Proceedings and Schedule of

5th IEEE International Conference on AI in Cybersecurity (ICAIC)

18 - 20 February 2026

University of Houston, 4800 Calhoun Rd, Houston, TX 77004

Proceeding Editors:

Hardik Gohel, Bishwajeet Pandey

Chair Message

On behalf of the organizing committee, it is our distinct honor and privilege to warmly welcome you to the 5th **International Conference on AI in Cybersecurity (ICAIC-2026)**. The conference will be held in **Hybrid Mode** from **February 18-20, 2026**, allowing participants to join us either in person in **Houston, USA**, or virtually from around the globe. ICAIC-2026 is dedicated to fostering innovative research and technical advancements in the fields of **artificial intelligence, cybersecurity, and emerging technologies**. This year, we are delighted to announce that we received an overwhelming response, with numerous high-quality submissions that showcase the depth and breadth of ongoing research in these domains. All accepted papers will be submitted to **IEEE Xplore**, and we aim for them to be indexed in **Scopus** and **DBLP**, further amplifying their global reach.

As conference chair, I am thrilled to host a diverse community of researchers, professionals, and thought leaders who are coming together to share their expertise and collaborate on the latest trends and challenges in AI and cybersecurity. We are committed to ensuring your experience in Houston is both memorable and enriching. We are confident that our carefully curated lineup of keynote speeches, technical sessions, and networking opportunities will provide invaluable insights and foster meaningful connections. Notably, ICAIC® 2026 marks the **35th conference hosted by Gyancity Research Consultancy**, in collaboration with our esteemed partner universities worldwide. We deeply value the trust and enthusiasm of our participants and look forward to making this event a rewarding experience for all. Next conferences in 2026 are following:

11th International Conference on Green Computing and Engineering Technologies (ICGCET ®)
01-02 April 2026 Liberty Central Saigon Riverside Hotel, Ho Chi Minh City, Vietnam
https://icgcet.org/

2026 IEEE Conference on Generative AI for Secure Systems (GAISS)
28-30 October 2026 Austin Texas USA https://gaiss.info/
https://conferences.ieee.org/conferences_events/conferences/conferencedetails/66401

Best wishes.

Prof Hardik Gohel, University of Houston, USA
Prof Bishwajeet Pandey, GL Institute of Technology and Management, India
Tel/Whatsapp: +1-786-376-5284, +91-74-28-640-820,

Email: gohelh@uhv.edu, dr.pandey@ieee.org

ICAIC-2026 AGENDA

18 February 2026 (OFFLINE (on-site))

Date: February 18, 2026 **Location:** University of Houston Student Center South, 4455 University Drive, Houston, TX 77204		
Time	**Event**	**Location**
8am to 8:30am	Breakfast	SC South Midtown
8:30am to 8:50am	Conference Inauguration	SC South Midtown
8:3am to 8:50am	Break for Breakout Session	
9:00am to 10am	**Keynote Speech by** 1. Mr. Amit Kumar Padhy, Adobe, USA 2. Dr. Siddhant Sharma, Holland & Hart LLP, USA 3. Mr. Tejas Pravinbhai Patel, Amazon, USA 4. Mr. Subhash Bondhala, Independent Researcher, United States	SC South Midtown
9:00am to 10am	**Keynote Speech by** 1. Mr. Sundeep Bobba, Southwest Airlines, USA 2. Mr. Ravi Sankar Susarla, Collabera, USA 3. Mr. Gopi Krishna Pamula, Mitchell Martin Inc, USA 4. Atish Kumar Dash Solutions Consultant, ADVANCE Solutions Corp, USA	SC South Astrodome
10am to 10:15am	**Break**	

10:15am to 12:30pm	Parallel presentations – track 1	SC South Midtown
10:15am to 12:30pm	Parallel presentations – track 2	SC South Astrodome
12:30pm to 2pm	**Lunch**	**Moody towers**
2pm to 3:45pm	Parallel presentations – track 3	SC South Midtown
2pm to 3:45pm	Parallel presentations – track 4	SC South Astrodome
3:45pm to 4:15pm	Award Ceremony of Day 1	SC South Midtown
5:00pm to 6:00pm	Networking Reception	TBA

Registered Listener:
Ms Vipra Gupta
Mr Ryan Perry
Mr Troy Dixon
Mr Sathish Pandurangan
Mr Edward French
Mr Komal Subhash More
Mr Hirokazu Hasegawa
Mr Masahito Kumazaki
Mr Masahito Kumazaki

Track 1 Session Chair:
Dr. Prasanna Ranjith Christodoss, Messiah University, USA

Paper Id	Paper Title	Presenter	
630	HCT-Net: A hybrid convolutional and transformer based network for corrosion detection in industrial images	Mohini Gohel North Carolina Agricultural and Technical State University, USA	mohinigohel@gmail.com
164	SC-GAN: A GAN-Based Data Augmentation Approach for Stablecoin Fraud Detection on Imbalanced Transaction Data	Mohan Sankaran, PayPal, USA	mohansankaran@ieee.org
342	Enhancing Trust in AI: Addressing Vulnerabilities and Ensuring Privacy with ML Security Protocols Across the Lifecycle	Ramkinker Singh Carnegie Mellon University Pittsburgh, USA	ramkinks@alumni.cmu.edu
177	Evaluating Jailbreak Vulnerabilities in LLMs: A Taxonomy and Comparative Analysis in Romance Fraud Scenarios	Parra Bautista, Yohn J,	yohn.parrabautista@famu.edu
193	Application of Federated Learning to Semantic Segmentation of Aerial Images	Yong-Lin Kuo, Ying-Wei Chuang National Taiwan University of Science and Technology, Taiwan	yl_kuo@yahoo.com
446	Hierarchical Attention Distillation for Real-Time Cyber Threat Detection and Mitigation in Large-Scale Networks	Ramkinker Singh Carnegie Mellon University Pittsburgh, USA	ramkinks@alumni.cmu.edu
499	Privacy-Preserving Federated Learning for Multi-Tenant CRM Systems	Nidhi Sharma Independent Researcher, USA	nidhi.sharmatechlead@gmail.com
610	Automated SIEM Detection Rule Translation System	Adelia Ibragimova, Epam Systems, USA,	adelina.30stm@inbox.ru

Track 2 Session Chair:

- **Mr. Shiva Krishna Kodithyala, Bread Financial, USA**

Paper Id	Paper Title	Presenter	
217	Reengineering Cybersecurity Processes with Generative AI: From Automation to Strategic Alignment	Mehrdad Sharbaf IEEE, USA	msharbaf@ieee.org
220	Identity Governance in DevSecOps: Automated Access Reviews for CI/CD Pipelines	Sunnykumar Kamani ,	sunnykumar.kamani @gmail.com
234	Combined Tempered MRG32k3a: A High-Quality and Reproducible Pseudo-Random Number Generator for AI in Cybersecurity	Hussein Alzoubi , German Jordanian University (GJU) Amman, Jordan	hussein.alzoubi@gju.edu.jo
251	GraphAE: Plant-informed Graph Autoencoder for ICS Anomaly Detection with SHAP-based Explanations	Vahid Heydari, Morgan State University, United States	vahid.heydari@morgan.edu
255	KyVul, an LLM Created C/C++ Vulnerability Dataset	Martin Carlisle Texas A&M University	carlislem@tamu.edu
284	Prompt Engineering vs Context Engineering: A Strategic Framework for Insurance AI Applications	Rakesh More A J Gallagher, USA	rakeshmore@gmail.com
287	Evaluating Adversarial Resilience of Deep Reinforcement Learning Algorithms for Network Intrusion Detection	Curtis Rookard Indian River State College, United States	crookard@irsc.edu
542	Quantifying Risk Reduction: AI-Driven Model for Automating Access Certification and Segregation of Duties (SOD) Controls	Sunnykumar Kamani	sunnykumar.kamani @gmail.com,

Track 3 Session Chair:

Ms Nidhi Sharma , Second Vice Chair, IEEE Long Island Section, NY, USA

Paper Id	Paper Title	Presenter
330	Assessing Agile Frameworks for Software Development Efficiency Industry Insights and Implementation with Jira Atlassian Agile Tools	Ravi Sankar Susarla Collabera, USA raviangirasa@ieee.org
339	Confidential and Attack-Resilient Edge LLM Serving for Multilingual Chatbots	Satya Karteek Gudipati Peritus Inc sskmaestro@gmail.com
343	Interval-Based Estimation of Generalization Accuracy in Supervised Learning via Confusion Matrix Resampling and Bayesian Inference	Amir Liron Texas State University, USA amir_liron@txstate.edu
345	AI-Driven Prediction of Flexural Properties in Kevlar/Carbon/Glass Fiber Hybrid Composites Using Random Forest Regression	**Sathish Pandurangan, Prasanna Ranjith Christodoss,**
346	Optimizing Customer Engagement through IoT Data Integration in CRM Ecosystems	Sathish Kumar Velayudam, Independent Researcher, USA sathish.velayudam@ieee.org
350	Conditional Password Generation Using a FiLM-Enhanced WGAN: A Controlled Comparison Against Standard GAN Baselines	Pranav Shetty, Sahyadri College of Engineering and Management, India studytimemail24@gmail.com
361	CityCopilot X: A Real Time Explainability Panel for Retrieval Augmented Generation (RAG)	SATYA KARTEEK GUDIPATI Publicis Groupe, USA sskmaestro@gmail.com
290	Resource Allocation of IoT-Based Edge Computing in Smart Cities	Abdullah Alqahtani Jazan University, Saudi Arabia, amqahtani@jazanu.edu.sa

Track 4 Chair: Mr. Linga Reddy Alva, IT Spin Inc, USA and Mr Sunnykumar Kamani, SoftSages Technology, USA

Paper Id	Paper Title	Presenter
461	Adaptive Loyalty Systems: A Reinforcement Learning Framework for Dynamic and Context-Aware Benefit Allocation	Tarun Kalwani, Independent Researcher Atlanta, GA USA tarun.kalwani17@gmail.com
472	Real-World Intrusion-Aware Zero Trust Architecture: An AI-Driven ASPM Framework Using CICIDS-2017 Network Attack Traffic	Subhash Bondhala Independent Researcher, United States subhashbondhala@gmail.com
381	Reliable Extraction of Cyber-Physical System Threat Knowledge with Multi-Run LLMs	Lara Habashy Mission Critical Cybersecurity DRDC – Valcartier Research Centre Quebec, Canada lara.habashy@drdc-rddc.gc.ca
386	A Survey on Machine Learning Applications for Operating System Fingerprinting	Siri Siqveland Embry-Riddle Aeronautical University, USA SIQVELAS@my.erau.edu
369	ChatXplain: Interpretable Explanations for Intelligent Assistants via Modular Rationale and Saliency	Naveen Anand Mishra Cigna, USA naveenmishra5@gmail.com
372	A Systematic Security Analysis of Model Context Protocol: Vulnerabilities, Exploits, and Mitigations	Theophilus Siameh Mississippi State University, USA theodondre@gmail.com
375	ENHANCING BREAST CANCER DIAGNOSIS USING LIGHTWEIGHT DEEP LEARNING MODELS	Jordan Mozebo, North Carolina A&T State University, USA jtmozebo@aggies.ncat.edu

19 February 2026 (OFFLINE (on-site))

Date: February 19, 2026 **Location:** University of Houston Student Center South, 4455 University Drive, Houston, TX 77204		
Time	**Event**	**Location**
8am to 8:30am	Breakfast	SC South Downtown
8:30am to 9:15 am	• **Keynote speech by** 1. Mr. Linga Reddy Alva, IT Spin Inc, USA 2. Mr. Shiva Krishna Kodithyala, Bread Financial, USA 3. Mr. Sreenivasulu Gajula, Principal Full-Stack Engineer, USA	SC South Downtown
9:15am to 10:00 am	**Keynote speech by** 1.Mr. Siva Teja Reddy Kandula, Senior Software Developer, USA 2. Mr. Sandeep Shivam, Associate Director, Touchless Lending platform, Tavant, USA 3. Mr. Srikanth Kavuri, Senior Software QA Automation Engineer, USA	SC South Downtown
10am to 10:15am	**Break**	
10:15am to 12:30pm	Presentations – track 5	SC South Downtown
12:30pm to 1:30pm	**Lunch**	**Moody towers**
1:30pm to 2:45pm	Presentations – track 6	SC South Downtown
2:45pm to 3:15pm	Award Ceremony of Day 2	SC South Downtown

3:15pm to 3:30pm	Conference Closing Ceremony	SC South Downtown

Track 5 Session Chair:
Mr. Sundeep Bobba, Southwest Airlines, USA

Paper Id	Paper Title	Presenter
326	Advancing Cybersecurity in Critical Infrastructure Systems via Machine Learning-Based Threat Detection and Mitigation	Siva Teja Reddy Kandula, Independent Researcher, sivateja.kandula@ieee.org
506	An Empirical Evaluation of Deep Neural Networks as Hash-Like Mappings Under Digital Signature Threat Models	Lynn Vonderhaar, Embry-Riddle Aeronautical University, USA vonderhl@my.erau.edu
510	A Self-Adaptive Red Teaming Framework for Vulnerability Profiling with Dynamic Attack Graphs and Retrieval Augumented Generation	Jothsna Praveena Pendyala Independent Researcher, Allen, TX, USA jothsnapraveena1421@gmail.com
529	FraudSentinel: Federated Multi-Agent Reinforcement Learning for Privacy-Preserving Cross-Marketplace Fraud Detection in Distributed E-Commerce Ecosystems	Tejas Pravinbhai Patel, Independent Researcher, IEEE, USA tejas.patel@ieee.org
533	TrustGraph: Federated Graph Neural Networks for Cross-Platform Trust and Fraud Propagation Analysis	Tejas Pravinbhai Patel, Independent Researcher, IEEE, USA tejas.patel@ieee.org
534	AgentSCO: A Multi-Layer Agentic Framework for Security Operations Automation	Joyjit Roy, KForce Inc, USA, joyjit.roy.tech@gmail.com
550	Zero-Shot Tokenizer Transfer for Targeted Attacks on Retrieval-Augmented Generation	Mark Spanier, Dakota State University, USA mark.spanier@dsu.edu
490	An Explainable Machine Learning Framework for Predicting Software Defects in Large-Scale Software Systems	Srikanth Kavuri SCDHHS, USA srikanthkavuri.research@gmail.com

566	Accelerating AI Maturity: A Framework for Reducing Development Lifecycles and Increasing Predictive Accuracy	Gopi Krishna Pamula Mitchell Martin Inc, USA gopipamula@gmail.com
567	Scalable Data Governance Through Engineering-Driven Quality and Consistency Controls	Gopi Krishna Pamula Mitchell Martin Inc, USA gopipamula@gmail.com
568	Anomaly Detection in Financial Payment Transactions Using Efficient Data-Driven Machine Learning Techniques	Sandeep Shivam, Tavant, USA, sandeep.shivam@ieee.org,
569	Securing the LLM Supply Chain: Analyzing Threats and Mitigation Strategies	Atish Kumar Dash ADVANCE Solutions Corp., USA atish.dash.7@gmail.com
573	The Adaptive AI SOC Agent – Moving Beyond Linear Playbooks with Cognitive Reasoning	Valentin Chichurov EPAM, USA valentinchichurov@gmail.com
575	Integrating AI into DevSecOps Pipelines for Secure Cloud Microservices Management	Atish Kumar Dash ADVANCE Solutions Corp., USA atish.dash.7@gmail.com

Track 6 Session Chair:
Mr. Tejas Pravinbhai Patel, Amazon, USA

PId	Paper Title	Presenter
368	AI-Assisted Detection of Malicious Changes in Infrastructure-as-Code for Secure DevOps	Shalini Sudarsan, DevOps Engineering Manager KinderCare Learning Companies Oregon, USA shallene.s@gmail.com
611	Digital Precognition: Teaching Transformers to Map Cyber Threats Before Analysts Can	Shane Waldrop Angelo State University, USA shanewaldrop123@gmail.com
621	Semantic Chunking for Triple Extraction from Cyber Threat Intelligence Reports	Shane Waldrop Angelo State University, USA shanewaldrop123@gmail.com
628	SARS-CoV-2 Classification Using Classical vs. Quantum Machine Learning: A Performance Comparison of SVM, QSVM, and Pegasos	Sthefanie J. G. Passo, The University of Texas at San Antonio, USA sthefanie.passo@utsa.edu
629	Beyond Single-Hop:Link Prediction Through Multi-Hop Reasoning In Malware Knowledge Graphs	Kibrom Bahlibi Angelo State University, USA kbahlibi@angelo.edu
634	Bridging the Black Box: Explainable Anomaly Detection for Critical Infrastructure Systems	William Mitchell Angelo State University, USA wmitchell9@angelo.edu
579	Integrating Customer Data Platforms (CDPs) with Experimentation Tools: A Framework for Optimizing Customer Journey Revenue	Gopi Krishna Pamula Mitchell Martin Inc, USA gopipamula@gmail.com
580	Hybrid Transformer and XGBoost Model for Federated IoT Intrusion Detection	Madhu Siddharth Suthagar Illinois Institute of Technology, USA madhusiddharths2@gmail.com
581	Architecting Agentic AI Systems with Multimodal Reasoning for Scalable Visual Pattern Recognition	Linga Reddy Alva IT Spin Inc, USA alvalingareddy@gmail.com
585	Integrating AI and Cloud to Advance Scalable, Secure, and Automated Information Management in Enterprises	Shiva Krishna Kodithyala Bread Financial, USA reachkodithyala@gmail.com

19 February 2026 (ONLINE (on Google Meet))

Meeting link: meet.google.com/nyh-cajt-bbk

Session 1 Chair: Mr. Prashant Vajpayee, Salesforce, USA pvajpayee.researchconnect@gmail.com

8:01-8:15 (TEXAS TIME)

Paper Presenter, Paper Id: Niranjan Pachaiyappan, tptniranjan@gmail.com, 219

Paper Title: Agentic Commerce: A Comprehensive Analysis of Cybersecurity Risks, Privacy Challenges, and Trust Mechanisms in Autonomous AI-Driven Marketplaces

8:16-8:30

Paper Presenter, Paper Id: Rajesh Vayyala, Harrisburg, NC, USA, vayyalarajesh@gmail.com, 470

Paper Title: A Time-Series and Machine Learning Framework for Forecasting GDP and Unemployment Using Global Economic Indicators (2020–2024)

8:31-8:45

Paper Presenter, Paper Id: Komaldeep Singh, Chandigarh University, India kdsguraya26@gmail.com, 210

Paper Title: Multilingual Image-to-Speech Conversion: Leveraging OCR and Language Translation for Enhanced Accessibility

8:46-9:00

Paper Presenter, Paper Id: Henry Cyril, T-Mobile, USA, henry.cyril.tech@gmail.com, 457

Paper Title: DeepNetDetect: A Deep Learning-Based Approach for Early Anomaly Detection in Network Traffic

9:01-9:15

Paper Presenter, Paper Id: Vaishali Vinay, Microsoft, USA vaishali.papneja@microsoft.com , **351**

