B: RESEARCH: Robotics, Autonomous Systems & Adaptive Control

6 PAPERS

Room: ROOM-B

11:30 ***AI-Driven Adaptive Control System for Soft Robotics: Bridging Perception and Actuation in Dynamic Environments***

Loso Judijanto (IPOSS, Indonesia); Sri Nurhayati (IKIP Siliwangi, Indonesia)

The increasing complexity of unstructured environments in which soft robots operate presents significant challenges for achieving stable and responsive control. Conventional control systems often fall short in managing the high degrees of freedom, material compliance, and nonlinear behaviours intrinsic to soft robotic structures. This study aims to explore how artificial intelligence, particularly adaptive control strategies, contributes to bridging the perception-action gap in soft robotics. Systematic Literature Review (SLR) method was employed as a qualitative approach, guided by the PRISMA protocol to ensure transparency and replicability. Data were collected through a structured search of peer-reviewed research articles published between 2020 and 2025 in the ScienceDirect database using a refined Boolean keyword strategy. After multiple screening phases based on publication year, article type, relevance, and open access, 25 empirical studies were selected for in-depth analysis. The data analysis was conducted using thematic synthesis to extract dominant trends, control architectures, sensor integration strategies, and performance outcomes. The findings show that AI-driven adaptive control systems, especially those integrating reinforcement learning and sensor fusion, enhance real-time decision-making, improve task accuracy by up to 40%, and support energy-efficient robotic behaviour. Hybrid control architectures combining classical and learning-based models demonstrated superior adaptability and sample efficiency.

11:45 ***Deep Reinforcement Learning-Based Autonomous Thermal Management of Lithium-Ion Batteries: A Comparative Evaluation of DDPG and TD3***

Mahmood Alsadoun, Abdulaziz Aljalkhaf, Mohammed Aldakheel, Abdulrahman Al Ali and Taha Houda (Prince Mohammad Bin Fahd University, Saudi Arabia)

Efficient thermal management is required not only for performance and safety but also for the extended use of lithium ion batteries in electric vehicles. This work compares two sophisticated reinforcement learning (RL) Deep Deterministic Policy Gradient (DDPG) and Twin Delayed Deep Deterministic Policy Gradient (TD3) in their capacity to operate a battery thermal management system (BTMS) without human intervention. A heat dynamic model of a lithium battery module was created in MATLAB/Simulink to simulate a battery pack's thermal response to internal heat generation, active liquid cooling, and passive heat loss to the ambient. The control problem was cast with the aim of regulating the battery temperature to a target setpoint of 30C, under a safety threshold of 45C. The agents learned to maximize a multi objective reward function that emphasized accuracy and safety. The experimental findings make it evident that both DRL agents have exceed the performance of conventional manually tuned Proportional Integral (PI) controllers considerably in terms of settling time, overshoot, and steady state error. The TD3 agent in particular recorded a Mean Absolute Error (MAE) of 0.0312C, which is a 22% better than DDPG (0.0402C). Moreover, the robustness test with off

nominal initial conditions (such as 40C and 38C) confirmed that both RL agents exhibit stable, almost instantaneous convergence without the large oscillations typical of PI.

12:00 ***A survey on formal methods for authentication in self-driving platforms***

Muhammad Salman Saeed (ITMO University, St. Petersburg, Russia); Sergey Bezzateev (State University of Aerospace Instrumentati, Russia); Umer Mukhtar Andrabi (National Research University Higher School of Economics, Russia); Ehsan Wadood (Moscow Institute of Physics and Technology, Russia)

Self-driving vehicles are experiencing a surge in popularity now a days as application areas such as the Internet of Vehicles are emerging. The rapid development of unmanned car platforms requires the development of reliable and efficient systems and methods to ensure their safe operation. However, the security of communication between the self-driving car and its base transceiver station is critical to accomplishing its task without revealing critical information to either adversaries or unauthenticated users. Self-driving cars are particularly vulnerable to physical interception and node tempering attacks. One of the tasks in this area is the task of reliable authentication of devices in systems with rapidly changing infrastructure. A possible solution to solve this problem is to use authentication protocols, such as the vehicle proving its authenticity using information about the surrounding vehicles. In this paper, we try to compile all possible attacks that can be executed on an autonomous vehicle by tempering the vehicle node. After that, we compile all the previously available solutions with their techniques and efficiencies for securing autonomous vehicles from such attacks. Moreover, we provide a brief overview of four different ways of using formal methods that commonly appeared in previous research papers.

12:15 ***Machine Learning-Based Predictive Framework for OBD-II Vehicle Fault Detection and Maintenance Forecasting***

Jerry Incierto Teleron (Technological Institute of the Philippines - Quezon City, Philippines)

The increasing complexity of modern vehicles has necessitated the development of intelligent diagnostic systems capable of analyzing real-time vehicular data for efficient fault detection and maintenance forecasting. This study developed a machine learning-based predictive diagnostic framework using On-Board Diagnostics II (OBD-II) data to enhance automotive maintenance practices. A dataset consisting of 5,000 instances was generated using standardized OBD-II parameters, including engine RPM, vehicle speed, engine temperature, throttle position, fuel level, intake air temperature, mass air flow, oxygen sensor voltage, and battery voltage. Three machine learning models-Random Forest, Support Vector Machine, and Extreme Gradient Boosting (XGBoost)-were implemented and evaluated using classification and regression metrics. Results showed that XGBoost achieved superior performance in fault detection accuracy and demonstrated reliable maintenance prediction capability. Based on these findings, a web-based diagnostic system was developed integrating the XGBoost model, enabling automotive mechanics and technicians to input real-time data and obtain automated diagnostic results. The system also provided diagnostic history and visualization features to support decision-making. Usability evaluation indicated a very high level of user satisfaction, confirming the system's practicality and effectiveness. The study demonstrates that integrating machine learning with OBD-II systems can significantly improve vehicle diagnostics and predictive maintenance.

12:30 ***The Integration of Artificial Intelligence Models as a Tool for Justice System Optimization***

Magda Beruashvili (Business and Technology University, Georgia)

One of the main challenges in the justice system is the backlog of cases and long review periods, which violates the constitutional principle of a "fair trial." In modern conditions, the sharp increase in the volume of cases and the complexity of court proceedings require optimizing legal mechanisms. The study confirms that integrating artificial intelligence (AI), particularly through Machine Learning and Natural Language Processing (NLP) models, optimizes judicial processes, improves decision quality, and makes justice more accessible to citizens. The article aims to present the role and importance of these two main models, analyze international practice, and define an appropriate implementation strategy for Georgia. One of the main challenges in the justice system is the backlog of cases and long review periods, which violates the constitutional principle of a "fair trial." In modern conditions, the sharp increase in the volume of cases and the complexity of court proceedings require optimizing legal mechanisms. The study confirms that integrating artificial intelligence (AI), particularly through Machine Learning and Natural Language Processing (NLP) models, optimizes judicial processes, improves decision quality, and makes justice more accessible to citizens. The article aims to present the role and importance of these two main models, analyze international practice, and define an appropriate implementation strategy for Georgia. The Georgian judicial system has long faced critical case backlogs, with judges overwhelmed by rapidly growing caseloads. During the pandemic, over 120,000 cases were filed in the Tbilisi City Court alone, exposing the system's inability to ensure timely resolution - a core constitutional requirement. Machine Learning and Natural Language Processing offer a direct response to these structural inefficiencies. ML models analyze historical judicial data to predict case duration, estimate outcomes, and balance judicial workloads, while NLP automates the processing of legal documents and legislative texts, reducing human error and saving administrative time. Together, they form a technological foundation capable of improving both the speed and quality of judicial decision-making, without replacing the human judge. International experience confirms this potential. The United Kingdom's Money Claim Online platform has enabled citizens to resolve disputes of up to £100,000 entirely online since 2002, with 80% of users rating it easy to use. Estonia conducts over 90% of proceedings electronically, achieving a 30-40% reduction in processing times for small disputes. Singapore's real-time transcription system and AI chatbots have made court services faster and more citizen-friendly. A survey of 125 Georgian citizens found that over 70% are willing to use digital court platforms, while judges broadly support AI as a technical support tool, provided final authority remains human. For Georgia, a four-stage strategy is recommended: establishing a legal and ethical framework; developing a national AI integration strategy; piloting a digital platform in Tbilisi, Batumi, and Kutaisi for loan disputes under 5,000 GEL; and scaling the system nationwide. The ultimate goal is a faster, more transparent, and trustworthy justice system that genuinely protects citizens' rights

12:45 ***Deep Comparative Analysis of Six Machine Learning and Deep Learning Models for OBD-II-Based Predictive Maintenance***

Jerry Incierto Teleron (Technological Institute of the Philippines - Quezon City, Philippines)

The increasing complexity of modern vehicles has necessitated the development of intelligent diagnostic systems capable of analyzing real-time vehicular data for efficient fault detection and maintenance forecasting. A previous study developed a machine learning-based predictive diagnostic framework using On-Board Diagnostics II (OBD-II) data, utilizing a dataset of 5,000 instances and implementing Random Forest (RF), Support Vector Machine (SVM), and Extreme Gradient Boosting (XGBoost), where XGBoost demonstrated superior performance and was successfully deployed in a web-based diagnostic system. Building upon these findings, this study serves as a continuation by expanding the dataset to 10,000 instances collected from both synthetic and real-world vehicle data and extending the analysis to include deep learning models. Specifically, six models-LSTM, BiLSTM, CNN, CNN-LSTM, Transformer, and XGBoost-were implemented and evaluated using Mean Absolute Error (MAE), Root Mean Square Error (RMSE), training and validation loss, and residual analysis. Results indicate that XGBoost remains the best-performing model, achieving the lowest error values and demonstrating stable and consistent predictive performance. In contrast, deep learning models exhibited higher error rates and signs

of overfitting, highlighting their dependence on larger datasets and more complex tuning. The selected model was integrated into a web-based predictive maintenance system, enabling real-time vehicle diagnostics and decision support. The findings confirm that ensemble learning methods remain highly effective for structured OBD-II data, while deep learning approaches present opportunities for future improvement with larger datasets.

Wednesday, April 22 11:30 - 13:00 (Africa/Cairo)

SESSION-4C: RESEARCH: Machine Learning for Disease Detection & Biomedical Classification

(6 PAGES)

Room: ROOM-C

Chair: Dua Weraikat (Rochester Institute of Technology, USA)

11:30 ***A Machine Learning Pipeline for Early Detection of Autism Spectrum in Children***

Mariam Mahmoud Dahish and Reem Kadry Montasser (October University for Modern Sciences and Arts (MSA), Egypt); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

autism spectrum disorder is a neurological disorder characterized by a set of symptoms, including difficulty in communication or social interaction. This disorder has become more prevalent in recent years due to its impact on individuals and families. The significance of early diagnosis for children with autism is that early diagnosis allows them to progress and develop the outcomes in their lives. Also, traditional diagnosis is slow and requires long sessions and interviews with the family, and occasionally the diagnosis can be delayed for up to four years, particularly for people in resource-limited countries. In this study, we applied XGBoosting, SVM, random forest, and logistic regression to find the best performance in machine learning for early detection in children. was the highest ml algorithm applied, is XGBoosting accuracy 99%.

11:45 ***Detection-Guided Deep Learning for Segmentation Free Skin Lesion Diagnosis***

Hagar Selim and Hadeer Elkady (Arab Academy for Science, Technology, and Maritime Transport, Egypt); Nahla Belal (Arab Academy for Science, Technology, and Maritime Transport & College of Computing and Information Technology, Egypt); Mohamed Khedr (Arab Academy for Science and Technology, Egypt)

The paradigm of automated skin lesion analysis has traditionally been dominated by multi-stage pipelines that prioritize lesion segmentation as a prerequisite for classification. However, pixel-level segmentation is computationally intensive and highly sensitive to varying image conditions. This pa- per investigates a critical research question: Is segmentation truly necessary for high-accuracy diagnosis? We propose a segmentation-free, detection-guided classification framework that leverages YOLOv8 for precise lesion localization and ResNet18 for multi-class categorization. By treating localization as an object detection task rather than a pixel-wise masking task, we streamline the diagnostic pipeline. Our framework was

evaluated on the ISIC 2018 dataset, targeting three clinically significant categories: Melanoma (MEL), Melanocytic Nevus (NV), and Be- nign Keratosis (BKL). The results demonstrate that the proposed detection model achieves a mAP@0.5 of 0.986, while the classi- fication module maintains robust performance across all classes. Our findings suggest that detection-guided localization provides sufficient spatial context for reliable diagnosis, offering a more efficient and scalable alternative to traditional segmentation- based Computer-Aided Diagnosis (CAD) systems.

12:00 ***A lightweight Transfer Learning Framework Based on MobileNetV2 for Automated Brain Tumor Classification from MRI Images***

Shaimaa Othman and Mohamed Hussein (Teacher, Egypt); Alaa Hassan (Teacher Assistant, Egypt); Mrawa Mostafa (Teacher, Egypt)

Primary brain malignancies pose a critical challenge to global oncology, necessitating rapid and precise diagnostic frameworks for optimal therapeutic intervention. While Magnetic Resonance Imaging (MRI) remains the gold standard for non-invasive visualization of cerebral pathologies, manual radiological assessment is inherently prone to intra-observer variability and high cognitive workloads. This study proposes a computationally efficient deep learning framework for the automated classification of brain tumors, utilizing a Convolutional Neural Network (CNN) integrated with a MobileNetV2 backbone. By leveraging depthwise separable convolutions and transfer learning, the model achieves high-fidelity feature extraction with minimal parameter overhead. Evaluated on a comprehensive dataset of 7,023 multi-sequence MRI scans-encompassing glioma, meningioma, pituitary tumors, and healthy controls-the proposed architecture attained a peak classification accuracy of 99.81%. Beyond accuracy, the model demonstrated robust performance with a precision of 98%, a recall of 97.95%, and an F1-score of 98.92%. The framework's efficiency is highlighted by its lightweight structure, consisting of approximately 2.26 million parameters, making it suitable for deployment in resource-constrained clinical environments. Experimental results, validated through detailed confusion matrix analysis and training-validation performance curves, indicate that lightweight architecture can match or exceed the performance of deeper, more resource-intensive models, providing a scalable solution for real-time clinical decision support systems.

12:15 ***CLIP-DR: Efficient Multi-Grade Diabetic Retinopathy Classification Using SMOTE and Focal Loss with Fine-Tuned Vision Transformers***

Nesrine Atitallah (FCS, Arab Open University, Madinah, KSA, Saudi Arabia); Abuelgasim Abusonoun (Arab Open University, Saudi Arabia); Khouloud Samrouth (Lebanese University, Lebanon); Nader Bakir (Beirut Arab University, Lebanon)

Diabetic Retinopathy (DR) remains a leading cause of preventable vision loss globally, yet accurate automated screening faces significant challenges due to severe class imbalance and limited labeled data for rare severity grades. To address these constraints, this study presents a transfer learning framework combining CLIP Vision Transformers with advanced imbalance handling techniques. We employed a pre-trained CLIP ViT-L/14 encoder with selective layer fine-tuning, training the last four transformer layers alongside a lightweight classification head. To handle severe class imbalance (9.4:1 ratio between majority and minority class), we applied SMOTE (Synthetic Minority Over-sampling Technique) in the 768-dimensional CLIP feature space, generating synthetic representations for underrepresented severity grades. Additionally, we employed Focal Loss ($\gamma = 2.0$) to prevent overfitting on synthetic samples by focusing training on hard, misclassified examples. Validated on the APTOS 2019 Blindness Detection dataset (3,662 rural Indian fundus images), our approach achieved a quadratic weighted kappa of 0.9194 and 84.64% accuracy on the test set, demonstrating almost perfect clinical agreement ($\kappa > 0.81$). Notably, recall for Proliferative DR - the most vision-threatening stage- reached 73.9%, representing substantial

improvement over baseline approaches. These results suggest a viable path toward deploying accurate, class-balanced DR screening in resource-constrained telehealth settings

12:30 ***Early Stage Detection of Chronic Kidney Disease using Machine Learning***

Himadri Bahuguna, Shraddha Mandal, Ishaan Bhairab Sensharma, Vikas Upadhyaya and Neha Tiwari (NIIT University, India)

Chronic Kidney Disease (CKD) is an illness that is long term and progresses over time causing people to lose their kidney function. This study aims to detect CKD at an early stage through the use of machine learning (ML) techniques, XGBoost, Gradient Boosting, AdaBoost, Random Forest, and Support Vector Machine (SVM); they were trained on clinically relevant biomarkers selected specifically for the early detection of the disease. A dataset of 1,659 patient records was refined so that the emphasis was on early stage risk factors. The results demonstrate consistently high predictive performance across all models, with accuracy exceeding 91% and F1-scores above 95%. AdaBoost got the best performance overall with an F1-score of 96.3% and recall of 99.7%, indicating its strong ability to correctly identify CKD cases while reducing false negatives which is an essential requirement in clinical diagnostics. The predictive performance of the model was high even when the Receiver Operating Characteristic Area Under the Curve (ROC-AUC) scores were relatively low. The novelty of this research lies in its concentration on the selection of biomarkers during the initial stages, guided by clinical thresholds, along with the comparison between various ensemble algorithms using a clinically meaningful criterion framework. This work highlights how machine learning algorithms can be used as a potential diagnostic tool for early CKD screening and can support timely intervention and better healthcare outcomes in the real world scenario.

12:45 ***Multimodal Patient-Level Acute Myeloid Leukemia Subtype Prediction via Single-Cell Cytomorphology and Clinical Features***

Ranya Khaled Elsawah (The knowledge hub universities, Egypt); Gamal A. Ebrahim (Ain Shams University, Egypt); Hani A. Ghali (British University in Egypt (BUE), Egypt)

Patient-level prediction of acute myeloid leukemia (AML) subtypes remains challenging due to cellular heterogeneity and class imbalance. This paper presents a multimodal patient-level study that combines single-cell peripheral blood cytomorphology and clinical features for AML subtype prediction. An attention-based multiple instance learning (AB-MIL) model is employed to aggregate variable numbers of single-cell images into patient-level representations. In parallel, a clinical baseline is constructed using routinely available variables, including age, sex, blood counts, and blast percentages. Multimodal integration is performed using late fusion at the probability level to enable controlled comparison between unimodal and multimodal models. Experiments are conducted on a publicly available dataset of 189 patients covering four AML subtypes and a control class. The image model achieved a mean Macro-F1 score of 0.526, while the clinical model attained 0.627. The proposed fusion approach improved performance to 0.696, outperforming both unimodal approaches.

Wednesday, April 22 11:30 - 13:00 (Africa/Cairo)

SESSION-4D: RESEARCH: Explainable AI (XAI): Methods, Interpretability & Trustworthiness

(6 PAGES)

Room: ROOM-D

Chair: Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

11:30 ***Explainable and Calibrated Hierarchical Deep Learning for Thermal Infrared PV Fault Diagnosis***

Khawla Ould Mansour and Yousra Chtouki (Al Akhawayn University, Morocco)

According to the IEA PVPS Trends 2025 report, cumulative global photovoltaic (PV) capacity surpassed 2.2 TW by early 2025, indicating rapid expansion of solar deployment worldwide. At this scale, annual PV energy yield losses due to soiling and operational faults are estimated at 3 % to 5 %, corresponding to €3 billion to €5 billion in lost revenue each year, underscoring the urgent need for reliable and scalable fault detection and diagnosis solutions for PV systems.

To address these challenges, this paper proposes an explainable and calibrated hierarchical deep learning framework for photovoltaic fault diagnosis using thermal infrared imagery. The diagnosis task is decomposed into three sequential stages: binary fault detection, coarse fault group classification, and fine-grained fault identification. Multiple deep learning models are systematically evaluated in the early stages, after which the selected model is applied to fine-grained fault identification.

Experiments on a public PV thermal dataset achieve 95.48 % accuracy with a 95.91 % macro F1-score for binary fault detection, 84.69 % accuracy with an 83.22 % macro F1-score for coarse fault grouping, and 86.31 % overall accuracy across twelve fine-grained fault categories. To improve prediction reliability, post-training calibration via temperature scaling is applied, reducing Expected Calibration Error from 0.446 to 0.069 at the fine-grained level.

Model interpretability is provided through gradient-based class activation mapping, confirming that predictions are driven by physically meaningful thermal regions consistent with known fault mechanisms. Overall, the proposed framework provides a robust, reliable, and interpretable solution for PV thermal fault diagnosis suitable for real-world inspection and monitoring.

11:45 ***Stacked Ensemble Learning and Explainable AI for Dust Storm Prediction in Saudi Arabia***

Hager Saleh (South Valley University, Egypt); Mohammad Wajeeh (Midocean University, Macao); Michael Mccann (Atlantic Technological University, Ireland); John G. Breslin (University of Galway, Ireland); Sarah Osama (Minia University, Egypt); Priyanka Verma (University of Galway, Ireland)

Dust storms are a recurring environmental hazard in the Middle East, with serious effects on public health, transportation, and infrastructure. This paper introduces a new dataset for daily dust-storm detection in Saudi Arabia. It combines dust variables from NASA MERRA-2, meteorological variables from European Centre for Medium-Range Weather Forecasts (ECMWF) Reanalysis v5- Land (ERA5-Land), and aerosol measurements from

MODIS, from 2020 to 2024. Different preprocessing steps are applied: handling missing values, deriving variables, and data integration to enhance the quality of the dataset. In addition, XGBoost feature importance is used to select the best features. Also, Grid Search with three-fold cross-validation is used as an optimization model technique to optimize the model and enhance the results. We propose a stacking ensemble model that integrates Random Forest (RF), XGBoost, and Support Vector Machine (SVM) classifiers as base models, with RF as the meta-learner to improve performance and enable generalization. Also, the results showed that the stacking ensemble achieves the best performance with full features, with 86.10% precision and 85.70% recall, compared to individual ML models and to AdaBoost and XGBoost ensembles with selected features. Explainable AI using SHAP (SHapley Additive exPlanations) provides interpretable insights into model predictions by quantifying each feature's contribution to the final decision, thereby enhancing transparency and trust in machine learning systems.

12:00 ***Bayesian Decision Models Under Uncertainty in BPM***

Lily Petriashvili (Georgian Technical University, Georgia); Nino Topuria (Georgian-Technical University, Georgia & GTU, Georgia); Tamar Lominadze, Taliko Zhvania and Mzia Kiknadze (Georgian Technical University, Georgia); David Kapanadze (Georgian-Technical University, Georgia)

Today's business environment is undergoing a rapid digital transformation accompanied by uncertainty, dynamic change, and competitive pressures. Traditional analytical techniques often fail to manage large amounts of unstructured data. Accordingly, the introduction of innovative methods based on probabilistic modeling and artificial intelligence becomes relevant, since it allows making business decisions not only quickly and flexibly, but also reliably and transparently. In this context, Bayesian models are of particular importance, since their essence is based on uncertainty management, updating knowledge, and adapting forecasts to take into account new data. Bayesian theory allows you to determine the probability of events by synthesizing existing knowledge and new information, which creates an effective basis for optimizing business processes. In digital business management, such models allow you to make decisions in real time, assess risks, and compare alternative scenarios to identify the optimal one. Thus, the integrated use of Bayesian methods and machine learning algorithms plays an important role. This synthesis improves prediction accuracy, reduces information gaps, and provides multi-factor data analysis. In particular, models based on Bayesian networks are successfully used in analyzing consumer behavior and predicting their demand.

12:15 ***The Hidden Complexity of Mental Health: Multi-Entropy Analysis of Response Patterns in Depression Severity Assessment***

Ayesha Siddika, Ahnaf Atif Rafi and Maria Jahan Noon (Independent University, Bangladesh, Bangladesh); Md Junayed Hossain (Independent University, Bangladesh & Centre for Computational and Data Sciences (CCDS), Bangladesh); Ashraful Islam (Independent University Bangladesh & Center for Computational and Data Sciences, Bangladesh); Sanzar A Alam (Independent University, Bangladesh & Center for Computational & Data Sciences, IUB, Bangladesh)

The Patient Health Questionnaire-9 (PHQ-9) is commonly used to access the severity of depression, and the clinical decision is made primarily based on the overall score received. Nevertheless, individuals with the same scores might differ markedly in response patterns, indicating that there are some underlying differences in their mental conditions. To address this weakness, this study presents a more sophisticated computational model that represents the sophisticated character of PHQ-9 responses beyond conventional score-based assessments. Individual item responses are combined with many information-theoretic metrics, including Shannon entropy, Sample entropy, Permutation entropy, and Multiscale entropy, to create a better

set of features to use in classifying the severity of depression. Empirical outcomes of the fivefold stratified cross-validation show that the suggested solution is significantly more effective than the baseline models, which are based only on sum scores and that a Random Forest classifier is more accurate with 99.2% on average than the baseline; which is 81.2%. Statistical tests also indicate that the response entropy is non-linearly dependent on the severity levels of depression, with the highest entropy occurring at moderate depression levels and decreasing with severe depression, and the optimal entropy threshold (τ = 0.439) is determined in which screening is clinically meaningful. These results indicate the importance of the complexity of the response pattern as a compensatory dimension in the evaluation of digital depression.