Paper Title: The Evolution of Agentic AI in Cybersecurity: From Single LLM Reasoners to Multi-Agent Systems and Autonomous Pipelines

9:16-9:30

Paper Presenter, Paper Id: Vaishali Vinay, Microsoft, USA vaishali.papneja@microsoft.com, 405

Paper Title: Jailbreaking Large Language Models: Techniques, Trends, Defenses, and Open Challenges

9:31-9:45

Paper Presenter, Paper Id: Michal Jaworski, TU Dublin, michal.jaworski11@gmail.com, 300

Paper Title: Hybrid Intelligence Endpoint Defense (HIED): A Data-Fusion Approach for Proactive Malware Detection

9:46-10:00

Paper Presenter, Paper Id: Henry Cyril, T-Mobile, USA, henry.cyril.tech@gmail.com, 488

Paper Title: DevSecOps-Driven Security Integration in the Software Development Lifecycle Using CI/CD Pipelines

10:01-10:15

Paper Presenter, Paper Id: Srikumar Nayak, Sr Member IEEE, Incedo Inc, USA, Srikumar.nayak2025@gmail.com, 459

Paper Title: FraudGNN: Self-Supervised Graph Neural Anomaly Detection for Real-Time Financial Fraud with Adversarial Robustness and Explainable Reasoning

10:16-10:30
Paper Presenter, Paper Id: Roudha Bin Redha, Zayed University, UAE, 301
Paper Title: Phishing URL Detection Using RNN – LSTM Models for Safer Web Browsing
10:31-10:45
Paper Presenter, Paper Id: Chaitanya Kulkarni, Oracle America Inc, USA, ckulkarni@ieee.org, 554
Paper Title: Cybersecurity Architecture for Cloud Database Infrastructure- A Case Study of Oracle Autonomous Database on Microsoft Azure
10:46-11:00
Paper Presenter, Paper Id: Rupam Priya, Manager - Marketing Analytics, USA, rupampriya.001@gmail.com, 537
Paper Title: The Architect of Advantage: How Robust Data Curation and Edge-Case Analysis Provides a Disproportionate Market Edge
11:01-11:15
Paper Presenter, Paper Id: Vidhubala J, SRM IST, India,vj0154@srmist.edu.in, 382
Paper Title: LightGBM-Based Multi-Class Botnet Detection Framework for IoT Networks
11:16-11:30
Paper Presenter, Paper Id: Hongmei Chi, hongmei.chi@famu.edu , 179
Paper Title: A Framework on Advancing Cybersecurity Education via Quantum Machine Learning

PARALLEL SESSION ON 19 FEBRUARY 2026

Meeting link: https://meet.google.com/hpb-gdhj-njh
Session 2 Chair: Mr Sanjoy Mukherjee, Cognizant, USA sanjoymukherjee302@gmail.com and Mr Anurag Ekkati, Principal Software Engineer, Palo Alto, USA

08:01-8:15
Paper Presenter, Paper Id: Sangharsh Agarwal Antioch, California – 94531, USA sangharshcs@gmail.com, 338
Paper Title: Causal-Contrastive Graph Neural Networks for Robust, Explainable Multi-Modal Intrusion Detection
08:16-8:30
Paper Presenter, Paper Id: Chandrashekhar Medicherla. chandrashekhar.medicherla@ieee.org, 633
Paper Title: CodeGraph- Malware Detection via Control Flow Graph Embeddings and Graph Neural Networks
08:31-8:45
Paper Presenter, Paper Id: Venkaiah Chirumavilla, DIGITAL SCRIPTS INC, USA, 556
Paper Title: Breaking the Monolith: A Data-First Strategy for Enterprise Microservices Migration with Measured Improvements in Scalability and Resilience
08:46-9:00
Paper Presenter, Paper Id: Venkaiah Chirumavilla, DIGITAL SCRIPTS INC, USA, 561
Paper Title: Strategic Unification: How Blazor is Redefining Enterprise Web & AI Stack
09:01-9:15
Paper Presenter, Paper Id: Manisha Guduri, Lawrence Technological University, USA, United States, manishaguduri@ieee.org, 178
Paper Title: Heart Disease Prediction Using Feature Reconstruction and Hybrid Deep Learning
09:16-9:30
Paper Presenter, Paper Id: Snehal Mehta, Wal Mart Associates Inc, USA snehal.mehta203@gmail.com, 500
Paper Title: AI-Driven experimentation for user experience: An Architectural Blueprint for Self-Optimizing iOS Experiences at Enterprise Scale
09:31-9:45
Paper Presenter, Paper Id: Jothsna Praveena Pendyala, Clark University, USA jothsnapraveena1421@gmail.com, 510
Paper Title: A Self-Adaptive Red Teaming Framework for Vulnerability Profiling with Dynamic Attack Graphs and Retrieval Augumented Generation
09:46-10:00
Paper Presenter, Paper Id: Mohith Reddy Patlolla Cyma Systems Inc, USA, mohithrpatlolla@gmail.com, 483
Paper Title: Intelligent Repair Bots for Power BI: An Agentic Automation Approach
10:01-10:15
Paper Presenter, Paper Id: Mohith Reddy Patlolla Cyma Systems Inc, USA, mohithrpatlolla@gmail.com, 484
Paper Title: Integrating Deep Learning with Power BI Admin APIs for Intelligent Data Governance

10:16-10:30
Paper Presenter, Paper Id: Mohith Reddy Patlolla Cyma Systems Inc, USA, mohithrpatlolla@gmail.com, 485
Paper Title: Agentic Metadata Cataloging for Power BI via LLM Scanners
10:31-10:45
Paper Presenter, Paper Id: Jigar Solanki, INCEDO INC, USA, jigarmahendrabhaisolanki@gmail.com , 546
Paper Title: Machine Learning Integration in Loan Decision Systems: Enhancing Kafka Based Workflows with Predictive Analytics
10:46-11:00
Paper Presenter, Paper Id: Jigar Solanki, INCEDO INC, USA, jigarmahendrabhaisolanki@gmail.com, 586
Paper Title: AI-Driven Multi-Cloud Strategies for Financial Services: Ensuring High Availability with AWS and Kubernetes
11:16-11:30
Paper Presenter, Paper Id: Shiva Kumara, T-Mobile, USA, reachkumaras@gmail.com, 438
Paper Title: A Lightweight Deep Learning Based Classification Models for Non-Human Identity Threat Detection
11:31-11:46
Paper Presenter, Paper Id: Varun Pandey, Independent Researcher, USA pandey.varun.087@gmail.com, 608
Paper Title: Secure and Compliant AI/ML-Based KYC: A Cybersecurity-Aware Architecture for Regulatory Identity Verification
11;46-12:00
Paper Presenter, Paper Id: Manoj Kumar, New York University, United States, **236**
Paper Title: AI-Driven DDoS Detection for Network Security: A Performance Analysis of Machine-Deep Learning Methods on Network Traffic Data

20 February 2026 (ONLINE (on Google Meet))

8:00-11:15 (TEXAS TIME)
Session 3:
Chair: Pranjal Sharma, Oracle, USA
Meeting Link: https://meet.google.com/reh-jwco-sfw

PAPER PRESENTER AND PAPER TITLE

8:01-8:15
Paper Presenter, Paper Id: Gaurav Dwivedi, Florida International University, Miami, FL, USA, gdwiv001@fiu.edu, 265
Paper Title: Detecting and Adapting to Normality Shifts in Learning-Based Security Anomaly Detection

8:16-8:30
Paper Presenter, Paper Id: Viswapriyan Ragupathy, Dr, Lewis Center, Ohio, USA viswapriyan.ragupathy@ieee.org, 296
Paper Title: Secure Agent-Based Architectures for Decentralized AI Identity Management: The DAIS Framework

8:30-8:45
Paper Presenter, Paper Id: Viswapriyan Ragupathy, Dr, Lewis Center, Ohio, USA viswapriyan.ragupathy@ieee.org, 297
Paper Title: Architectural Resilience in AI-Driven Decision Systems under Adversarial Conditions

8:46-9:00
Paper Presenter, Paper Id: Ankit Gupta, Shilpi Mittal, 145
Paper Title: AI-Driven Quantum Cryptography

9:01-9:15
Paper Presenter, Paper Id: Shilpi Mittal, Ankita Gupta, mittalshilpi@hotmail.com, 175
Paper Title: Neuromorphic Architectures for Infrastructure-Scale AI: Design, Equations, and Synthetic Benchmarks

9:16-9:30
Paper Presenter, Paper Id: Sanjoy Mukherjee, Cognizant, USA sanjoymukherjee302@gmail.com , 519
Paper Title: Early Detection and Prediction of Depression Based on Data-Driven Machine Learning Techniques in Mental Healthcare

9:31-945
Paper Presenter, Paper Id: Joseph Childress, jrc093@latech.edu,143
Paper Title: A Review of User Account Anomaly Detection and Insider Threat Detection Techniques

9:46-10:00
Paper Presenter, Paper Id: Mohanakrishnan Hariharan, Apple Inc. Austin, TX, USA, m_hariharan@apple.com, 373
Paper Title: Semantic Mastery: Enhancing LLMs with Advanced Natural Language Understanding

10:01-10:15
Paper Presenter, Paper Id: Mohanakrishnan Hariharan, Apple Inc. Austin, TX, USA, m_hariharan@apple.com, 374
Paper Title: Reinforcement Learning Integrated Agentic RAG for Software Test Cases Authoring

10:16-10:30

Paper Presenter, Paper Id: Sakshyam Ghimire, Minnesota State University, Mankato, USA sakshyam.ghimire@mnsu.edu, 570

Paper Title: AI-Assisted Zero Trust Architecture for Continuous Risk Assessment of Programmable Logic Controllers in Food Processing Infrastructure

10:31-10:45

Paper Presenter, Paper Id: Gajendra Babu Thokala, Independent Researcher, IEEE, USA, 528

Paper Title: Machine Learning–Based Fault Prediction in Large- Scale Distributed Systems

10:46-11:00

Paper Presenter, Paper Id: Md Maruf Hassan Southeast University, Bangladesh ancssf@gmail.com, 599

Paper Title: Identity-Spoofing Prompt Injection: Empirical Analysis of Instruction Leakage in Custom GPTs

11:01-11:15

Paper Presenter, Paper Id: Kateryna Babii, Eleks Inc., USA katrusyabb@gmail.com, 571

Paper Title: AI-Guided Adaptive Compression for Secure and Efficient Web Resource Delivery

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 4:

Session Chair: Shiva Kumara, T-Mobile, USA, reachkumaras@gmail.com,

Meeting Link: https://meet.google.com/kkv-bduq-wdz

8:01-8:15

Paper Presenter: Kateryna Savenko, catherineborykina@gmail.com 609

Paper Title: Client Side AI Based Intent Verification for Defending Against CSRF Attacks in Modern Web Applications

8:16-8:30

Paper Presenter, Paper Id: Saisuman Singamsetty, saisuman.singamsetty@gmail.com, 199

Paper Title: Next-Generation Digital Twin Analytics in Smart Manufacturing using Cross-Layered Edge Cloud Deep Learning

8:31-8:45

Paper Presenter, Paper Id: Vineeth Sai Narajala, 314

Paper Title: Enterprise-Grade Security for the Model Context Protocol (MCP): Frameworks and Mitigation Strategies

8:46-9:00

Paper Presenter, Paper Id: Vineeth Sai Narajala, 315

Paper Title: Agent Name Service (ANS): A Universal Directory for Secure AI Agent Discovery and Interoperability

09:01-09:15

Paper Presenter, Paper Id: Venkatesh Prabu Parthasarathy PROPHECY CONSULTING INC venkateshprabu2003@gmail.com, 322

Paper Title: Agentic AI Integration for Process Automation in MSMEs

09:16-09:30

Paper Presenter, Paper Id: Mitesh Patel, Uday Korat IEEE Senior Member USA mitesh.rf@gmail.com, 304

Paper Title: Swarm Optimization Algorithm-Enhanced Clustering Techniques for Reliable Wireless Sensor Networks Communication

9:31-9:45

Paper Presenter, Paper Id: Manjunath Venkatram, manjunath.venkatram@thoughtdata.com, 148

Paper Title: Unknown Threat Detection using AI, ML and DL methods

9:46-10:00

Paper Presenter, Paper Id: Sanjoy Mukherjee, Cognizant, USA sanjoymukherjee302@gmail.com, 504

Paper Title: Machine Learning-Enabled Classification of Diabetes Using Clinical via Lifestyle Health Data

10:01-10:15

Paper Presenter, Paper Id: Claudia Larramendi-Ferras, Florida International University Miami, USA, clarrame@fiu.edu, 252

Paper Title: IARS: Low-Information Area Recognition for Steganography

10:16-10:30

Paper Presenter, Paper Id: Santosh Kumar Kotakonda, Independent Researcher, USA santoshkumarkotakonda84@gmail.com, 482
Paper Title: Benchmarking Container Orchestration Platforms: A Comparative Analysis of Kubernetes and AWS ECS for Stateful Microservices
10:31-10:45
Paper Presenter, Paper Id: Santosh Kumar Kotakonda, Independent Researcher, USA santoshkumarkotakonda84@gmail.com, 480
Paper Title: A Cloud-Native Framework for the Petabyte-Scale Purge of Regulated Data: Achieving Compliance and Performance in the Financial Domain
10:46-11:00
Paper Presenter, Paper Id: Ritesh Ruparel, CSG, USA ritesh.ruparel@yahoo.com, 584
Paper Title: Prompting for LLM Security and RAG: A Survey from Zero-Shot to Automatic Prompt Optimization (APO) and Prompt-Injection Defenses
11:01-11:15
Paper Presenter, Paper Id: Oleh Polishchuk, Sofo-Group, USA oleg.polischuks@gmail.com , 572
Paper Title: AI-Based Cross-Layer Vulnerability Management for Cloud-Native Systems
11:16-11:30
Paper Presenter, Paper Id: Arbaz Surti Independent Researcher, United States, arbaz.m.surti@gmail.com, 204
Paper Title: Bridging the Gap: Using GPT-4 to Summarize and Explain Static Analysis Warnings at Scale

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 5:

Session Chair: Varun Pandey Independent Researcher, USA
pandey.varun.087@gmail.com

Meeting Link: meet.google.com/vgb-hkfh-fxo

8-8:15

Keynote Address by

Paper Presenter, Paper Id: Sarvapriya Tripathi, Florida International University Miami, FL, USA, strip001@fiu.edu, 266

Paper Title: Designing Effective Quantum Generators: A Comparative Study of Variational Ansätze in Hybrid QGANs

8:16-8:30

Paper Presenter, Paper Id: Madhusudan Bangalore Nagaraja, eSystems Inc, USA madhunagaraja@ieee.org, 617

Paper Title: Interpretable Generative AI for Predictive Project Risk and Success Analytics

8:31-8:45

Paper Presenter, Paper Id: Damodhara Reddy Palavali, Software Engineer, Aubrey, USA damodharapalavali@gmail.com, 618

Paper Title: AI-Driven Claims Adjudication: Optimizing Healthcare Systems with Machine Learning and Deep Neural Networks

8:46-9:00

Paper Presenter, Paper Id: Aisvarya Adeseye University of Turku, Finland aisvarya.a.adeseye@utu.fi, 473

Paper Title: Iterative Verification and Batch Processing for Enhancing Accuracy and Confidence Computation in LLM-Based Phishing Email Detection

9:01-9:15

Paper Presenter, Paper Id: Aisvarya Adeseye University of Turku, Finland aisvarya.a.adeseye@utu.fi, 474

Paper Title

9:16-9:30

Paper Presenter, Paper Id: Aisvarya Adeseye University of Turku, Finland aisvarya.a.adeseye@utu.fi, 619

Paper Title: LLM-Assisted Codebook Development for Cybersecurity Interviews with Enhanced Accuracy and Reduced Hallucination

9:31-9:45

Paper Presenter, Paper Id: Phaneendra Yerra, Bank of America, USA phani.net4@gmail.com, 540

Paper Title: Quantum-Resilient Secure Framework for Agentic LLM Workflows in E-Commerce and FinTech Systems

9:46-10:00

Paper Presenter, Paper Id: Humberto Goncalves, New York Institute of Technology, USA hdasilva@nyit.edu 637

Paper Title: Toward Deployable Disinformation Defense: Benchmarking Lightweight Transformers on FakeNewsNet

10:01-10:15

Paper Presenter, Paper Id: Humberto Goncalves, New York Institute of Technology, USA hdasilva@nyit.edu, 638
Paper Title: Explainable AI for Cloud Intrusion Detection: A User Study of SHAP and LIME in AWS GuardDuty

10:16-10:30

Paper Presenter, Paper Id: Dr. Mousumi Munmun, Metro State University, USA mousumi.munmun@metrostate.edu, 358
Paper Title: Leveraging Artificial Intelligence to Predict Unfunded Loans in Peer-to-Peer Lending Platforms

10:31-10:45

Paper Presenter, Paper Id: Naresh Kshetri, Rochester Institute of Technology, USA kshetrinaresh@gmail.com , 549
Paper Title: dataRLsec: Safety, Security, and Reliability With Robust Offline Reinforcement Learning for DPAs

10:46-11:00

Paper Presenter, Paper Id: Andrew Wheeler, Tennessee Technological University, USA amwheeler43@tntech.edu, 385
Paper Title: Graphene: Leveraging Transformers with Control Flow Modalities for Malware Detection

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 5:

Session Chair: Mr Biky Chouhan, Chandigarh University, India dr.chouhan@ieee.org

Meeting Link: https://meet.google.com/ahw-ckgk-ujz

7:45-8:00

Keynote Speech by Mr. Dharnisha Narasappa Senior Network Architect at Versa Networks, USA

8:01-8:15

Paper Presenter, Paper Id: Jayesh Soni, Florida International University, United States, jsoni@fiu.edu, 274

Paper Title: LLM-Based Decision Making Framework for Autonomous Drone Navigation

8:16-8:30

Paper Presenter, Paper Id: Jayesh Soni, Florida International University, United States, jsoni@fiu.edu, 275

Paper Title: Federated Transformer Model for Water Contamination Detection in Distributed Monitoring Systems

8:31-8:45

Paper Presenter, Paper Id: Vatsal Mavani, CVS Health Inc, USA vatsalkishorbhai.mavani@gmail.com, 541

Paper Title: Codebase Aware Generative Agents for the SDLC: Automating Documentation, Dependency Analysis and Test Generation

8:46-9:00

Paper Presenter, Paper Id: Vatsal Mavani, CVS Health Inc, USA vatsalkishorbhai.mavani@gmail.com, 543

Paper Title: An Autonomous Governance Framework for Generative AI: Real-Time PII Redaction and Compliance in LLM-Driven Data Pipelines

09:01-09:16

Paper Presenter, Paper Id: Xiantian Zhou, California State University, East Bay, USA, 364

Paper Title: Scalable Graph-Based Detection of Fraud Rings in Large-Scale Networks

09:16-9:30

Paper Presenter, Paper Id: Prasanthi Sreekumari, University of Louisiana at Monroe, USA sreekumari@ulm.edu, 555

Paper Title: HopeBot: An AI-Powered Mental Health Chatbot Built to Support College Students