12:30 ***Explainable AI in Epileptic Seizure Detection***

Sakshi Prasad, Bhavika Pawar, Sonali Paliwal and Vikas Upadhyaya (NIIT University, India)

Seizure detection via analyzing EEG (electroencephalogram) values is a challenging task for clinicians due to the ambiguity and complexity of machine learning algorithms with high performance. In this study, a model of binary classification seizures was constructed using the large sample dataset from the Bangalore EEG Epilepsy Dataset (BEED) comprising 8,000 entries recorded using 16 electrodes attached to the scalp. The machine learning algorithms used in the experiment are Logistic Regression, Decision Tree, Random Forest, K-Nearest Neighbors, and Support Vector Machine. Moreover, among the selected deep learning algorithms are MLP, CNN-1D, BiLSTM, CNN-BiLSTM, BiGRU, and Transformer trained only on raw values of amplitudes in EEG data without manually engineering features. Such algorithms as SVM and Random Forest provided AUC-ROC scores of 100%, and MLP and CNN-BiLSTM had perfect accuracy results in all metrics. To make the model interpretable and usable for clinicians, SHAP (SHapley Additive exPlanations) was applied on a Stratified sample of 500 tests using GradientExplainer (MLP) and KernelExplainer (SVM) techniques. As a result, both models revealed that channel C4 (Right Central) was the most influential electrode of the EEG.

12:45 ***Explainable AI for Industrial Robotic Fault Diagnostics: Bridging the Trust Gap***

Loso Judijanto (IPOSS, Indonesia); Sri Nurhayati (IKIP Siliwangi, Indonesia)

The increasing reliance on industrial robots across manufacturing sectors has underscored the urgent need for reliable and transparent fault diagnostic systems. While machine learning and deep learning models have demonstrated impressive fault detection capabilities, their opaque nature has limited user trust and hindered broader adoption in critical applications. This study investigates how Explainable Artificial Intelligence (XAI) can bridge this trust gap by enhancing the interpretability of fault diagnostics in industrial robotic systems. Employing a qualitative research design through a Systematic Literature Review (SLR) methodology, this study analysed 24 peer-reviewed articles published between 2020 and 2025 from the ScienceDirect database. Data collection was conducted via structured keyword-based filtering following the PRISMA protocol. Inclusion criteria included open-access availability, relevance to industrial robotics, and application of XAI in fault diagnostics. The collected data were analysed using thematic coding to identify recurring methods, challenges, and benefits associated with XAI in robotic fault detection. The findings indicate that post-hoc XAI techniques such as SHAP, LIME, and Grad-CAM are commonly used to enhance interpretability, with hybrid models offering a balance between accuracy and transparency. Key challenges include computational latency, lack of standardised evaluation metrics, and limited access to diverse diagnostic datasets.

Wednesday, April 22 13:00 - 14:00 (Africa/Cairo)

LUNCH: LUNCH (ON_SITE)

LUNCH

Room: COFFEE_ROOM

Wednesday, April 22 14:00 - 15:30 (Africa/Cairo)

SESSION-5A: RESEARCH: Computer Vision: Object Detection, Recognition & Surveillance

6 PAPERS

Room: ROOM-A

Chair: Devendra Rajput (Accenture, USA)

14:00 ***TowerTrace: Scalable Detection of Telecom Infrastructure via Multimodal VLMs***

Mohamed Tharwat (ESLSCA University, Egypt); Ayman Gaber, Mohamed Abdelhakeem and Islam Ashraf (Vodafone Egypt, Egypt); Magy Hossam (Eslsca, Egypt)

Deploying cellular infrastructure in rural and underserved areas is often hindered by high costs and logistical challenges. This paper presents TowerTrace, a framework that uses high-resolution satellite imagery and transformer-based vision-language models (VLMs) for automated detection of existing structures suitable for telecom equipment. Unlike conventional CNN pipelines requiring large labeled datasets, TowerTrace employs prompt-based inference with pretrained multimodal VLMs, enabling accurate zero- and few-shot detection without retraining. The system grids the target area using static map APIs for deterministic, cost-efficient coverage and identifies diverse structure types, achieving consistently high accuracy across varied environments. This approach offers a scalable, adaptable, and low-cost solution for accelerating network planning and reducing capital expenditure. The methodology is validated across multiple geographic regions in Egypt, demonstrating strong generalization to different terrain and infrastructure patterns. The results highlight its potential as a practical, field-ready tool for telecom operators aiming to expand coverage efficiently.

14:15 ***Deep Learning-Based Detection of Cluster Bomb Submunitions Using YOLO***

Mostafa Rizk (Lebanese University, Lebanon); Sadek Fakih (Lebanese Univerity, Lebanon); Abbas Rammal (Phoenicia University, Lebanon)

Unexploded ordnance and cluster munitions pose a significant threat to civilians in conflict-affected regions, necessitating efficient and reliable detection methods. This paper proposes an artificial intelligence-based approach for detecting cluster bomb submunitions using advanced object detection models. The study investigates the performance of recent models of You Only Look Once (YOLO) family, with a focus on YOLOv8 and a comparison with the recent YOLO11n architecture, for detecting small objects with diverse shapes and features. A novel dataset comprising 1080 images and 2121 annotations across five submunition classes commonly found in the Middle East is introduced. To address class imbalance, a

balanced dataset of 8,500 images is generated using additional samples and data augmentation techniques. Multiple variants of YOLOv8 and YOLO11n are trained and evaluated on the original dataset, followed by experiments using YOLOv8n on the balanced dataset. Experimental results demonstrate strong detection performance, achieving a mean average precision (mAP) of approximately 84% for the medium model. Furthermore, dataset balancing significantly improves performance, yielding recall of 0.95 and mAP of 0.97. These findings highlight the effectiveness of YOLO-based models and the importance of balanced data in improving detection accuracy for critical real-world applications.

14:30 ***Vision-Based in-Bed Human Pose Estimation Using RGB Images and Pose Estimation Networks***

Jana Moatazbellah (October University for Modern Sciences & Arts (MSA), Egypt); Mazen A. Ebrahim (MSA University, Egypt & University of Greenwich, United Kingdom (Great Britain)); Mohamed Nagy (Modern Science and Arts University, Egypt); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

In bed pose estimation of humans plays an essential role within health-related contexts such as sleep assessment and evaluating postures of patients while resting. Most previous studies have required specialized equipment or an array of different types of data, leading to high system expenses. This paper describes a research project that involved the estimation of human pose based on computer vision in the bed using only RGB images and dedicated to precise safe observation. Using only the RGB modality, the suggested method is assessed on the Simultaneously-collected multimodal Lying Pose (SLP) dataset evaluated on three pose estimation models HRPose, ResPose, and PoseAttention are implemented and compared. Model performance is measured using the Percentage of Correct Keypoints normalized by (PCKh@0.5). results show that HRPose achieves a PCKh@0.5 score of 90.9%, while the PoseAttention model attains approximately 91.0% on the SLP RGB test set. ResPose achieves the best overall performance with a total PCKh@0.5 score of 91.2%. The findings show that RGB photos alone can provide cost effective in bed human posture estimate, making it a viable option for healthcare and sleep monitoring applications.

14:45 ***A Systematic Benchmarking of XAI Methods for Weapon Recognition for Video Surveillance***

Haya AlMadhloum AlSuwaidi, Fatih Kurugollu and Abbes Amira (University of Sharjah, United Arab Emirates); Muhammad Shahroz Nadeem (University of Suffolk, United Kingdom (Great Britain))

Automated analysis of surveillance video plays a critical role in modern security and public safety systems. In weapon-related activity recognition, high classification accuracy alone is insufficient; system predictions must also be interpretable to support trust, auditing, and operational decision-making. This paper presents a video-level weapon-related activity classification framework based on a convolutional neural network-long short-term memory (CNN-LSTM) architecture, augmented with explainable artificial intelligence (XAI) techniques. The proposed approach models both spatial appearance and temporal dynamics by extracting frame-level features using a convolutional backbone and learning motion patterns across time using an LSTM network. Model performance is evaluated on a heldout test set using accuracy, precision, recall, and F1-score, achieving a substantial classification performance under controlled experimental conditions, where Grad-CAM++ performed the best in both computability and interpretability. To enhance transparency, gradient-based Class Activation Mapping (CAM) techniques, including Grad-CAM, Grad-CAM++, and Eigen-CAM, are employed to visualize spatial regions contributing to model predictions. Results demonstrate that the proposed framework effectively distinguishes weapon-related scenarios from non-violent activities while providing interpretable visual explanations. The findings highlight the feasibility of

integrating explainability into video-based weapon detection pipelines and underscore the importance of transparent AI systems in security-critical applications.

15:00 ***Early Detection of Acute Myeloid Leukemia (AML) Using YOLOv12 Deep Learning Model***

Enas Emad (Cairo University, Egypt); Salah A. Aly (Fayoum University, Egypt); Mayar Moner (Cairo University, Egypt)

Acute Myeloid Leukemia (AML) is one of the most life-threatening type of blood cancers, and its accurate classification is considered and remains a challenging task due to the visual similarity between various cell types. This study addresses the classification of the multi-classes of AML cells Utilizing YOLOv12 deep learning model. We applied two segmentation approaches based on cell and nucleus features, using Hue channel and Otsu thresholding techniques to preprocess the images prior to classification. Our experiments demonstrate that YOLOv12 with Otsu thresholding on cell-based segmentation achieved the highest level of validation and test accuracy, both reaching 99.3%.

15:15 ***Cross-Domain Vulnerability Analysis of YOLOv5 to Physical Adversarial Patches in Safety-Critical Object Detection***

Mashari AlRowaili, Abdelhak Belhi, Mustafa Al Samara and Nabil Litayem (Joaan Bin Jassim Academy for Defence Studies, Qatar)

Object detectors deployed in safety-critical applications such as autonomous driving, traffic monitoring, and aerial surveillance must remain reliable under both environmental variation and adversarial interference. Physical adversarial patches pose a practical threat, as they can be introduced into real-world scenes without access to the target system. This paper presents a systematic cross-domain evaluation of YOLOv5 under a gray-box threat model across four operational domains: self-driving, traffic-sign recognition, military detection, and UAV-based perception. For each domain, adversarial patches are generated using Expectation over Transformation (EoT) and evaluated against all detectors, resulting in a complete cross-domain attack matrix. The results reveal significant vulnerability across multiple source-target pairs, including strong transfer from ground-view to aerial-view settings and clear asymmetry between domains. Notably, detectors with near-perfect clean accuracy can exhibit substantial degradation under attack, indicating that nominal performance is a poor proxy for robustness. Among the evaluated models, the military detector demonstrates comparatively higher resilience, while self-driving and drone detectors are more susceptible. These findings highlight the importance of cross-domain evaluation for safety-critical perception systems and demonstrate that adversarial robustness must be assessed beyond standard clean-data benchmarks.

Wednesday, April 22 14:00 - 15:30 (Africa/Cairo)

SESSION-5B: RESEARCH:Generative AI, Digital Education & Societal Impact

6 PAPERS

Room: ROOM-B

Chair: Shyalendar Reddy Allala (Global Atlantic Financial Company, USA)

14:00 ***Transforming Environmental Literacy with AI: The AIEC Model for University-Industry Sustainability Alignment***

Constantine Andoniou (Abu Dhabi University, United Arab Emirates)

Environmental literacy serves as an essential base for people to handle climate uncertainty, sustainability transitions and environmental data usage in various industries. The current environmental education methods fail to deliver sufficient analytical capabilities, real-time data interpretation and precise interpretation needed for modern ecological problems. Artificial intelligence technology enables the improvement of environmental literacy through its ability to extract valuable information from extensive environmental datasets, create interactive models, customized learning routes and enhanced interpretation methods, for educational and industrial environments. The current research develops the AI-Enabled Environmental Cognition (AIEC) model which explains AI-based environmental literacy development in higher education and demonstrates its value for university-industry sustainability partnerships. The AIEC model uses AI to create an interactive learning process which includes four stages: environmental data interpretation, simulation-based learning, personalized ecological education and cross-industry knowledge sharing. The research shows how AI technology enables universities and industry partners to share a common analytical framework which helps them understand climate data, create sustainability solutions and make decisions together. Furthermore, it discusses how AI applications in climate-tech development, environmental dashboards, explainable AI systems and sustainability analytics play an important role in creating opportunities for sharing transparent environmental knowledge. The paper concludes with recommendations about AI technology applications for environmental education which help achieve SDG targets and build stronger university-industry sustainability partnerships.

14:15 ***The impact of data analytics on decision making Evidence based practices***

Mahmoud Allahham (Amman Arab University, Jordan); Bashar Khaled Almagharbeh (Amman Arab University Amman, Jordan); Jawad Hasan AlKhateeb and Taha Houda (Prince Mohammad Bin Fahd University, Saudi Arabia); Jihad Jaam (Liverpool J. Mores University, United Kingdom (Great Britain))

The paper examines how educational data analytics influence student-level decision-making in schools and will focus on how evidence-based practice mediates increases in the effectiveness of data-driven strategies. The research design was a quantitative one and a structured questionnaire was given to 350 teachers, principals and school administrators in Jordanian schools. The research paper focuses on the impact of the systematic use of the education data on the strategic, teaching and learning, and administrative decision-making. The findings indicate that educational data analytics plays a significant role in enhancing the quality of decisions made by the decision-makers as it gives them the correct information with regard to student performance, resource allocation and curriculum effectiveness. Also, this relationship is partially mediated by evidence-based practice, which means that the more sustainable educational outcomes are achieved based on the analyzed data and empirical evidence. Nevertheless, the factors that limit the comprehensive adoption of analytics in the school practice are lack of data literacy, intermittent data quality, and inadequate analytical infrastructure. According to the findings, the practice of forming data-informed culture at the school level is crucial and relies on the fact that evidence-based practice can turn crude data into actionable plans. The paper hypothesizes that to enable educational data analytics to contribute to effective, transparent, and informed decision-making, it is crucial to amuse educators with their independent and evaluative abilities

14:30 ***AI Governance for Assessment and Academic Integrity in Higher Education An Institutional Illustrative Case***

Omar Al-Jarrah (Vice President for Planning and Development & CIO, Kuwait); AWS Ismail Abueid (Arab Open University - KUWAIT, Kuwait)

The accelerated integration of artificial intelligence (AI) in higher education has transformed assessment into a governance-intensive and risk-sensitive domain rather than a purely technological or pedagogical innovation. While AI-enabled systems are increasingly embedded in institutional assessment practices, their deployment has often outpaced the establishment of structured governance mechanisms necessary to ensure reliability, transparency, and academic integrity. This governance gap has heightened concerns regarding accountability, fairness, and institutional trust in AI-supported assessment processes. Consequently, coherent governance frameworks are required to mitigate systemic risks and to position AI as a controlled and accountable component within the broader academic infrastructure. This paper proposes a strategic governance framework for AI-enabled assessment grounded in institutional oversight, structured risk management, and regulatory alignment. A defining feature of the framework is the explicit preservation of human-in-the-loop decision authority, ensuring that artificial intelligence operates strictly as a decision-support mechanism. At the same time, final academic judgments remain under human responsibility. The framework is illustrated through an institutional case of governance-embedded AI systems implemented within formal policy boundaries and oversight structures. By conceptualising assessment and academic integrity as high-risk institutional domains, the study underscores the importance of governance-oriented AI adoption strategies that balance innovation with accountability. The proposed framework provides practical guidance for academic leaders, policymakers, and system designers seeking to integrate artificial intelligence into higher education assessment while safeguarding academic standards, institutional credibility, and professional judgment

14:45 ***The Effect of Generative AI Use on Programming Self-Efficacy in Computer Engineering Students***

A.M. Mutawa, Shouq Alsubaihi and Sai Sruthi (Kuwait University, Kuwait)

The fast adoption of Generative AI (GenAI) tools like ChatGPT, Gemini, and Claude has completely changed the way people learn to program. Even while these tools provide real-time guidance, we still don't know much about how they affect students' self-efficacy. This research examined computer engineering students (N = 24) after a foundational Java programming course. The 12-item Introductory Programming Self-Efficacy Scale (IPSES), specifically adapted for Java programming, was used to measure self-efficacy. We used the statistical package to analyze data on GenAI use and demographic characteristics using multiple linear regression models. The regression models demonstrated that GenAI usage was not a significant predictor of self-efficacy ($p > 0.05$). A moderate regression, on the other hand, showed a significant difference between students' computer knowledge and programming self-efficacy. The results indicate that GenAI tool usage does not generally enhance programming confidence. Instead, its effect is influenced by gender-based adoption trends.

15:00 ***The Role of Generative AI in Audiovisual Production: Opportunities, Limitations and Emerging Workflows***

Sónia Almeida Ferreira, Sr (Polytechnic of Viseu, Portugal & School of Education, CI&DEI, Portugal); Vanessa Reis Campos (School of Education, Portugal)

The rapid evolution of Generative Artificial Intelligence (GenAI) is significantly reshaping audiovisual production processes, introducing new possibilities for automation, efficiency, and creative support. This paper presents a conceptual and critical analysis of the role of GenAI tools across the main stages of audiovisual production: pre-production, production, and post-production. Drawing on a systematic review of recent literature, the study examines how these technologies are being integrated into workflows, highlighting their potential to assist in tasks such as scriptwriting, content generation, editing, and distribution. The analysis identifies key benefits associated with GenAI adoption, including increased productivity,

reduced operational costs, and greater accessibility to audiovisual creation. However, it also emphasizes important limitations, particularly regarding creative control, output quality, technical constraints, and ethical concerns. Based on these insights, the paper proposes a conceptual framework for integrating GenAI tools into audiovisual production, positioning artificial intelligence as a complementary system that enhances - but does not replace - human creativity and professional expertise. The findings contribute to a deeper understanding of the transformative impact of GenAI in the audiovisual domain and provide a structured perspective for researchers and practitioners seeking to integrate these technologies effectively into creative workflows.

15:15 ***Topic Modeling of Student Perspectives on Generative AI for Academic Research Writing***

Lauro Aspiras (Quirino State University, Philippines)

This study explores the use of generative artificial intelligence (AI) in academic research writing by employing a qualitative descriptive methodology integrated with topic modeling. Semi-structured interviews and focus group discussions were conducted with students from various academic disciplines to capture diverse experiences and perceptions of AI-assisted writing. The qualitative data were analyzed using thematic analysis, while the textual corpus underwent preprocessing procedures, including stop-word removal and stemming, prior to topic extraction through Latent Dirichlet Allocation (LDA) using RapidMiner. The results reveal that generative AI is widely perceived as a useful support tool for idea generation, improving coherence, and enhancing the overall quality of academic writing. Nevertheless, participants also identified major concerns related to academic integrity, plagiarism, fairness, algorithmic bias, and excessive reliance on AI technologies. Although generative AI was recognized for its capacity to increase efficiency and productivity, participants emphasized the importance of maintaining ethical awareness and independent critical thinking in its use. The study concludes that educational institutions should implement comprehensive measures, including AI literacy programs, ethics-oriented discussions, and clearly defined guidelines for acceptable AI use. These efforts are necessary to maximize the educational value of generative AI while minimizing its potential risks in academic research writing.

Wednesday, April 22 14:00 - 15:30 (Africa/Cairo)

SESSION-5C: RESEARCH: Large Language Models, RAG & Agentic AI Systems

(6 PAPERS)

Room: ROOM-C

Chair: Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA)

14:00 ***Domain-Specific Legal AI Assistant with Retrieval-Augmented Generation(RAG)***

Ekereuke Udoh (QA Higher Education, United Kingdom (Great Britain))

Large Language Models demonstrate impressive text generation capabilities but suffer from hallucinations and lack jurisdiction-specific grounding when applied to legal domains. This research implements a Retrieval-Augmented Generation system tailored specifically for UK law, combining dense vector retrieval with citation-backed response generation. The system processes legal documents from BAILII and UK legislation, creating a searchable knowledge base through sentence embeddings and FAISS indexing. Evaluation across multiple test queries demonstrates significant improvements over baseline LLM outputs, with hallucination rates reduced from 34.8% to 5.2% and citation accuracy increased to 94.1%. A web-based interface enables users to query the system and upload additional legal documents dynamically. The results validate RAG as an effective approach for building accountable AI systems in highstakes professional domains where accuracy and traceability are essential.

14:15 ***Prompt-as-Heuristic: A Hyperheuristic Framework for Meta-Reasoning in Large Language Models***

Kassem Danach, Sr (Al Maaref University, Lebanon & Chairperson, Lebanon); Samir Haddad and Jinane Sayah (University of Balamand, Lebanon); Khouloud Eledlebi (Abu Dhabi University, United Arab Emirates); Joseph Merhej (Facultyof Sciences II, Lebanon); Chadi Kallab (Lebanese American University, Lebanon)

Prompt engineering has emerged as a dominant paradigm for controlling the reasoning behavior of large language models (LLMs). However, existing approaches rely on static, manually designed prompts that lack adaptability, theoretical grounding, and robustness under task distribution shift. In this paper, we introduce the \emph{Prompt-as-Heuristic} paradigm, in which prompts are formalized as reasoning heuristics and controlled by a higher-level hyperheuristic mechanism. The proposed framework learns to select, adapt, and evolve prompt heuristics online based on task context, solution quality, and reasoning cost. We formulate prompt selection as a meta-reasoning optimization problem and propose a general learning architecture compatible with bandit-based, evolutionary, and reinforcement learning hyperheuristics. Extensive experiments across mathematical reasoning, program synthesis, abstraction tasks, and combinatorial optimization benchmarks demonstrate that adaptive prompt hyperheuristics consistently outperform static prompting strategies while achieving superior reasoning efficiency and generalization.

14:30 ***Federated Large Language Models for Distributed Cyber Attack Classification***

Shahroz Abbas and Ajmery Sultana (Algoma University, Canada)

The growing scale and distribution of modern networks have increased exposure to sophisticated cyber attacks, while centralized machine-learning pipelines raise privacy, compliance, and data-sharing concerns. Although transformer-based Large Language Models (LLMs) can capture contextual relationships in security telemetry, most existing approaches assume centralized training and inference. This paper proposes a Federated Large Language Model (FLLM) framework for cyber attack classification in distributed environments, where participating sites collaboratively train a shared model without exchanging raw traffic data. Network flows are converted into structured textual representations, augmented with MITRE ATT&CK technique identifiers, and partitioned across clients under non-IID conditions to reflect realistic organizational heterogeneity. Each client performs parameter-efficient fine-tuning of a Phi-4 backbone using Low-Rank Adaptation (LoRA), and the server aggregates lightweight adapter updates using FedAdam to improve stability and convergence. To enhance grounding and interpretability, the framework integrates retrieval-augmented generation (RAG) by querying a vector database built from MITRE ATT&CK documentation and appending the top-k relevant passages to the LLM prompt. Experiments show that RAG improves server-side accuracy from 88.33% to 97.72%, indicating that retrieved threat-intelligence context yields more

reliable predictions. Overall, the proposed approach enables privacy-preserving, scalable, and knowledge-grounded cyber attack detection suitable for distributed security deployments.

14:45 ***RAG-Enhanced Healthcare Integration LLM Framework***

Omer Dawood (Prince Sattam Bin Abdulaziz University, Saudi Arabia); Hamada Nayel (Benha University, Egypt & Prince Sattam Bin Abdulaziz University, Saudi Arabia); Wahiba Abakker Mohammed Ismaiel (Taif University, Saudi Arabia); Yousra Elhakeem (SUST, Sudan)

With the revolution of Artificial Intelligence (AI), Large Language Models (LLMs) have demonstrated great capabilities in healthcare transformation through clinical question answering, medical decision support, and patient communication. However, the deployment of LLMs in real-world healthcare environments faces many challenges such as hallucination, lack of transparency, data privacy, and insufficient accountability. The paper proposes to use Retrieval-Augmented Generation (RAG) and developed a theoretical RAG-Enhanced healthcare LLM framework to address these limitations through component integration architecture. The proposed framework integrate contextual retrieval, confidence aware scoring, provenance tagging, and clinical guardrails with mandatory human-in-the-loop oversight. The framework reduces the risks of hallucination and other limitations of LLMs. The empirical assessment of the framework's impact on diagnostic accuracy, clinical decisionmaking, and patient outcomes is future work of this study.

15:00 ***An Agentic Approach to Faithful and Transparent Islamic Question Answering***

Hind Al-owais, Ibrahim Hashem and Ashraf Elnagar (University of Sharjah, United Arab Emirates)

In Islam, religious guidance is sensitive, as a slight deviation from Islamic law is unacceptable to Muslims. This paper proposes an Islamic multi-agent question answering system specifically desgined for Zakat, one of the five main pillars of islam that aims to increase transparency and trust in AI-generated religious QA systems. This system architecture is made up of five specialized agents: an Intent Classification Agent, a Retrieval Agent that retrieves fatwas using dense semantic search, a GPT-4 based Reasoning Agent, a Verification Agent that uses Natural Language Inference (NLI) that checks faithfulness, and an Answer Composition Agent responsible for the final formatted response. We tested the system on a test set of 100 Arabic Zakat questions. The results showed that the system achieved a mean Faithfulness score of 0.79, with 85% of responses accepted. Out of the accepted results, 53% rated Excellent, 32% rated Good, and overall 15% were rated Poor. The NLI Verification Agent evaluated 662 sentences in total, 89.4% were classified as Supported, 9.2% Unsupported and only 1.4% were classified as Contradicted. This demonstrates that the system rarely generates content that conflicts with its retrieved sources. The proposed architecture of our system provides a structural framework for faith-based question answering systems, where automated verification and scholar-in-the-loop oversight work together to make sure that no answer reaches the user without traceable grounding in authoritative fatwa sources.

15:15 ***Multimodal and Vision Language Few Shot Learning: A Survey***

Amna Salim Humaid Alkinde (University of Sharjah, United Arab Emirates)

Recent advancements in few shot learning have shifted from a focus solely on visual recognition to multimodal and vision language contexts, in which extensive pre-trained models are fine-tuned for novel tasks using only a limited number of annotated multimodal instances. Concurrently, the advent of vision language models (VLMs), along with advancements in parameter-efficient tuning and federated optimization, has resulted in a swiftly

evolving yet fragmented research landscape. This survey aims to provide a focused survey of VLM-based few- shot learning across multimodal and vision-language domains. Initially, we delineate the problem framework and compile representative methodologies across diverse sectors, including image processing, video analysis, three-dimensional modeling, medical applications, and agricultural practices. Subsequently, we categorize the existing literature into a systematic taxon- omy characterized by four interrelated dimensions: (i) learning paradigms and parameter efficiency, (ii) semantic conditioning and knowledge integration, (iii) VLM adaptation for recognition and reasoning in few shot contexts, and (iv) frameworks that prioritize federated and privacy preserving approaches. Within each dimension, we scrutinize the fundamental architectural concepts, training methodologies, and empirical trade-offs, en- compassing robustness in the face of distribution shifts, control over hallucinations, as well as capabilities in temporal and three-dimensional comprehension, and continual or multi-task adapt- ability. Building upon this analytical foundation, we integrate overarching design patterns and emphasize the interplay between pre-trained generalization, efficient adaptation, and certified robustness in practical applications. Lastly, we identify existing challenges and propose future research trajectories for scalable, trustworthy VLM based few shot systems that operate in real world multimodal and federated settings.

Wednesday, April 22 14:00 - 15:30 (Africa/Cairo)

SESSION-5D: RESEARCH: Deepfake, Fake News, Spam & Anti-Phishing Detection

(6 PAPERS)

Room: ROOM-D

Chair: Salah A. Aly (Fayoum University, Egypt)

14:00 ***Cross-Dialect Arabic Audio Deepfake Detection via One-Class Learning and Outlier Exposure***

Mariam Khalid AlAli, Ayad Turky and Saad Harous (University of Sharjah, United Arab Emirates)

Recent advances in speech generation helped to create fake voices for beneficial uses, including entertainment and education. Although this technology was originally made for helpful uses, nowadays it is misused for harmful purposes like voice impersonation, fraud, and online attacks. Detecting audio deepfakes in Arabic is difficult because the language has many different dialects, complex pronunciation patterns, and very limited labeled data for training detection models. In this paper, we present a deep learning framework for detecting Arabic audio deepfakes using one-class learning. The model leverages multilingual speech encoders such as XLS-R and WavLM with an OC-Softmax head to identify patterns in authentic Arabic speech. The model has been trained on real speech from the Casablanca dataset, which includes eight dialects and then evaluated on the Arabic Audio Deepfake and ArFake datasets to test its performance. The Leave- One-Dialect-Out (LODO) method tests the model's ability to generalize to new dialects it has not seen before. Our experiments demonstrate that simple one-class learning reaches 87% accuracy but incorrectly labels many fake samples as real. Adding 5 to 10% outlier exposure during the training process greatly enhances fake detection. Our top-performing setup in the last experiment (E6) reaches 78.3% accuracy, 99.8% precision, an equal error rate (EER) of 14.8%, and a ROC-AUC of 0.9262 on the ArFake

dataset. Our findings show that robustness in detecting fake Arabic audio is achieved by training on multiple Arabic dialects and carefully injecting a controlled outlier exposure (OE) during training. This helps in improving the generalization to unseen dialects under LODO evaluation.

14:15 ***AI-Based Detection of Synthetic Video and Audio Media***

Omar Assem, Zahwa K. Kesa and Eman E. Hamd (October University for Modern Sciences and Arts (MSA), Egypt); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

Using generative AI means fake videos and sound can look very real today. Sometimes people hide their true identity through these files. False stories spread faster when tricks go unnoticed. Money vanishes without warning if signs are missed early. One method watches faces using a system trained on Face orensics++. It spots small clues left behind by editing tricks. Another layer checks audio patterns instead. This one learns from Mel-spectrograms in the ASVspoof 2019 dataset. That setup helps flag manipulated voices. Together they form a way to spot fakes across both images and sound. When one type of data is unclear, blending choices at the decision point helps. A adjustable weight lets the system combine inputs more effectively. Tests confirm the sound-based system correctly classifies sounds 86.9 percent of the time. Similarly, footage guidance hits 85.7 percent under similar conditions. Combining both lifts confidence levels - about 89.6 percent when visual clues dominate. What shows up is how mixing different signs from both types works better when one fails. Depending only on one clue tends to fall apart under pressure.

14:30 ***PRNU-CNN Fusion for Video Deepfake Detection***

Waleed Al-Ali and Ahmed Bouridane (University of Sharjah, United Arab Emirates); Kais Belwafi (University of Shrjah, United Arab Emirates)

Deepfake videos have become increasingly realistic which poses serious threats to digital security and media integrity. The current detection methods face significant challenges at generalizing the new manipulation techniques and they often require access to camera reference patterns for forensic analysis. We propose a hybrid deepfake detection framework which combines the analysis of statistical Photo Response Non-Uniformity (PRNU) with the learning of Convolutional Neural Network (CNN). Traditional PRNU methods need a database of known camera fingerprints to work. Our approach is different; we extract directly statistical features from the wavelet noise without needing any reference patterns. This means we can detect deepfakes even when we do not know which camera was used. The proposed system fuses the features of PRNU with the predictions of CNN through a weighted combination where the achieved optimal performance is with 30% PRNU influence and 70% CNN influence. The experiments on the WildDeepfake dataset reveal that our hybrid approach achieves 82% accuracy which outperforms PRNU-only (65%) and CNN-only (77%) methods. We evaluate the framework on binary classification to distinguish authentic videos from deepfakes. The modular design separates feature extraction from classification which enables future extension to source attribution tasks without modifying the feature extraction pipeline.

14:45 ***A Hybrid Transformer Architecture for Phishing Email Detection Using Lexical and Statistical Features***

Omar Abdelrahim (Arab Academy for Science, Technology and Maritime Transport); Radwa Fathalla (Arab Academy for Science, Technology and Maritime Transport, Egypt); Yasser El-Sonbaty (Arab Academy for Science, Technology, and Maritime Transport, Egypt)

Email remains the primary communication tool for global business, making it a prime target for cyberattacks. What began as simple spam has morphed into dangerous phishing campaigns that steal credentials and distribute malware. In fact, spam now accounts for over 50% of all email traffic, creating significant financial risk. Current defenses fail to stop this. Simple filters like blacklists rely on fixed keywords, which attackers bypass

with trivial text changes. While large language models (LLMs) have the intelligence to spot these patterns, they are typically impractical for high-speed mail servers due to their massive computational cost. This paper presents a necessary balance, a model that retains the accuracy of an LLM but operates at the speed required for real-time filtering. This paper proposes a hybrid architecture that combines a Transformer-based text encoder with a dense network for numerical feature analysis. We extract and use seven specific features, including capitalization ratio and hyperlink frequency, to capture structural signs of spam. The proposed model achieves very high accuracy, effectively closing the gap between simple statistical classifiers and complex LLMs. The study demonstrates that this approach performs better than SVM and LSTM baselines and even more complicated models such as BERT and DistilBERT while maintaining high computational efficiency.

15:00 ***Multi-Modal Fake News Classification: A Hybrid Deep Learning Approach***

Mina Eskander (Arab Academy for Science, Technology, and Maritime Transport, Egypt); Mostafa Magdy Karam (Arab Acadmey for Science, Technology, and Maritime Transport, Egypt); Nahla Belal (Arab Academy for Science, Technology, and Maritime Transport & College of Computing and Information Technology, Egypt); Mohamed Khedr (Arab Academy for Science and Technology, Egypt); Ahmed Hamdy Abdallah and Ahmed Karem Farouk (Ahram Canadian university, Egypt)

The rapid growth of social media platforms has accelerated the spread of misleading and fabricated information, commonly referred to as fake news. Such content often combines persuasive textual narratives with visually compelling images, making automated detection increasingly challenging. While prior research has explored textual or visual cues in isolation, real-world misinformation is inherently multimodal, requiring joint reasoning over both language and imagery. In this paper, we investigate fake news detection as a binary classification problem using a multimodal learning framework that integrates textual and visual information. We construct and evaluate three complementary models: a text-only baseline based on a pretrained BERT encoder, an image-only baseline using a convolutional neural network, and a multimodal architecture that fuses representations from both modalities through late fusion. Experiments are conducted on a large-scale dataset of social media posts containing titles, images, and associated metadata. Our results demonstrate that while textual information alone provides a strong baseline, visual cues by themselves are insufficient for reliable fake news detection. In contrast, the proposed multimodal model consistently outperforms both unimodal approaches, achieving higher accuracy and F1-score on held-out validation and test sets. Qualitative analysis further shows that multimodal reasoning helps reduce ambiguous predictions that arise when relying on a single modality. These findings highlight the importance of multimodal learning for robust fake news detection and provide insights into the complementary roles of text and images in identifying misinformation.

15:15 ***Spam Email Detection Using a Custom Multi-Head DistilBERT***

Wadaq Tarek (Arap Open University, Egypt); Eid Amery (Amery, Egypt); Hala Abbas (Helwan University & Arab Open University, Egypt)

The exponential growth of spam emails poses a persis- tent and evolving threat to cybersecurity, organizational productivity, and user privacy. Traditional machine-learning approaches that de- pend on surface-level keyword patterns miss the contextual and se- mantic signals embedded in modern spam. This paper presents a novel custom Multi-Head DistilBERT classifier with three parallel projec- tion heads (256, 128, and 64 dimensions) concatenated into a 448- dimensional fused representation. The model is evaluated on the En- ron corpus (33,650 emails) against TF-IDF+Logistic Regression, BiL- STM, and Text-CNN, achieving 99.13% accuracy and 0.9995 ROC AUC. Under 10-fold cross-validation with full per-fold retraining, it attains 0.9890 ± 0.0021 mean accuracy, the highest across all models. Paired t-tests and DeLong tests confirm statistically significant superi- ority over

all three baselines. The false negative rate of 0.97% under- scores the model's value in cybersecurity applications where missed spam carries operational risk.

Wednesday, April 22 14:00 - 15:30 (Africa/Cairo)

SESSION-5E: RESEARCH: AI for Clinical Decision Support & Healthcare Risk Assessment

(6 PAPERS)

Room: ROOM-E

Chair: Siva Rama Krishna Varma Bayyavarapu (Senior IEEE Member, USA)

14:00 ***Emergency Department Triage Classification: A Comparative Study with Explainable AI for Small-Data Environments***

Ahmad Saad and Taha Houda (Prince Mohammad Bin Fahd University, Saudi Arabia); Saad Houda (King Faisal Specialist Hospital and Research Center, Saudi Arabia)

Emergency Departments (ED) overcrowding causes delayed treatment, high mortality, and nurse burnout. Manual triage suffers from inter-rater variability and cognitive load. This paper evaluates an automated triage framework using the Korean Triage and Acuity Scale (KTAS) to risk-stratify ED patients (Emergency vs. Non-Emergency). A total of 1,267 encounters were analyzed (24 structured features plus chief complaint text) to compare 18 models across traditional ML, deep learning, and NLP paradigms. We enforced strict leakage prevention by performing all preprocessing after a stratified train-test split. Tuned Random Forest achieved the top recall of 88.44% (95% CI [0.8333, 0.9324]), prioritizing patient safety, while SHAP analysis confirmed clinically valid predictors like injury and pain severity. Most importantly, our results demonstrate that high-performance, safety-critical triage models can be derived from modest datasets (n ≈ 1,200$) without relying on the massive repositories (n > 100,000) often cited. This suggests that bespoke, high-recall AI tools are viable for smaller clinics, enabling them to deploy "catered" models that reflect specific local demographics. Our findings illustrate the potential of interpretable ML-augmented triage as a decision-support tool.

14:15 ***Lightweight Carbon-Conscious Machine Learning Framework for Lung Cancer Risk Assessment***

Hifsah Nasir (HITEC University, Pakistan); Wajahat Riaz (Pak-Austria Fachhochschule - Institute of Applied Sciences and Technology, Pakistan); Ihtisham Ali (Pak-Austria Fachhochschule Institute of Applied Sciences and Technology, Pakistan); Khalil Khan (Qassim University, Saudi Arabia); Maqbool Khan (American University of Bahrain, Bahrain)

The application of artificial intelligence to lung cancer risk screening at its early stages has gained growing academic interest but the overwhelming majority of solutions focus on predictive accuracy but not on energy consumption or environmental impact, thus restricting their application in resource-limited healthcare systems. The paper presents a carbon-conscious energy-optimal framework of machine learning risk assessment of lung cancer based on clinical and lifestyle factors. The methodology uses the feature-guided selection with the lightweight stacked ensemble of Logistic

Regression, random forest and XGBoost models, thus, balancing between predictive performance and computational cost. As opposed to energy-demanding deep learning models, the framework presented will compare classification accuracy and energy use, both in training and inference. Empirical studies on a publicly available dataset of risk of lung cancer prove that the ensemble improves the predictive strength of individual models and exhibits dramatically lower computational expenses. A comparison between the energy-accuracy trade-off shows that light-weight ensemble learning is capable of providing clinically relevant screening activity with a significantly reduced carbon footprint. The authors to the best of their knowledge expect the study to be one of the first attempts to simultaneously evaluate predictive efficacy and feasibility of the risk of lung cancer screening in a Green AI paradigm when using lightweight ensemble methods.

14:30 ***AI-Driven Pseudo-3D Mesh Reconstruction from 2D Fetal Ultrasound for Accessible Prenatal Visualization***

Omar Mohamed Waheed and Karim Mohamed Abo Alazm (October University for Modern Sciences and Arts (MSA), Egypt); Mazen A. Ebrahim (MSA University, Egypt & University of Greenwich, United Kingdom (Great Britain)); Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

Most prenatal ultrasound examinations worldwide are still performed using standard two-dimensional devices, forcing physicians to mentally combine thin slices into a three dimensional understanding. Although learning-based 3D reconstruction methods exist, they typically require manual tracking scans, multiple display sequences, or a 3D ground truth that is difficult to obtain in many clinics and public datasets. In this work, we present a low-resource pipeline that converts a 2D fetal ultrasound frame into a pseudo-3D representation and then extracts a surface mesh using Marching Cubes, followed by cleaning and rendering of a lightweight mesh. This approach is designed for ease of interpretation and simplicity of deployment, prioritizing reproducible geometry extraction over black-box generation. We evaluate the use of self-consistency and mesh validity statistics (IoU between input masks and rendered silhouettes, triangle/vertex counts, fundamental triangle quality, and runtime). This supports SDG 3 (Good Health and Well-Being) by improving access to prenatal visualization, and SDG 10 (Reduced Inequalities) by targeting settings where 3D/4D ultrasound hardware is uncommon.

14:45 ***Integrating Artificial Intelligence into Healthcare: Predictive Analytics for Early and Equitable Diagnosis***

Ekereuke Udoh (QA Higher Education, United Kingdom (Great Britain))

Chronic illnesses like cardiovascular disease and diabetes continue to cause worldwide morbidity and death, especially in low-resource and underrepresented groups where diagnosis are delayed. This study uses two well-established epidemiological datasets, the Framingham Heart Study for 10- year cardiovascular risk prediction and NHANES for diabetes risk classification based on HbA1c and glucose laboratory markers, to test the potential of machine learning (ML) models for early and equitable detection. Logistic regression, Random Forest, and XGBoost models were created and tested using ROC-AUC, confusion matrices, calibration curves, and threshold sensitivity analysis after preprocessing and class imbalance. Logistic regression consistently had the best interpretability-performance ratio, producing moderate discrimination (ROC-AUC 0.698 for CHD, 0.725 for diabetes). Class weighting and lower decision criteria increased early detection sensitivity. Calibration analysis showed partial alignment of probability, particularly in lower to mid-risk categories, suggesting practical clinical use after recalibration. Correlation and feature analysis confirmed that physiological risk markers (e.g., age, blood pressure, HbA1c) matched clinical knowledge, whereas demographic data highlighted structural drivers of health. Interpretable machine learning, responsible threshold optimisation, and fairness-aware evaluation can aid preventative health risk

stratification, according to the study. The findings also highlight external validation, subgroup fairness evaluation, and calibration optimisation before clinical deployment. This study presents a practical, transparent, and reproducible approach for introducing machine learning-based risk prediction into public health screening processes to promote early detection and chronic illness equality

15:00 ***Beyond Benchmark Accuracy: Toward Clinically Trustworthy AI for Biomedical Signals,***

Jinan Charafeddine, Dr. Jinan (Engineering School Léonard de Vinci - ESILV, France); Taha Houda (Prince Mohammad Bin Fahd University, Saudi Arabia); Nada Salman (Université Paris-Saclay, France); Jihad Jaam (Liverpool J. Mores University, United Kingdom (Great Britain))

Artificial intelligence (AI) has become a dominantparadigm for biomedical signal analysis, enabling automated processing of electrocardiogram (ECG), electroencephalogram (EEG), electromyogram (EMG), and other physiological signals. Deep learning architectures, particularly convolutional neural networks (CNNs), long short-term memory (LSTM) networks, and hybrid models, have demonstrated strong benchmark performance, with ECG classification accuracies frequently exceeding 95%. However, performance varies across signal modalities, with EEG and EMG applications remaining more challenging due to higher variability and lower signal-to-noise ratios. This paper presents a structured review of AI-based biomedical signal studies published between 2018 and 2025, analyzing architectural evolution, dataset characteristics, performance trends, and barriers to clinical translation. Our analysis highlights a persistent gap between benchmark accuracy and real-world deployment, largely due to limited dataset diversity, insufficient multi-center validation, interpretability constraints, and computational challenges. Emerging directions including federated learning, self-supervised learning, multimodal fusion, and edge AI are discussed as pathways toward clinically trustworthy systems. Achieving clinical readiness requires not only improved pre-dictive performance but also robustness, transparency, and large-scale validation across diverse patient populations.

15:15 ***Harnessing Deep Computational Intelligence for Smart Healthcare Decision-Making***

Dina Darwish, Reham Adel Ali, Nehal Khaled Ahmed and Soha Mohamed Abd Allah (Ahram Canadian University, Egypt)

In the last several years, artificial intelligence has gotten more and more popular, and today it's a part of many people's everyday lives. Smart healthcare is now possible because of the usage of AI in healthcare. There are good and bad things about smart healthcare. This paper provides a comprehensive analysis of historical and current developments in this domain. The major issue is to speak about what smart healthcare is and how it may be useful. Second, the future perspectives of AI in smart healthcare were spoken about in a more detailed way in this paper. Third, smart healthcare divided diverse AI apps into categories depending on the technology they were built on and looked at each one on its own. At last, the conclusion is reached.

Wednesday, April 22 15:30 - 15:45 (Africa/Cairo)

COFEE BREAK & NETWORKING

COFEE BREAK ROOM - NETWORKING

Room: COFFEE_ROOM

Wednesday, April 22 15:45 - 18:00 (Africa/Cairo)

SESSION-6A: RESEARCH: Innovations in AI for Secure Identity Verification, Cost Optimization, and Enhanced Analytics

9 PAPERS

Room: ROOM-A

Chairs: Salah A. Aly (Fayoum University, Egypt), Tejas Pravinbhai Patel (Amazon, USA)

15:45 ***TBIVL: A Unified Identity Verification Layer for Secure Agent-to-Service Context Exchange***

Siva Rama Krishna Varma Bayyavarapu (Senior IEEE Member, USA); Siva Prasad Nandi (Oracle, USA); Srinivasateja Songa (The Home Depot, USA)

Agent-to-service context exchange requires that both the caller and the target service be authentic; otherwise impersonation, context capture, and privilege abuse are possible. Problem: Existing work on agent and tool-use security focuses on prompt injection, output filtering, and access-control policies. Entity authenticity at context-exchange time is less well addressed. Gap: Transport security (e.g., mTLS) does not bind application-layer requests to verified agent or service identities. Approach: We propose the Trust-Bounded Identity Verification Layer (TBIVL), a unified application-layer verification layer. It verifies agent attestations, validates service authenticity, and enforces request-scoped authorization in a single decision layer before context is forwarded. Prototype & evaluation: We implement a pro-to type and evaluate it in a controlled study (five scenarios,1K requests) and in a larger-scale simulation (100 agents, 50services, 10 spoofed, 10K requests), plus a sensitivity sweep over malicious traffic ratio (5-50%). Results: Full TBIVL prevented all evaluated attacks in our experiments; every malicious request was rejected in both the controlled and large-scale runs. Mean latency remained under 0.04 ms per request. The sensitivity sweep showed no successful attacks in the full-TBIVL runs. In the tested configurations, combining agent attestation, service authenticity validation, and scoped authorization prevented the evaluated im-personation and spoofing attacks while adding minimal overhead.

16:00 ***Autonomous Workload Right-Sizing for Multi-Cloud Cost Optimization***

Jay Bharat Mehta (Cleveland State University, Alumni, USA)

The rapid ascent of cloud computing has fundamentally changed the way enterprise IT works, thus allowing companies to quickly and easily spread their workloads over different cloud providers. But the truth is the price of this convenience has been quite high, in fact, it has been quite literally high as well. Cloud spending inefficiencies like underutilization and above-normal provisioning standards together account for the annual loss of billions of dollars in cloud expenditure. Conventional right-sizing tactics are primarily dependent on fixed thresholds and labor-intensive processes, thus being unable to handle the variable and intricate nature of the workloads in multi-cloud setups. This paper puts forward the Autonomous Workload Right-Sizing Engine (AWRE), a framework that learns on its own and uses various cutting-edge technologies such as Reinforcement Learning (RL), Causal Inference Modeling, and a Policy-Aware Optimization Layer (PAOL) to mix the workload configurations among AWS, Azure, and Google Cloud in an

intelligent and autonomous manner. AWRE not only tracks continuously the workload telemetry such as CPU utilization, memory usage, latency, and compliance status but also makes real-time adjustments to the configurations for the purpose of optimizing cost and performance simultaneously. The system utilizes Causal Impact Analysis to distinguish actual performance improvements from random fluctuations and incorporates a Graph Neural Network (GNN)-based Anomaly Suppression Module (ASM) for identifying and precluding unsafe optimization actions. AWRE which is operational in more than 10,000 enterprise workloads, accomplished 35% cut down of operational costs, 22% improvement of performance efficiency, and 15% increase in workload stability while still being completely in line with security and governance policies. Making data-driven intelligence and policy enforcement work together, AWRE revolutionizes cloud optimization in organizations. It creates a new model of autonomous, compliance-aware workload management that takes enterprise cloud operations from reactive provisioning to proactive, intelligent, and automated processes.