09:31-9:46

Paper Presenter, Paper Id: Harsh Patel, UST Global Inc, USA harshpatelv1009@gmail.com, 624

Paper Title: Optimizing Production Variance and Yield Reporting through Cross-Modular Integration in SAP S/4HANA

9:46-10:00

Paper Presenter, Paper Id: Harsh Patel, UST Global Inc, USA harshpatelv1009@gmail.com, 626

Paper Title: A Hybrid Methodology for SAP Implementations: Blending ASAP with Agile Principles for Enhanced Flexibility and Stakeholder Engagement

10:01-10:16

Paper Presenter, Paper Id: Devisharan Mishra, Amazon Web Services, 562

Paper Title: Software Engineering Challenges in the Deployment of Generative AI Models at Scale

10:16-10:30

Paper Presenter, Paper Id: Vaishnavi Gudur, Microsoft, USA gudur.vaishnavi@gmail.com,144

Paper Title: Leveraging Generative AI And Reinforcement Learning For Autonomous Cyber Threat Mitigation

10:31-10:45

Paper Presenter, Paper Id: Utham Kumar Anugula, Independent Researcher, Atlanta, USA, mailuthamkumar@gmail.com, 153

Paper Title: A GenAI-Powered Cybersecurity Mesh for Real-Time Risk Detection in Digital Payments

10:46-11:00

Paper Presenter, Paper Id: Hariharan Velu Microsoft, USA hariharanvelu.research@gmail.com, 250

Paper Title: Enhancing Movie Recommendation Systems Through Explainable Machine Learning Models in the Entertainment Industry

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 6:

Session Chair: Mr. Tejas Pravinbhai Patel, Amazon, USA

Meeting Link: meet.google.com/tuz-peag-hfa

8:01-8:15

Paper Presenter, Paper Id: Yulia Bobkova, Vancouver Island University, Canada, Ajay.Shrestha@viu.ca, 491

Paper Title: Comparative Insights: A Multigroup Analysis of Privacy Management Across Youths, Parents, Educators, and AI Professionals in AI Applications

8:16-8:30

Paper Presenter, Paper Id: Molly Campbell , Vancouver Island University, Canada, Ajay.Shrestha@viu.ca, 517

Paper Title: Age-Differentiated Pathways to Privacy Protection in Smart Voice Assistants: A Multigroup PLS-SEM Study of Youth

8:31-8:45

Paper Presenter, Paper Id: Vaishnav Anand, The Athenian School, vaishnavanand90@gmail.com, 221

Paper Title: A Domain Shift Evaluation Protocol for Fake Satellite Image Detection: CNN Superiority and the Enhanced Model Paradox

8:46-9:00

Paper Presenter, Paper Id: Vaishnav Anand, The Athenian School, vaishnavanand90@gmail.com, 622

Paper Title: Quad-Stream Deepfake Detection: Combining Spatial and Frequency Domain Analysis for Robust Video Authentication

9:01-9:16

Paper Presenter, Paper Id: Ilia Sedoshkin, Whead, USA, i@wehead.com, 578

Paper Title: AI-Driven Frontend Cybersecurity for Real-Time Phishing and Threat Detection in ReactJS and Swift Applications

9:16-10:30

Paper Presenter, Paper Id: Pranav Gangwani, Florida International University, pgangwan@fiu.edu, 355

Paper Title: Non-Intrusive Machine Learning-Based Anomaly Detection for Heterogeneous Embedded Platforms

9:31-9:45

Paper Presenter, Paper Id: Sohrab Farooq, University of Salford, UK, S.Farooq8@edu.salford.ac.uk, 645

Paper Title:
Adaptive Honeypots using Reinforcement Learning Algorithms DQN and DDQN in Ensemble Framework: A Systematic Literature Review

9:46-10:00

Paper Presenter, Paper Id: Md. Sharif Hassan, Taylor's University, Malaysia, mdsharif.hassan@taylors.edu.my, 347

Paper Title: Cybersecurity in Blockchain-based fintech platforms: Social Perceptions and Adoption Intention

10:01-10:16
Paper Presenter, Paper Id: Sheikh Thanbir Alam, University Malaysia Perlis, Malaysia, sk.tamim56@gmail.com, 348
Paper Title: Security Challenges and Attack Vectors in Modern Steganography
10:16-10:30
Paper Presenter, Paper Id: Sheikh Thanbir Alam, University Malaysia Perlis, Malaysia sk.tamim56@gmail.com 363
Paper Title: Reverse Entropy Spiral Deep Neural Steganography for Secure Medical Ultrasonogram Videos
10:31-10:45
Paper Presenter, Paper Id: Abdullah Al Siam, Daffodil International University, Bangladesh, abdullah35-462@diu.edu.bd 367
Paper Title: Explainable Machine Learning for Malware Detection: A SHAP-Based LightGBM Framework
10:46-11:00
Paper Presenter, Paper Id: Deepika Bhatia, NVIDIA, United States reachdeepikabhatia@gmail.com, 422
Paper Title: MBIST++: An Adaptive March Algorithm Generator for Memory Test Coverage Enhancement in Post-Silicon Validation
11:01-11:15
Paper Presenter, Paper Id: Deepika Bhatia, NVIDIA, United States reachdeepikabhatia@gmail.com, 423
Paper Title: DFT AI: Machine Learning--Guided Test Point Insertion for Pre-Silicon Debug and Post-Silicon Diagnosis in Complex VLSI Systems

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 7:

Session Chair: Ms Krishnaveni Palanivelu, SVP, Citi, USA

Meeting Link: https://meet.google.com/oja-mjce-zen

8:01-8:15

Paper Presenter, Paper Id: Sashidhar Chinthala, IEEE, USA sashichs@ieee.org, 192

Paper Title: Adversarial Machine Learning in Cybersecurity: Attacks, Defenses, Robustness, and Explainability

8:16-8:30

Paper Presenter, Paper Id: Dilip Patel, Uber, USA dilip.patel.cal@gmail.com, 306

Paper Title: Integrating Price Elasticity and Reinforcement Learning: A Data-Driven Framework for Strategic E-commerce Pricing

8:31-8:45

Paper Presenter, Paper Id: Akshay Mittal University of the Cumberlands, USA, 331

Paper Title: Privacy-Driven Cloud AI: Federated Learning and Zero-Trust for Secure Multi-Domain Collaboration

8:46-9:00

Paper Presenter, Paper Id: Qixin Deng, 344

Paper Title: Lanyard Policy Tracker: A Secure, Privacy-Aware Student Compliance System for K–12 Environments

9:01-9:16

Paper Presenter, Paper Id: Rajgopal Devabhaktuni, Independent Researcher, Atlanta, USA, devabhaktuni.rajgopal@gmail.com, 349

Paper Title: Enclave-Driven Tokenization: Reducing PCI DSS Scope in Cloud-Native Checkout Systems

9:16-10:30

Paper Presenter, Paper Id: Madhusudan Sharma, Vadigicherla Integra LifeSciences Corp, USA, reachmadhusv@gmail.com 469

Paper Title: Supply Chain with Sixth Sense Agentic AI

9:31-9:45

Paper Presenter, Paper Id: Rahul Cherekar, Chewy, USA, rahul.cherekar@gmail.com, 395

Paper Title: Improving Temporal Ordering with External Data

9:46-10:00

Paper Presenter, Paper Id: Dileep Jain, FedEx, USA, jdileep41@gmail.com, 514

Paper Title: Artificial Intelligence (AI)-Driven Network Traffic Anomaly Detection for IT Infrastructure Security

10:01-10:16

Paper Presenter, Paper Id: Rajani Kant, University of the Cumberlands, USA rajani.kant@ieee.org, 521

Paper Title: AI-Based Cybersecurity in Healthcare: A Data-Driven, Governance-Aware Framework for Secure Clinical Systems

10:16-10:30

Paper Presenter, Paper Id: Vishwa Lakhnakiya, Tata Consultancy Services Inc, USA vishwalakhnakiya1628@gmail.com, 524

Paper Title: Future-Proofing Cloud Infrastructure: AI/ML-Driven Automation for Predictive Cloud Operations

10:31-10:45

Paper Presenter, Paper Id: Vishwa Lakhnakiya, Tata Consultancy Services Inc, USA vishwalakhnakiya1628@gmail.com , 525

Paper Title: Automating Multi-Cloud Deployments at Scale Using an Advanced GitOps Framework

10:46-11:00

Paper Presenter, Paper Id: Suresh Pairu Subramanyam, Avanade, USA sureshpairusubramanyam@gmail.com, 530

Paper Title: AI-Driven Data Architecture: Building Intelligent Analytics Platforms with Azure and Python

11:01-11:15

Paper Presenter, Paper Id: Suresh Pairu Subramanyam, Avanade, USA sureshpairusubramanyam@gmail.com, 532

Paper Title: DevOps and CI/CD Maturity in Large-Scale Organizations: A SonarQube and Jenkins Approach

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 8:

Session Chair: Milankumar Rana, MIlanrana@ieee.orga

Meeting Link: https://calendar.app.google/15r7fWRAUsZ1smsb6

8:01-8:15

Paper Presenter, Paper Id: Ashish Pokhrel,Vivint, USA, ashishbdm90@gmail.com, 548

Paper Title: Evaluating Modern Software Design Trends for Efficient and Maintainable Application Development

8:16-8:30

Paper Presenter, Paper Id: Venkata Thej Deep Jakkaraju, Gamestop, USA thejdeep.j@outlook.com, 552

Paper Title: Predictive Green FinOps: Joint Optimization of Cost, Carbon, and Reliability in AI-Intensive Clouds

8:31-8:45

Paper Presenter, Paper Id: kailash.thiyagarajan@ieee.org, 259

Paper Title: Observability in Large-Scale Multi-Agent Ecosystems: Coordination, Emergence, and Failure Modes

8:46-9:00

Paper Presenter, Paper Id: Rutul Desai, Kunai Inc, USA, rutuldesai193@gmail.com, 596

Paper Title: A Secure Biometric Authentication Framework for iOS with Cross-Platform Extensions: Addressing OS-Level Vulnerabilities and Enhancing Real-Time Protection

9:01-9:16

Paper Presenter, Paper Id: Rutul Desai, Kunai Inc, USA, rutuldesai193@gmail.com, 595

Paper Title: Optimizing AI-Driven Mobile Applications on iOS: A Comparative Analysis of Core ML and TensorFlow Lite for Cross-Platform Performance

9:16-9:30

Paper Presenter, Paper Id: Aldo Hernandez-Suarez, Instituto Politécnico Nacional, Mexico, alhernandezsu@ipn.mx , 574

Paper Title: BRUJO: Baseline-calibrated Risk from Unified Joint Observations for MITRE ATT\&CK-based Time Series Forecasting in Security Operations Centers

9:31-9:45

Paper Presenter, Paper Id: Mohamed Gebril, George Mason University, USA, mgebril@GMU.EDU, 376

Paper Title: A Comprehensive Assessment Tool for Prompt Injection Attacks in Large Language Models (LLMs)

9:46-10:00

Paper Presenter, Paper Id: Vamshi Krishna Pamula, Artek Solutions Inc, USA vamshikpamula@gmail.com 582

Paper Title: The Medallion Architecture in Practice: A Framework for Building Scalable and Governed Data Lakehouses on Microsoft Fabric

10:01-10:16

Paper Presenter, Paper Id: Shravya Bussari, HCLTech Inc, USA shravyabussari@gmail.com 583

Paper Title: RAGSec: Retrieval-Augmented Generation for Cybersecurity Threat Intelligence in Enterprise Networks

10:16-10:30

Paper Presenter, Paper Id: Vamshi Krishna Pamula, Artek Solutions Inc, USA vamshikpamula@gmail.com 589

Paper Title: From Science to Production: A Systematic Framework for Operationalizing and Governing Machine Learning Model

10:31-10:45

Paper Presenter, Paper Id: Tarek Mahmud Texas A&M University-Kingsville, USA, Tarek.Mahmud@tamuk.edu 447

Paper Title: TabNet-IDS: A TabNet-Driven Tabular Deep Learning Framework for Intrusion Detection Systems

10:46-11:00

Paper Presenter, Paper Id: Avdesh Mishra, Texas A and M University-Kingsville, USA avdesh.mishra@tamuk.edu

Paper Title: Enhancing Intrusion Detection with Image-Based CNN and CTGAN Synthetic Oversampling

11:01-11:15

Paper Presenter, Paper Id: Shujaatali Badami, Independent Researcher, USA shujaatali@ieee.org, 359

Paper Title: AI-Optimized VLSI Architecture for Energy Efficient and Sustainable IoT Systems

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 9:

Session Chair: Bhushan Chavan, Director of engineering Identity access management at MRI, USA

Meeting Link: https://meet.google.com/ujn-kgbh-nrj

8:01-8:15

Paper Presenter, Paper Id: Mohamed Gebril, George Mason University, USA, mgebril@GMU.EDU, 377

Paper Title: Multi-Layered Defense Proxy for Home Assistant utilizing Large Language Models

8:16-8:30

Paper Presenter, Paper Id: Khrystyna Terletska, Datadog, USA khrystynaterletska1@gmail.com, 603

Paper Title: Secure Execution of Post Inference Scripts in AI Driven Vector Search Pipelines

8:31-8:45

Paper Presenter, Paper Id: Sabarinathan Govindaraj, PWC, USA gsabari_89@hotmail.com, 604

Paper Title: Triangulating Digital Forensics to Detect Friendly Fraud and Abuse at Scale: A Cloud-Native, Agentic AI Framework

8:46-9:00

Paper Presenter, Paper Id: Mohammad Masum, San Jose State University, USA, mohammad.masum@sjsu.edu, 605

Paper Title: HRIT: A Human-Readable Framework for Phishing URL Detection using Large Language Models

9:01-9:16

Paper Presenter, Paper Id: Jonathan Roy, Université du Québec à Chicoutimi, Canada jonathan.roy@uqac.ca 607

Paper Title: Toward Context-Aware Alert Classification in Security Operations Centers Using LLMs

9:16-10:30

Paper Presenter, Paper Id: Samir Abood, Prairie View A&M University, Prairie View, TX, USA, siabood@pvamu.edu, 516

Paper Title: AI-Driven Cybersecurity for SCADA-Integrated Microgrids: A Real-Time Detection Framework

9:31-9:45

Paper Presenter, Paper Id: Bireswar Banerjee, VISA, USA, bireswar.infosys@gmail.com, 613

Paper Title: AI-Powered Economic Digest and A Composite Index for National Financial Well-being: An AI-Driven Approach to Quantifying United States Financial Health with the USFHI

9:46-10:00

Paper Presenter, Paper Id: Jyoti Kunal Shah, ADP, USA, thejyotishah83@ieee.org, 614

Paper Title: Unified AI Framework for Real-Time Customer Churn Prediction Using Behavioral and Event-Log Anomaly Signals

10:01-10:16

Paper Presenter, Paper Id: Md Muntasir Hossain, Lamar University, USA mhossain54@lamar.edu, 616

Paper Title: A Dual-task Prediction Model for Starlink Maritime Performance

10:16-10:30
Paper Presenter, Paper Id: Jothsna Praveena Pendyala, Clark University, USA
jothsnapraveena1421@gmail.com, 620
Paper Title: ZT-ICAS: A Zero-Trust Integrity-Constrained Framework for Agentic Vulnerability Scanning
10:31-10:45
Paper Presenter, Paper Id: Ranjith Kumar Vanaparthi, Ventois Inc, USA, ranjithkumarvanaparthi21@gmail.com, 631
Paper Title: From UIKit to SwiftUI: A Quantitative Analysis of Migration Strategies for Large-Scale Financial Applications
10:46-11:00
Paper Presenter, Paper Id: Ranjith Kumar Vanaparthi, Ventois Inc, USA, ranjithkumarvanaparthi21@gmail.com, 632
Paper Title: Architecting for Privacy and Performance: A Federated Learning Framework for On-Device Financial AI in Mobile Applications
11:01-11:15
Paper Presenter, Paper Id: Sherif Abdelhamid Virginia Military Institute Lexington, VA, USA, abdelhamidse@vmi.edu , 337
Paper Title: TrustNet: A Hybrid Machine Learning and LLM-Based Multi-Agent System for Scam Website Detection

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 10:

Session Chair: Nixal Patel, LG Electronics North America, New Jersey, United States

Meeting Link: https://meet.google.com/swn-yfvh-qpg

8:01-8:15

Paper Presenter, Paper Id: Hareendra Sura, Coupang, USA, Hareendrasura@gmail.com
642

Paper Title: Proactive Discrepancy Detection at Scale: A Continuous Validation Framework for Distributed Data Systems

8:16-8:30

Paper Presenter, Paper Id: Hareendra Sura, Coupang, USA, Hareendrasura@gmail.com, 643

Paper Title: Forecasting and Overcommit Pipelines for Cloud Storage: A Model for Persistent Disk Capacity Management

8:31-8:45

Paper Presenter, Paper Id: Varun Singh, hritesh.yadav@ieee.org, 646

Paper Title: Securing 5G/6G Networks: Handover, Slicing, and QoS Security Issues

8:46-9:00

Paper Presenter, Paper Id: Datta Snehith Dupakuntla Naga, Teladoc Health Inc, USA, dndattasnehith1989@gmail.com, 650

Paper Title: Intelligent Test Data Automation: A Python-Based Framework for Deterministic and Scalable Software Testing

9:01-9:16

Paper Presenter, Paper Id: Keshav Kumar, Gyancity Research Lab, India keshav@gyancity.com, 686

Paper Title: Multi-Frequency Implementation and Timing Analysis of LAES on Spartan-7 and Kintex-7 FPGAs

9:16-10:30

Paper Presenter, Paper Id: Devanand Patil, Sant Gajanan Maharaj College of Engineering, Mahagaon, India, 2505C40027@sru.edu.in, 721

Paper Title: Resource allocation in 6G Networks:A comprehensive review of Machine Learning Approaches

9:31-9:45

Paper Presenter, Paper Id: Ramsha Ali, University of Bradford, UK, ramshaali88@yahoo.com, 615

Paper Title: A Federated Learning Framework for the IoMT System to Promote the Personalised Detection of COPD in a Non-Clinical Environment – A Pilot Study

9:46-10:00

Paper Presenter, Paper Id: Anshul Sharma, Independent Researcher, USA
anshul.sharma@ieee.org, 360

Paper Title: Enhancing Safety and Performance in Intelligent Vehicular Networks Using Edge-Based Explainable AI Models

10:01-10:16

Paper Presenter, Paper Id: Moutaz
Alazab moutaz.a@oryx.edu.qa, 370

Paper Title: TransCall: A Transformer-Driven Framework for Zero-Day Malware Detection Using System Call Sequences

10:16-10:30

Paper Presenter, Paper Id: Abdullah AlSiam, abdullah35-462@diu.edu.bd, 371

Paper Title: Real-Time Multi-Source Threat Intelligence Fusion Using Large Language Mode

10:31-10:45

Paper Presenter, Paper Id: Advait Patel, Independent Researcher, United States advaitpa93@gmail.com, 489