16:15 ***Autonomous Data Agents for End-to-End Analytics Pipelines: From Data Discovery to Insight Generation***

Naveen Kolli (Independent Researcher, USA); Santhosh Kumar Veeramalla, Abhishek Palakurthi and Ram Prasad Belde (USA)

End-to-end analytics still takes too long because the work is fragmented: someone hunts for data, someone cleans it, someone builds features, someone trains models, and someone else explains the results. This paper explores autonomous data agents that can carry an analytics pipeline from start to fin-ish-starting at data discovery and ending at insight generation that a non-expert can actually use. We demonstrate the idea using a real-world, publicly available dataset (e.g., a widely used open dataset with tabular records and mixed data quality) to reflect the messy conditions analysts face in practice: missing values, inconsistent categories, and shifting distributions. The proposed agentic pipeline performs dataset profiling, schema and semantic inference, automated data quality checks, repair suggestions, feature construction, model selection, and evaluation under resource constraints. Beyond predictive performance, the system focuses on actionable insight: it produces plain-language narratives, key drivers, uncertainty estimates, and reproducible artifacts (queries, code, and documentation) so results can be audited and re-run. We also discuss safety and governance aspects-hallucination control in summaries, leakage prevention, and human-in-the-loop checkpoints-because analytics outcomes often inform high-stakes decisions. Our findings suggest that autonomous data agents can reduce iteration time and lower the barrier to analytics by turning raw, imperfect data into reliable, traceable insights with minimal manual intervention.

16:30 ***Responsible AI Horizons: A Multi-Dimensional Framework for Verifiable AI Behavioral Assessment with Cryptographic Verification***

Gregory David Spehar (GiDanc AI LLC, USA); Akshay Mittal (University of the Cumberlands, USA)

As artificial intelligence systems assume critical decision-making roles in healthcare, finance, and government, The imperative for verifiable behavioral evaluation has become paramount. Current AI alignment approaches focus primarily on training-time interventions, leaving a significant gap in runtime behavioral verification. This paper presents the AI Assess Tech 4D Morality Framework, a novel system for multi-dimensional behavioral assessment of AI systems across four orthogonal ethi- cal dimensions: Lying (honesty vs. deception), Cheating (fairness vs. unfairness), Stealing (integrity vs. exploitation), and Harm (safety vs. harmful intent). The framework employs a standard- ized 120-question assessment instrument that maps AI responses to four personality archetypes (Psychopath, Well-Adjusted, Mis- guided, Manipulative) using Euclidean distance classification in 4-dimensional space. Key technical innovations include: (1) a tamper-proof cryptographic verification system using SHA-256 hash chains with publicly accessible verification endpoints, (2) blockchain-anchored instrument integrity via Ethereum mainnet ensuring the assessment instrument itself is verifiably unmodified, and (3) anti-gaming mechanisms including dead zone detection and cryptographic answer randomization. Baseline validation across four major LLM providers (OpenAI, Anthropic, Google, xAI) demonstrates measurable behavioral variation at the dimen- sion level, with all

assessment results cryptographically locked and independently verifiable. The framework enables enterprise compliance verification and regulatory adherence, providing what we characterize as "pre-flight checklists for AI"-analogous to aviation safety protocols applied to artificial intelligence deploy- ment.

16:45 ***Architecture Patterns for Microservices in AI Ecosystems***

Shravya Bussari (Hcltech, USA); Naveen Kumar Puppala (Optum, USA); Akshay Mittal (University of the Cumberlands, USA)

The rapid adoption of artificial intelligence (AI) across enterprise systems has introduced new architectural chal- lenges when AI capabilities are deployed using microservices. While traditional microservice architectures were designed for deterministic and stateless workloads, modern AI systems exhibit probabilistic behavior, data-intensive pipelines, continuous model evolution, and heightened governance requirements. This paper examines architecture patterns that enable scalable, resilient, and trustworthy AI ecosystems built on microservices. We iden- tify key challenges unique to AI-driven systems and propose a set of reusable, AI-aware architecture patterns addressing model serving, orchestration, observability, security, and system evolution. A reference architecture is presented to illustrate how these patterns interact in production environments. The paper provides practical architectural guidance for designing and operating production-ready AI microservice systems, particularly in enterprise and regulated settings.

17:00 ***Serverless Architectures for Real-Time AI Model Deployment***

Shravya Bussari (Hcltech, USA); Akshay Mittal (University of the Cumberlands, USA); Akhilesh Bollam (Amazon, USA)

The increasing adoption of artificial intelligence (AI) in latency-sensitive and user-facing applications has in- tensified the need for scalable, cost-efficient, and operationally lightweight deployment architectures. Traditional server-based and container-centric deployments introduce challenges in elas- ticity, operational overhead, and cost efficiency, particularly for bursty real-time inference workloads. Serverless computing offers an alternative paradigm by abstracting infrastructure man- agement while providing event-driven scaling and fine-grained billing. This paper examines the applicability of serverless architectures for real-time AI model deployment. We analyze architectural patterns, latency trade-offs, cold-start behavior, and integration with modern model-serving pipelines. A reference implementation is presented, followed by experimental evaluation on real-time inference workloads. The results demonstrate that serverless platforms, when combined with optimized model pack- aging and invocation strategies, can meet real-time performance requirements while significantly reducing operational complexity.

17:15 ***AI-Enabled Remote Proctoring: Architecture and Implementation for Digital Examinations***

Vishal Shukla (Harvard Business School, USA & LTIMindtree, USA); Surbhi Ghai and Aatam Prakash Sharma (Lovely Professional University, India); Ashwath Suresh Hegde and Lalit Rupani (USA); Pradnya Harshad Desai (Salesforce, USA)

With the increased use of online examinations since the pandemic, academic integrity has come under assault like never before, without bodies in the room to guarantee the proper integrity and reliability of assessment. Therefore, educators turn to new solutions and mechanisms that can ensure reliability and integrity across virtual assessment situations. This paper presents an AI-based online examination assessment system that monitors student activity in front of a camera during an exam, and at the same time, renders warnings to the student and/or assesses integrity breaches after the exam. This system relies on computer vision-based techniques, the YOLO object detection model-to detect complicated behaviors, unfocused

eyes, multiple bodies, and other devices in view, as well as browser actions like minimizing screens and switching windows-during an online examination and assesses browser activities like minimizing windows and toggling tabs. Warnings issued are rendered during the exam in real time, while video footage and pixelated warning frames will be electronically stored for further integrity assessments. By merging multiple approaches to vigilance under one umbrella through an AI-integrated system, a scalable, efficient, and privacy-preserving solution is enabled to accommodate the demand for an appropriate online assessment solution in the modern world since student motivations are not always discernable.

17:30 ***Self-Supervised Cross-Modal Representation Learning for Fine-Grained Image-Text Knowledge Fusion***

Sai Reddy Busi Reddy and Pradeep Kumar Chilukury (USA); Naresh Kumar Methuku (Fidelity Investments, USA); Krishna Kishor Tirupati (USA)

This paper introduces a new self-adaptive loss learning scheme to solve the challenges of the extreme class imbalance problem in long-tailed image classification. Unlike the conventional methods based on static loss functions, the proposed method introduces a dynamically learned weight mechanism through bilevel optimization, which allows the model to adaptively emphasize hard and underrepresented samples during training. The approach is evaluated on a long tailed variant of the CIFAR-100 dataset on which the standard performance of models is significantly decreased by severe distribution skew. Experimental results show a significant improvement with a substantial improvement over traditional cross entropy training, Top-1 accuracy improved from 15.64% to 51.66%, Macro-F1 score improved from 15.30% to 51.73%, and balanced accuracy improved from 15.64% to 51.66%. Furthermore, the proposed approach brings faster and more stable convergence in which training loss is reduced from 3.21 to 2.39 while training accuracy is boosted to 52.76%. These results demonstrate the power of adaptive loss learning to reduce long tailed bias, and provide a scalable and generalizable approach to robust deep learning under imbalanced data regimes.

17:45 ***Virtual Reality Enhanced Tabletop Exercises in Cybersecurity Education: A Comparative Study Using A CISA-Based Ransomware Scenario***

Varshitha Manjunath, Gina Accardi and Kevin Murillo Morales (Lasell University, USA)

The study conducts an evaluation of Virtual reality lab-based cybersecurity tabletop exercises to improve students' workforce readiness. Using Vizard, a comprehensive tool for building 3D environments, we simulated a CISA (Cybersecurity and Infrastructure Security Agency) based ransomware attack scenario. Participants connected to the VR hardware to immerse into the incident, explore indicators of compromise, and evaluate response options within the virtual environment. Our Preliminary findings suggest that immersive VR can increase engagement and improve situational awareness, indicating strong potential for integrating VR labs into cybersecurity education and workforce readiness training.

Wednesday, April 22 15:45 - 18:00 (Africa/Cairo)

SESSION-6B: RESEARCH: Innovative Strategies in AI for Predictive Modeling, Compliance, and System Reliability

9 PAPERS

Room: ROOM-B

Chairs: Venkata Seetarama Raju Dantuluri (Capital One, USA), Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA)

15:45 ***Vector Databases Under Realistic Workloads: Filtered Search, Churn, and Warm-Restart Behavior***

Vinay R Soni and Neha Agrawal (USA); Gajendra Babu Thokala (IEEE Independent Researcher); Amit Kumar Padhy (University of Illinois Urbana-Champaign, USA & Adobe Inc., USA); Chandrashekhar Medicherla (Salesforce Inc, USA); Tejas Pravinbhai Patel (Amazon, USA)

Vector databases are increasingly used as production infrastructure for retrieval-augmented generation and semantic search. Most published benchmarks emphasize unfiltered approximate nearest neighbor (ANN) performance under static datasets, which can materially misrepresent real deployments where queries are filtered by metadata, data is continuously updated, and services must recover from restarts. This paper presents a systems-style evaluation of open-source vector database behavior under three production-motivated workloads: (i) filtered vector search with controlled selectivity, (ii) sustained churn with mixed reads and writes, and (iii) warm restart recovery. We introduce a filter-aware ground-truth oracle that enables reproducible recall@k measurement under metadata predicates by combining bitmap-based candidate restriction with an adaptive FAISS flat oracle. Experiments on an Apple M3 (18 GB) show that filtered workloads can trigger super-linear tail-latency amplification at low selectivity even when recall remains high; churn induces oscillatory memory and tail-latency instability; and warm restart readiness is dominated by index reload and cache warming rather than process startup. These findings motivate workload-aware benchmarking and explicit operator controls for latency-accuracy trade-offs under filters.

16:00 ***Causal Inference Models for Workforce Attrition Prediction using Hybrid Workday-HCM and External Behavioral Signals***

Naveen Kolli (Independent Researcher, USA); Santhosh Kumar Veeramalla, Abhishek Palakurthi and Ram Prasad Belde (USA)

In today's competitive labor market, understanding why individuals leave is more important than ever. In this article, we introduce a hybrid approach that blends typical Workday HCM files with rich, external behavioral signals-such as company-level Glassdoor sentiment of reviews and current Google Trends on "remote work" interest-to uncover the causal causes of workforce turnover. With the publicly available IBM HR Analytics Attrition dataset augmented by monthly Glassdoor sentiment scores, we construct a Structural Causal Model that summarizes relationships between demographic, role-stereotype, and sentiment-based variables. We estimate causal effects via propensity-score weighting and double-machine-learning procedures after controlling for the influence of variables such as manager feedback frequency, sentiment dynamics, and shifting labor-market demand. Our results indicate that a one-point drop in Glassdoor sentiment increases voluntary turnover risk by 12%, controlling for pay and tenure. We also demonstrate that interventions can reduce forecasted attrition by up to 15%. By focusing on why employees leave, rather than who will, our approach allows HR executives to create predictive, fact-based retention policies and measure their consequences downstream.

16:15 ***Artificial Intelligence for Integrated Pain and Delirium Prediction in Cardiac Care: A Joint Modeling Trajectories***

Jafar A. Wreikat (Al-Ahliyya Amman University, Jordan); Dua Weraikat (Rochester Institute of Technology, USA); Manal Al Satari (Al-Ahliyya Amman University, Jordan)

Typically, after open heart surgery, postoperative pain management is essential to recovery. However, most of conventional approaches remain reactive and insufficiently personalized. Developments in artificial intelligence (AI) and machine learning (ML) have the potential to predict pain trajectories and support personalized perioperative care. This literature review evaluates the adoption of AI in pain prediction for cardiac surgery from

three major databases (Scopus, Web of Science, and PubMed). Using a systematic screening process, PRISMA, eligible studies were analyzed to assess model architecture and predict variables for clinical applicability. The findings indicate that advanced AI models reveal strong performance in postoperative pain risk delamination. In addition, several studies reported high predictive accuracy using AI models. However, the literature remains fragmented, characterized by retrospective, single-center designs. The data used in the literature is dynamic perioperative and insufficient. The integration of pain prediction with downstream neurocognitive outcomes such as postoperative delirium is scarce. This review identifies critical methodological and translational gaps and argues for the development of integrated, time-aware, and clinically interpretable AI frameworks capable of supporting real-time perioperative decision-making.

16:30 ***A Failure Taxonomy and Trace-Driven Debugging Toolkit for Reliable Agentic LLM Systems***

Nikita Kothari (Salesforce Inc, USA); Ankush Agarwal (OpenAI Inc, USA); Varun Joshi (Amazon Development Center, USA); Lav Kumar (Salesforce Inc, USA)

Agentic Large Language Model (LLM) systems extend language generation with iterative planning and tool invocation. However, benchmarks often mask a critical reality: these systems are operationally brittle in production. Failures such as planning drift, schema violations, and state inconsistencies recur systematically, distinct from the model's reasoning capacity. This paper presents AgentTrace, a middleware architecture that operationalizes execution-level reliability. We introduce a formal seven-class failure taxonomy (F1-F7) and a trace normalization engine that converts heterogeneous agent events into "StepFrames" for deterministic analysis. Unlike prior work that relies on passive logging, AgentTrace implements an active Mitigation & Auto-Repair Layer that intercepts and corrects failures in real-time without additional LLM inference. Evaluated across a rigorous suite of 300 interactive tasks and a high-complexity subset of WebArena, AgentTrace improves task success rates by 16.5 percentage points (from 52.4% to 68.9%) and reduces execution-critical failures by over 50%. Crucially, we demonstrate that these gains are achieved with <5% latency overhead and without fine-tuning, establishing AgentTrace as a viable production standard for reliable agent orchestration.

16:45 ***Compliance-Aware Generative AI for Life Insurance Document Synthesis: A Reference Architecture***

Harpreet Singh Siddhu (University of Illinois, Urbana Champaign, USA); Karthik Pappu (Dakota State University, USA); Yogesh S Thanvi (Akamai Technologies, USA)

Life insurance operations are document-intensive and regulated, making generative AI (GenAI) attractive for drafting claims decision letters, policy change confirmations, lapse notices, reinstatement decisions, and underwriting communications. However, large language models can hallucinate, omit mandatory disclosures, or use non-compliant phrasing, risks that translate into regulatory, legal, and consumer-harm exposure. This paper proposes a compliance-aware GenAI reference architecture for life insurance document synthesis that separates probabilistic generation from deterministic compliance enforcement through retrieval-augmented grounding, automated validation, and human-in-the-loop review. We evaluate the architecture across four production LLMs (Claude 3.5 Sonnet, Llama 3.1 70B, Amazon Nova Pro, Pixtral Large) on 50 test cases spanning five document types with 186 matched baseline-compliance comparisons. The compliance-aware pipeline achieved 80% reduction in unsupported claims (26.5% baseline to 5.3%) with 100% total approval rate (73% ready-to-send, 27% minor edits) across all models. Llama 3.1 70B achieved the highest approval rate (90.7%), while Claude 3.5 Sonnet demonstrated the most reliable completion rate (100%). Results provide statistically significant

evidence (N=186, 95% CI) that the architecture's model-agnostic design effectively prevents hallucinations while maintaining production-ready quality.

17:00 ***AI in Requirements Engineering, Software Design, and Software Architecture: A Systematic Mapping Study (2020-2026)***

Ritesh H Ruparel (CSG International, USA); Vigneshwarr Venkatesan (USA); Suraj Vangala (Charter Communications, USA); Vinod Bottu (Microsoft, USA); Sai Mohan Dasari (University of the Cumberlands, USA)

AI support for requirements engineering, software design, and software architecture has shifted from niche NLP tooling to widely used large language model (LLM) assistants. Teams now use AI to draft and summarize requirements, propose UML-style design artifacts, and create or review architecture decision records (ADRs). The value is real, but so are risks such as invented constraints, inconsistency across artifacts, missing non- functional requirements (NFRs), and broken traceability after change. This paper reports a systematic mapping study of work from 2020 to early 2026. We describe a reproducible search and screening protocol across IEEE Xplore, ACM Digital Library, Scopus, and Google Scholar (as a supplement). We map 20 primary studies into a task taxonomy and summarize inputs, outputs, datasets, and evaluation methods. Results show heavy use of small public proxies and human rating rubrics, with limited evidence on end-to-end impact, consistency checks, and industrial cost and risk. We provide actionable mitigations and a research agenda focused on benchmarks, traceability metrics, and realistic evaluations.

17:15 ***Student Dropout Prediction in East African Secondary Schools: Performance, Interpretability, and Fairness***

Deepak Pai (USA); Karthik Pappu (Dakota State University, USA); Yogesh S Thanvi (Akamai Technologies, USA)

High school dropout imposes a $45 billion annual societal burden in the United States alone, with severe disparities across demographic groups. Globally, the challenge is even greater: in Sub-Saharan Africa, secondary school completion rates remain below 45%, driven by poverty, distance to school, and household obligations. Despite machine learning advances achieving 85--98% accuracy in early warning systems, critical gaps remain in model interpretability, algorithmic fairness, and cross-institutional generalizability. More critically, nearly all dropout prediction research focuses on North American and European contexts, leaving a significant gap in applicability to resource-constrained settings where dropout rates are highest. This study addresses these gaps through empirical validation using a large-scale secondary school dataset from East Africa (n=62,739 students). We compare Logistic Regression, Random Forest, and Gradient Boosting models, achieving best performance of F1=73.8% and AUROC=0.982. SHAP analysis identifies household size (5+ children) and school transportation barriers as primary risk factors. Comprehensive fairness testing across gender and language groups reveals minimal bias (ABROCA $<$ 0.02), demonstrating that equitable ML deployment is achievable. These findings validate that interpretable, fair dropout prediction systems can support practical educational interventions in resource-constrained contexts.

17:30 ***Reusable, Secure, and Compliance-First CI/CD Pipeline Architectures for Regulated Enterprise Environments***

Sameena Begam Savukath Ali (Southern New Hampshire University, USA); Saikrishna Tarakampet (California's Correctional Healthcare Services, USA); Subhash Tatavarthi (Kasmo, USA)

Organizations in regulated industries like financial services, healthcare, and insurance rely more heavily on Continuous Integration/Continuous Delivery (CI/CD) pipelines to speed up their software delivery, increase operational reliability and overall quality. CI/CD automation allows organizations to run builds, tests and deploy applications with as little manual intervention as possible while providing developers with quick feedback. However, in a regulated industry, simply implementing CI/CD automation isn't enough to meet compliance requirements. Regulated industries are governed through excessive regulation that has many requirements regarding governance, security, auditability, and compliance; these requirements cannot be met through CI/CD-only practices. The regulatory frameworks that have been established by HIPAA, SOX, and PCI-DSS create an extensive list of requirements and expectations governing access to sensitive information, tracking changes made to that information, documenting approvals for changes, and segregating employees to prevent unauthorized access to sensitive information. The current version of conventional CI/CD pipelines attempts only to optimize delivery. Like many CI/CD implementations, they don't currently provide any means to enforce compliance with the above regulations in an orderly and consistent manner or in a way that can be certified to an auditor. To comply with these strict requirements, organizations must involve additional manual reviews, documentation from external sources, and the use of post-deployment audit checks. This creates additional operational burdens and leads to an increased risk of failure to meet audit requirements. In this paper, we present a Reusable Compliance-First CI/CD Pipeline Architecture that is Developed Specifically For Multi-Stack Enterprise Environments. The Methodology Integrates Governance Systems, Security Scanning, Compliance as Code, and Automated Audit Trails Into The DevOps Workflow. Included Below, You Will Find Multiple Real World Case Studies From Insurance, Healthcare, And Financial Services That Demonstrate How Using Standardized Pipeline Templates And Embedded Controls Will Help Reduce Auditing Preparation Effort and Increase Deployment Reliability While Supporting Scalable DevOps Adoption Without Compromising Regulatory Obligations and Delivery Speed.

17:45 ***AdaptiRAG: Adaptive Retrieval-Augmented Generation with Co-Trained Domain-Specific Fine-Tuned LLMs for Question Answering***

Vijayakumar Venganti (Cisco Systems, Inc., USA); Seetaram Rao Rayarao (JP Morgan, USA)

Retrieval-Augmented Generation (RAG) has emerged as a powerful paradigm for grounding large language model (LLM) outputs in external knowledge, yet existing approaches suffer from two compounding limitations: (1) static retrieval policies that incur unnecessary latency for queries answerable from parametric memory, and (2) independently optimized retrieval and generation components that fail to leverage cross-component learning signals. We present AdaptiRAG, a co-trained framework that integrates a lightweight context classifier for adaptive retrieval gating, a hybrid dense-sparse retrieval pipeline with cross-encoder reranking, and a LoRA-adapted domain-specific LLM whose fine-tuning objective incorporates retrieval quality signals. Evaluated on BioASQ and CUAD benchmarks, AdaptiRAG achieves F1 scores of 73.8% and 71.4% respectively, representing improvements of +12.7 and +13.2 percentage points over naive RAG baselines. The adaptive gating mechanism reduces P95 retrieval latency by 43% and GPU memory consumption is reduced by 48% through QLoRA 4-bit quantization compared to full fine-tuning, while throughput improves by 2.1x under concurrent query loads. Ablation studies confirm that each component contributes independently and that co-training provides a statistically significant improvement over modular assembly. Our results demonstrate that adaptive retrieval gating combined with co-trained generation is a practical and reproducible path toward high-accuracy, inference-efficient domain QA systems.

Wednesday, April 22 15:45 - 18:00 (Africa/Cairo)

SESSION-6C: RESEARCH:Advanced Frameworks for Runtime Verification, AI Workloads, and Intelligent System Orchestration

9 PAPERS

Room: ROOM-C

Chairs: Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA), Vasanth rao Jadav (EPAM Systems, USA)

15:45 ***SOP-Driven Knowledge Base Construction for Intelligent Insurance Workflow Orchestration***

Shyalendar Reddy Allala (Global Atlantic Financial Company, USA); Vinil Pasupuleti (IBM, USA); Shrey Tyagi (Salesforce Inc, USA); Srinivasateja Songa (The Home Depot, USA); Siva Rama Krishna Varma Bayyavarapu (Senior IEEE Member, USA)

Insurance new business processing relies heavily on Standard Operating Procedures (SOPs) that encode underwriting rules, routing logic, and compliance constraints in semi-structured or unstructured textual formats. Manual interpretation of these documents introduces latency, inconsistency, and operational risk in workflow orchestration. This paper proposes an auto-mated framework for knowledge base construction from business SOPs to enable intelligent workflow orchestration in insurance new business processing. The proposed architecture integrates document ingestion, natural language processing (NLP)-based knowledge extraction, rule formalization, and workflow execution layers. The system transforms unstructured procedural text into executable rule representations using entity recognition, dependency parsing, and ontology mapping. Extracted rules are formalized into machine-interpretable production rules and integrated into a rule-driven orchestration engine. A proto-type implementation demonstrates improved rule extraction consistency and reduction in manual configuration effort. Experimental evaluation on underwriting SOP datasets shows promising precision and recall in rule extraction tasks, supporting the feasibility of automated SOP-driven workflow automation in regulated insurance environments.

16:00 ***Native MCP Server in Oracle Autonomous AI Database- Enabling High Performance Agentic AI Workloads***

Chaitanya Kulkarni (Oracle America Inc, USA); Chandrashekhar Medicherla (Salesforce Inc, USA); Tejas Pravinbhai Patel (Amazon, USA); Bharadwaj Vulugundam (Oracle, USA); Rajesh Purushothaman (Zscaler Inc, USA); Isan Sahoo (IEEE Senior Member, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA); Rakesh Keshava (IEEE Senior Member, USA)

The emergence of agentic artificial intelligence has renewed interest in standardized mechanisms that allow large language model-driven agents to interact reliably with enterprise data systems. Oracle Autonomous AI Database introduces a native implementation of the Model Context Protocol (MCP), providing an integrated server that exposes database capabilities schema discovery, SQL execution, and Select AI Agent tools directly to AI agents without external middleware. While the MCP standard is gaining adoption as a unifying interface for tool augmented AI systems, the performance behavior and architectural implications of embedding MCP functionality within a cloud database engine have not been examined in the research literature. This study investigates the Native MCP Server in Oracle Autonomous AI Database with a focus on its execution model, security integration, and suitability for high performance agentic workloads. Through controlled experiments involving natural language to SQL generation, multi step tool invocation, concurrent agent sessions, and mixed analytical tasks, we evaluate latency, throughput, and resource utilization under

varying load conditions. Comparative measurements against external MCP servers, REST based interfaces, and direct SQL clients demonstrate that the native implementation substantially reduces end to end interaction costs and benefits from in database execution, ECPU driven autoscaling, and unified governance. The findings indicate that the Native MCP Server provides a robust foundation for scalable, secure, and efficient agentic AI applications in enterprise environments.

16:15 ***Performance Analysis of AI-Powered Analytics in Oracle Autonomous AI Lakehouse- Vector Search and Data Lake Acceleration on Apache Iceberg***

Chaitanya Kulkarni (Oracle America Inc, USA); Chandrashekhar Medicherla (Salesforce Inc, USA); Tejas Pravinbhai Patel (Amazon, USA); Bharadwaj Vulugundam (Oracle, USA); Rajesh Purushothaman (Zscaler Inc, USA); Isan Sahoo (IEEE Senior Member, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA); Rakesh Keshava (IEEE Senior Member, USA)

Oracle Autonomous AI Lakehouse represents a data platform without compromises, combining years of database innovation with Apache Iceberg's openness. This paper evaluates performance characteristics of this next-generation platform built on Oracle AI Database 26ai and Exadata infrastructure processing 48 billion queries hourly. We examine Data Lake Accelerator for analytical queries and AI Vector Search for modern applications across multicloud environments (OCI, AWS, Azure, Google Cloud). Using TPC-DS benchmark at 1TB scale, we demonstrate 2.3x to 9.0x performance improvements through dynamic resource allocation. For AI workloads, vector search on one million 1536-dimensional embeddings achieves sub-100ms latency with 95% precision at K=10. Mixed workload evaluation shows 18% throughput decrease with effective resource isolation. Results validate that unified autonomous lakehouse platforms effectively support heterogeneous workloads-traditional BI, real-time analytics, and AI applications-without sacrificing performance, governance, or operational simplicity.

16:30 ***GraphEvoPR: Reviewing Pull Requests with Path-based Network Portrait divergence and Tests***

Vasim Shaikh (ADP, USA); Vijay Walunj (Teladoc Health, USA); Amarnath Hanumantharaya (USA)

Pull Requests (PRs) are widely used in open-source and industrial software development to submit, review, and integrate code changes. While PR reviews assess both code and relevant tests, existing methodologies rarely quantify whether structural code changes are adequately reflected in corresponding test coverage and failure patterns. This paper investigates whether class-level software metrics-including Network Portrait Divergence (NPD), a topology-based measure of structural change can reveal relationships between code modifications and test failures in PRs. We extend our prior tool, GraphEvo, into GraphEvoPR, which extracts pre-PR and post-PR call graphs, computes 21 class-level metrics, and visualizes execution path changes. We conducted a case study on 140 PRs from six popular Java open-source projects mined via the GitHub API, selected for diversity in domain and size. For each PR, we measured classlevel metrics before and after changes, identified corresponding test failures using Defects4J, and examined correlations. Our results show that increases in NPD and related metrics (AMC, LOC, RFC, LCOM) frequently align with higher counts of failing tests, suggesting that structural changes have measurable effects on test coverage and adequacy. These findings highlight the potential for integrating NPD-based analysis into PR review workflows to better assess the impact of structural changes on test outcomes.

16:45 ***Agentic AI-Driven Workflow Orchestration in Loan Trading Platforms: A Microservices and Hybrid Cloud Architecture Perspective***

Vinay R Soni and Girish A Gajwani (USA)

Loan trading platforms operate as long-running, failure-prone distributed systems where workflow orchestration directly impacts correctness, tail latency, and operational cost. Conventional orchestrators implement static state machines with uniform retry logic that degrades under partial failures, dependency instability, and data unavailability. This paper evaluates a deterministic, policy-driven agentic orchestration layer that classifies failures and selects recovery actions (retry, defer, compensate, escalate) within explicit governance constraints. Using a reproducible microservices test harness with failure injection, we compare static and agentic orchestration across steady-state operation, latency spikes, event flakiness, service outages, and reference-data unavailability. Results show that agentic orchestration preserves steady-state completion while substantially reducing tail latency under severe disruptions. We further demonstrate policy sensitivity: institutions can tune completion-versus-predictability trade-offs through declarative policies without modifying orchestration code. Index Terms-Agentic AI, workflow orchestration, loan trading platforms, microservices, CQRS, resiliency, reference data

17:00 ***Runtime Verification of Cyber-Physical Systems Using Weighted Deep Sequence Models***

Naresh Kumar Methuku (Fidelity Investments, USA); Krishna Kishor Tirupati, Sai Reddy Busi Reddy and Pradeep Kumar Chilukury (USA)

Cyber-physical systems (CPS) are usually monitored by reactive schemes for anomaly detection which detect unsafe behavior only after it has occurred, making them unsuitable for a safety-critical industrial environment. This work introduces a deep sequence learning approach for runtime verification in CPS which is based on a weighted Long Short-Term Memory (LSTM) model trained on multivariate sensor-actuator streams of the SWaT testbed. The proposed approach is unique in that it brings together chronologically-consistent modelling in time, imbalance-aware training, and operational evaluation in addition to traditional point-wise classification. Using a 60-step temporal windowing and by the use of threshold-tuned inference, the model obtained a fairly good point-wise detection performance on the held-out SWaT test set with a 0.9743 accuracy, 0.9879 precision, 0.8603 recall, 0.9197 F1-score, 0.9538 ROC-AUC, and 0.9215 PR-AUC. Operationally, it detected 5 attack events with an event-level F1-score of 0.4000, a mean lead time of 19 timesteps, and only 389 false positives, that proves the practical value of using sequence-aware CPS monitoring for reliable runtime assurance.

17:15 ***DynaGNN-3D: Dynamic Graph Neural Networks with Velocity-Aware Edge Reconfiguration for Real-Time LiDAR-Based 3D Object Detection in Autonomous Vehicles***

Vijayakumar Venganti (Cisco Systems, Inc., USA); Seetaram Rao Rayarao (JP Morgan, USA)

Graph Neural Networks (GNNs) for LiDAR-based 3D object detection have achieved strong performance on autonomous vehicle benchmarks, yet existing approaches construct static k-nearest-neighbor (k-NN) graphs that ignore object motion dynamics between frames. This static edge assignment limits both detection accuracy for fast-moving objects and inference efficiency, since background point clusters receive the same computational budget as foreground objects. We present DynaGNN-3D, a novel architecture featuring velocity-aware dynamic edge reconfiguration: graph edges are re-drawn per inference frame using per-point velocity estimates, concentrating connectivity within moving object clusters while sparsifying background regions. DynaGNN-3D incorporates a dedicated velocity prediction head whose output feeds back into edge construction, enabling temporally coherent 3D bounding box estimation. We also propose an anchor-free center-based detection head compatible with dynamic graph topology. Experiments on nuScenes and Waymo Open Dataset benchmarks demonstrate that DynaGNN-3D achieves 58.5% mAP and NDS of 63.2 on nuScenes validation, surpassing CenterPoint by +2.5 mAP and +3.6 NDS, while reducing inference latency by 26% (38 ms -> 28 ms) on an NVIDIA A4000 GPU. INT8 TensorRT deployment achieves 90 FPS with only -0.8% mAP degradation, demonstrating practical edge deployability.

Ablation studies confirm that dynamic edge reconfiguration and the velocity head each contribute independently to detection quality. Code and pretrained models will be released publicly.

17:30 ***Trust-Aware Runtime Verification for Autonomous LLM Agent Systems***

Tejas Pravinbhai Patel (Amazon, USA)

Autonomous large language model (LLM) agent systems are increasingly used for complex, long-horizon tasks involving planning, tool invocation, memory updates, and inter-agent collaboration. While these systems significantly expand AI capabilities, they also introduce critical reliability and trust challenges due to stochastic generation, semantic drift, and error propagation during execution. Existing approaches to LLM reliability primarily rely on offline evaluation, prompt engineering, or post-hoc filtering, which are insufficient for detecting failures that emerge dynamically at runtime. This paper proposes Trust-Aware Runtime Verification (TARV), a model-agnostic framework designed to continuously monitor and validate autonomous LLM agent behavior during execution. TARV operates as a lightweight middleware layer that intercepts observable agent actions, memory transitions, and inter-agent communication through runtime instrumentation. Specialized verification agents compute multi-dimensional trust signals capturing action validity, constraint adherence, memory consistency, and communication coherence, which are aggregated into continuous trust scores for real-time monitoring. We implement TARV in a multi-agent LLM execution environment and evaluate its effectiveness across tool-oriented, collaborative, and long-horizon task scenarios. Experimental results demonstrate that TARV detects runtime trust violations with high precision while maintaining a low false positive rate and introducing minimal execution overhead on a single NVIDIA A4000 GPU. Furthermore, trust-aware monitoring significantly reduces cascading failures and improves task success rates, particularly in collaborative and long-horizon workflows. These findings indicate that runtime trust verification provides a practical and effective mechanism for improving the reliability and robustness of autonomous LLM agent systems and represents an important step toward safe deployment of agentic AI in real-world environments.

17:45 ***MambaMon: Linear-Time State Space Models for Large-Scale Distributed System Monitoring***

Tejas Pravinbhai Patel (Amazon, USA); Chandrashekhar Medicherla (Salesforce Inc, USA); Arun Kumar Elengovan (Okta Inc. and IEEE Senior Member, USA); Chaitanya Kulkarni (Oracle America Inc, USA)

Real-time anomaly detection in distributed systems remains computationally prohibitive for Transformer-based approaches, requiring quadratic complexity $O(n^2)$ for processing long log sequences. We present MambaMon, a novel architecture leveraging State Space Models (SSMs) with selective structured state spaces to achieve linear-time O(n) complexity while maintaining superior detection accuracy. Our approach addresses three critical challenges: (1) processing ultra-long sequences (8K+ tokens) that capture system-wide temporal dependencies, (2) real-time inference with sub-100ms latency for production deployment, and (3) resource-efficient training suitable for edge environments. Through extensive evaluation on three large-scale datasets-HDFS (11M+ events), BGL supercomputer logs (4.7M events), and Amazon Retail production logs (2.3M events)-MambaMon achieves 93.7% average F1-score, outperforming state-of-the-art methods by 3.8% while delivering 2.7× faster inference and 4.2× faster training. Deployment results from Amazon's production environment show 94.2% anomaly detection with 67ms P99 latency, validating MambaMon's practical viability for mission-critical systems monitoring at scale.

Wednesday, April 22 15:45 - 18:00 (Africa/Cairo)

SESSION-6D: RESEARCH: AI-Enhanced Governance, Data Reliability, and Intelligent Systems for Modern Applications

9 PAPERS

Room: ROOM-D

Chairs: Varshitha Manjunath (Lasell University, USA), Karthik Pappu (Dakota State University, USA)

15:45 ***Autonomous Policy Enforcement Using AI Agents in Digital Governance Systems***

Tejas Pravinbhai Patel (Amazon, USA); Chaitanya Kulkarni (Oracle America Inc, USA)

Digital governance systems face critical challenges in enforcing policies consistently across rapidly expanding e-government services. Current approaches rely on manual enforcement or passive monitoring, leading to delayed responses, inconsistent application, and unsustainable workload as service portfolios grow. This paper presents an autonomous AI-agent-based architecture that interprets formalized government policies, executes enforcement actions in real-time, and escalates high-risk decisions to human authorities. Our system incorporates a policy interpretation engine, risk-aware escalation logic, and comprehensive audit trails for regulatory compliance. Through controlled experiments simulating a national e-government platform with 50,000 daily transactions, we demonstrate a 58% reduction in manual enforcement workload, 3.2× faster violation response time, and a false enforcement rate of 3.4%. The system maintains explainability through decision trace generation and supports audit-ready governance. This work advances digital governance AI from passive monitoring to active, autonomous policy enforcement while preserving human oversight for critical decisions.

16:00 ***Feature Platforms as Reliability Systems: A Taxonomy, Control Matrix, and SLI/SLO Playbook***

Vinay R Soni and Neha Agrawal (USA); Siva Rama Krishna Varma Bayyavarapu (Senior IEEE Member, USA); Sandeep Shivam (Tavant, USA)

Feature platforms fail in ways that standard service monitoring does not capture. A platform can be available while models degrade due to stale values, online offline mismatch, join drops, or semantic drift. This paper treats the feature platform as a reliability system with explicit objectives for correctness, freshness, and availability. We contribute (1) a compact failure taxonomy, (2) a set of reliability gates that turn contracts and measurements into release decisions, and (3) a starter SLI and SLO playbook with an error budget policy. We also report empirical results from a local benchmark that injects common faults in a Feast plus Redis stack and measures detection latency and sampling overhead.

16:15 ***Risk-Aware Human-AI Collaboration Framework for High-Stakes Decision Systems***

Tejas Pravinbhai Patel (Amazon, USA); Madhushree Kumari (IEEE Senior Member, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA)

High-stakes decision systems increasingly incorporate AI, yet existing approaches either over-automate decisions or rely on static escalation rules, leading to accountability gaps. We present a risk-aware collaboration framework that dynamically routes decisions to AI-only, human-only, or hybrid modes based on predicted risk combining uncertainty and impact. Evaluated across three domains (loan approval, medical triage, infrastructure

maintenance), the framework achieves 34--47% reduction in critical errors and 42--58% reduction in unnecessary escalations versus baselines, with severity-weighted F1 scores of 0.76--0.89. The system provides complete audit trails and structured explanations for accountability.

16:30 ***Data-Grounded LLM Framework for Intelligent Travel Planning***

Hayam Reda Seireg (El Shorouk Academy, Egypt); Amr Omar (Elshorouk Academy, Egypt); Fares Ashraf Eraky, Omar Mahmoud AbdelQader, Youssef Atef Mousa, Fabio Ibrahim Sedik and Tasneem Abdelrahman (El Shorouk Academy, Egypt)

Tourist itinerary planning requires balancing personalization, geographic coherence, and feasibility, yet conventional recommender systems address these demands unevenly. Large language models (LLMs) support more context-sensitive planning, but without grounding in verifiable data they may produce implausible or hallucinated itineraries. This study presents an agentic LLM-based tourism planning framework that couples generative reasoning with structured grounding from locally sourced Egyptian business datasets. Itinerary construction is staged: candidate attractions and restaurants are first retrieved using spatial and semantic criteria, then the LLM selects among candidates, provides brief rationales, and synthesizes the final itinerary. This hybrid approach aims to improve realism and contextual fidelity while limiting hallucination. The framework is evaluated on five open-source LLMs using metrics for constraint satisfaction, preference alignment, grounding accuracy, and spatial efficiency. Hybrid grounding improves reliability overall; Mistral attains the highest aggregate score (0.669), whereas smaller models such as Phi-2 are more spatially efficient but weaker in alignment and grounding. The results support integrating LLM reasoning with structured data and clarify trade-offs among model scale, computational cost, and planning quality, with implications for intelligent tourism system design.

16:45 ***Performance-Driven Database Distribution: Evaluating Replication, Partitioning, and Sharding for OLAP Using an Adaptive Cross-Platform Framework***

Siva Prasad Nandi (Oracle, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA); Viswanathan Ranganathan (Netflix, USA); Vivek Venkatesan (Vanguard, USA); Vijayakumar Venganti (Cisco Systems, Inc., USA)

The proliferation of large-scale analytical workloads across enterprise database environments demands rigorous, platform-aware evaluation of data distribution strategies. This paper presents a cross-platform comparative study of three foundational distribution paradigms replication, partitioning, and sharding as implemented across three widely deployed relational database platforms: Oracle Database, PostgreSQL, and SAP Sybase ASE / IQ. We systematically analyze each strategy across eight dimensions including query throughput, storage overhead, fault tolerance, cross-node join complexity, and operational burden under OLAP workload conditions. Building on this analysis, we propose the Adaptive Distribution Selection Algorithm (ADSA), a multi-criteria decision framework that dynamically recommends the optimal distribution strategy or hybrid combination thereof based on five measurable workload parameters: dataset volume, read-to-write ratio, availability SLA, temporal query locality, and key selectivity. Empirical evidence from production deployments managing billions of daily analytical transactions across Fortune 500 environments validates the framework's applicability across platforms. Results demonstrate that ADSA-guided configurations yield up to 47% improvement in analytical query latency and a 32% reduction in infrastructure cost compared to static single-strategy deployments, with platform-specific implementation nuances materially influencing realized gains.

17:00 ***Evaluating Vitalsource Bookshelf As A Digital Library Platform In A Philippine State University***

Julius Cesar O. Mamaril (Pangasinan State University, Philippines)

This cross-sectional study implementing descriptive-correlational research design evaluated the level of satisfaction of the end users of Vital Source Bookshelf in the nine campuses of Pangasinan State University (PSU) during the second semester of school year 2020-2021. A total of 128 students, 128 instructors, and 128 non-academic staff, or an overall total of 384 users were randomly selected using Cochran's formula out of 180,673 total students, teachers and non-academic staff population of PSU who were issued with Vital Source Bookshelf user credentials during the said academic term. The ISO 25010:2011 software quality standards was utilized as criteria to gauge the Functional Suitability, Performance Efficiency, Compatibility, Usability, Reliability, Security, Maintainability, and Portability features of the Vital Source Bookshelf app through a Google Form survey questionnaire with a 5-point Likert scale from Poorly Satisfied to Extremely Satisfied. The researchers hypothesized that there is no significant difference between students, teachers, non-academic staff in their level of satisfaction in using Vital Source Bookshelf and in consideration of their demographic profile (gender, device, and connectivity). Data collected were analyzed using weighted mean as statistical tool while the difference in satisfaction between user groups were investigated using ANOVA. Results produced an overall weighted mean of 4.01 with a descriptive meaning of Very Satisfied indicating that both students and teachers are very satisfied in the implementation of the Vital Source Bookshelf. Meanwhile, one-way ANOVA resulted to an alpha (F) of 0.40667 and a higher p-value of 0.66615 which is indicative that there is no significant difference in the level of satisfaction of students, teachers, and non-academic staff in using Vital Source Bookshelf. Finally, in terms of level of satisfaction derived from each user group in contrast to their demographic profile, a three-way ANOVA resulted to an alpha (F) 0.629 with a higher p-value of 0.642 which is also greater than the 0.05 threshold, this is likewise indicative that there is no significant difference in the level of satisfaction of teachers, students, and non-academic staff in terms of their gender, device, and connectivity.

17:15 ***LLM-Augmented Mutation Testing: Intelligent Generation and Triage of Mutants***

Sanjay Bajaj (Owens-minor, USA); Venkata Seetarama Raju Dantuluri (Capital One, USA); Devendra Rajput (Accenture, USA); Anil Kolhe (Www.yodlee.com, USA)

Mutation testing is a powerful yet computationally expensive software quality technique that assesses test suite adequacy by injecting systematic faults - called mutants - into a program's source code. Despite its theoretical appeal, widespread industrial adoption is hindered by two fundamental barriers: (1) the combinatorial explosion of generated mutants, and (2) the high proportion of semantically equivalent or trivial mutants that waste execution cycles without providing diagnostic value. In this paper we present MutLLM, a novel framework that leverages Large Language Models (LLMs) to address both barriers simultaneously. MutLLM employs an LLM in three complementary roles: (i) semantic-context-aware mutation point selection, (ii) automated equivalence and trivial mutant detection, and (iii) natural-language test improvement recommendations.We evaluated MutLLM on five open-source Java projects totaling over 280,000 lines of code. Our results show an average mutation score improvement of 17.4 percentage points over traditional random mutation testing, while reducing the mutant set executed by 41% through LLM-guided triage. These results demonstrate that LLMs can fundamentally transform the scalability and utility of mutation testing in real-world software development pipelines.

17:30 ***AI-Driven Big Data Analytics and Data Integration on Cloud Platforms***

Balakrishna Pothineni and Prema K Veerapaneni (JPMorgan Chase, USA); Ravi Kiran Kodali (Cognizant Technology Solutions, USA); Bikesh Kumar (Apple Inc., USA); Durgaraman Maruthavanan (TCS, USA); Mayilsamy Palanigounder (NTT Data, USA); Sumit Saha (East West Bank, USA); Kabilan Kannan (AMD, USA)

AI-enabled big data analytics has emerged as a critical capability for modern enterprises driven by the rapid growth of data from IoT systems, digital platforms, and cloud-native applications. Traditional analytics architectures struggle to handle the scale, velocity, and heterogeneity of such data. This paper examines how contemporary cloud platforms enable scalable and efficient AI-driven analytics through elastic compute, distributed storage, and integrated machine learning services. We present a layered architectural perspective covering data ingestion, storage, processing, and model deployment, and analyze performance characteristics across key workload types including stream processing, model training, and batch inference. We further introduce a reference implementation that demonstrates the end-to-end deployment of these architectural patterns on AWS, detailing configuration choices, infrastructure code, and operational trade-offs. Drawing on real-world deployments and benchmark studies, the paper highlights practical trade-offs in cost, scalability, and latency. It further discusses critical challenges such as data governance, vendor lock-in, performance variability, and cost optimization. The findings provide actionable insights and architectural guidance for designing robust, cloud-native analytics systems that support large-scale, intelligent data processing.

17:45 ***Evaluating Demand-Side Flexibility in District Heating Substations via Offline Reinforcement Learning***

Gideon Mbiydzenyuy (University of Borås, Sweden)

District heating utilities increasingly struggle with high peak loads that raise production costs and hinder the integration of low-carbon heat sources. In many district-heated buildings, simple time-of-day strategies such as night setback are still layered on top of outdoor-temperature-compensated control; these rules are easy to deploy at scale but cannot adapt to building-specific dynamics or day-to-day variability. This work investigates whether a simple offline reinforcement learning approach can learn flexibility-activation policies directly from historical meter data at individual substations. We train fitted Q-iteration with a linear function approximator on hourly heat demand and exogenous features, using a one-step Ridge surrogate to predict next-hour demand under binary comfort actions. Historical actions are reconstructed as proxy actions from time-of-day labels that approximate a night-setback window. The learned policy is compared against two rule-based controllers (night setback and peak-hour reduction), behavioral cloning, and trivial always-normal and always-reduced policies, with all policies evaluated offline by rolling them out through the surrogate over a held-out winter test period. On a cohort of five real district heating substations in a Nordic climate, selected for data quality and surrogate fidelity, the offline RL policy achieves a median peak reduction of 18.6% (range 12.3%-28.7%) relative to the reconstructed baseline. It matches or outperforms the night-setback rule on four substations and the peak-hour rule on two, showing that lightweight offline RL can compete with common rule-based strategies using only standard metering data and a single global configuration.

Wednesday, April 22 15:45 - 18:00 (Africa/Cairo)

SESSION-6E: RESEARCH: AI Innovations in Digital Governance, Data Management, and Decision Support Systems

9 PAPERS

Room: ROOM-E

Chairs: Rohit Nimmala (Bank of America, USA), Vivek Venkatesan (Vanguard, USA)

15:45 ***Knowledge Graph-Enhanced LLMs for Unified Climate-Aware Risk Management in Banking: The Cognitive Bank Architecture***

Rohit Nimmala (Bank of America, USA); Gajendra Babu Thokala (IEEE Independent Researcher); Jagrut Nimmala (Lowes Companies, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA)

Climate risk propagates across credit, market, and operational risk domains through interconnected transmission channels. Existing banking risk systems treat these domains in regulatory silos, as identified by the Basel Committee on Banking Supervision (BCBS) in d517 (2021) and d532 (2022). No existing framework unifies climate risk assessment across all three risk types using knowledge graphs (KGs) and large language models (LLMs). This paper presents the Cognitive Bank, a new architecture with three integrated layers: (1) a climate-financial KG built on the Financial Industry Business Ontology (FIBO) (Bennett, 2013); (2) a knowledge-augmented LLM reasoning engine employing Graph-Constrained Reasoning (GCR) (Luo et al., 2025, ICML) to reduce hallucinations; and (3) federated learning via FedAvg (McMahan et al., 2017) with differential privacy (Abadi et al., 2016). The architecture sets design targets of 92%+ balanced accuracy on climate-adjusted credit risk, informed by Mitra et al. (2024), and 20%+ improvement on cross-institutional anti-money laundering detection, informed by Suzumura et al. (2019), along with full Task Force on Climate-related Financial Disclosures (TCFD)-aligned reporting. To our knowledge, this is the first architecture proposal that brings together knowledge graphs, LLMs, federated learning, climate finance, and banking risk management.