Paper Title: Secure Prompt Engineering Patterns for Cloud LLM Agents

PARALLEL SESSION ON 20 FEBRUARY 2026

8:00-11:15 (TEXAS TIME)

Session 11:

Session Chair: Shalini Sudarsan,DevOps Engineering Manager KinderCare Learning Companies Oregon, USA shallene.s@gmail.com

Meeting Link: https://meet.google.com/mzn-dpfy-xiy

8:01-8:15

Paper Presenter, Paper Id: Abid Ahmad, Southeast University, Bangladesh, abidzufi@gmail.com, 599

Paper Title: Identity-Spoofing Prompt Injection: Empirical Analysis of Instruction Leakage in Custom GPTs

8:16-8:30

Paper Presenter, Paper Id: Avdesh Mishra,Texas A and M University-Kingsville, USA avdesh.mishra@tamuk.edu, 587

Paper Title: GXMalDetect: A Hybrid GA–XGBoost Architecture for Malware Detection Using Static Image Features

8:31-8:45

Paper Presenter, Paper Id: Vitaly AndrejeusAngelo State University, USA vandrejeus@angelo.edu, 623

Paper Title: Automated Validation and Repair of Knowledge Graph Triples for Cyber Threat Intelligence

8:46-9:00

Paper Presenter, Paper Id: Aman Singh, Dakota State University, USA aman.singh@trojans.dsu.edu, 404

Paper Title: Mid-Generation Jailbreaks in Open-Source LLMs Using a Pause-and-Edit Attack

9:01-9:15

Paper Presenter, Paper Id: MD WAHIDUR RAHMAN Texas A&M University-Kingsville, USA wahidtuhin0@gmail.com, 313

Paper Title: Privacy-Preserving Federated Deep Learning for Nomophobia Risk Prediction from Smartphone Usage Logs

9:16-9:30

Paper Presenter, Paper Id: Madhuri Margam, Director, Software Engineer, USA, madhurimargam2@gmail.com, 531

Paper Title: A Secure and Scalable Cloud-Based Banking Service Model Leveraging AI and Advanced Cyber Security

9:31-9:45

Paper Presenter, Paper Id: Jeevana Swaroop Kalapala, Northern Arizona University, USA, jk2396@nau.edu, 602

Paper Title: Packing Induced Bias in Deep Learning Malware Classifiers: A Systematic Experimental Study

9:46-10:00

Paper Presenter, Paper Id: Sabitha Muppuri, Independent Researcher, USA, sabitha594@gmail.com 335

Paper Title: Generative Temporal Diffusion Models for Early Prediction of Cloud Service Degradation

10:01-10:15

Paper Presenter, Paper Id: Ajay Kumara Makanahalli Annaiah University of North Carolina Wilmington, USA makanahalliannaiaha@uncw.edu, 388

Paper Title: Exploring the Capabilities of LLMs in Binary Decompilation and Deobfuscation

10:16-10:30

Paper Presenter, Paper Id: BhuLakshmi Makkena, Mastercard, USA, Bhulakshmimakkena9@gmail.com, 307

Paper Title: Interpretable Advanced Machine Learning Models for Early Identification of Health Insurance Claim Fraud Detection

10:31-10:45

Paper Presenter, Paper Id: BhuLakshmi Makkena, Mastercard, USA, Bhulakshmimakkena9@gmail.com, 308

Paper Title:Harnessing Natural Language Processing (NLP) and Generative AI Techniques for Social Media Sentiment Analysis with Text Classification

10:46-11:00

Paper Presenter, Paper Id: Aruun Kumar, Amazon Web Services, USA aruunkumar@gmail.com 389

Paper Title: Multi-Layered Security Framework for Financial AI Solutions - PISA

ICGCET'2015: 1st International Conference of Gyancity at Dubai, UAE

RTCSE'16: 2nd International Conference of Gyancity at Kuala Lumpur, Malaysia

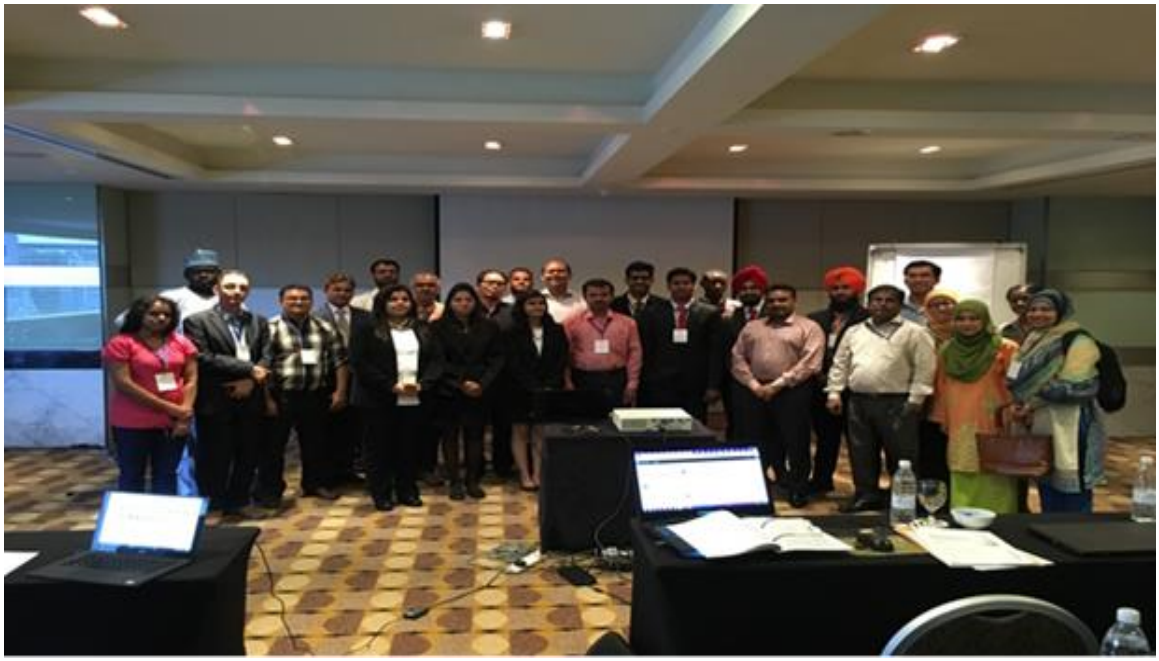

ICGCET'2016: 3rd International Conference of Gyancity at Aalborg University, Esbjerg, Denmark

Institut i Esbjerg samler forskere fra hele verden

DEL Af Edmund Jacobsen 15. august 2016 kl. 05:31

40 forskere og studerende fra hele verden samles på Institut for Energiteknik, Aalborg Universitet Esbjerg, i tre dage i denne uge, når der afvikles en international konference, der handler om at gøre computerteknologi mere grøn.

D.M. Akbar Hussain, lektor ved Institut for Energiteknik på Aalborg Universitet Esbjerg, har sammen med en kollega fra Indien arrangeret konferencen International Conference on Green Computing and Engineering Technologies.

Det er planen, at disse konferencer skal afvikles i Esbjerg hvert andet år – ganske enkelt fordi Institut for Energiteknik i Esbjerg er internationalt anerkendt.

RTCSE'17: 4th International Conference of Gyancity at Kuala Lumpur, Malaysia

IMCES'17: 5th International Conference of Gyancity at Kuala Lumpur, Malaysia

ICGCET'2018: 6th International Conference of Gyancity at Limerick, Ireland

RTCSE'2018: 7th International Conference of Gyancity at Bangkok, Thailand

ICGCET'18: 8th International Conference of Gyancity at Aalborg University, Esbjerg, Denmark

RTCSE'2019: 9th International Conference of Gyancity at Univeristy of Hawaii, USA

IMCES'2019:10th International Conference of Gyancity at Port Louis, Mauritius

ICGCET'2019: 11th International Conference of Gyancity at Casablanca, Morocco

RTCSE'2020: 12th International Conference of Gyancity at University of Hawaii, USA

IMCES'2020: 13th International Conference by Gyancity at Jakarta, Indonesia

ICGCET'2020: 14th Conference by Gyancity at St Petersburg, Russia

RTCSE'2021: 15th International Conference of Gyancity at University of Hawaii, USA

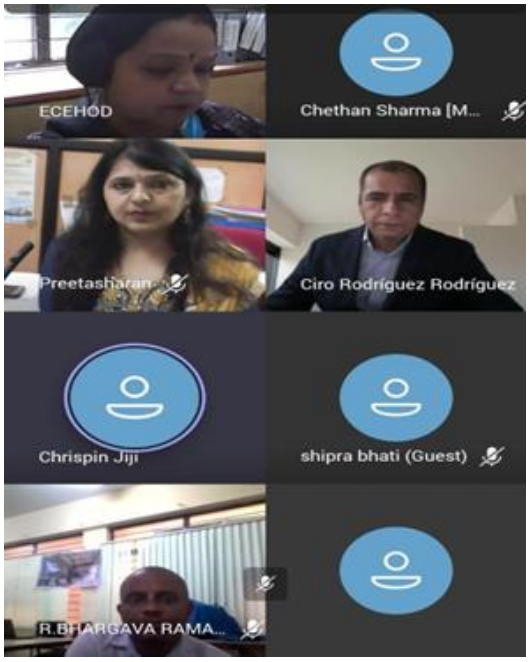

BMESS'2021: 16th Virtual Conference by Gyancity

IMCES'2021: 17th International Conference by Gyancity at Yarsi University, Indonesia

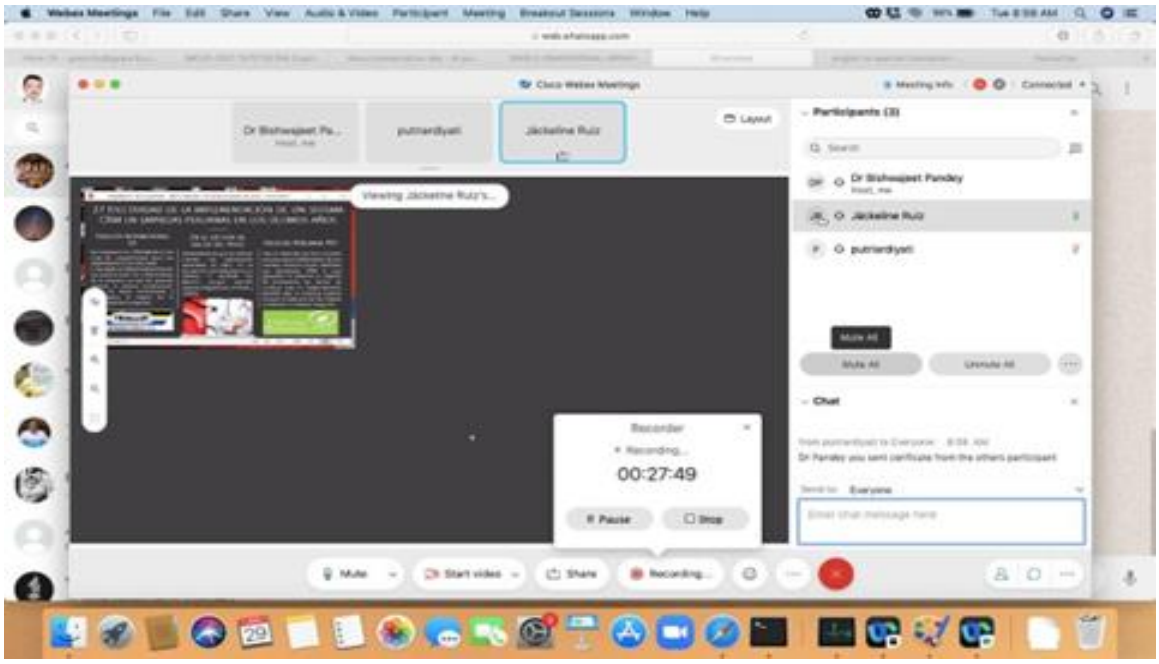

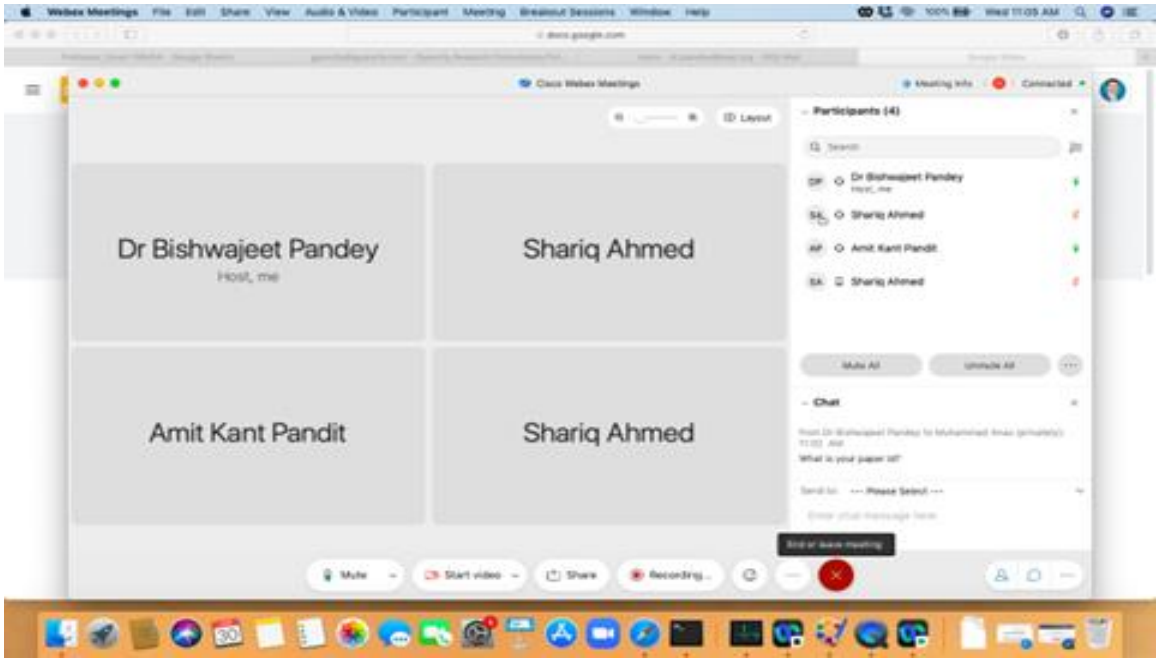

ICGCET'2021: 18th International Conference by Gyancity at National University of Federico Villareal, Lima, Peru

Evento se dará el 22 y 23 de septiembre. Foto: difusión

La República
larepublica_pe
ediciondigital@glr.pe

16 Set 2021 | 12:40 h
Actualizado el 16 de Setiembre 2021 | 12:40 h

Este 22 y 23 de septiembre se realizará la 7ª Conferencia Internacional sobre Tecnologías de Ingeniería y Computación Ecológicas 2021 (ICGCET-2021) y la 13ª Conferencia Internacional en Inteligencia Computacional y Redes de Comunicación 2021 (CICN 2021), eventos que tendrán como sede a la Universidad Villareal (UNFV).

Juan Alfaro, rector de la UNFV, será el encargado de inaugurar los referidos certámenes, el miércoles 22 a las 10.00 a. m. Previamente, Akbar Hussain, de la Universidad Aalborg de Dinamarca, será el encargado de brindar las palabras de bienvenida.

La ICGCET-2021 presentará las investigaciones de diferentes áreas de la ciencia y la tecnología, y proporcionará una plataforma para que investigadores y científicos de todo el mundo intercambien y compartan sus experiencias y resultados de investigación.

La República

ÚLTIMAS NOTICIAS POLÍTICA ECONOMÍA SOCIEDAD MUNDO DEPORTES ESPECTÁCULOS RED

EN VIVO - Emmy 2021: sigue aquí la premiación a lo mejor de la TV y el streaming

NOTAS DE PRENSA

Conferencias internacionales se desarrollarán en Universidad Villarreal

Cada evento contará con la participación de destacados expertos de la investigación.

ICGCET'2021: 18th International Conference by Gyancity at National University of Federico Villareal, Lima, Peru

RTCSE'2022: 19th International Conference of Gyancity at University of Hawaii USA

BMESS'2022: 20th International Conference by Gyancity at Bath Spa University UAE

ICAIC'2022: 21st International Conference by Gyancity at University of Houston-Victoria, USA

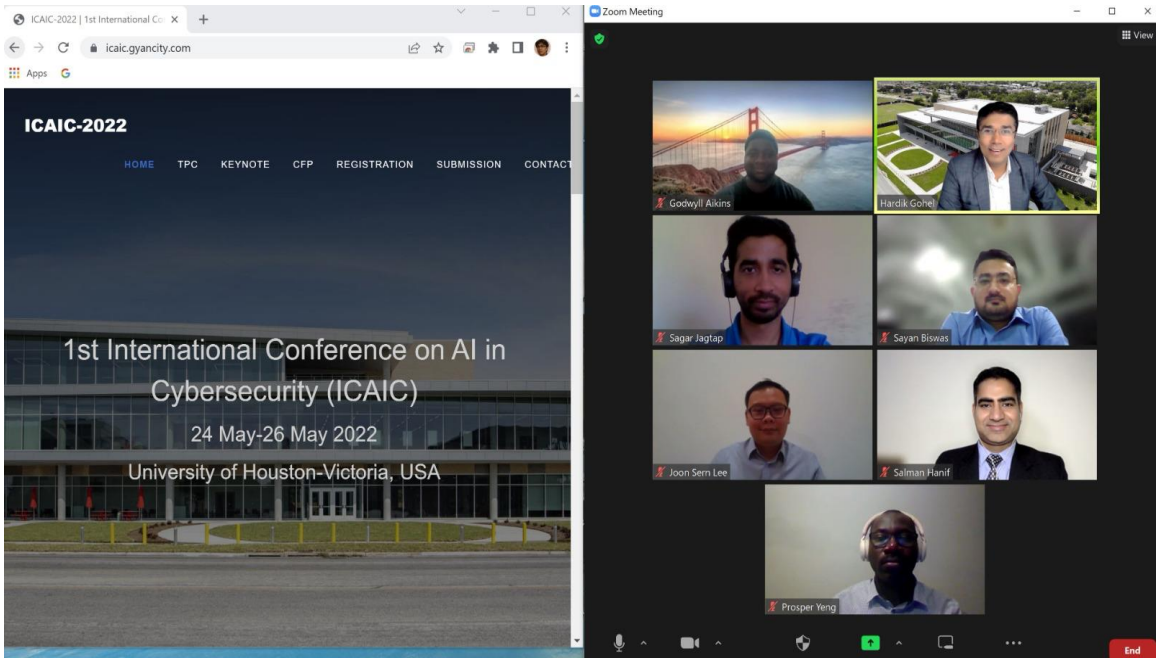

IMCES'2022: 22nd International Conference by Gyancity at Aalborg University, Esbjerg, Denmark

ICGCET'2022 GROUP PHOTO: 23rd International Conference of Gyancity at Mauritius

RTCSE'2023 GROUP PHOTO: 24th International Conference of Gyancity at University of Hawaii USA