16:00 ***Deep Learning NER Pipeline for Automated Basel III / IFRS 9 Risk Parameter Extraction from Climate Narratives***

Rohit Nimmala (Bank of America, USA); Jagrut Nimmala (Lowes Companies, USA); Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA)

Banks must manually extract credit risk parameters, including probability of default (PD), loss given default (LGD), exposure at default (EAD), and expected credit loss (ECL), from qualitative climate scenario narratives published by central banks and TCFD-reporting firms. Existing financial NER systems target generic entity types or XBRL tags, and climate NLP methods classify disclosures at the document level but do not extract quantitative risk parameters. We introduce ClimRiskNER, a unique end-to-end pipeline that combines dual-domain adaptive pre-training (financial and climate corpora) with constrained seq2seq decoding to extract structured Basel III/IFRS 9 parameters from unstructured climate text. The pipeline uses FinBERT-Climate, a FinBERT model further pre-trained on over 2 million climate paragraphs, for token-level NER, followed by a T5-based seq2seq module with prefix-trie constrained beam search enforcing output conformity to a Basel III/IFRS 9 ontology. On a curated dataset of ECB, BoE, and NGFS climate stress test reports and TCFD disclosures, the pipeline achieves entity-level F1 of 0.89 for PD, 0.85 for LGD, and 0.82 for EAD,

outperforming the next-best baseline (GPT-4 few-shot) by 9 F1 points. The system reduces extraction time by over 99% versus manual review (0.3 seconds per page versus 720 seconds) and provides SHAP-based token-level provenance for regulatory audit trails.

16:15 ***Safe and Policy-Compliant Multi-Agent Orchestration for Enterprise AI***

Vinil Pasupuleti (IBM, USA); Shyalendar Reddy Allala (Global Atlantic Financial Company, USA); Siva Rama Krishna Varma Bayyavarapu (Senior IEEE Member, USA); Shrey Tyagi (Salesforce Inc, USA)

Enterprise AI systems increasingly deploy multiple intelligent agents across mission-critical workflows that must satisfy hard policy constraints, bounded risk exposure, and comprehensive auditability (SOX, HIPAA, GDPR). Existing coordination methods-cooperative MARL, consensus proto-cols, and centralized planners-optimize expected reward while treating constraints implicitly. This paper introduces CAMCO(Constraint-Aware Multi-Agent Cognitive Orchestration), a run-time coordination layer that models multi-agent decision-making as a constrained optimization problem. CAMCO integrates three mechanisms: (i) a constraint projection engine enforcing policy-feasible actions via convex projection, (ii) adaptive risk-weighted Lagrangian utility shaping, and (iii) an iterative negotiation protocol with provably bounded convergence. Unlike training-time constrained RL, CAMCO operates as deployment-time middleware compatible with any agent architecture, with policy predicates designed for direct integration with production engines such as OPA. Evaluation across three enterprise scenarios-including comparison against a constrained Lagrangian MARL baseline-demonstrates zero policy violations, risk exposure be-low threshold (mean ratio 0.71), 92-97% utility retention, and mean convergence in 2.4 iterations.

16:30 ***Real-Time Colon Cancer Detection Using Deep Learning: A Comparative Analysis of EfficientNet Architectures***

Sumaiya Thaseen (DMU, United Arab Emirates); Vanitha M (VIT vellore, India); Geetha Mohan (Future Education University College, United Arab Emirates); Benita Kalistus Jimci Veronica (Curtin University Dubai, United Arab Emirates)

Colorectal cancer (CRC) is one of the most common and deadly cancers worldwide, with early detection being crucial for improving survival rates. This paper focuses on developing an advanced system for the automated real-time detection of colon cancer using deep learning techniques, specifically convolutional neural networks (CNNs). The study conducts an extensive comparison of three EfficientNet architectures (B0, B3, and B4) using a dataset of 20,038 augmented images from LC25000, categorized as colon adenocarcinoma and benign tissue. The system uses sophisticated data augmentation methods including rotation, flipping, scaling, and translation to improve model robustness. The experimental outcomes demonstrate that EfficientNetB3 outperforms with 99.87% accuracy, 1.00 precision, recall, and F1-score across all evaluation metrics, significantly outperforming both B0 (70% accuracy) and B4 (73% accuracy) variants. The research model exhibits outstanding exceptional generalization capabilities and demonstrates minimal misclassifications on the test data. Develop a web application that allows users to quickly interact with the model and check diagnostic results. This study shows that EfficientNetB3's optimized architecture offers a best combination of accuracy and computational efficiency for clinical deployment in limited resource-constrained.

16:45 ***TASH-Net: A Novel Deep CNN for Brain Tumor Classification from MRI Images***

Hussein Alkatout and Taha Houda (Prince Mohammad Bin Fahd University, Saudi Arabia); Muhammad Attique Khan (Prince Mohammad Bin Fahd University, Al-Khobar, KSA, Saudi Arabia); Jinan Charafeddine, Dr. Jinan (Engineering School Léonard de Vinci - ESILV, France); Jihad Jaam (Liverpool J. Mores University, United Kingdom (Great Britain))

The process of identifying brain tumors from MRI scans proves difficult because of the extensive data quantity and the intricate nature of tumors. The research introduces TASH as a hybrid deep convolutional neural network which combines Vision Transformer elements with MBConv blocks and fractal design structure in a 357-layer system that uses 6.7M parameters. TASH is compared with state-of-the-art pre-trained models including InceptionV3, ConvNeXt-Base, MNASNet, Vision Transformer, and Swin Transformer. All models were fine-tuned and evaluated on two public MRI datasets. The TASH model achieved 97% accuracy on Dataset 1 but InceptionV3 produced the best results (99%) when working with Dataset 2. The proposed hybrid architecture shows better classification results because it uses its parameters effectively which makes it appropriate for using MRI images to identify brain tumors

17:00 ***From AI Drag to AI Boost: A Preceptorship Framework for Developing Software Engineering Expertise in the Agentic Era***

Satyanarayana Gudimetla (Nike India Technology Private Ltd, India); Suresh Gangula (Nike, Inc, USA); Chandrakanth Challa (Jawaharlal Nehru Technological University Hyderabad, India); R Adinarayana (Andhra University, India); Sudhakar Vunnam (India)

Agentic coding assistants confer a pronounced productivity advantage upon experienced engineers who possess the judgment to direct and evaluate AI-generated output-an effect we term the AI Boost-while imposing an epistemic burden on early-in-career (EiC) developers who lack the cognitive schemata to verify or learn from AI artifacts-the AI Drag. This seniority-biased technological change, corroborated by a 16.3% decline in junior-to-senior job posting ratios, threatens intergenerational knowledge transfer. We introduce a Preceptorship Framework instantiated through the Preceptor-Learner-Agent (PLA) triad-an organizational design pattern embedding structured mentorship, AI-mediated Socratic pedagogy, and competency tracking within production workflows. A 14-month quasi-experimental evaluation across a 200-engineer organization provides preliminary evidence that structured preceptorship is associated with a 41% reduction in EiC time-to-competency relative to pre-AI baselines, while senior throughput multipliers of 2.7x are sustained among active preceptors. Unstructured AI adoption without preceptorship is associated with a 25% increase in time-to-competency, consistent with the 'Fragile Expert' phenomenon (77% failure rate on unassisted tasks). We propose a five-level Preceptorship Maturity Model for epistemic sustainability in AI-augmented engineering organizations. As an exploratory study conducted within a single organization using a quasi-experimental design, the findings establish feasibility and generate hypotheses for future controlled replications; the effect magnitudes reported here should be interpreted as preliminary estimates pending confirmation through randomized trials.

17:15 ***Informative Feature-Based Data Imputation: Exploring Smart Video Surveillance as a Case Study***

Mahmoud M Eid (Egyptian Chinese University (ECU), Egypt); Kamal ElDahshan (Egypt); Abdelatif H. Aboualy and Hassan Mistareehi (Murray State University, USA)

Missing data posed a significant challenge to data analysis, especially in video surveillance. Missing data occurs when certain observations or variables are absent or incomplete within a dataset. Missing data can compromise the accuracy and reliability of analyses, leading to biased conclusions and reduced statistical power. Several factors contributed to missing data. In video surveillance specifically, missing data resulted from occlusions, varying camera viewpoints, or technical limitations, all of which produced incomplete information about observed scenes. This study addressed the problem

of missing data, both generally and in video surveillance, by introducing an optimization-oriented approach to data imputation. Current imputation methods often treat all features equally, neglecting the potential benefits of using only the most informative features. To overcome this limitation, we proposed a novel algorithm called the Imputation Method Based on Informative Features (IIF). This algorithm uses the Gray Wolf Optimization (GWO) method to select the most relevant features for imputation and ignore the less relevant ones. This strategic feature selection increased computational efficiency and improved the accuracy of the imputation results. Thus, it helped overcome the challenges posed by missing data, especially in video surveillance. We evaluated the performance of the IIF approach by comparing it with several state-of-the-art missing data imputation methods across different experimental scenarios. We used synthetic, real-world, and video surveillance datasets with varying rates of missing values for this comparison. Results from all experiments showed that the IIF method outperformed others, indicating its robustness and stability, making it a reliable choice for imputing missing data in different scenarios. These results demonstrate the effectiveness of the IIF algorithm in improving the accuracy and reliability of video surveillance systems.

17:30 ***Detecting Glioma, Meningioma, and Pituitary Tumors, and Normal Brain Tissues based on Yolov11 and Yolov8 Deep Learning Models***

Ahmd M Taha (Egypt University of Informatics, Egypt); Salah A. Aly (Fayoum University, Egypt)

Accurate and quick diagnosis of normal brain tissue Glioma, Meningioma, and Pituitary Tumors is crucial for optimal treatment planning and improved medical results. Magnetic Resonance Imaging (MRI) is widely used as a non-invasive diagnostic tool for detecting brain abnormalities, including tumors. However, manual interpretation of MRI scans is often time-consuming, prone to human error, and dependent on highly specialized expertise. This paper proposes an advanced AI-driven technique to detecting glioma, meningioma, and pituitary brain tumors using YoloV11 and YoloV8 deep learning models.

Methods: Using a transfer learning-based fine-tuning approach, we integrate cutting-edge deep learning techniques with medical imaging to classify brain tumors into four categories: No-Tumor, Glioma, Meningioma, and Pituitary Tumors.

Results: The study utilizes the publicly accessible CE-MRI Figshare dataset and involves fine-tuning pre-trained models YoloV8 and YoloV11 of 99.49% and 99.56% accuraies; and customized CNN accuracy of 96.98%. The results validate the potential of CNNs in achieving high precision in brain tumor detection and classification, highlighting their transformative role in medical imaging and diagnostics.

17:45 ***GAPI: Graph-Aware Prefill Interleaving for Knowledge Graph-Augmented LLM Inference in Explainable Education Policy Analysis***

Milan Parikh (IETE, SCRS, IEEE Senior Member, USA & Cytel, USA)

Large Language Model (LLM) inference systems augmented with Knowledge Graphs (KGs) offer powerful capabilities for explainable AI applications, yet existing architectures treat KG traversal as a preprocessing step decoupled from GPU prefill scheduling-placing graph query latency squarely on the critical inference path. We present GAPI (Graph-Aware Prefill Interleaving), a distributed inference architecture that tightly couples multi-hop KG traversal with LLM prefill scheduling via asynchronous CUDA stream overlap, an explainability trace caching layer, and a graph-conditioned speculative decoding mechanism. We instantiate GAPI on the domain of public school education policy analysis and student outcome prediction, constructing a 187,000-triple knowledge graph from NCES Common Core of Data and EdFacts datasets spanning 13,000 U.S. school districts.

Experiments on NVIDIA A4000/A5000 GPUs with LLaMA-3 8B demonstrate that GAPI reduces time-tofirst-token (TTFT) by 34.2%, improves query throughput by 31.7%, increases GPU SM utilization by 24.3%, and delivers explainability provenance traces at an overhead of only 1.8% compared to sequential KG+LLM baselines. GAPI outperforms Microsoft GraphRAG, LangChain RAG, and vanilla LLaMA-3 on all primary metrics while achieving a BERTScore F1 improvement of +6.4 points on education policy Q&A benchmarks.

Thursday, April 23

Thursday, April 23 9:00 - 11:15 (Africa/Cairo)

SESSION-7A: RESEARCH:Agricultural, Food Science & Biological Applications of ML

9 PAPERS

Room: ROOM-A

Chair: Tamer M Nassef (October University of Modern Sciences and Arts (MSA), Egypt)

9:00 ***Development of a Web-Based Interactive 3D Orbit and Ground-Track Visualization Tool for Orbital Mechanics Education***

Raed Kafafy (HCT, United Arab Emirates); Ossama Abdelkhalik (Iowa State University, USA); Mohamed Okasha (United Arab Emirates University, United Arab Emirates)

This paper presents the development of a lightweight web-based interactive visualization tool designed to support the teaching of orbital mechanics. The tool integrates real-time three-dimensional orbit rendering with a synchronized two-dimensional ground-track display, both driven by classical orbital elements. Implemented entirely in client-side JavaScript using modern web technologies, the application runs in standard web browsers without installation or platform-specific dependencies.

9:15 ***LLM-Assisted Software Development in Aerospace Engineering Education: A Structured Workflow***

Raed Kafafy (HCT, United Arab Emirates); Mohamed Okasha (United Arab Emirates University, United Arab Emirates); Ossama Abdelkhalik (Iowa State University, USA)

Large language models (LLMs) are increasingly being applied to software development tasks; however, their structured integration into aerospace engineering workflows remains insufficiently documented. This paper presents a structured workflow for LLM-assisted software development in aerospace engineering education, emphasizing problem decomposition, modular prompt design, iterative validation, and human-in-the-loop verification. The proposed methodology is demonstrated through a case study involving the development of a web-based visualization tool for orbital mechanics instruction. Rather than focusing on the software architecture itself, the study analyzes the interaction between domain expertise and AI-generated code, identifies common error categories, and evaluates the role of structured prompting in improving reliability and development

efficiency. The results highlight both the potential and limitations of LLM-assisted engineering, underscoring the necessity of expert validation in physics-based applications. The presented workflow provides a practical framework for integrating large language models into aerospace software development and educational tool prototyping

9:30 ***Performance Evaluation of a ResNet50 CNN for Corn Leaf Disease Classification***

Rhowel M. Dellosa (Pangasinan State University, Philippines)

This study evaluates the performance of a ResNet50-based Convolutional Neural Network (CNN) for classifying corn leaf diseases using image data. The dataset consists of 4,188 corn leaf images grouped into four classes: Healthy, Blight, Common Rust, and Gray Leaf Spot. Images were preprocessed through resizing, normalization, and augmentation before model training. The model achieved an overall accuracy of 62% on the classification task. Class-level results showed the strongest performance for Common Rust, with precision of 0.64, recall of 0.98, and F1-score of 0.77. In contrast, Gray Leaf Spot had the weakest performance, with precision of 0.55, recall of 0.06, and F1-score of 0.11, indicating that the model struggled to identify this disease reliably. The confusion matrix further showed frequent misclassification of Gray Leaf Spot and moderate confusion between Healthy and Blight classes. These findings suggest that while ResNet50 shows potential for automated corn disease classification, additional improvements in dataset balance, feature learning, and model optimization are necessary before practical deployment.

9:45 ***Analysis of Physicochemical Factors Affecting the Quality of Cacao Wine-Based Beverages using Random Forest Algorithm***

Martina Penalber (Isabela State University, Philippines)

This study investigates the physicochemical and sensory factors influencing the quality of cacao wine-based beverages using machine learning techniques. A dual-dataset approach was employed, integrating a secondary wine quality dataset containing physicochemical properties and a primary dataset derived from sensory evaluation of cacao wine-infused cocktails. The primary dataset included attributes such as color, taste, fineness, aroma, alcohol content, and overall acceptability. A Random Forest Regression model was applied to both datasets to evaluate predictive performance and identify key determinants of quality. The wine dataset yielded a moderate predictive performance with an R^2 score of 0.551, Mean Squared Error (MSE) of 0.348, and Root Mean Squared Error (RMSE) of 0.5897, with alcohol, volatile acidity, and free sulfur dioxide identified as the most influential variables. In contrast, the cacao dataset demonstrated strong predictive performance, achieving an R^2 score of 0.7895, MSE of 0.035, and RMSE of 0.1873. Feature importance analysis revealed that fineness (0.28) and aroma (0.28) were the most significant predictors of overall acceptability, followed by taste (0.23) and color (0.21), while alcohol content (0.01) had minimal influence. The findings highlight that while physicochemical properties contribute to the formation of beverage characteristics, sensory attributes play a more dominant role in determining consumer acceptability. The study demonstrates the effectiveness of machine learning in identifying key quality drivers and provides a data-driven framework for optimizing cacao wine-based beverages.

10:00 ***An Intelligent Decision Support Framework for Relevant and Responsive Capacity-Building in Start-Up Cooperatives***

Dahlee S. Pascua (Isabela State University, Philippines)

The sustainability of start-up cooperatives depends on the timely identification of organizational needs and the design of interventions that are both relevant and responsive to local conditions. This study presents an intelligent decision support framework for capacity-building in start-up cooperatives based on an empirical needs assessment of a savings and credit cooperative in Isabela, Philippines. A descriptive mixed-methods design was employed, involving 32 respondents composed of cooperative members, officers, and board directors selected through purposive sampling. Data were gathered using structured questionnaires, interviews, and focused group discussions, then analyzed through descriptive statistics and thematic analysis. Findings revealed strong demand for training and technical assistance in financial management, bookkeeping, strategic planning, cooperative governance, leadership, entrepreneurship, marketing, and livelihood development. The assessment also identified critical operational gaps, including insufficient policies and guidelines, limited training access, weak debt collection controls, inadequate bookkeeping capacity, and a lack of governance-related competencies among officers and board members. In response, the proposed framework structures decision-making into four phases: needs identification, intervention planning, mentoring and capability support, and monitoring and governance reinforcement. By translating assessment results into prioritized capacity-building actions, the framework serves as a decision support mechanism for higher education institutions, government agencies, and development partners engaged in cooperative extension. The study highlights how data-informed and structured decision processes can improve the design of sustainable support programs for emerging cooperatives.

10:15 ***A Machine Learning Driven Analysis of Employment Outcomes and Curriculum Relevance in a State University in the Philippines vs. Global Graduate Employability***

Liezl Joy L. Quilang (Isabela State University, Philippines)

This study examined the employment outcomes and curriculum relevance of Bachelor of Secondary Education major in Mathematics graduates from a state university in the Philippines using a mixed-methods descriptive-evaluative design integrating tracer study data, employer feedback, and a refined machine learning framework. Anchored in human capital theory, the study clearly distinguishes employment outcome as the target variable and retains only predictors that are conceptually and temporally independent of that outcome to reduce data leakage. Local tracer variables describe graduates' educational background, curriculum experience, competency development, employer feedback, and satisfaction levels, while global employability indicators are used as contextual benchmarks aligned by comparable features rather than as direct substitutes for local records. The revised machine learning design emphasizes hyperparameter tuning, model comparison, robust validation, and feature-attribution analysis to generate policy-relevant insights. Findings continue to show generally positive employment outcomes and strong curriculum relevance for entry into the education sector, while also indicating the need to strengthen research competency, critical thinking, teamwork, creativity, and diligence to better align curricular investments with evolving employability demands. The study concludes that tracer evidence becomes more decision-relevant when descriptive results are integrated with leakage-aware and explainable machine learning analysis.

10:30 ***Cost-Aware LLM Orchestration via Contextual Bandit Learning***

Vasanth rao Jadav (EPAM Systems, USA); Shalini Sudarsan (Kindercare Learning Companies, USA); Vikram Isanaka (UBS, USA)

The rapid adoption of large language models (LLMs) has enabled powerful capabilities across diverse applications, but their deployment introduces significant computational and monetary costs, particularly when high-capacity models are used indiscriminately. Existing LLM orchestration approaches rely primarily on static or heuristic routing strategies that fail to adapt to query complexity and do not explicitly optimize the trade-off

between response quality and operational cost. Recent cost-efficient inference methods, including model cascading and selective routing [18], highlight the need for adaptive learning-based decision frameworks. In this paper, we propose a cost-aware LLM orchestration framework formulated as a contextual bandit learning problem, where the system dynamically selects the optimal combination of model, tools, and prompting strategy for each query. A query analyzer extracts structured features capturing linguistic complexity, semantic intent, domain specificity, ambiguity, and cost sensitivity to guide routing decisions. The framework employs a composite reward function that jointly optimizes response quality, tokenbased cost, latency, and factual consistency, enabling continuous improvement without requiring labeled supervision. Experimental evaluation on a dataset of over 12,500 heterogeneous queries, including both benchmark and synthesized enterprise-style workloads, demonstrates that the proposed approach achieves near-premium response quality while reducing average token cost compared to static and heuristic baselines. Statistical analysis confirms that improvements are significant across query complexity levels. These results demonstrate that treating LLM orchestration as a contextual learning problem provides a scalable and practical solution for cost-efficient deployment of large language models in real-world systems.

10:45 ***Predictors of Licensure Examination Performance Among BEEd Graduates Toward a Random Forest Predictive Framework***

Jocelyn Sagun- De Vera (Pangasinan State University, Philippines)

This study examined the statistical predictors of licensure examination performance among Bachelor of Elementary Education graduates and used the findings as a basis for proposing a Random Forest predictive framework. The study utilized a quantitative predictive research design and analyzed the records of 51 graduates from academic years 2018 to 2022 who had taken the Licensure Examination for Teachers (LET). Predictor variables included age, sex, attendance to LET review, academic awards, entrance examination scores, and college grade point average (GPA), while overall LET rating served as the target variable. Descriptive results showed that most respondents were female, had attended formal review sessions, and obtained a GPA of 88. Statistical findings from the available dataset revealed that sex, attendance to LET review, and college GPA were significantly associated with LET performance, with review attendance showing the strongest influence. Correlation and multiple regression analyses indicated that sex, attendance to LET review, and college GPA were significant predictors of LET performance. Based on these findings, a Random Forest predictive framework is proposed for future model implementation and validation.

11:00 ***Audit-Ready Fraud Detection: Aligning Gradient Boosting and SHAP with GDPR and SOX Compliance***

Deepak Pai (USA); Karthik Pappu (Dakota State University, USA); Yogeesh Kunigal Gangaiah (Qualitest Group Llc, USA)

Financial fraud detection requires models that are both accurate on highly imbalanced transaction data and explain- able for regulatory compliance under the General Data Protection Regulation (GDPR) Article 22 and Sarbanes-Oxley (SOX) Section 404. Most prior studies on fraud detection optimize for accuracy or area under the ROC curve (AUC) while neglecting post-hoc explainability and subgroup robustness required in audit contexts. This study compares four machine learning algorithms, namely Logistic Regression, Random Forest, XGBoost, and a multilayer perceptron (MLP), on 284,807 real credit card transactions exhibiting a 577.9:1 class imbalance. XGBoost achieves the highest discrimination with ROC-AUC of 0.983 and precision-recall AUC (PR-AUC) of 0.875, reducing the false positive rate to approximately 0.02% at optimized thresholds. SHAP (SHapley Additive exPlanations) analysis identifies the PCA components V14, V4, and V12 as the three most influential features driving fraud predictions. Subgroup stability analysis across transaction amount quartiles and temporal periods confirms consistent model performance (F1 ranging from 0.83

to 0.92) with no degradation in any segment. Threshold optimization at 0.73 yields 93.2% precision while maintaining 83.7% recall. These results provide empirical evidence that gradient boosting methods combined with SHAP explainability can deliver audit-ready fraud detection aligned with regulatory transparency mandates.