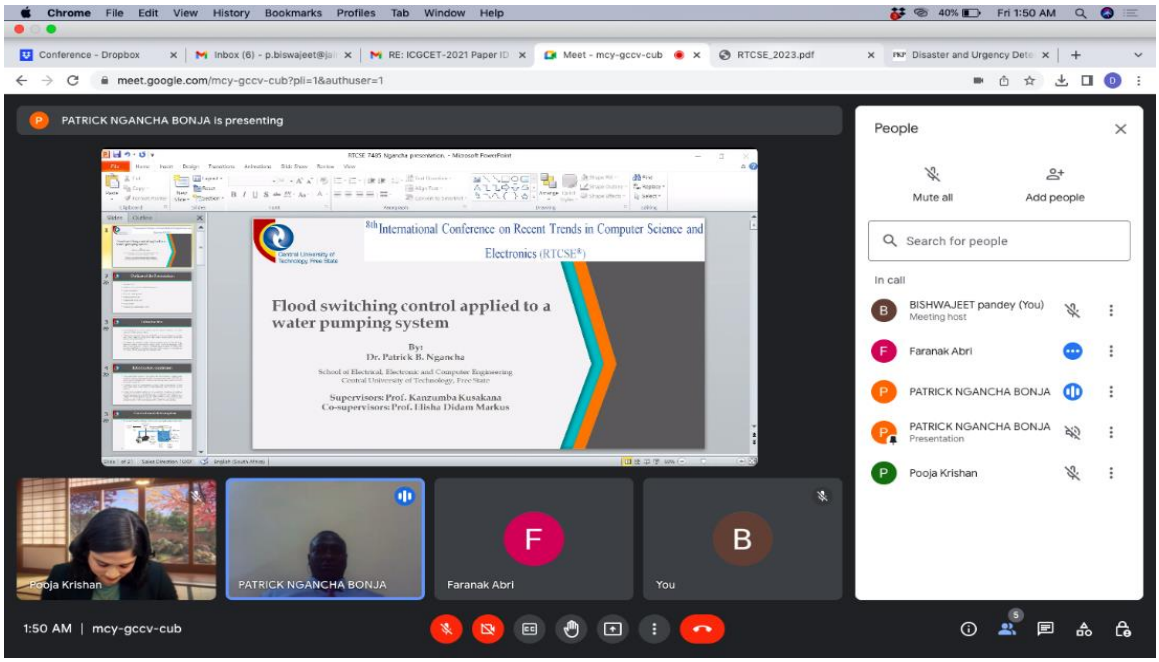

ICAIC'2023 GROUP PHOTO: 25th International Conference of Gyancity at University of Houston-Victoria, USA

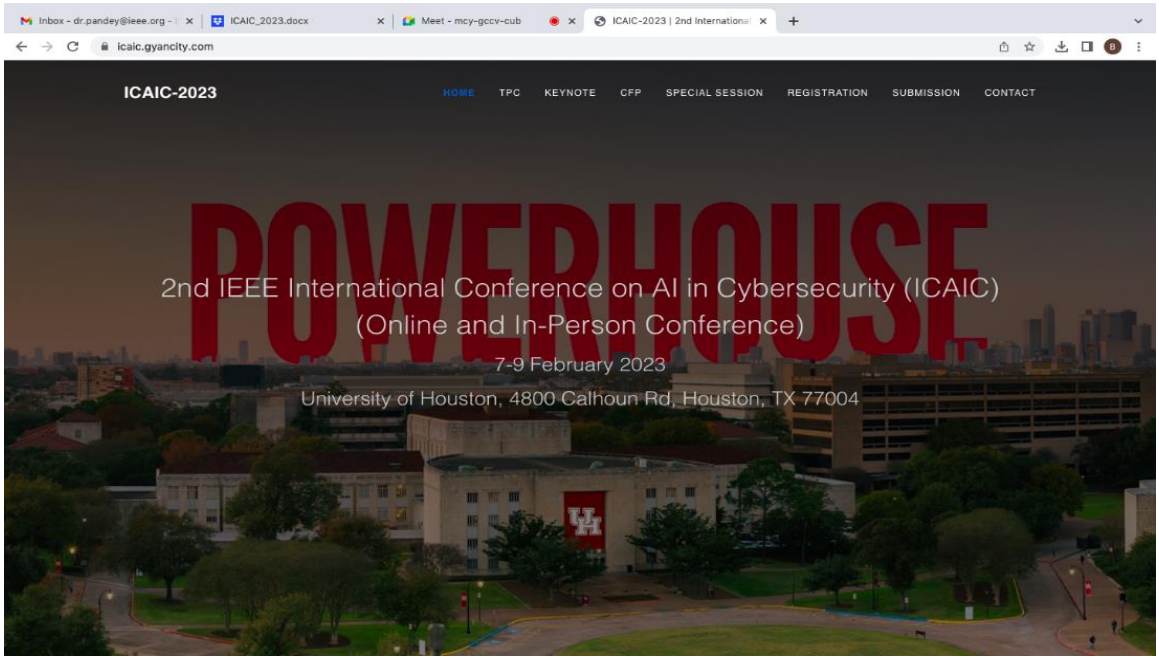

BMESS'2023 GROUP PHOTO: 26th International Conference of Gyancity at Bath Spa University, UAE

IMCES'2023: 27th International Conference by Gyancity at Yarsi University, Jakarta, Indonesia

ICGCET'2023: 28th International Conference by Gyancity at Cape Town South Africa

ICAIC'2024: 29th International Conference by Gyancity at University of Houston, USA

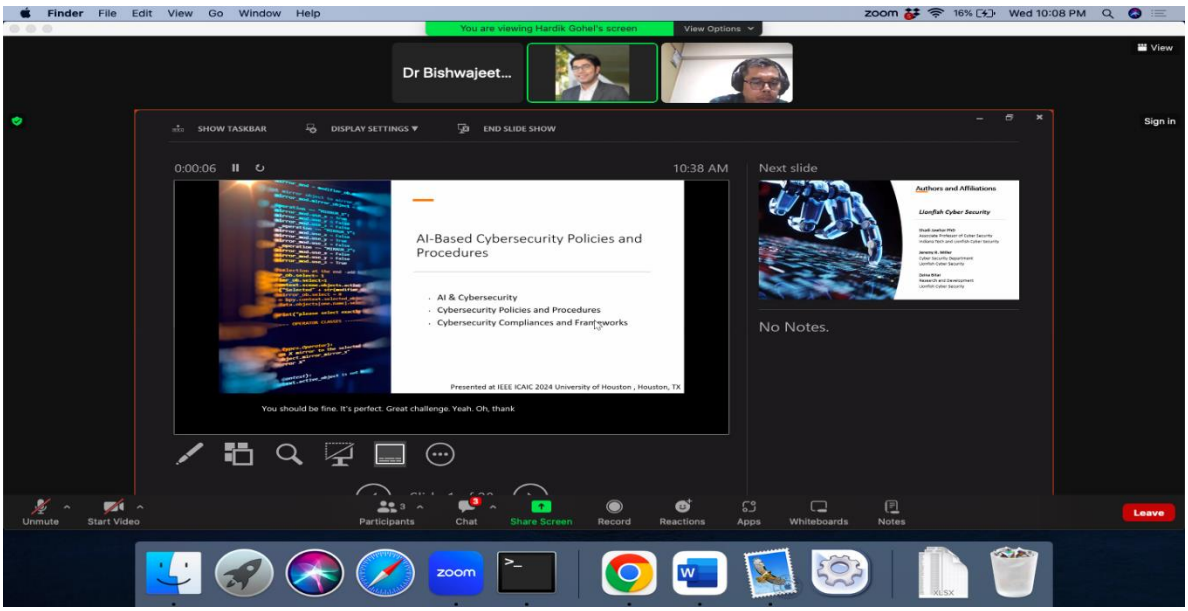

IMCES'2024 and BMESS'2024: 30th and 31st International Conference by Gyancity at Bath Spa University UAE

ICGCET'2024: 32nd International Conference by Gyancity at Sea Cliff Resort, Zanzibar

ICAIC'2025: 33rd International Conference by Gyancity at University of Houston

RTCSET'2025: 34th International Conference by Gyancity at Yarsi University, Jakarta and University Technology of Bandung, Indonesia

Abstract of Paper Accepted in ICAIC-2026

143

A Review of User Account Anomaly Detection and Insider Threat Detection Techniques

Joe Childress
Louisiana Tech University, USA

ABSTRACT

Modern information systems can be subjected to malicious attacks from the compromise of a valid user's account credentials, or even from the malicious use of an internal authorized user. These types of attacks can be extremely difficult to detect. This paper reviews a range of literature on a variety of user account compromise and user behavior anomaly detection techniques. The methods proposed in the literature covered here include both supervised and unsupervised approaches.

Keywords: Machine learning, cybersecurity, supervised learning, unsupervised learning, user anomaly detection

Abstract of Paper Accepted in ICAIC-2026

144

Leveraging Generative AI And Reinforcement Learning For Autonomous Cyber Threat Mitigation

Vaishnavi Gudur
Independent Researcher, IEEE, USA
gudur.vaishnavi@gmail.com
Shiva Kumar Ramavath
University of the Cumberlands, USA
r92shivakumar@gmail.com

ABSTRACT

Cybersecurity is up against attackers that are getting smarter and faster than traditional defenses. This paper proposes a new AI-driven cybersecurity framework that leverages machine learning (ML) and generative artificial intelligence (GenAI) to enhance threat detection, automate incident response, improve anomaly detection, and bolster zero-day exploit prevention. The framework solves some of the greatest concerns with AI-based security systems today, such as too many false positives, difficulty detecting new attacks, sluggish response times, and difficulty integrating. It does this by employing a structure that blends generative models (such GANs and VAEs), sophisticated deep learning, and large language model (LLM) methods. We explain how synthetic data generation and generative modeling help find unusual behavior and unknown threats before they happen, and how reinforcement learning and automation make it easier to take quick action to contain them. A proof-of-concept simulation shows that the framework could greatly improve the accuracy of detection and the speed of response compared to standard solutions. We also talk about how the framework works with tools that are already in place (like SIEM and SOAR) and keeps data private. The results show that using GenAI can make a cyber defense that is more flexible and smart. This work introduces a multi-layer GenAI-driven framework that integrates anomaly detection, LLM-based reasoning, and reinforcement learning to automate cybersecurity operations and reduce false positives.

Keywords: Cybersecurity, Generative AI, Large Language Models, Anomaly Detection, Intrusion Detection Systems, Explainable AI, Reinforcement Learning, Threat Simulation, AI for Security, Zero-Day Attack Detection.

Abstract of Paper Accepted in ICAIC-2026

145

AI-Driven Quantum Cryptography

Shilpi Mittal
Tyson Foods Inc., USA
mittalshilpi@hotmail.com
Ankit Gupta,
Exeter Finance LLC, USA
ankitgupta_ag@outlook.com

ABSTRACT

Integrating artificial intelligence in quantum cryptography is another masterstroke in secure communication, since other cryptographic models have flaws exposed in the quantum computing world. This paper aims to analyze how AI can be used simultaneously to enhance quantum cryptographic algorithms, such as Quantum Key Distribution (QKD), through their predictive and adaptive attributes in combating challenges like quantum noise, resource management, and real-time threat identification. Quantum cryptography removes the practical barriers of scalability, error correction, and critical management typically applied to AI-driven systems, thereby enhancing the security of communications against both classical and quantum threats. The study also demonstrates how AI can improve the development of cryptographic systems by utilizing machine learning algorithms, neural networks, and reinforcement learning frameworks. It investigates the effectiveness of these techniques in noise reduction, adapting to dynamic protocols, and implementing self-turnstile mechanisms in response to emerging threats. The paper also explores the potential applications of AI-based quantum cryptography in critical sectors, including finance, healthcare, defense, and telecommunications, where secure information sharing and system reliability are paramount. However, several issues are raised that cut across the technical and operational aspects, including, but not limited to, computational overhead, interpretability, and ethical issues. This study emphasizes the combination of operational and system approaches in creating standardized models and recruiting lightweight forms of AI to achieve easy accessibility, efficiency, and optimal security.

Keywords: artificial intelligence (ai), quantum cryptography, quantum key distribution (qkd), ai-driven optimization, secure communication, scalability, noise mitigation, resource management, hybrid cryptographic systems, anomaly detection

Abstract of Paper Accepted in ICAIC-2026

148

Unknown Threat Detection using AI, ML and DL methods

Manjunath Venkatram

ThoughtData, USA

ABSTRACT

Cyber threats in any organization is a constant concern for the IT teams. Many tools available in the market only detect known threats and many uses fixed signatures which hackers and attackers are unlikely to reuse. Any rule-based engine deployed is ineffective given the nature of manifestation of threat attacks hackers use these days. If a system can learn the network over time from what is expected and what is not expected it can provide better detection of unusual traffic patterns in the network which are threat oriented. This research explores a new way of detecting unknown threats which are happening for the first time in a IT network by using machine learning or behavior analysis of historical data and also certain probabilistic functions to qualify a cyber threat. It depends of historical behavior learning from an IT network where behavior patterns are learnt over time and used to detect anomalous conditions and qualifying them further as potential threats by different means.

Keywords: Cyberthreats, Deep packet inspection, DPI, AI, ML, Deep Learning, Zero-day exploits, Behavior Analysis

Abstract of Paper Accepted in ICAIC-2026

153

A GenAI-Powered Cybersecurity Mesh for RealTime Risk Detection in Digital Payments

Utham Kumar Anugula, Vijayanand Ananthanarayanan

ABSTRACT

The surge in digital payment systems has introduced significant challenges in securing real-time financial transactions against sophisticated cyber threats. This paper presents a GenAI-powered cybersecurity mesh architecture designed to detect and mitigate risks dynamically within decentralized digital payment ecosystems. Our approach integrates transformer-based deep learning models, explainable AI (XAI) mechanisms, and decentralized threat telemetry across a microservices-driven architecture. The system supports autonomous policy enforcement and real-time fraud detection using SHAP-explainable risk scores aggregated from user behavior, device fingerprinting, and transaction metadata. Deployed within a zero-trust security model, the mesh adapts to adversarial patterns through self-improving AI agents. Experimental results on simulated fintech datasets show over 94% detection accuracy with sub-200ms latency, demonstrating viability for high-throughput environments. The proposed architecture advances the state-of-the-art in real-time payment security and paves the way for scalable GenAI-driven threat response infrastructures in financial services.

Keywords: GenAI, Cybersecurity Mesh, Digital Payments, Explainable AI, Transformer Models, Real-Time Risk Detection, Zero-Trust Architecture, Financial Security

Abstract of Paper Accepted in ICAIC-2026

164

SC-GAN: A GAN-Based Data Augmentation Approach for Stablecoin Fraud Detection on Imbalanced Transaction Data

Mohan Sankaran, PayPal, USA, mohansankaran@ieee.org
Nagaraju Jooluri, Incode Technologies, j.nagaraju@gmail.com
Srimaan Yarram, Coupa, srimaan.yarram@gmail.com
Balasundaram Subbusundaram, Walmart, zelabala1@gmail.com
Balaji Sundareshan, Tumeke Inc, kss.balaji1@gmail.com

ABSTRACT

Stablecoins are becoming more common in the Fin-Tech (Financial Technology) ecosystem as they keep their value stable and combines easily with decentralized finance inherentin the financial technology ecosystem due to their price stability and simplicity of integration in decentralized finance (DeFi),cross-border payments, and automated trading systems. However, the same characteristics that propel utility transaction speed, pseudonymity, and automation through smart contracts have also made them vulnerable to financial manipulation. Tactics such as wash trading, spoofing, and pump-and-dump schemes have become more prevalent, compromising market integrity significantly. However, major technical challenge in detecting these fraudulent activities and behaviors, especially under conditions of extreme class imbalance even the legitimate transactions vastly outnumber fraudulent ones.This paper introduces SC-GAN, a conditional Generative Adversarial Network that addresses the scarcity of fraudulent samples by synthesizing realistic blockchain-based fraud instances.The model conditions on key financial and transactional features native to blockchain systems, enabling the generation of highfidelity synthetic data. We then compare SC-GAN with traditional oversampling techniques like SMOTE and Borderline SMOTE , on a variety of supervised classification models. Our experiments on a real-world stablecoin transaction dataset show with the help of SC-GAN improves both the F1 Score and overall accuracy that is resulting in more efficient detection of rare but crucial fraudulent transactions. This approach alos provides the possibility for stronger fraud prevention methods and risk management policies within FinTech platforms

Keywords: Stablecoins, FinTech, Blockchain Analytics, Fraud Detection, GANs, Data Augmentation, Imbalanced Classification, SMOTE, Borderline-SMOTE, DeFi.

Abstract of Paper Accepted in ICAIC-2026

175

Neuromorphic Architectures for Infrastructure-Scale AI: Design, Equations, and Synthetic Benchmarks

Shilpi Mittal
Tyson Foods Inc., USA
mittalshilpi@hotmail.com
Ankit Gupta,
Exeter Finance LLC, USA
ankitgupta_ag@outlook.com

ABSTRACT

Neuromorphic computing offers brain-inspired, event-driven processing with fine-grained parallelism and near-memory computation, making it an attractive architectural direction for large-scale AI in IT infrastructure. We present a concise architecture-focused study anchored in synthetic benchmarks and datasets (e.g., N-MNIST, DVS-like streams), quantifying advantages for sparse, streaming workloads. We formalize spiking neuron dynamics, give an energy-to-solution model, and report fully reproducible synthetic results and figures (spike rasters, energy per inference, latency/throughput, node-level integration). We outline an integration path for heterogeneous nodes (CPU+GPU+NPU), identify software tooling for reproducible evaluation, and summarize open challenges in programmability, scaling, and benchmarking. The paper targets Track 6 (AI in IT Infrastructure) by emphasizing system-level fit, energy proportionality, and deployability in data centers.

Keywords: Neuromorphic computing, spiking neural networks, event-driven architectures, energy efficiency, heterogeneous HPC/IT infrastructure, synthetic benchmarks.

Abstract of Paper Accepted in ICAIC-2026

177

Evaluating Jailbreak Vulnerabilities in LLMs: A Taxonomy and Comparative Analysis in Romance Fraud

Yohn Jairo Parra Bautista, Hongmei Chi, Richard Alo
Florida A&M University
yohn.parrabautista@famu.edu, hongmei.chi@famu.edu, richard.alo@famu.edu
Vinny Lima
Purdue University
vlima@purdue.edu

ABSTRACT

Romance fraud conversations provide high-stakes contexts where jailbreak prompts can elicit policy-violating guidance from large language models (LLMs). We present a taxonomy-driven benchmark and comparative evaluation across three leading LLM families using 60 romance-themed jailbreak prompts. Prompts are auto-tagged into 12 strategy primitives---e.g., roleplay persona, multi-agent mimicry, policy evasion, system override, format smuggling, translation obfuscation, emotional coercion, cot coaxing, and no refusal clause, and summarized into five higher-level categories: persona roleplay, safety bypass, system injection, coaxing frame, and obfuscation. We measure Attack Success Rate (ASR), refusal rate, and latency, and provide pairwise confusion/effect size analyses by strategy and category. Results show substantial variance in susceptibility: Gemini achieved the highest ASR (0.983; refusal 0.017; ~1.08 s), OpenAI models were moderately vulnerable (ASR 0.533; refusal 0.467; ~4.08 s), and Claude was most conservative (ASR 0.250; refusal 0.750; ~5.75 s), each evaluated on n=60 prompts. High-yield attacks frequently combined persona-roleplay with explicit refusal-suppression clauses or multi-agent mimicry. The proposed taxonomy and measurement workflow surface model-specific blind spots and enable targeted mitigation testing, including strengthened system-prompt hardening, refusal-sandbox detection, and filters for emotionally coercive or obfuscated requests.

Keywords: fraud detection, jailbreaking, LLMs, psychological traits, social media scam

Abstract of Paper Accepted in ICAIC-2026

178

Heart Disease Prediction Using Feature Reconstruction and Hybrid Deep Learning

Saba Shamsher
Zillow, United States
Sandeep Thota
Oracle Inc, United States
Hari Suresh Babu, Narendra Chennupati
Frontier, United States
Shivakrishna Deepak Veeravalli
Benchling, United States
Manisha Guduri
Lawrence Technological University, USA

saba.shamsher@gmail.com , sandeepthota@ieee.org ,
gummadiharisureshbabu@gmail.com , chennupati.be@gmail.com ,
Shivakrishnadeepakv@gmail.com , manishaguduri@ieee.org

ABSTRACT

Cardiovascular disease is the leading cause of death worldwide, so creating new, quick, and accurate diagnostic techniques is essential. To improve feature categorization, this study proposes a novel hybrid deep learning architecture combining the strengths of sparse autoencoders (SAEs), multilayer perceptrons (MLPs), and convolutional neural networks (CNNs). After the initial steps, including categorical encoding and normalization, are complete, the model generates precise feature representations from the UCI Heart Disease dataset. This collaborative architecture enables the collection of substantial abstract data while preserving vital physiological activities. In terms of accuracy (95.1%), precision (94.6%), recall (93.8%), and F1-score (94.2%), the model outperforms all other machine learning methods, including Learning Vector Quantification and Logistic Regression with Boruta Feature Selection. According to the statistics, the proposed CNN-MLP + SAE architecture is a reliable and effective method for early identification of cardiac problems.

Keywords: Cardiovascular disease, UCI Dataset, CNN, SAE

Abstract of Paper Accepted in ICAIC-2026

179

A Framework on Advancing Cybersecurity Education via Quantum Machine Learning

Celestina Kolog, Hongmei Chi
Florida A&M University, USA
celestina1.kolog@famu.edu, chi7356@gmail.com
Jie Yan, Bowie State University, USA, jyan@bowiestate.edu
Xian Mallory
Florida State Universuty, USA
xfan2@fsu.edu

ABSTRACT

https://www.glbitm.orgQuantum computing and artificial intelligence (AI) are two powerful technologies that are beginning to shape how we solve problems in cybersecurity and education. As cyber threats grow more complex, there's a growing need for students to understand how these advanced tools work and how to apply them in real-world situations. This paper presents a simple but effective framework for designing hands-on labs that teach students about quantum machine learning (QML), which combines the strengths of quantum computing and AI. These labs are meant to help students learn by doing solving practical cybersecurity challenges while building technical skills in coding, quantum circuits, and machine learning models. We also reviewed key research from experts in the fields of quantum computing, cybersecurity, and education. Their insights helped us decide how to organize the topics and where they should appear in a learning program or thesis. The goal is to make the learning process more interactive, scalable, and relevant to the digital world today. This approach does not just teach students new technologies; it gives them the opportunity to experience how these tools can be used to defend against real cyber threats.

Keywords: quantum machine learning, cybersecurity, experiential learning, machine learning, cybersecurity education

Abstract of Paper Accepted in ICAIC-2026

192

Adversarial Machine Learning in Cybersecurity: Attacks, Defenses, Robustness, and Explainability

Sashidhar Chinthala,
IEEE, USA sashichs@ieee.org

ABSTRACT

Most modern cybersecurity systems, including intrusion detection, malware classification, anomaly detection, and authentication, utilize machine learning. However, ML models are susceptible to adversarial machine learning (AML): designed inputs or modified training data that lead to false results without raising a red flag or damaging the model. This survey encompasses recent research (2020-2025) on AML in cybersecurity, categorizing attacks and defenses into a single taxonomy, examining the relationship between robustness and explainability, comparing and contrasting methods across domains, and identifying open problems and actionable directions in research and practice. We also underscore practical limitations (e.g., maintaining malware functionality or realistic traffic) and the significance of standardized benchmarks, lifecycle defenses, and explainable techniques that are both informative and resilient to abuse.

Keywords: Adversarial machine learning; cybersecurity; adversarial attacks; adversarial defenses; robustness; explainable AI; intrusion detection; malware

Abstract of Paper Accepted in ICAIC-2026

193

Application of Federated Learning to Semantic Segmentation of Aerial Images

Yong-Lin Kuo, Ying-Wei Chuang
National Taiwan University of Science and Technology,
Taipei, Taiwan
yl_kuo@yahoo.com, home5o5o0089@gmail.com

ABSTRACT

This paper presents a federated learning algorithm, which is based on the framework of the federated learning with personalization layers. Besides, the batch normalization and meta learning are integrated to alleviate the effects of heterogeneity between various clients' datasets. The proposed algorithm is applied to semantic segmentation of the aerial images from Inria Aerial Image Dataset, and the intersection over union (IoU) is used as an indicator to show the performances of the proposed algorithm.

Keywords: federated learning, semantic segmentation, aerial image

Abstract of Paper Accepted in ICAIC-2026

199

Next-Generation Digital Twin Analytics in Smart Manufacturing using Cross-Layered Edge Cloud Deep Learning

Saisuman Singamsetty, Sudheer Singamsetty, S Satyanarayana
American Unit Inc, San Antonio, Texas, USA.
saisuman.singamsetty@gmail.com
EDNA Technology Consulting Limited, Ontario, Canada.
sudheersingamsetty99@gmail.com
Department of AI & ML, Malla Reddy University, Hyderabad, India.
drssnaiml1@gmail.com

ABSTRACT

The emergence of digital twins is transforming smart manufacturing by enabling real-time system monitoring, predictive maintenance, and optimization of industrial processes. Despite these advantages, traditional cloud-centric analytics face limitations such as high latency, bandwidth constraints, and reduced responsiveness, which hinder deployment in dynamic factory environments. To address these challenges, a cross-layered edge–cloud framework integrated with deep learning is proposed. The architecture employs long short-term memory (LSTM) networks to model temporal dependencies in sensor data, enabling accurate prediction of equipment degradation. Edge devices are responsible for local data acquisition, preprocessing, and lightweight LSTM-based inference, while the cloud layer is utilized for large-scale model training, digital twin synchronization, and advanced analytics. A bidirectional communication strategy ensures that updated cloud models are periodically transferred back to edge nodes, supporting continuous refinement of predictive intelligence while maintaining low-latency inference. The framework is validated using publicly available datasets, including the AI4I 2020 predictive maintenance dataset. Experiments show that our proposed learning model can improve the performance of the cloud manufacturing platform in real-life applications efficiently. Finally, the integration of deep learning with cross-layered edge–cloud computing establishes a scalable and adaptive foundation for digital twin analytics in next-generation smart manufacturing.

Keywords: Digital Twin, Smart Manufacturing, Edge Cloud Computing, Deep Learning, LSTM, Industry 4.0.

Abstract of Paper Accepted in ICAIC-2026

204

Bridging the Gap: Using GPT-4 to Summarize and Explain Static Analysis Warnings at Scale

Arbaz Surti

Independent Researcher, United States

arbaz.m.surti@gmail.com

ABSTRACT

Static analysis tools such as SonarQube are widely adopted in industry to detect bugs, code smells, and security vulnerabilities. However, their effectiveness is often hindered by developer fatigue caused by large volumes of low-priority or cryptic warnings. We investigate how GPT-4, a Large Language Model (LLM), can assist in summarizing and remediating static analysis results. We propose a prompt-based pipeline that converts SonarQube output into rich, contextual guidance using LLMs. By analyzing over 2,700 real-world Java warnings from the Apache Commons Lang project, we compare GPT-4 responses across rule types and severity levels. Our study shows that GPT-4 consistently provides actionable advice that improves upon default tool outputs. We categorize the nature of the guidance, quantify verbosity and structure, and evaluate alignment with developer expectations. This work demonstrates the promise of human-AI collaboration in static analysis triage, extending earlier efforts on LLM-assisted software quality. Our methodology and artifacts are publicly available to support reproducibility and further research.

Keywords: Static analysis, SonarQube, large language models, GPT-4, software quality, prompt engineering, developer productivity, code triage

Abstract of Paper Accepted in ICAIC-2026

210

Multilingual Image-to-Speech Conversion: Leveraging OCR and Language Translation for Enhanced Accessibility

Aryan, Komaldeep Singh, Harkaran Singh, Aakash Rampal
Department of CSE
Chandigarh University
Mohali, India
rohillaaryan0911@gmail.com, kdsguraya26@gmail.com, balakaran32@gmail.com, rampalaakash@gmail.com

ABSTRACT

As digital content is growing rapidly and there is a need for communication that is non-discriminatory, the conversion of textual information from pictures to speech has become a very important matter. This paper introduces a groundbreaking Image-to-Speech Converter which combines Optical Character Recognition (OCR) with language translation and text-to-speech (TTS) synthesis. The system gets the text from the image with the help of a modern OCR model, changes it into the language required, and makes the sound output of very good quality. So the users have access to the information in the language they like. The performance of the system was tested through experiments that measure the accuracy of the system in extracting text, the correctness of the translation, and the clarity of the speech. The intended device has a huge potential in implementing the concept of the technological revolution such as in the field of education, the use of languages, and communicative interaction across languages, thus, creating less distance between the two types of information which are visual and auditory.

Keywords: Image-to-Speech, OCR, Optical Character Recognition, Multilingual Translation, Text-to-Speech, Assistive Technology, Cross-Lingual Communication.

Abstract of Paper Accepted in ICAIC-2026

217

Reengineering Cybersecurity Processes with Generative AI: From Automation to Strategic Alignment

Mehrdad Sharbaf
IEEE, USA
msharbaf@ieee.org

ABSTRACT

Generative AI (GenAI) is rapidly transforming enterprise operations, yet its potential in cybersecurity remains underutilized beyond automation. This paper presents a strategic framework for integrating GenAI into cybersecurity process reengineering, moving from reactive, compliance-driven models to proactive, quality-managed systems. We explore GenAI's capabilities in translating complex frameworks, enhancing stakeholder alignment, and accelerating continuous improvement. Through technical applications and real-world examples, we demonstrate how GenAI can elevate cybersecurity maturity, optimize workflows, and foster strategic resilience.

Keywords: Generative AI, Cybersecurity Process Reengineering, Quality Management, Maturity Assessment,

Abstract of Paper Accepted in ICAIC-2026

219

Agentic Commerce: A Comprehensive Analysis of Cybersecurity Risks, Privacy Challenges, and Trust Mechanisms in Autonomous AI-Driven Marketplaces

Niranjan Pachaiyappan
Visa Inc, USA
tptniranjan@gmail.com

ABSTRACT

Agentic commerce represents a paradigm shift in digital marketplaces where autonomous AI agents execute transactions with minimal human intervention. This paper analyzes cybersecurity risks and privacy challenges through systematic examination of OWASP's 15-category threat taxonomy, regulatory compliance frameworks, and real-world implementations. Traffic from AI browsers increased 4,700% year-over-year, while the global market projects growth from $7.55 billion (2025) to $199.05 billion (2034). Critical vulnerabilities include prompt injection (73% of deployments), credential exposure (100% of SDKs), and multi-agent cascading failures. We evaluate Zero Trust architectures, federated learning, and explainable AI mechanisms, providing actionable recommendations for secure agentic commerce deployment.

Keywords: agentic commerce, autonomous AI agents, cybersecurity risks, prompt injection, data poisoning, Zero Trust architecture, GDPR compliance, privacy-preserving AI, federated learning, explainable AI (XAI), multi-agent systems, OWASP agentic security, identity management, threat modeling, AI governance.

Abstract of Paper Accepted in ICAIC-2025

220

Identity Governance in DevSecOps: Automated Access Reviews for CI/CD Pipelines

Sunnykumar Kamani
SoftSages Technology, USA

ABSTRACT

The critical security issue within CI/CD pipelines and digital cloud infrastructure is the increase in Non-Human Identities (NHIs) in the system. This paper proposes a conceptual approach towards the integration of Identity Governance and Administration (IGA) in the DevSecOps lifecycle. This paper illustrates comprehensive methodology for auto identification, mapping ownership, and periodic access certification for all cloud and CI/CD resources. This approach enhances security posture, streamlines audit and compliance processes and improves significant operational efficiency.The framework proposed in this paper reduces attack surface by implementing with existing tools, HashiCorp Vault, GitLab, Jenkins. This paper also describes future directions that may emanate from the infusion of artificial intelligence and machine learning toward an ultimate goal of a self-governing security ecosystem.

Keywords: CI/CD security, cloud security posture management (CSPM), DevSecOps, Identity Governance and Administration (IGA), Non-Human Identity (NHIs)

Abstract of Paper Accepted in ICAIC-2026

221

A Domain Shift Evaluation Protocol for Fake Satellite Image Detection: CNN Superiority and the Enhanced Model Paradox

Vaishnav Anand
The Athenian School

Ivan Rodriguez
Brown University
vaishnavanand90@gmail.com, ivan_felipe_rodriguez@brown.edu

ABSTRACT

The rapid rise of AI-generated satellite imagery introduces new risks for cybersecurity, national security, and geospatial intelligence. We propose a domain shift evaluation protocol for assessing fake satellite image detection models under realistic deployment conditions. Our protocol evaluates models trained on one set of synthetic imagery against completely different generation methods, simulating operational scenarios where novel synthesis techniques may be used for deception or cyber-manipulation. Using this protocol, we benchmark multiple architectures including CNNs, Vision Transformers, and CLIP models. Key findings include: (1) CNNs demonstrate superior domain shift robustness, with all CNN models achieving ≥99.9% accuracy compared to 96.7%-99.8% for Vision Transformers, (2) CLIP models, despite perfect test accuracy, exhibit severe operational degradation (0.44–0.68% detection rates), and (3) the enhanced model paradox, where training CLIP models with 64% more diverse data leads to 36% worse performance due to over-specialization on the combined training distribution rather than learning generalizable detection principles. Both CLIP models show concerning overconfidence (95–98%) in incorrect predictions. These results establish CNNs as the preferred architecture for satellite image authentication while highlighting critical evaluation methodology limitations that our proposed protocol addresses.

Keywords: Domain shift, fake image detection, CNN, Vision Transformer, CLIP, satellite imagery, evaluation protocol, cybersecurity, national security, geospatial intelligence

Abstract of Paper Accepted in ICAIC-2026

234

Combined Tempered MRG32k3a: A High-Quality and Reproducible Pseudo-Random Number Generator for AI in Cybersecurity

Hussein Alzoubi ,

Senior Member, IEEE Department of Computer Science, German Jordanian University (GJU) Amman, Jordan Email: hussein.alzoubi@gju.edu.jo

ABSTRACT

The random number generator (RNG) is a central component of any cybersecurity system as well as to artificial-intelligence agents. A strong pseudo random number generator should have certain traits such as being statistically robust, reproducible, and computationally efficient in order to be reliably used in critical tasks such as cryptographic operations, training and testing of machine learning models, and simulations. In this paper, we propose a novel RNG based on the well-known L'Ecuyer's MRG32k3a generator. We call the proposed RNG combined tempered MRG32k3a as it combines two independent and multiple recursive pseudo random number generators using a carefully designed tempering layer. The combined tempered MRG32k3a provides highly desirable properties in addition to reproducibility and simplicity, including longer effective period, better statistical uniformity, enhanced spectral characteristics, and improved bit-level diffusion. These traits are critical to AI and cybersecurity. To illustrate the effectiveness of the combined tempered MRG32k3a, experiments at multi-terabyte scales have been carried out using PractRand and TestU01. Results show that the combined tempered MRG32k3a passes all TestU01 SmallCrush tests successfully and shows no anomalies in over than four terabytes regarding PractRand reproducibility. Our results are compared with the state-of-the-art random number generators. We also make our implementation publicly available for researchers and practitioners.

Keywords: pseudo-random number generator (PRNG), reproducibility, tempering, cybersecurity, PractRand

Abstract of Paper Accepted in ICAIC-2026

236

AI-Driven DDoS Detection for Network Security: A Performance Analysis of Machine-Deep Learning Methods on Network Traffic Data

Maunik shah, Google, United States
Manoj Kumar, New York University, United States
maunik.shah27989@gmail.com, mk10141@nyu.edu

ABSTRACT

The proliferation of the use of internet-enabled devices has considerably raised the probability of DDoS attacks, which may overwhelm systems with malicious information and impair network security. It processes a real-time DDoS traffic dataset using ML and DL models in an effort to design an AI-based system of an efficient DDoS detection. The evaluation of the methodology may include accuracy, precision, recall, F1-score, cross-validation, and real-time performance analysis and implies a large amount of preprocessing, feature engineering, and model training using classifiers: SVM (Support Vector Machine), KNN (K-Nearest Neighbors), ND (Naive Bayes), RF (Random Forest), Gradient Boosting (GB), and DNN (Deep Neural Network) and LSTM (Long Short-Memory) models. Results revealed that DNN had the highest accuracy of 93.50 outperforming traditional models and other modern models, including Decision Tree, BiLSTM, BERT, and conventional Naive Bayes models with the best computational performance in real-time application. Study also used interpretability methods like permutation importance, SHAP and LIME to ensure model transparency and actionable insights. The proposed method was demonstrated to be more accurate in detection and balanced than the current methods, which will offer a powerful, scalable, and interpretable solution to dynamic network systems.

Keywords: Cyber Security, DDoS Attack, Network Traffic, Network Security, AI, Machine Learning, Deep Learning.

Abstract of Paper Accepted in ICAIC-2026

250

Enhancing Movie Recommendation Systems Through Explainable Machine Learning Models in the Entertainment Industry

Hariharan Velu
Microsoft, USA
hariharanvelu.research@gmail.com

ABSTRACT

The entertainment industry has realized the role of movie recommendation system in enhancing user experience and finding content because of the proliferation of digital content exponentially. Recommendation systems can propose films to viewers based on the effective use of user behavior, preference and watching history data. The primary goal of this project is to come up with a recommendation system integrating a correlation-based method and sentiment analysis. The IMDb 50K dataset of 25,000 positive and negative reviews was taken as a subset and comprehensive data preprocessing methods such as tokenization, lemmatization, stop words elimination and text pre-processing using Count Vectorizer were used. The reason the model uses the XG Boost algorithm to classify sentiments is because it demonstrates better performance and efficiency. The goal is to assess the suggested models' efficacy by analyzing their ROC-AUC curve, recall, accuracy, precision, and F1-score, as well as their model explainability. Interpretability tools like SHAP and LIME were applied to model predictions to be transparent. Experimental results demonstrate that XG Boost achieved 87.07% accuracy. In comparison, hybrid recommendation model, collaborative filtering, CNN, GRU, RNN, Naive Bayes and Decision Tree classifiers demonstrated lower performance across all metrics. This hybrid model offers high levels of predictive performance and explainability in its decision-making process, and helps fill the explanation-performance gap. The proposed system is notable since it allows making sentiment-conscious recommendations, based on correlations, and can provide a sound, explainable basis for promoting more personalized recommendations in any real-life entertainment system.