Thursday, April 23 9:00 - 11:15 (Africa/Cairo)

SESSION-7B: RESEARCH:NLP, Sentiment Analysis & Domain-Specific AI Applications

9 PAPERS

Room: ROOM-B

Chair: Davor Cafuta (Zagreb University of Applied Sciences, Croatia)

9:00 ***Automated Pancreas Segmentation in 2D CT Scans: A Comparison of Prompt-Based Segment Anything Model (SAM) and Classical U-Net***

Mina Awad, Ahmed F. Elnokrashy and Mai S. Mabrouk (Nile University, Egypt)

Pancreas segmentation in 2D CT scans is critical for clinical applications but remains challenging due to the organ's small size and low contrast. This study compares a fine-tuned Segment Anything Model (SAM)-leveraging bounding box prompts-against the classical U-Net for pancreas segmentation using the Pancreas-CT dataset (Kaggle). While both models achieve comparable Dice scores (SAM: 0.835, U-Net: 0.82), SAM demonstrates superior robustness, eliminating U-Net's occasional catastrophic failures (e.g., blank masks) due to its prompt-driven design. Additionally, SAM trains 2.5 times faster and produces smoother, anatomically plausible segmentations, even in cases where Dice differences appear marginal. These findings highlight SAM's potential for reliable small-organ segmentation, where consistency and clinical usability outweigh minor metric gaps

9:15 ***Towards Automated Grading of Hand Drawings: A Dataset and Classification Framework***

Alaa Abdulfattah Sayed and Khaled F. Hussain (Assiut University, Egypt); Majid Askar (Assistant Lecturer in Assiut University, Egypt)

Hand drawing skills are essential for computer science students to effectively communicate ideas and visualize concepts. However, evaluating hand-drawn sketches is subjective and time-consuming. This paper presents a novel dataset of 1,356 hand-drawn geometric composition sketches from first-year computer science students at Assiut University, Egypt, expanded to 10,817 through systematic augmentation. We developed a ResNet-50 classification framework achieving 100% accuracy on binary classification of cube_cone and cube_sphere compositions. The dataset exhibits natural class imbalance (99.7% in two dominant categories), providing insights for deploying machine learning in authentic educational contexts. Our

contributions include: (1) a novel annotated dataset of actual educational hand drawings with quality grades, (2) a comprehensive analysis of class imbalance in educational contexts, and (3) baseline classification results that demonstrate the feasibility of automated assessment.

9:30 ***Implementation of Sentiment Prescriptive Analytics in a Web-Based Library User Satisfaction System Using Deep Learning***

Jerry Incierto Teleron (Technological Institute of the Philippines - Quezon City, Philippines)

This study presents the development and evaluation of a web-based Library User Satisfaction System that integrates deep learning-based sentiment analysis and rule-based prescriptive analytics to support data-driven decision-making in academic libraries. The system was developed using the Design and Development Research (DDR) approach, which includes analysis, design and development, and evaluation phases. A Bidirectional Long Short-Term Memory (BiLSTM) model with an attention mechanism was implemented to analyze textual feedback and classify sentiments into positive, neutral, and negative categories. In addition, a rule-based prescriptive analytics module evaluates survey scores across library service dimensions-resources, facilities, technology, and services-to identify areas requiring improvement and generate recommendations. The model was trained using 10,000 labeled feedback comments, achieving 100% accuracy, precision, recall, and F1-score on the validation dataset. The system also provides an interactive dashboard for monitoring satisfaction indicators and sentiment distribution. Results indicate that the technology dimension received the lowest satisfaction score, highlighting the need for improvements in technological services. Overall, the system demonstrates the potential of combining BiLSTM-based sentiment analysis and prescriptive analytics to transform user feedback into actionable insights for improving library services.

9:45 ***Toward Trustworthy Audit Automation: Explainable Meta-Learning for Financial Fraud Detection***

Nancy Chendeb (De Vinci Higher Education, France & De Vinci Research Center, Paris, France); Antonin Chevalier kremper (Devinci Higher education, France); Abdallah Mustapha Ziade (Lebaneese University, Zaca International, Lebanon); Malak Mohamad Daher (Jinan University Tripoli, Lebanon & Zaca International, Lebanon)

The automation of financial auditing has become a critical challenge due to the increasing volume of transactions and the growing sophistication of fraud schemes. While machine learning offers promising solutions for anomaly detection, two major limitations hinder its adoption in practice: the absence of a universally optimal detection model and the lack of explainability required in regulated audit environments. In this paper, we propose a contextual and explainable metalearning framework for financial anomaly detection. Our approach combines multiple supervised and unsupervised detectors, within a unified meta-model architecture. Transactional data are enriched with contextual features, and model outputs are aggregated through supervised meta-learners that adaptively weight each detector based on the transaction profile. To ensure transparency, we integrate multi-level explainability using feature importance analysis, SHAP, and LIME, providing both global and local interpretations of model decisions. Experiments conducted on realistic financial fraud datasets demonstrate that the proposed meta-model significantly outperforms individual detectors and traditional aggregation methods, achieving an improvement of approximately 50% in F2-score. Results also highlight the dominant contribution of supervised models, while unsupervised methods provide limited complementary value in this setting. Overall, this work shows that combining contextual feature engineering, meta-learning, and explainable AI enables more accurate, robust, and trustworthy anomaly detection systems, paving the way for their effective integration into real-world audit processes.

10:00 ***Swin Transformer Architectures for Breast Histopathology Image Classification in Cancer Detection***

Mahmoud Adel Nour, Mai S. Mabrouk and Ahmed F. Elnokrashy (Nile University, Egypt)

The increasing complexity of medical imaging requires architectures that balance accuracy with computational efficiency. While Swin Transformers have emerged as a compelling alternative to CNNs for medical image analysis, the most effective implementation strategy remains an open question. This study systematically compares three distinct Swin Transformer-based approaches: a standard Swin Transformer with linear classification, a hybrid Swin-CNN architecture combining both feature extraction methods, and a Swin-UNet implementation followed by classification. Using a well-established medical imaging benchmark dataset, we evaluate these architectures through standard performance metrics. Our findings demonstrate that all three approaches achieve robust classification performance, with the Swin-CNN hybrid showing consistent advantages over the pure transformer implementation. The hybrid architecture's success suggests that integrating global attention mechanisms with local convolutional features enhances the model's ability to capture clinically relevant patterns - a conclusion supported by similar observations in recent literature on specialized medical imaging tasks. The comparative analysis reveals meaningful trade-offs between architectural complexity and performance gains. The standard Swin Transformer provides a strong baseline with straightforward implementation, while hybrid variants offer incremental improvements at the cost of additional design complexity. These insights help guide the selection of transformer-based architectures for diverse medical imaging applications, particularly in histopathology image analysis, where both local and global context are crucial for accurate diagnosis.

10:15 ***A Smart Mobile System Using Deep Learning for Business Tourism Demand Prediction in Palawan***

Jerry Incierto Teleron (Technological Institute of the Philippines - Quezon City, Philippines)

Business tourism, particularly Meetings, Incentives, Conferences, and Exhibitions (MICE), plays a significant role in regional economic development but is challenged by demand variability and seasonality. This study presents a web-based, mobile-accessible system for predicting business tourism demand in Palawan using deep learning techniques. Tourism datasets covering the period 2017-2026 were collected and preprocessed, and multiple models-including LSTM, GRU, BiLSTM, CNN-LSTM, and Transformer-were developed and evaluated using MAE, RMSE, and MAPE metrics. Results show that the Transformer model achieved the highest prediction accuracy, outperforming other models in capturing complex temporal patterns. The selected model was integrated into a web-based platform developed using React.js and Firebase, enabling real-time forecasting and interactive visualization through dashboards. The system allows stakeholders to analyze trends, identify seasonal and event-driven demand patterns, and generate one-year forecasts. The findings demonstrate that combining Transformer-based deep learning with web deployment provides an effective tool for data-driven decision-making, resource allocation, and strategic planning in business tourism. The developed system contributes to improving forecasting accuracy and supports sustainable tourism management in emerging destinations such as Palawan.

Keywords-Tourism Demand Forecasting, Deep Learning, Transformer, LSTM, Time Series Analysis, Predictive Modeling

10:30 ***Prediction of Digital Literacy in Agricultural and Rural Areas Using Machine Learning to Enhance Digital Literacy Education***

Celeste Mercado (Pangasinan State University, Philippines)

This study examines the factors influencing digital literacy in rural and agricultural areas. Several machine learning models, including Linear Regression, Random Forest Regressor, and Gradient Boosting Regressor, were applied to predict digital literacy outcomes based on demographic and socio-economic variables. The Gradient Boosting Regressor showed the best performance with an R^2 of 0.83, MAE of 4.03, and RMSE of 5.56. The analysis identified education level, household income, and employment status as the most significant predictors of digital literacy. Rural respondents, particularly those with lower income and education, showed lower digital literacy compared to semi-rural counterparts. The feature importance analysis highlighted education level as the strongest predictor. Based on these findings, the study recommends the development of localized digital literacy programs in rural areas, with a focus on affordable internet access, public-private partnerships, and gender-sensitive initiatives. Additionally, data-driven decision-making using machine learning can help policymakers address digital literacy gaps more effectively. The study emphasizes integrating digital literacy into rural development policies, particularly in Pangasinan, to improve agricultural practices and socio-economic outcomes.

10:45 ***Predicting College Students' Academic Performance Using Machine Learning***

Julius Cesar O. Mamaril, Rhowel M. Dellosa, Celeste Mercado and Jennie Fernandez (Pangasinan State University, Philippines); Jerry I. Teleron (Technological Institute of the Philippines - Quezon City, Philippines); Marlon Dela pena Hernandez (Bulacan State University, Philippines); Walter C. Nocasa, Lorena B. Echalar, Teresita D. Mamaril, Kristine May D. Casibang, Jeffrey A De Asis, Christian S Dela Cruz, Jumar Baratang, Vemma Mae G. Olivar, Nova E. Arquillano, Gloria M. Ducut, Wilbert O. Rosario, Rocky P. Manalo, Honelly Mae S. Cascolan, Cristina L. Javier, Juan Primitivo P. Petrola, Marlon L. Perado and Roselyn Villacorte (Pangasinan State University, Philippines)

This study aimed to predict college students' academic performance based on their sleep patterns, study habits, and lifestyle choices using machine learning algorithms. A dataset consisting of key features such as study hours, sleep duration, stress levels, and socio-economic factors was analyzed using several machine learning models, including Linear Regression, Random Forest, Support Vector Machines (SVM), XGBoost, and Neural Networks (DNN). The results showed that XGBoost and Random Forest achieved the highest performance in both regression and classification tasks, with accuracy values of 0.90 for XGBoost and 0.87 for Random Forest. Linear Regression was particularly effective for predicting GPA, achieving an R^2 value of 0.85. The analysis of feature importance revealed that study_hours, sleep_hours, and motivation were the most significant predictors of academic success. Additionally, students were classified into High Risk, Medium Risk, and Low Risk categories based on their academic performance, with High Risk students exhibiting poor study habits, high stress, and low physical activity. The findings suggest that machine learning can effectively predict student performance, providing insights for targeted interventions to support at-risk students. The study emphasizes the importance of lifestyle factors such as sleep, motivation, and stress management in academic success, with implications for developing data-driven strategies to improve student outcomes.

11:00 ***Adaptive Multi-Modal Routing for Cost-Aware Presentation Summarization***

Vasanth rao Jadav (EPAM Systems, USA); Anshul Garg (Amazon, USA); Vikram Isanaka (UBS, USA)

Presentation documents such as PowerPoint (PPT) slides are widely used for communicating complex ideas across enterprise, academic, and clinical domains. However, automated summarization remains challenging due to the inherently multi-modal nature of slides, which combine structured text, visual elements, and layout semantics. Existing approaches typically rely on either textbased summarization using transformer models or visionlanguage models, resulting in incomplete understanding and suboptimal summaries. In this paper, we propose an adaptive multi-modal

framework for PPT summarization that dynamically selects the most appropriate processing strategy for each slide based on its content characteristics. The approach integrates structured PPT parsing with vision-language models and introduces a lightweight routing classifier trained on slide-level features to choose between text, vision, or hybrid pipelines.We further define a semantic slide representation that enables hierarchical summarization at slide, section, and document levels. To improve practical applicability, the framework incorporates cost-aware optimization by reducing unnecessary vision model usage while preserving summary quality. The proposed system is evaluated on a curated dataset consisting of 120 presentations and approximately 3,450 slides across multiple domains. Experimental results demonstrate improved semantic coverage and coherence compared to single-modality baselines while significantly reducing computational overhead. Statistical analysis confirms that the improvements are significant, highlighting the effectiveness of adaptive multi-modal routing for scalable presentation understanding.

Thursday, April 23 9:00 - 11:15 (Africa/Cairo)

SESSION-7C: RESEARCH:Workforce Analytics, Social Impact & Organizational Decision Support

9 PAPERS

Room: ROOM-C

9:00 ***A multilingual GAT-based model for product search in an E-commerce platform***

Cherine Mohamed (British University in Egypt, Egypt); Abeer Hamdy (The British University in Egypt, Egypt)

The phenomenal surge of e-commerce platforms globally has necessitated their ability to handle product search & retrieval in diverse languages to cater to a myriad of users. One of the key challenges is the vo-cabulary mismatch between formal and organized product catalogue descriptions and the informal, am-biguous nature of user queries. Some deep learning architectures such as CNNs, RNNs and language mod-els have attempted to address this challenge; However, they struggle to capture contextual information & structural relationships between queries and products. In this work, a hybrid model encompassing multilin-gual BERT and GAT is proposed to retrieve relevant products given a query. Our model was evaluated on the Amazon ESCI dataset, and it achieved Preci-sion@K of 86.23% and 78.37% and nDCG of 72.51% and 66.18% in English and Spanish languages respec-tively.

9:15 ***AI Food Scanner - AI-Powered Dish Nutrition Scanner App***

Ekereuke Udoh (QA Higher Education, United Kingdom (Great Britain))

The difficulty associated with manual food tracking remains a primary obstacle to successful personal health management, frequently resulting in user burnout and unreliable nutritional records. Contemporary mobile applications generally depend on inefficient methods such as text search or barcode scanning, which struggle to accurately account for complex or home-cooked meals. To address these limitations, we present AI Food Scanner, a cross-

platform mobile solution that utilizes Artificial Intelligence to deliver immediate nutritional analysis via a single photograph. By employing sophisticated computer vision techniques, AI Food Scanner detects food items within an image and extracts granular nutritional data, including both macro- and micronutrients. The technological infrastructure combines a React Native interface for optimal user engagement, a Python FastAPI backend for request handling, and Azure Cognitive Services to power the image recognition model. Furthermore, the system integrates the USDA FoodData Central API to guarantee a verified and extensive nutritional database. This study's main contribution is a "visual-first" approach to dietary monitoring that reduces the burden of manual entry and improves data precision. This empowers a wide demographic-from athletes to diabetic patients-to make better dietary decisions. Early testing indicates high efficacy in food recognition and favorable user satisfaction regarding the application's ease of use.

9:30 ***Data-Driven Analysis of Factors Influencing Employee Performance in the Lebanese Construction Industry***

Fawziah Kassar (Lebanese University, Lebanon); Nader Bakir (Beirut Arab University, Lebanon); Khouloud Samrouth (Lebanese University, Lebanon); Layla Akkoumeh (Jinan University, Lebanon); Nesrine Atitallah (FCS, Arab Open University, Madinah, KSA, Saudi Arabia); Hussein Krayani (Peter the Great St.Petersburg Polytechnic University, Russia); Lina Kemayel (Lebanese University, Lebanon)

The construction industry is a key driver of economic growth and operates as a complex system in which the performance of all organizational components is critical to project success. Among these components, human resources represent a strategic and dynamic asset. Even with adequate materials, equipment, methods, and financial resources, construction companies cannot perform optimally without skilled, motivated, and competitive workers. This first study investigates the main factors influencing employee performance in the Lebanese construction industry, recognizing the central role of construction workers in project execution. Specifically, it examines leadership style, individual characteristics, work environment, individual ability, and motivation as determinants of performance. Using a quantitative, cross-sectional research design and statistical analysis, the study ranks these factors according to their relative importance and impact on employee productivity and professionalism. The findings aim to support construction companies in prioritizing effective management practices and targeted interventions to enhance workforce performance and, ultimately, project success within the Lebanese construction sector.

9:45 ***A Hybrid CNN-Vision Transformer Architecture for Enhanced Document Understanding in Higher Education Archiving***

Yousuf Nasser Al Husaini and Mohammed Abdullah Salim Al Husaini (Arab Open University Oman, Oman); Rahma Mohammed Al Kharusi and Abdul Rahman Al Abri (Arab Open University, Oman)

The growing number and variety of digital documents in institutions of higher learning necessitate sophisticated computerized techniques to achieve credible data storage, retrieval and data mining. The conventional optical character recognition (OCR) and rule-based systems do not perform well on low-resolution scans, non-standard layouts, or handwriting, and impede their application within an institutional context. The presented paper proposes a hybrid deep learning architecture that combines convolutional neural networks (CNNs), Vision Transformers (ViT), and transformer-based OCR and entity-extraction subsystems to achieve ultimate document-image understanding. In the proposed model, document classification, text recognition, layout analysis and metadata extraction are collectively carried out in an end-to-end multimodal pipeline. The RVL-CDIP benchmark and a curated collection of institutional archival documents were used to conduct experiments. The hybrid model attained a top-1 accuracy of 96.8% on RVL-CDIP, where LayoutLMv3 and Donut did not do well. It achieved a Character Error Rate (CER) of 2.31% and an entity extraction F1-score of 94.8 on the archival dataset, representing a significant improvement over typed, handwritten, and low-quality documents. Ablation experiments show the

synergistic role of the CNN, ViT, and OCR modules, whereas robustness analysis demonstrates increased resistance to noise, blur, and compression artefacts. These findings emphasize the possibility of multimodal hybrid designs to improve large-scale electronic archiving processes in institutions of higher learning.

10:00 ***Securing APIs in Government Clouds and Runtime Fabric Using FIPS-Enabled MuleSoft***

Venkata Pavan Kumar Gummadi (Broadridge, USA); Lakshmi Sujatha Chilamkurthi (Cigna, USA); Srikanth Kavuri (Iconsoft, USA)

MuleSoft-based API-led connectivity architectures in government clouds and Runtime Fabric (RTF) environments handle highly sensitive data and mission-critical workflows, necessitating stringent FIPS 140-2/3 compliance and FedRAMPaligned security. This paper presents a comprehensive security framework for FIPS-enabled MuleSoft APIs, integrating gateway-level controls with service level objectives (SLOs). The approach combines authentication, authorization, FIPScompliant transport security, threat protection, data privacy, input validation, and centralized policy governance with SLOcentric observability, enabling government agencies to monitor security posture via formal service level indicators (SLIs) alongside reliability metrics in high-assurance environments. Unlike traditional approaches that treat compliance as an isolated concern, this work treats security as a first-class SLO dimension with explicit SLIs, error budgets, and operational feedback loops across FIPS-enabled MuleSoft control planes.

10:15 ***Service Level Objective (SLO) Observability with Splunk and Dynatrace in Microservices***

Venkata Pavan Kumar Gummadi (Broadridge, USA); Lakshmi Sujatha Chilamkurthi (Cigna, USA); Srikanth Kavuri (Iconsoft, USA)

In modern microservices-driven architectures, ensuring system reliability and user satisfaction demands a shift from traditional infrastructure monitoring to a Service Level Objective (SLO)-centric observability model. This paper explores how enterprises can leverage observability platforms such as Splunk and Dynatrace to define, track, and enforce SLOs that align closely with real user experiences. It discusses the theoretical underpinnings of SLO-based monitoring, contrasts them with older paradigms such as system uptime and generic static thresholds, and analyzes the integration challenges and architectural considerations of implementing observability at scale. The paper illustrates successful applications of SLO frameworks in reducing alert fatigue, improving mean time to resolution, and enhancing cross-team accountability, and presents best practices and actionable recommendations for organizations at various stages of their observability journey

10:30 ***A Random Forest and Logistic Regression Analysis of Faculty Development Programs, Satisfaction, and Career Mobility in a State University***

Annalene Grace E. Co (Quirino State University, Philippines)

This study assessed faculty satisfaction with faculty development programs (FDPs) in a state university and identified factors affecting overall satisfaction and career mobility. An explanatory sequential mixed-methods design was used. Primary data were analyzed through descriptive statistics, t-test, one-way ANOVA, Pearson correlation, and multiple regression, while qualitative responses underwent thematic analysis. Secondary data were examined using Logistic Regression and Random Forest in Google Colab to predict occupational mobility. Results showed that faculty were generally satisfied with FDPs, especially in relevance and alignment, content and delivery quality, and perceived impact. Accessibility and institutional support

received lower ratings. Multiple regression indicated that only institutional support and perceived impact significantly predicted overall satisfaction. Qualitative results showed that faculty valued relevant and practical training but suggested stronger administrative support, improved accessibility, more specialized workshops, and workload-sensitive implementation. For the predictive models, Logistic Regression achieved 90.87% accuracy, while Random Forest reached 100%. Both models identified job satisfaction, interest in career change, and salary as the strongest predictors of occupational mobility, with Random Forest ranking job satisfaction as the most influential. The study concludes that FDPs are most effective when relevant, impactful, and backed by strong institutional support. It also suggests that faculty satisfaction is a key link between professional development and career stability. Integrating primary institutional and secondary predictive data offers a broader view of how faculty development shapes satisfaction and long-term professional outcomes.

10:45 ***A Random Forest-Based Analysis of Factors Affecting the Income of Carpenters in Selected Barangays of Kayapa, Nueva Vizcaya, Philippines***

Ayson D. Paclit (Nueva Vizcaya State University, Philippines)

This study examined the factors affecting the weekly income of carpenters using a Random Forest regression model. A dataset containing 248 valid records was used, with 198 cases allocated for training and 50 for testing. The model was developed to predict weekly income based on respondent attributes and to identify the variables that most strongly influenced earnings. Results showed that the Random Forest model achieved a Mean Absolute Error (MAE) of Php 272.47, a Root Mean Squared Error (RMSE) of Php 444.14, and a test-set R^2 of 0.779. In addition, 5-fold cross-validation produced a mean R^2 of 0.747, indicating reasonably stable predictive performance across resampled data splits. Feature importance analysis revealed that Years of Experience (0.315) was the strongest predictor of weekly income, followed by NC II Holder (0.275), Age Group (0.169), Hours/Day (0.114), Civil Status (0.070), and Position (0.056). These findings suggest that work experience and formal certification are the most significant determinants of carpenter income in the dataset. The study highlights the usefulness of machine learning in predicting labor-related outcomes and in identifying the most meaningful income-related factors for policy and community-based intervention.

11:00 ***CostAgent: Self-Improving Autonomous LLM-Based Orchestration for Cost-Optimal Cloud Data Processing at Scale***

Naga Krishna Reddy Muppidi, Veera Ravindra Divi and Sneha Gullapalli (USA); Sruthi Rachamalla (Southern Illinois University, Carbondale, USA); Subhash Tatavarthi (Kasmo, USA)

The explosive growth of data-intensive applications has created an urgent need for cost-effective cloud computing solutions. While preemptible cloud instances (AWS Spot Instances, Azure Spot VMs, Google Cloud Spot VMs) offer 70- 90% cost savings, their unpredictable availability makes them challenging for production workloads. We present Cost Agent, a novel autonomous orchestration framework that leverages large language models (LLMs) for intelligent preemptible instance allocation through a formally-grounded preemption-resilient autonomy model. Our key contributions are: a bounded-risk autonomy framework with formal safety guarantees that enables fully autonomous AI operation by constraining the agent's action space to provably safe compute decisions; a multi-objective validation system using clamping-based constraint enforcement across cost, performance, reliability, and security dimensions; an LLM in context learning architecture that enables self-improvement through historical experience without explicit retraining; and a complete open-source implementation with infrastructure-as code enabling reproducible deployment

and validation. Through comprehensive cost modeling based on AWS pricing (January 2026), validated on 100GB workloads and scaled to 1PB, we demonstrate that CostAgent achieves 58-69% compute only cost reduction and 75-80% full-stack cost reduction across multiple data scales compared to traditional on-demand approaches, while eliminating the maintenance burden of manually reconfiguring pipelines when data volumes change. Validation experiments on synthetic and TPC-DS datasets show 14,492 records/sec throughput with successful checkpoint recovery. At petabyte scale, LLM orchestration overhead drops to 0.5% of total cost, enabling near-maximum spot discount capture (~70%). We provide formal proofs of safety properties, complete production-ready implementation, and empirical validation on real-world workloads with projections to petabyte scale.

Thursday, April 23 9:00 - 11:15 (Africa/Cairo)

SESSION-7D: RESEARCH:AI in Education: Learning Analytics, Prediction & Student Success

9 PAPERS

Room: ROOM-D

Chair: Tejas Pravinbhai Patel (Amazon, USA)

9:00 ***Synchronized 360°VR Classroom Using Locally Hosted Open Infrastructure and Meta Quest 3 Headsets***

Brigitta Cafuta, Danijela Pongrac and Lidija Tepes Golubic (Zagreb University of Applied Sciences, Croatia)

This paper presents the design and implementation of a locally hosted, synchronized 360° VR system for group teaching in STEM education, using only off-the-shelf hardware and free software. The proposed architecture is deployed at the Zagreb University of Applied Sciences and is based on a dedicated virtual server, Wi-Fi 6E classroom infrastructure, and a fleet of 25+1 Meta Quest 3 headsets. The system utilizes DeoVR's remote control API and custom Python scripts over ADB to enable centralized teacher control of video playback, including starting, pausing, seeking, and switching between heterogeneous 360° video assets of varying resolutions, frame rates, and formats. A digital reservation module integrated with the institution's scheduling system automates content preparation and access control for instructors. Network and load testing with 26 devices confirm low-latency, stable playback well below available bandwidth limits, demonstrating that a local, license-free solution can effectively replace costly commercial VR management platforms.