Keywords: Movie Recommendation System, Sentiment Analysis, IMDb Reviews, XG Boost, Natural Language Processing (NLP), Explainable AI.

Abstract of Paper Accepted in ICAIC-2026

251

GraphAE: Plant-informed Graph Autoencoder for ICS Anomaly Detection with SHAP-based Explanations

Vahid Heydari, Kofi Nyarko

Morgan State University, United States
vahid.heydari@morgan.edu, kofi.nyarko@morgan.edu

ABSTRACT

Network timing in industrial control systems (ICS) distorts cross-tag relationships and inflates false alarms. We propose a lightweight graph autoencoder (GraphAE) that encodes plant structure: nodes are tags and edges follow P&ID couplings, with cross-phase links up-weighted 2.5x to capture inter-stage propagation. The model is trained on normal data and scores windows by reconstruction error. On the Secure Water Treatment (SWaT) testbed (N=44), we evaluate GraphAE against a non-graph AE-LSTM and the graph-centric GDN under identical preprocessing. All methods are compared at a matched true-positive rate of 0.90 on A2, and we report AUC-ROC and false-positive rate (FPR) both overall and inside near-prop windows (a +/- 10-step band around positives). At this operating point, GraphAE attains AUC-ROC = 0.861 versus 0.870 for AE-LSTM and 0.841 for GDN. FPR (overall) is 0.507 for GraphAE, 0.533 for AE-LSTM, and 0.689 for GDN. In near-prop regions, FPR is 0.556 for GraphAE, 0.625 for AE-LSTM, and 0.641 for GDN. SHAP analyses with a compact A1 background (2,048 samples) highlight the tags driving high scores and support operator triage. These results show that modeling inter-signal structure reduces false alarms in networking-sensitive regimes without meaningful inference overhead. Public SWaT/WADI releases do not include instrument-level P&IDs; our edges are derived from tag-name heuristics and stage structure (no P&ID graph was used for the reported results).

Keywords: industrial control systems (ICS), anomaly detection, graph autoencoder, graph neural networks (GNN), explainable AI (XAI), SHAP, SCADA, PLC

Abstract of Paper Accepted in ICAIC-2026

252

LIARS: Low-Information Area Recognition for Steganography

Claudia Larramendi-Ferras, Alexander Perez-Pons, Gustavo Chaparro-Baquero

Electrical and Computer Engineering Florida International University Miami, USA
clarrame@fiu.edu

ABSTRACT

Steganography conceals information within digital media, making its detection a crucial challenge in modern digital forensics. This paper introduces a framework that integrates machine learning to automatically identify near uniform regions in images, to then embed data in them, and to validate the presence of such concealed content through a controlled forensic pipeline. The proposed workflow segments images into patches, classifies them as high-information or low-information using a convolutional neural network, and embeds a textual payload exclusively within low-entropy zones. A validation process based on difference-map and optical character recognition module are used to extract and confirm the hidden text. Experimental results demonstrate that the proposed work can reliably detect and reconstruct hidden payloads even when pixel intensity modifications are as subtle as ±5 RGB units, imperceptible to the human eye. The results highlight the potential of low-information modeling and adaptive AI segmentation as effective tools for explainable and forensic-oriented steganalysis.

Keywords: Steganography, Steganalysis, Machine Learning, Image Segmentation, Low-Information Areas, Digital Forensic

Abstract of Paper Accepted in ICAIC-2026

255

KyVul, an LLM Created C/C++ Vulnerability Dataset

Bryson Brown, Martin Carlisle
Texas A&M University
Department of *Computer Science and Computer Engineering*
College Station, Texas
bryson.brown@tamu.edu, carlislem@tamu.edu

ABSTRACT

Software vulnerabilities have long been a significant issue, serving as the root cause of many cyber-attacks. The use of artificial intelligence to identify these vulnerabilities is of great interest to both researchers and businesses; however, the lack of high-quality data has constrained much of the research in this area. Most existing datasets are constructed by labeling code that has already been written by humans. In this context, KyVul stands out as the first large language model (LLM) dataset focused on C/C++ vulnerabilities, boasting an impressive overall accuracy rate of 97%. This level of accuracy surpasses that of most existing benchmark datasets. When combined with the quality-controlled dataset PrimeVul for training, the area under the precision-recall curve (AUPRC) experienced an average increase of 47.5%, while the area under the receiver operating characteristic curve (AUROC) increased by an average of 1.2%. Both enhancements were found to be statistically significant, as confirmed by an unpaired t-test. This dataset effectively showcases the potential of LLMs to generate large quantities of both vulnerable and non-vulnerable code.

Keywords: vulnerabilities, datasets, LLM

Abstract of Paper Accepted in ICAIC-2026

265

Detecting and Adapting to Normality Shifts in Learning-Based Security Anomaly Detection

Gaurav Dwivedi, Alexander Perez-Pons

Department of Electrical and Computer Engineering, College of Engineering and Computing Florida International University Miami, FL 33174 USA

Email: gdwiv001@fiu.edu, aperezpo@fiu.edu

ABSTRACT

Learning-based security applications are faced with many issues such as concept drift, which regularly leads to frustrations and decreased performance of the models as they get stale. It is a common assumption of these applications that both training and deployment models are identical, or the close-world assumption. Zero-positive anomaly detection is one of the most significant tasks in the security domains as it is resistant to the drift of abnormal behavior when trained without abnormal data. However, when normality shifts, this immunity results in more serious effects. The normality shift for zeropositive anomaly identification has received little attention in previous works, which have mostly concentrated on the drift concept of anomalous behavior. This paper promotes a general framework, to close this gap for deep learning-based anomaly detection in security applications. It recognizes, elaborates, and adjusts to normality shift in practice. By detecting shifts in unsupervised manner, eliminating the requirement for manual labeling, and improving adaption performance by distributionlevel tackling helps in surpassing earlier methods. Three securityrelated anomaly detection applications are used to demonstrate the effectiveness of the model in various practical experiments with real-world long-term data. Results show that model offers superior normality shift adaption performance with reduced labeling overhead.

Keywords: Nomality Shifts, Learning, Anomaly Detection

Abstract of Paper Accepted in ICAIC-2026

266

Designing Effective Quantum Generators: A Comparative Study of Variational Ansätze in Hybrid QGANs

Sarvapriya Tripathi, Aakarsh Etar, Jayesh Soni, Himanshu Upadhyay
Department of Electrical & Computer Engineering
Florida International University
Miami, FL, USA
strip001@fiu.edu, aetar@fiu.edu, jsoni@fiu.edu, upadhyay@fiu.edu

ABSTRACT

Quantum Generative Adversarial Networks (QGANs) combine quantum circuit based generative models with classical discriminators to exploit quantum advantages in generative modeling. This work investigates a hybrid QGAN architecture with a quantum-enhanced generator and a classical discriminator, applied to the Modified National Institute of Standards and Technology (MNIST) dataset. We implement and benchmark three different quantum circuit ansätze, namely a basic entangler circuit, a simple two-design circuit and a strongly entangling circuit, with each QGAN model being constructed using six layers of the respective ansätze. To manage the 28x28 image complexity, the generator employs a patch wise scheme, dividing each image into small patches generated by separate sub-circuits. The performance of each ansatz is evaluated in terms of training stability and output image quality. Our results demonstrate that the choice of quantum circuit ansatz significantly impacts the QGAN's ability to learn the distribution of handwritten digits. In particular, the strongly entangling ansatz achieves the best performance at the cost of stability and image fidelity among the three yielding recognizable digit images. While no quantum advantage was demonstrated, this study provides comparative analysis and insights into designing effective quantum generators for hybrid QGANs on image data, serving as a benchmark for future developments in quantum-enhanced generative models.

Keywords: ***Quantum Generative Adversarial Network (QGAN), Hybrid Quantum-Classical GAN, Variational Quantum Circuit, Ansatz, Quantum Machine Learning***

Abstract of Paper Accepted in ICAIC-2026

274

LLM-Based Decision Making Framework for Autonomous Drone Navigation

Mirza Aarish Baig, Brad Alvarez, Richard Lage, Jayesh Soni, Himanshu Upadhyay
Florida International University, USA

ABSTRACT

Drones are getting more and more common in today's world. However, maneuvering them is still a challenge for most people. It becomes difficult to scale up the number of drones deployed if each drone needs an operator. With Large Language Models (LLMs) gaining popularity, it may be possible to utilize them to simplify the control scheme of a drone and to assign them as drone operators. This paper presents an LLM driven framework for autonomous drone navigation that integrates perception, reasoning, and control within a unified architecture. The system operates in a high-fidelity Unreal Engine 5 simulation and combines real-time visual processing using the YOLOE-11-L detector with language-based reasoning for decision-making. The LLM interprets visual encodings, maintains contextual awareness, and generates flight actions executed through the PX4 control stack. To evaluate the framework, a representative mission, “Land on a couch”, is performed in a realistic 3D environment reconstructed from real-world imagery. Experimental results demonstrate that the LLM can effectively interpret sensor data, reason about spatial goals, and plan safe trajectories without domain-specific training, highlighting the potential of language-based reasoning for embodied aerial autonomy.

Keywords: LLM, VLM, Autonomous, Drone simulator, Foundation Model

Abstract of Paper Accepted in ICAIC-2026

275

Federated Transformer Model for Water Contamination Detection in Distributed Monitoring Systems

Jayesh Soni, Raja Kumar, Himanshu Upadhyay
Florida International University, USA

ABSTRACT

Safety of water distribution systems is a serioussocietal health concern. The emergence of IoT sensors has facilitated real time monitoring, yet this creates huge time-series data that is high-dimensional and generates massive amounts of data that are often isolated across various utilities for privacy and security reasons. Such fragmentation of data prevents the creation of effective anomaly detection models. We introduce a privacy-sensitive federated learning (FL) system to water contamination detection based on a Transformer-based architecture. Our solution is based on the FedAvg algorithm that allows several, distributed water utilities to jointly learn a strong global model, without ever exchanging their raw, sensitive sensor data. The self-attention module of the Transformer is particularly useful when attempting to detect intricate spatio-temporal relationships and long-range connections among various water quality metrics (e.g., pH, turbidity, chlorine), which tend to be antecedents of contamination occurrences. We show that the federated Transformer method shows high levels of detection performance, validating its potential to develop a generalized and resilient system of anomaly detection and strictly adhering to the data privacy requirements necessary to critical infrastructure.

Keywords: Federated Learning, Transformer, Anomaly Detection, Water Contamination, Time-Series, Privacy, SelfAttention

Abstract of Paper Accepted in ICAIC-2026

284

Prompt Engineering vs Context Engineering: A Strategic Framework for Insurance AI Applications

Rakesh More
A J Gallagher, USA
rakeshmore@gmail.com

ABSTRACT

The rapid evolution of artificial intelligence in the insurance sector has led to transformative changes in underwriting, claims processing, risk management, customer service, and compliance. In particular, the engineering of input prompts and contextual information for large language models (LLMs) has emerged as a critical factor in achieving reliable, explainable, and efficient AI systems. This paper compares two prominent approaches: prompt engineering, which leverages task instructions and demonstration examples (e.g., zero-shot and few-shot techniques), and context engineering, which systematically structures diverse sources of information into dynamic and modular inputs for LLMs. Drawing upon extensive research on context engineering methodologies alongside practical insights from insurance AI applications, we propose a strategic decision framework to guide stakeholders in selecting the most suitable engineering method for specific use cases in underwriting, fraud detection, claims automation, and regulatory compliance. The framework is substantiated with empirical metrics, visualization of decision flows, and detailed case studies illustrating performance, transparency, and explainability improvements. Our findings underscore that while prompt engineering offers simplicity and ease-of-deployment for less complex tasks, context engineering provides enhanced performance, stability, and explainability for high-stakes insurance applications.

Keywords: artificial intelligence, insurance technology, prompt engineering, context engineering, retrieval augmented generation, risk management, regulatory compliance

Abstract of Paper Accepted in ICAIC-2026

287

Evaluating Adversarial Resilience of Deep Reinforcement Learning Algorithms for Network Intrusion Detection

Curtis Rookard
Indian River State College, United States

ABSTRACT

Recent advancements in the field of cybersecurity has allowed for greater defensive technologies. The application of machine learning and artificial intelligence in cybersecurity has allowed for greater detection of cyber threats, especially in the field of network intrusion detection. Reinforcement learning, a prominent subfield in machine learning, aims to train an intelligent agent through providing rewards in response to stimuli in an environment. These advances, however, are often subject to adversarial attacks. In this study, we propose the application of a reinforcement learning-based network intrusion detection system by utilizing a deep recurrent Q-network for threat classification. We then implement adversarial machine learning attacks using the Fast Gradient Sign Method and Basic Iterative Method against our intrusion detection system. Our results indicate that the adversarial attacks were able to reduce the effectiveness of our intrusion detection systems on certain performance metrics by as much as 25%.

Keywords: Reinforcement Learning, Q-Learning, Deep QNetworks, Recurrent Neural Networks, Adversarial Machine Learning, Adversarial Attacks, Network Intrusion Detection Systems

Abstract of Paper Accepted in ICAIC-2026

290

Resource Allocation of IoT-Based Edge Computing in Smart Cities

Abdullah Alqahtani Jazan University, Saudi Arabia

amqahtani@jazanu.edu.sa

ABSTRACT

The rapid proliferation of Internet of Things (IoT) devices is accelerating the development of smart cities, where billions of sensors and actuators collaborate to enhance safety, efficiency, and sustainability. Applications such as intelligent transportation, smart grids, healthcare, and surveillance generate massive, heterogeneous data streams that demand timely processing and reliable decision-making. Traditional cloud computing provides scalable storage and analysis but struggles to meet latency and bandwidth requirements for mission critical tasks. Edge computing addresses these limitations by bringing computation closer to data sources, enabling low-latency, energy-aware operations. However, resource allocation in smart city systems remains a critical challenge due to limited computing power, memory, spectrum, and energy, constrained by device capabilities and shared infrastructure. This paper proposes a novel multi-stage framework for IoT-based edge computing that integrates predictive workload modeling, proportional fairness distribution, and adaptive Lagrangian optimization. The framework ensures scalable and fair allocation of heterogeneous resources while minimizing energy consumption and latency. Experimental results demonstrate a significant reduction in latency (62 ms vs. 80–95 ms), increased throughput (182 Mbps vs. 130–150 Mbps), improved energy efficiency (39 J vs. 48–55J), and enhanced system utilization (92% vs. 71–78%). Reliability improves to 96%, while fairness and security ratings reach 9 on a normalized scale, outperforming existing approaches. By incorporating federated learning and blockchain, the framework enhances privacy, accountability, and resilience. Overall, the proposed model establishes a robust, scalable, and sustainable approach for smart city IoT resource allocation, advancing both technological efficiency and social responsibility.

Keywords: Adaptability, Edge computing, Energy efficiency, Fairness, IoT resource allocation, Latency optimization, Reliability, Scalability, Smart cities, Sustainability.

Abstract of Paper Accepted in ICAIC-2026

296

Secure Agent-Based Architectures for Decentralized AI Identity Management: The DAIS Framework

Viswapriyan Ragupathy
IEEE Senior Member
8753 Olenbrook Dr, Lewis Center, Ohio – 43035, USA
viswapriyan.ragupathy@ieee.org

ABSTRACT

This study proposes the multi-layered Decentralised Autonomous Identity System (DAIS) for safe, scalable, and policy-aware identity management for autonomous AI agents. DAIS uses distributed ledgers, trust registries, cryptographic credential processing, and adaptive governance principles to authenticate and authorise without central authority. For security and operational reliability, the architecture uses efficient verification techniques, Byzantine fault-tolerant consensus, hierarchical credential state management, and behaviour-aware trust scoring. The mathematical definitions of credential validation, revocation propagation, and consensus guarantees show how DAIS accomplishes constant-time verification and quick revocation dissemination over distributed systems. A 50-agent testbed, PBFT-based registry duplicates, and enterprise-grade hardware were used to evaluate DAIS to OAuth 2.0 and OpenID Connect under identical settings. DAIS regularly outperforms centralised solutions in latency, revocation speed, computational overhead, interoperability, and scalability, improving verification time by fourfold and revocation responsiveness by sixteen-fold. These findings show that DAIS is a durable, high-performance, and future-ready identification infrastructure for large autonomous agent ecosystems.

Keywords: decentralized identity, autonomous agents, trust registry, cryptographic verification, Byzantine consensus, scalability, revocation, interoperability

Abstract of Paper Accepted in ICAIC-2026

297

Architectural Resilience in AI-Driven Decision Systems under Adversarial Conditions

Viswapriyan Ragupathy
IEEE Senior Member
8753 Olenbrook Dr, Lewis Center, Ohio – 43035, USA
viswapriyan.ragupathy@ieee.org

ABSTRACT

The Self-Healing AI Architecture shows that resilient cognitive layers improve multi-agent AI systems' stability, autonomy and dependability. The method decreases recovery latency, cross-agent failure propagation and accuracy retention across failure scenarios with minimal operating overhead. Structured recovery patterns such as Semantic Checkpointing, Cognitive Rollback, Quarantine Isolation and Adaptive Consensus Reset enable systematic reasoning-level degradations, and empirical validation across 20 cooperative agents proves the architecture's scalability and enterprise suitability. These findings emphasise cognitive-layer robustness for reliable agentic AI. Several architectural upgrades are planned. In predictive fault models, time-series analysis and anomaly detection may enable proactive recovery before reasoning fails. Reinforcement learning–based optimisation may choose experience-based recovery approaches under changing situations. Vision, reinforcement learning, and symbolic agents enhance application beyond LLM-based agents. Cognitive-layer, infrastructure-level resilience, and formal verification may improve large-scale deployment accuracy and safety.