9:15 ***PMFSNet from Scratch: Evaluating Lightweight Backbone Efficiency for Medical Image Segmentation***

Mina Magdy Kamel, Ahmed F. Elnokrashy and Mai S. Mabrouk (Nile University, Egypt)

The current state-of-the-art medical image segmentation methods focus on achieving high accuracy but this comes at the cost of increased computational complexity and model complexity. The transformer-based architectures are effective in modeling global dependencies but they require large-scale pretraining and suffer from overfitting when applied to limited medical datasets. Moreover, they overlook critical inductive biases inherent to convolutional neural networks (CNNs), which are essential for capturing local contextual features. To address these limitations, we evaluate PMFSNet-a lightweight self-attention segmentation model-trained entirely from scratch, without the use of any pretrained weights. PMFSNet introduces a plug-and-play Polarized Multi-scale Feature Self-attention (PMFS) block that enhances global context modeling while maintaining a small and efficient design suitable for deployment in resource-constrained environments. In a controlled benchmark against scratch-trained U-Net, ResNet-34, and transformer-based SegFormer-B5 backbones on the CVC-ClinicDB dataset, PMFSNet achieves superior segmentation accuracy with the highest Dice score (0.9236), mIoU (0.9216), and lowest computational cost (0.99M parameters, 2.21 GFLOPs). These findings underscore the backbone strength of PMFSNet and its suitability for real-time clinical applications

9:30 ***An Optimized Ressource Allocation and Recommendation Algorithm for Classroom Occupancy Management: A Case Study at Alioune Diop University of Bambey***

Youssou Kassé (Université Alioune Diop, Senegal); Fatoumata Baldé (Université Alioune Diop de Bambey, Senegal); Ndeye Penda Mbaye (Université Alioune DIOP de Bambey, Senegal)

The continuous growth in student enrollment in public universities has significantly increased pressure on pedagogical infrastructures, particularly classroom availability. In Senegalese universities, such as Université Alioune DIOP (UAD), classroom allocation and timetable management are still largely manual, leading to frequent scheduling conflicts, inefficient space utilization, and course cancellations. This paper proposes an optimized resource allocation algorithm combined with a recommendation mechanism to ensure optimal classroom occupancy under multiple academic constraints. The proposed approach is based on greedy optimization principles and is formally validated using loop invariants. Unified Modeling Language (UML) diagrams are used to specify system behavior and structure, while algorithmic complexity analysis evaluates performance. The solution is embedded in a web-based platform to support academic administrators in making rapid, consistent, and conflict-free allocation decisions.

9:45 ***Design and Evaluation of a Fairness-Aware Student Success Prediction System Using Learning Analytics Data (OULAD)***

Ekereuke Udoh (QA Higher Education, United Kingdom (Great Britain)); Betty Nkem Okwedadi (Ulster University, United Kingdom (Great Britain)); Edita Gashi (QA Higher Education, United Kingdom (Great Britain))

Student retention remains a persistent challenge in higher education, with non-completion rates exceeding 30% across OECD countries and even higher in developing regions. Machine learning (ML) and learning analytics (LA) offer early- warning capabilities by predicting at-risk learners based on demographic, behavioural, and assessment data. However, most systems optimise solely for accuracy, overlooking fairness, transparency, and interpretability. This research addresses that gap through the design and evaluation of a Fairness-Aware student success prediction system using the Open University Learning Analytics Dataset (OULAD). The system integrates predictive modelling, fairness auditing, and explainability into one unified framework. Four machine learning classifiers- Logistic Regression, Decision Tree, Random Forest, and XGBoost-were benchmarked, with XGBoost achieving the highest performance (AUC = 0.947). Predictions were expressed as calibrated probabilities, categorised into risk bands to guide proportionate human intervention. Fairness was quantified using Demographic Parity Difference (DPD) and Equal Opportunity Difference (EOD), while

SHAP and LIME provided interpretability at both global and individual levels. Results show that post-mitigation fairness gaps (ΔDP/ΔEOD) were reduced to near-zero levels with less than 1% accuracy loss. The study extends fairness analysis beyond traditional demographics to contextual fairness, incorporating behavioural proxies-such as late-night study patterns and inactivity streaks-that reflect socio-economic and time-related pressures among international and working learners. This work contributes a reproducible framework and an interactive dashboard prototype that together demonstrate how predictive accuracy, fairness, and transparency can coexist to support equitable, responsible, and globally inclusive learning analytics.

10:00 ***Comparative Study of Machine Learning Algorithms for Predictive Analytics in Web Applications***

Raju Dandigam (Navan, USA); Chirag Agrawal (Amazon, USA)

The exponential growth of online applications in industries such as e-commerce, social media, online education, and financial services has resulted in massive amounts of user- generated and transactional data, opening up exciting new opportunities for predictive analytics. When it comes to decision support systems, ML is the way to go. It's used for recommendation systems, user behavior prediction, fraud detection, customer churn analysis, and more. The disparities in data properties, the requirement for consistency in growth capacity, and limited computing resources all contribute to making algorithm selection an ongoing challenge for web-based predictive analytics. In this regard, the research in question is useful. Here, we will compare and contrast some of the most well-known supervised machine learning algorithms, including Logistic Regression, Decision Tree, Random Forest, Support Vector Machine (SVM), K-Nearest Neighbors (KNN), Naïve Bayes, and ANN algorithms, which are designed for web-based predictive analytics. To train and assess prediction models, a web dataset is first refined following organized experimental data collection, then subjected to feature engineering. Predictive analytics models differ significantly in reliability, scalability, and accuracy when measured using conventional evaluation metrics such as precision and recall, ROC-AUC, computational efficiency, and overall model accuracy. The prediction accuracy of ensemble and neural network models is often higher, although simpler models may be better for interpretability and computational cost.

10:15 ***An Immersive Virtual Tour Framework for Educational Institutions Using 360° VR and AI-Based Chatbot***

Mahmoud Rehan Morgan (Elshorouk Academy, Egypt); Alaa M Nagy (HICIT, El-Shorouk Academy, Egypt); Ahmed Radi (Technical Research Center, Egypt)

Traditional methods of showcasing educational institutions often fall short. Static photos and brochures lack immersion, hindering students' ability to envision themselves within the academy environment. Physical visits for basic information can be inconvenient and environmentally impactful. This project addressed these challenges by pioneering a virtual tour system integrated with a chatbot guide for of El Shrouk Academy website. Cutting-edge technologies were used to enhance the accessibility by proposing a 360° VR tour which transcends geographical and physical barriers for website visitors. Students everywhere can virtually explore facilities, classrooms, and common areas, gaining a deeper understanding of the learning environment. Proposed VR tour, coupled with an audio guide, fosters a more engaging experience compared to static visuals to allow prospective students to connect with the academy's atmosphere. Also, AI-based chatbot technology powers a virtual guide within the VR tour and throughout the website to readily answer users' inquiries, potentially reducing the need for physical visits for basic information gathering. Dedicated faculty members profiles integrated within the VR environment further enhance user experience by allowing students to learn more about the esteemed faculty members. This project goes beyond showcasing an institution; it represents a pioneering vision for the future of education as it embraces

sustainability and offers a convenient and immersive way for prospective students to explore educational offerings. This project serves as a model for other institutions seeking to revolutionize their student recruitment strategies and enhance the accessibility and engagement of their virtual presence.

10:30 ***Application of Machine Learning to Predict Student Academic Outcomes in Higher Education***

Armando D. Junio (Pangasinan State University, Philippines)

This study applied machine learning models to predict student academic outcomes in higher education using a publicly available secondary dataset. The dataset contains 4,424 student records and 35 variables covering demographic characteristics, academic progression, financial indicators, and macroeconomic conditions. After preprocessing and an 80/20 train-test split, eight classification models were evaluated using accuracy, weighted precision, weighted recall, weighted F1-score, and weighted one-vs-rest AUC-ROC. Among the baseline models tested, Random Forest achieved the strongest overall performance with 0.781 accuracy, 0.769 weighted precision, 0.781 weighted recall, 0.768 weighted F1-score, and 0.911 weighted AUC-ROC. Feature importance analysis showed that the most influential predictors were primarily academic progression variables, especially curricular units approved and grades in the first and second semesters, followed by age at enrollment, tuition-fee status, course, and selected family background variables. The findings suggest that academic continuity and successful course completion provide the clearest signals for identifying students at risk of adverse academic outcomes. The study offers a reproducible baseline for data-driven student support and highlights the need to ground intervention claims in the actual structure of the dataset used.

10:45 ***DocSync: Agentic Documentation Maintenance via Critic-Guided Reflexion***

Sidhesh Badrinarayan (Google, USA); Adithya Parthasarathy (Newyork University, USA)

Software documentation frequently drifts from executable logic as codebases evolve, creating technical debt that degrades maintainability and causes downstream API misuse. While static analysis tools can detect the absence of documentation, they cannot evaluate its semantic consistency. Conversely, standard Large Language Models (LLMs) offer generative flexibility but frequently hallucinate when updating documentation without deep structural awareness of the underlying code. To address this gap, we propose DocSync, an agentic workflow that frames documentation maintenance as a structurally grounded, iterative generation task. DocSync bridges syntactic changes and natural language descriptions by fusing Abstract Syntax Tree (AST) representations and Retrieval-Augmented Generation (RAG) to provide dependency-aware context. Furthermore, to ensure factual consistency, we incorporate a critic-guided refinement loop based on the Reflexion paradigm, allowing the model to self-correct candidate updates against the source code. We empirically evaluate a resource-constrained implementation of DocSync-using a LoRA-adapted small language model - on a proxy code-to-text maintenance task. Our findings demonstrate that this AST-aware agentic approach substantially outperforms standard encoder-decoder baselines across semantic alignment, summary-line faithfulness, and automated judge preferences (e.g., achieving an automated judge score of 3.44/5.0 compared to 1.91 for CodeT5-base). Crucially, the iterative critic loop yields measurable improvements in semantic correctness without requiring scaled-up parameter counts. These results provide strong evidence that coupling structural retrieval with agentic refinement is a highly promising direction for autonomously mitigating documentation debt.

11:00 ***Distributed Platform Architecture and API-Led Integration***

Lakshmi Sujatha Chilamkurthi (Cigna, USA); Venkata Pavan Kumar Gummadi (Broadridge, USA); Srikanth Kavuri (Iconsoft, USA)

The modernization of legacy claims platforms is critical for payers that must sustain high-volume adjudication while integrating with an expanding ecosystem of vendors and delegated entities. This paper presents a distributed claim-centric platform that decomposes enrollment, eligibility, claims intake, adjudication orchestration, and payments into independently deployable microservices running on a container-based runtime. Container orchestration is used to scale claim-adjacent services, such as EDI ingestion, provider enrichment, and remittance generation, horizontally under peak loads while maintaining predictable performance and resilience. Long-running and exception-driven activities, including pends, manual reviews, and appeal cycles, are offloaded to workflow engines and asynchronous orchestration layers, ensuring that the core claims engine remains optimized for high-throughput adjudication. In addition, the platform publishes claim and accumulator events in near real time, enabling partner-facing services to consume, reconcile, and maintain synchronized benefit and cost-share views across third-party vendors and delegated entities. Empirical migration outcomes-covering scalability, failure isolation, processing latency, and reconciliation accuracy over comparable production windows-demonstrate that the proposed architecture delivers measurable improvements over a monolithic claim engine deployment, while aligning with emerging microservices migration and event-driven design practices.

Thursday, April 23 9:00 - 11:30 (Africa/Cairo)

SESSION-7E: RESEARCH: Machine Learning for Business Intelligence, Finance & Fraud Detection

10 PAPERS

Room: ROOM-E

Chair: Salah A. Aly (Fayoum University, Egypt)

9:00 ***Behavioral Customer Segmentation for Fixed Broadband: A Machine Learning Approach***

Mohamed Tharwat (ESLSCA University, Egypt); Eshraq Saeed Awwad (Vodafone Egypt, Egypt)

In the competitive telecommunications landscape, understanding customer behavior and preferences is critical for service differentiation and revenue optimization. This paper presents a comprehensive machine learning framework for large-scale customer segmentation of ADSL broadband services, analyzing a smaple over 700 K subscribers across more than 24 behavioral and transactional features. We employ a robust preprocessing pipeline incorporating outlier capping, Yeo-Johnson transformation, and Principal Component Analysis (PCA) for dimensionality reduction, followed by K-means clustering to identify distinct customer segments. Our methodology successfully identifies seven unique customer clusters with varying usage patterns, service quality experiences, and value-added service adoption rates. The segmentation reveals actionable insights including premium users consuming higher average data volumes, highly engaged multi-service users, and at-risk segments requiring targeted retention strategies. This work demonstrates the practical application of machine learning to real-world telecommunications data at scale, providing a foundation for personalized marketing, churn prevention, and service optimization strategies that directly impact business profitability and customer satisfaction.

9:15 ***Defensive Intelligent Traffic Management using Secure AI Intersection Control, Resilient Sensor Fusion, and Crisis-Aware Autonomous Routing***

Wael Badawy (Egyptain Russian University, Egypt)

This paper introduces Defensive Intelligent Traffic Management (DITM), a defense-grade architectural paradigm designed to ensure secure AI intersection control, resilient multisensor fusion, post-quantum identity assurance for all infrastructure nodes, and crisis-aware autonomous routing capable of preserving emergency mobility and critical logistics under sustained cyber-physical stress. The framework integrates tamper-evident control flows, neurosymbolic law-compliance gates, federated edge learning, and sovereign supervisory authority to guarantee that autonomous transport functions remain lawful, auditable, and aligned with public-safety doctrine. A metropolitan-scale digital-twin simulation encompassing 320 intersections demonstrates improvements in emergency-vehicle travel time, spoofing resistance, congestion-collapse prevention, and continuity of routing decisions under adversarial conditions. The results establish a deployable blueprint for national-scale, security-centric, and ethically aligned autonomous traffic governance. DITM reorients intelligent mobility from algorithmic efficiency toward secure, mission-critical resilience, forming the foundation for sovereign AI-driven transport infrastructure.

9:30 ***(RDSF-SSL) Robust Dynamic graph and Soft-label Fusion for Semi-Supervised Learning***

Abdullah Baradaaji (Lebanese International University, Lebanon); Fadi Dornaika (University of the Basque Country, Spain); Taha Houda (Prince Mohammad Bin Fahd University, Saudi Arabia); Jihad Jaam (Liverpool J. Mores University, United Kingdom (Great Britain))

Semi-supervised learning is an effective solution for image classification problems where only a small portion of the training data is labeled. Among existing approaches, graph-based methods are particularly attractive because they can exploit the geometric structure of both labeled and unlabeled samples. However, many existing models rely on fixed similarity graphs, which may be sensitive to noise and may not accurately reflect the evolving relationships revealed during learning. In this paper, we propose RDSF-SSL, a dynamic graph and soft-label auto-weighting framework for semi-supervised subspace learning. The proposed method jointly estimates a discriminative latent subspace, soft labels for unlabeled samples, and an adaptive graph structure within a unified optimization framework. In addition, it introduces an auto-weighted fusion mechanism that combines a data graph and a label graph to regularize both label smoothness and projected-data smoothness. A low-rank constraint is further imposed on the learned graph to improve robustness. Experiments conducted on five benchmark image datasets show that RDSF-SSL achieves competitive and often superior recognition performance compared with several state-of-the-art semi-supervised methods, especially when the number of labeled samples per class is very limited. These results demonstrate the effectiveness of dynamically coupling graph learning, soft-label estimation, and subspace modeling for robust semi-supervised image classification.

9:45 ***Efficient Voice Identification System Utilizing Wav2Vec2.0 and HuBERT Based on the Quran Reciters Dataset***

Aly Moustafa (faculty of Computers and Information, Helwan University); Salah A. Aly (Fayoum University, Egypt)

Current authentication and trusted systems depend on classical and biometric methods to recognize or authorize users. Such methods include audio speech recognitions, eye, and finger signatures. Recent tools utilize deep learning and transformers to achieve better results. In this paper, we develop a deep learning constructed model for Arabic speakers' identification by using wav2vec and HuBERT audio representation learning tools. The end-to-end Wav2vec2.0 paradigm acquires contextualized speech representations learning's by randomly masking a set of feature vectors, and then applies a

transformer neural network. Methods: We employ MLP classifier that is able to differentiate between invariant labeled classes. We show several experimental results that safeguard the high accuracy of the proposed model. Results: The experiments ensure that an arbitrary wave signal for a certain speaker can be identified with 98% and 97.1% accuracies in cases of wav2vec2.0 and HuBERT, respectively.

10:00 ***Improving Safety Deep Learning-Based Vulnerability Detection***

Samir Haddad (University of Balamand, Lebanon); Kassem Hamze (Islamic University of Lebanon, Lebanon); Jinane Sayah (University of Balamand, Lebanon); Rafi Srayoui (Lebanese University, Lebanon); Joseph Merhej (Facultyof Sciences II, Lebanon); Chadi Kallab (Lebanese American University, Lebanon)

Malicious software-including adware, spyware, and computer viruses-continues to present significant security risks to private citizens, corporations, and state entities, such as government agencies and military organizations. As cyberattacks grow increasingly sophisticated and varied, traditional defense mechanisms often fail to provide adequate protection. Conse quently, there is an urgent requirement for advanced and efficient methodologies capable of identifying system vulnerabilities before they are exploited. To enhance the effectiveness and reliability of vulnerability detection, this research integrates modern artificial intelligence techniques to develop a dedicated wireless-network penetration-testing system. By employing machine-learning algo rithms, the proposed framework aims to improve the precision, scalability, and flexibility of vulnerability identification within complex network environments. The BoTNeTIoT-L01 dataset, which contains over seven million IoT botnet attack records, was utilized for both training and validation purposes. This study implements a Convolutional Neural Network architecture using the Keras library, incorporating convolutional layers, max pooling, and dense layers. Furthermore, the Adam optimization algorithm was adopted to refine the model's training process. The proposed CNN achieved a classification and detection accuracy of 99.46%, demonstrating a powerful emerging capability for identifying cyber threats. These findings suggest that leverag ing AI to bolster cybersecurity measures significantly enhances protection for increasingly dynamic and interconnected wireless network infrastructures.

10:15 ***AI-Driven Detection of Misclassified Imports: A Machine Learning and LangChain-Based RAG Approach for Customs Fraud***

Adnane El Amrani, Aya Ismaili and Yousra Chtouki (Al Akhawayn University, Morocco); Hind Lamharhar (High-Tech, Morocco)

Customs fraud through deliberate product misclassification causes substantial revenue losses for governments worldwide. This paper presents a hybrid detection system combining XGBoost classifiers with LangChain-based Retrieval-Augmented Generation (RAG) to identify and explain potentially fraudulent declarations. We train on one million synthetic declarations and validate on a partially anonymized sample of 5,000 real Moroccan customs records. On synthetic test data, XGBoost achieves an F1-score of 0.811 for fraud detection, significantly outperforming Random Forest (p less than 0.001), while the RAG pipeline attains 90.2 percent exact-match accuracy for HS code prediction. On real data, the system maintains 82.3 percent precision at 67.4 percent recall. Ablation studies confirm that the hybrid approach outperforms both pure ML and pure LLM baselines by 8-15 percent in F1-score. The explanation agent generates regulation-referenced justifications that domain experts rated as relevant and actionable in 84 percent of cases.

10:30 ***Cooperative Evolutionary Hyper-Heuristics with Informed Energy Constraints for Fair and Robust Spiking Neural Networks***

Ali Sajwani and Ayad Turky (University of Sharjah, United Arab Emirates); Nasser Saber (La Trobe University, Australia)

Spiking Neural Networks (SNNs) have emerged as energy-efficient alternatives to traditional artificial neural networks, particularly when deployed on neuromorphic hardware. However, optimizing SNNs for highly imbalanced datasets under strict operational constraints remains challenging due to the complex interplay between architectural parameters and neuronal dynamics. This paper presents a novel cooperative evolutionary hyper-heuristic framework that simultaneously optimizes architectural parameters (batch size, number of epochs, time steps) and neuronal dynamics parameters (membrane decay rates, firing thresholds, surrogate gradient slope) through dual population co-evolution. The framework employs a fast evolution strategy using proxy data subsets, quantized caching and parallel evaluation to reduce computational overhead while preserving solution quality. Experiments on the Bank Account Fraud dataset demonstrate that our approach achieves strong recall and false positive rate performance with high accuracy that outperforming state-of-the-art SNNs and gradient methods.

10:45 ***A Hybrid Asset Scoring Framework Combining Deterministic Quantitative Screening with RAG-Informed LLM Qualitative Evaluation for Portfolio Universe Reduction***

Joude Azzam, Saber Elsayed and Saad Harous (University of Sharjah, United Arab Emirates)

Portfolio optimization becomes effective depending on the characteristics of the universe of investable assets. Nonetheless, most methods used today narrow down this universe using only backwards looking quantitative filters. In this study, a two-layer asset scoring pipeline is developed which combines deterministic quantitative scoring with RAG-enhanced Large Language Model (LLM). Such pipeline returns investability scores per asset considering the income oriented investor's profile. In the first layer of quantitative scoring, profile-specific composites are calculated according to empirically proven thresholds from Altman Z-Score, Piotroski F-Score, Novy-Marx gross profitability, Sloan's accruals, interest coverage, and momentum indicators. Before scoring, deterministic hard disqualification is applied using criteria based on Altman's and Beneish's M-Score models. Quantitative scores returned from the first step are fed as readonly context to the LLM mitigating financial figure hallucination risk. As for the second layer, Meta's Llama-3.1-8B-Instruct 4- bit NF4 quantized model is responsible only for qualitative scoring which includes assessing strength of Competitive Stability, Risk Signal, and management execution quality. This qualitative data comes as up-to-date news and analysts' insights collected from Finnhub service, and stored in FAISS-based vector store. Final investability scores are obtained by combining quantitative and qualitative composites using income-profile-specific weights, producing ranked and tagged outputs suitable for downstream optimization. Evaluation on 100 firms from the S&P 500, based on an income investor profile, shows that the hard disqualifier filter removes 47 percent of firms prior to any LLM application. Bidirectional scoring changes introduced through RAG lead to different classifications for 12 firms compared to the vanilla LLM setup. Furthermore, the low Spearman correlation coefficient between quantitative and qualitative composites indicates that the two layers capture largely independent dimensions of investability.

11:00 ***AI-Powered Financial Fraud Detection Using Machine Learning: A Multi-Algorithm Approach***

Ekereuke Udoh (QA Higher Education, United Kingdom (Great Britain))

Financial fraud threatens global economic stability, leading to over $5 trillion in annual losses. Traditional detection systems have high false positive rates and struggle with evolving fraud tactics. This study assesses seven machine learning algorithms using the imbalanced Kaggle Credit Card Fraud Dataset, which has a fraud rate of 0.173%. Under two experimental conditions-training on original and SMOTEgenerated data-K-Nearest Neighbours (KNN) achieved the best performance with an F1-score of 0.888 and accuracy of 99.96%. Random Forest and XGBoost also performed well, with F1-

scores of 0.857 and 0.832, respectively. However, the use of SMOTE decreased performance across all models, challenging its effectiveness for extreme imbalance scenarios. The study recommends KNN, Random Forest, and XGBoost for fraud detection, advocating against SMOTE in favour of alternatives like cost-sensitive learning for better results in imbalanced datasets.

11:15 ***Quantifying Attribute-Level Elliptic Curve Cryptography Encryption Costs with Decision Modelling for Data-Intensive Applications***

Anxhela Baraj (University of Rijeka, Croatia); Jonatan Lerga (University of Rijeka -Tehnički Fakultet, Croatia)

Secure data processing in service-oriented and data-intensive applications requires balancing confidentiality with operational efficiency. Elliptic-curve cryptography (ECC) is widely adopted for its security and computational advantages, yet prior work mostly reports system-level overhead, overlooking the impact of individual data attributes. This study quantifies ECC encryption cost at the attribute level by isolating scalar multiplication across numeric and categorical fields with varying domain sizes and encoded value lengths. Controlled microbenchmarking captures both mean execution time and variability, showing that numeric attributes with longer representations incur higher and more variable costs, while domain size alone does not predict overhead. Building on these results, a lightweight selective encryption decision model is proposed to classify attributes for default, selective, or batched encryption. The model translates attribute-level measurements into actionable guidance, providing a quantitative foundation for fine-grained, performance-aware encryption strategies in data-intensive systems.