Keywords: Adversarial resilience, AI security, robust architecture, design-science methodology, resilient decision systems, anomaly detection, redundant inference, DevSecOps, trustworthy AI

Abstract of Paper Accepted in ICAIC-2026

300

Hybrid Intelligence Endpoint Defense (HIED): A Data-Fusion Approach for Proactive Malware Detection

Michal Jaworski, Bahareh Pahlevanzadeh
(TU Dublin)
michal.jaworski11@gmail.com, Bahareh.Pahlevanzadeh@tudublin.ie

ABSTRACT

Hybrid Intelligence Endpoint Defense (HIED) fuses Endpoint Detection and Response (EDR) telemetry with Cyber Threat Intelligence (CTI) to improve malware detection using only existing enterprise telemetry and publicly available threat feeds. We collect synthetic and real Sysmon logs from Windows 10 endpoints and enrich them with AlienVault Open Threat Exchange (OTX) indicators. We then compare four fusion strategies—feature-level (early), decision-level (late) with two pipelines, late fusion with two classical models, and a combined approach—using Random Forest, XGBoost, and Local Outlier Factor. Unlike prior EDR–CTI "integrations" that apply post-hoc lookups, we inject CTI at the feature level and empirically quantify early vs late fusion. Early fusion performs best, improving Matthews Correlation Coefficient (MCC) from 0.94 (EDR-only) to 0.98, recall from 87.58% to 99.60%, and reducing false positive rate from 0.09% to 0.02% on a held-out validation set. We attribute these gains to CTI anchoring the decision boundary with external threat context, which mitigates alert fatigue. Aligned with ICAIC'2026 Tracks 8 (AI in Cyber Security), 12 (AI in Malware Analysis/Digital Forensics), and 6 (AI in IT Infrastructure), HIED shows what more can be done with standard Sysmon logs and open CTI—no new sensors required. We discuss deployment considerations and limitations, and outline future work on real-time fusion, multi-feed CTI, and broader malware coverage.

Keywords: AI in Cyber Security; Malware Analysis; Digital Forensics; IT Infrastructure; Endpoint Detection and Response; Cyber Threat Intelligence; Data Fusion; Sysmon; FastText

Abstract of Paper Accepted in ICAIC-2026

301

Phishing URL Detection Using RNN – LSTM Models for Safer Web Browsing

Roudha Bin Redha, Maitha Alsammak, Mohamed Almourad
Zayed University, UAE
r.binredhaa@gmail.com, M80009456@zu.ac.ae, basel.Almourad@zu.ac.ae

ABSTRACT

This position paper, conducted by Master students at Zayed University investigates phishing URL detection using deep learning, with a focus on RNN–LSTM models. The paper presents a systematic literature review of recent IEEE and ACM studies to analyze the strengths, limitations, and comparative performance of traditional machine learning, standalone deep learning, and hybrid architectures in URL-based phishing detection. The review identifies key limitations in current models, including high computational cost, lack of interpretability, and challenges in real-time deployment on resource-constrained platforms such as browsers and mobile clients. Building these findings, the paper identifies a research gap in lightweight, adaptive detection frameworks that can balance accuracy and efficiency for practical deployment in browsers and mobile clients. Instead of reporting experimental model results, the paper establishes a conceptual framework, clarifies research gaps, and justifies the need for further experimentation. This work, therefore, lays out the foundation for future experimental development and evaluation of a lightweight, adaptive phishing detection model suitable for deployment in resource-constrained client environments.

Keywords:

Abstract of Paper Accepted in ICAIC-2026

304

Swarm Optimization Algorithm-Enhanced Clustering Techniques for Reliable Wireless Sensor Networks Communication

Mitesh Patel, Uday Korat
IEEE Senior Member
USA
mitesh.rf@gmail.com, ukorat@gmail.com

ABSTRACT

The Wireless Sensor Networks (WSNs) are ubiquitous in the smart city, industrial surveillance, and environmental research literature but continue to have the fundamental challenges of inconsistent power supply, ineffective data management, and connectivity issues. The clustering of WSNs is a robust approach to control the topology of networks, allocate loads in the similar way, and extend the life of the networks. However, it is NP-hard to find a good cluster and the conventional approaches cannot always provide a good balance between power consumption and network survivability. This study presents a Hybrid Ant Colony Optimization-Particle Swarm Optimization (ACO-PSO) model that not only identifies Cluster Heads (CHs) in an intelligent manner, but also identifies optimal routing paths, and thus guarantees efficient energy management. The proposed ACO-PSO algorithm used very low energy (EC) of 3.920 J and high packet delivery ratio (PDR) of 97.94 when it was used through extensive simulations (50 nodes). The effectiveness of this approach is much higher compared to existing approaches, including ACO, PSO-ECSM, SSA, and MIGJOA. The framework offers a uniform coverage at various magnitude levels of the network, end-to-end latency is minimized substantially, and the scalability and reliability of swarm optimization are demonstrated in prospective resource-constrained wireless sensor network utilizations.

Keywords: *Wireless Sensor Networks (WSNs), Swarm Intelligence (SI), Cluster Head (CH) Selection, Energy Efficiency, Network Reliability, Optimization.*

Abstract of Paper Accepted in ICAIC-2026

306

Integrating Price Elasticity and Reinforcement Learning: A Data-Driven Framework for Strategic E-commerce Pricing

Dilip Patel
Group Product Manager, Uber Risk Management
Uber, USA
dilip.patel.cal@gmail.com

ABSTRACT

The retail business is an important aspect towards growth and general development. The growing competition among retailers makes customer acquisition and retention strategies increasingly significant. These strategies are based on proper pricing of products which has a direct impact on customer loyalty and revenue. The research paper examines reinforcement learning (RL) methods to dynamic e-commerce pricing based on the publicly available Retail Price Optimization dataset. Two RL-based pricing agents were Deep Q-Network (DQN) agent to price real-time dynamically and Monte Carlo Tree Search (MCTS) agent to price sequentially. The MSE, RMSE, MAE, MAPE, and R^2 measures were used to assess model performance, and reward analysis was employed to evaluate consistency and stability. Findings indicate that the DQN agent exhibits better predictive accuracy and consistent pricing behavior, as demonstrated by a high R-squared value of 98.74, a minimum reward variance of 8.24, and an RMSE of 8.24. Ablation research integrating price elasticity information and RL agents also enhance accuracy and minimize price deviation. The tests of statistical significance demonstrate that ensemble predictions are more efficient than single-agent techniques. The adaptability of RL-based agents especially DQN to the dynamic market conditions is better than traditional methods like XGBoost, multiple linear regression, and decision trees. The paper introduces a strong, data intensive framework in the e-commerce pricing that gives practical strategies grounded on elasticity, profitability optimization as well as the ability to make real time, competitive pricing decisions based on the product categories.

Keywords: Dynamic pricing, reinforcement learning, Deep Q-Network (DQN), Monte Carlo Tree Search (MCTS), e-commerce, price optimization, price elasticity.

Abstract of Paper Accepted in ICAIC-2026

313

Privacy-Preserving Federated Deep Learning for Nomophobia Risk Prediction from Smartphone Usage Logs

Md Wahidur Rahman[1], Mehdi Hasan[2], Mais Nijim[1], Muhammad Armughan Ul Haq[1], and Fawaz Ali Mohammed[1]
[1]Department of Electrical Engineering and Computer Science, Texas A&M University–Kingsville, Kingsville, TX 78363, USA
[2]Division of Computer Science, Mathematics, and Science, St. John's University, New York, USA
Corresponding author email: md_wahidur.rahman@students.tamuk.edu

ABSTRACT

Nomophobia, the fear or discomfort of being without a smartphone or connectivity, is increasingly common in young people and is linked to anxiety, poor sleep, and reduced well-being. Existing detection methods rely mainly on self-report scales or centralized machine learning, which limits real-time use and raises privacy concerns. This study proposes a privacy-preserving federated deep learning framework for nomophobia risk prediction based on real smartphone usage logs. We expand a public Kaggle dataset from 1{,}000 to 50{,}000 synthetic users, apply a unified preprocessing and percentile-based risk labeling scheme, and use particle swarm optimization to select informative usage features. Multiple federated models (1D CNN, LSTM, BiLSTM, and a GNN-style network) are trained with FedAvg over ten simulated clients and are compared against centralized baselines including SVM, RF, KNN, DT, XGB, and NB. The CNN-based federated model achieves an accuracy of 83.79\%, an F1-score of 0.8642\%, and an MCC of 0.6912\% on the unseen test set, while centralized RF reaches 94.15\% accuracy and a 0.8720 MCC. These experimental results indicate that on-device federated learning can approach centralized performance while keeping raw usage data local, providing a practical path toward continuous, privacy-aware nomophobia monitoring.

Keywords: Federated learning; privacy-preserving deep learning; nomophobia risk prediction; smartphone usage logs.

Abstract of Paper Accepted in ICAIC-2026

314

Enterprise-Grade Security for the Model Context Protocol (MCP): Frameworks and Mitigation Strategies

Vineeth Sai Narajala, Meta, USA abcvineeth.sai@gmail.com
Edan Habler, Cisco, USA edan.habler@gmail.com

ABSTRACT

The Model Context Protocol (MCP), introduced by Anthropic, provides a standardized framework for artificial intelligence (AI) systems to interact with external data sources and tools in real-time. While MCP offers significant advantages for AI integration and capability extension, it introduces novel security challenges that demand rigorous analysis and mitigation. This paper builds upon foundational research into MCP architecture and preliminary security assessments to deliver enterprise-grade mitigation frameworks and detailed technical implementation strategies. Through systematic threat modeling and analysis of MCP implementations and analysis of potential attack vectors, including sophisticated threats like tool poisoning, we present actionable security patterns tailored for MCP implementers and adopters. The primary contribution of this research lies in translating theoretical security concerns into a practical, implementable framework with actionable controls, thereby providing essential guidance for the secure enterprise adoption and governance of integrated AI systems.

Keywords: Model Context Protocol (MCP), AI Security, Zero Trust Architecture (ZTA), Tool Poisoning, Defense-in-Depth, Operational Security, Secure AI Integration, API Security, AI Governance

Abstract of Paper Accepted in ICAIC-2026

315

Agent Name Service (ANS): A Universal Directory for Secure AI Agent Discovery and Interoperability

Ken Haung, OWASP, USA, kenhuangus@gmail.com
Vineeth Sai Narajala, Meta, USA abcvineeth.sai@gmail.com
Edan Habler, Cisco, USA edan.habler@gmail.com
Akram Sheriff, Cisco, USA, sheriff.akram.usa@gmail.com

ABSTRACT

The proliferation of AI agents requires robust mechanisms for secure discovery. This paper introduces the Agent Name Service (ANS), a novel architecture based on DNS addressing the lack of a public agent discovery framework. ANS provides a protocol-agnostic registry mechanism that leverages Public Key Infrastructure (PKI) certificates for verifiable agent identity and trust. The architecture features several key innovations: a formalized agent registration and renewal mechanism for lifecycle management; DNS-inspired naming conventions with capability-aware resolution; a modular Protocol Adapter Layer supporting diverse communication standards (A2A, MCP, ACP, etc.); and precisely defined algorithms for secure resolution. We implement structured communication using JSON Schema and conduct a comprehensive threat analysis of our proposal. The result is a foundational agent directory service protocol addressing the core challenges of secure discovery and interaction in multi-agent systems, paving the way for future interoperable, trustworthy, and scalable agent ecosystems.

Keywords: Agent Name Service (ANS), Agentic AI, Service Discovery, Public Key Infrastructure (PKI), Interoperability, Secure DNS, Formal Methods, Multi-Agent Systems (MAS).

Abstract of Paper Accepted in ICAIC-2026

322

Agentic AI Integration for Process Automation in MSMEs

Venkatesh Prabu Parthasarathy
PROPHECY CONSULTING INC
venkateshprabu2003@gmail.com

ABSTRACT

This study investigates the integration of agentic AI-driven process automation in India's micro, small, and medium enterprises (MSMEs) to enhance productivity, efficiency, and competitiveness. Employing a mixed-methods approach, the research combines qualitative insights from interviews and surveys with quantitative analysis of operational metrics, including accounting accuracy, compliance consistency, and task efficiency. Findings reveal strong receptivity toward AI adoption across generational and sectoral lines, with formalized businesses demonstrating higher readiness due to access to digital infrastructures such as GST APIs and e-invoicing systems. Agentic AI systems—capable of autonomous decision-making, predictive analysis, and adaptive learning—show potential to optimize accounting, inventory management, customer service, and regulatory compliance. The study also introduces an AI-augmented digital guide that facilitates formalization and AI literacy among MSME owners. Results indicate that agentic AI adoption can transform Indian MSMEs into data-intelligent, self-optimizing organizations, fostering sustainable growth, innovation, and global competitiveness while bridging gaps in digital literacy and operational efficiency

Keywords: Agentic AI, Process Automation, MSMEs, India, Accounting Automation, GST Compliance, AI Adoption

Abstract of Paper Accepted in ICAIC-2026

326

Advancing Cybersecurity in Critical Infrastructure Systems via Machine Learning-Based Threat Detection and Mitigation

Siva Teja Reddy Kandula
Independent Researcher
sivateja.kandula@ieee.org

ABSTRACT

Cyberattacks represent a serious risk to any company, but especially those involved in critical infrastructure (CI). Modern cyberattacks are more complex, multi-vectored, and unpredictable, making the job of cyber security risk managers (CSRM) more difficult. To manage these dangers and lessen the impact, critical infrastructure must have an additional layer of protection. The goal of Cyber Threat Intelligence (CTI) is to prevent cyber risks by providing evidence-based information about these risks. This study introduces a DL-based IDS that employs the hybrid convolutional neural network (CNN) + bi-directional Long short-term memory (Bi-LSTM) model, to increase detection accuracy when complex network traffic is present. The UNSW-NB15 dataset is utilized to train and assess the model, which includes 49 features and several types of attacks. To prepare data for analysis, one must eliminate duplicates, deal with missing values, and use Min-Max normalization. To provide a robust evaluation, the dataset is separated into 20% testing and 80% training, and the key features are chosen using feature significance analysis. The propose hybrid CNN+Bi-LSTM model achieved 99.78% accuracy, 98.67% precision, 99.99 % recall and 98.89% F1score and outperformed more conventional variants of DNN, DBN and SVM models. Minimal overfitting and excellent generalization are confirmed with the confusion matrix and learning curves. This shows how DL can improve the capabilities of IDS against modern-day cyberattacks.

Keywords: *Cybersecurity, Cyber Threat Detection, Intrusion Detection System (IDS), Network Attacks, Deep Learning, UNSW-NB15 Dataset.*

Abstract of Paper Accepted in ICAIC-2026

330

Assessing Agile Frameworks for Software Development Efficiency Industry Insights and Implementation with Jira Atlassian Agile Tools

Ravi Sankar Susarla
Collabera, USA
raviangirasa@ieee.org

ABSTRACT

Software development projects often struggle with predictability, resource allocation and resolving issues in a timely manner, leading to inefficiencies in Agile workflows. This study addresses these challenges through data mining of caused data from Jira Atlassian software to mix applied analytics for simulation and process modeling and optimization of the Agile Scrum process. Project issues were extracted from the Jira platform. Using these collected data a program application in Python with the use of pandas, NumPy, Plotly, and Seaborn was utilized to analyze the data. Descriptive statistics, the correlation between features, and Kaplan-Meier methods were used to modelling tickets. Results revealed issues taken an average time of approximate 1.82 hours and lead time of 0.07 days. Furthermore, discrete-event simulation was utilized to model the project issues of development, testing, and project review workflows. Monte Carlo simulations were utilized to improve planning with respect to estimating sprints and catching areas of potential problems. Interactive visualizations supported all analyses and included network graphs and Sankey diagrams. This result indicates that analyzing historical data from Jira and simulation using discrete event and Monte Carlo approaches, provided opportunities to improver sprint predictability, workflow inefficiencies, and data-driven decisions for efficient Agile project management.

Keywords: Agile Software Development, Python Simulation, Jira Analytics, Resource Management, Sprint Planning, Monte Carlo Simulation, Workflow Optimization.

Abstract of Paper Accepted in ICAIC-2026

331

Privacy-Driven Cloud AI: Federated Learning and Zero-Trust for Secure Multi-Domain Collaboration

Akshay Mittal
University of the Cumberlands, USA
Prashanthi Matam, Vasanth Rao Jadav
Independent Researcher
Karthik Pappu
Dakota State University, USA
amittal18886@ucumberlands.edu, prashumatam@gmail.com, karthik.pappu@trojans.dsu.edu, jadavvasanth@gmail.com

ABSTRACT

This study introduces GWO-FLNet, a privacy-preserving cloud AI framework that combines Federated Learning (FL), Zero-Trust Architecture (ZTA), and the Grey Wolf Optimizer (GWO) to enable secure and accurate collaboration across distributed, multi-domain systems. The primary objective is to improve model accuracy while ensuring data privacy in scenarios where client datasets are statistically heterogeneous and domainspecific. The proposed method employs a CNN-BiLSTM hybrid architecture to capture spatial-temporal patterns, trained locally at each client without sharing raw data. To address inter-client variability and training imbalance, GWO dynamically tunes the learning rate, BiLSTM units, and personalized aggregation weights during the federated training process. Secure authentication and access control are enforced using ZTA principles across all clients. Experimental results demonstrate significant improvements over conventional FL approaches (e.g., FedAvg, FedProx). GWO-FLNet achieved an average increase of 7.8% in Dice Similarity Coefficient (DSC), 8.3% improvement in Intersection over Union (IoU), and 6.4% higher F1- score, while reducing communication overhead by 15%. These results highlight the framework's robustness and scalability in environments characterized by decentralized, non-uniform data distributions, making it well-suited for real-time, privacy-sensitive cloud deployments.

Keywords: Federated Learning, Zero-Trust Architecture, Grey Wolf Optimizer, Privacy-Preserving Machine Learning, Data Heterogeneity, Cloud AI Security, CNNBiLSTM

Abstract of Paper Accepted in ICAIC-2026

335

Generative Temporal Diffusion Models for Early Prediction of Cloud Service Degradation

Sabitha Muppuri
Antioch, California - 94531, USA
sabitha594@gmail.com

ABSTRACT

Cloud computing infrastructure depends majorly on predictive analytics to assure reliability. However, current machine learning systems find difficulties in forecasting service degradation due to the presence of noisy and rapidly shifting telemetry. To overcome these issues, this paper introduces a generative temporal diffusion model (GTDM) for predicting future cloud performance trajectories by reversing a learned corruption process. The proposed framework produces a probabilistic distribution describing potential future system states rather than issuing single, point-in-time predictions. This assists a timely detection of critical issues such as resource exhaustion and unsteady services. The proposed model incorporates temporal self-attention, conditioning on the diffusion step, and a Transformer-style refinement process to remove noise. These components allow it to capture dependencies across long time spans and accurately reconstruct plausible trajectories, even when significant portions of the input data are corrupted. Critical evaluation of the proposed approach confirms and validates that the proposed approach surpasses current techniques based on the models such as LSTM, TCN, and standard transformers. Overall, the findings from the proposed model suggest its effectiveness in operational cloud prediction.

Keywords: cloud computing; temporal diffusion; transformer; LSTM; reliability

Abstract of Paper Accepted in ICAIC-2026

337

TrustNet: A Hybrid Machine Learning and LLM-Based Multi-Agent System for Scam Website Detection

Sherif Abdelhamid, James Bangura, Katelyn Redlinger, Gunnar Romsland
Virginia Military Institute Lexington, VA, USA
abdelhamidse@vmi.edu, bangurajs26@vmi.edu, redlingerke27@vmi.edu, romslandgc26@vmi.edu

ABSTRACT

The increase in the number of phishing and scam websites adds a continuous threat to online users. In this research, we present TrustNet, a hybrid AI platform that integrates traditional machine learning (ML) models with large language models (LLMs). We used different machine learning algorithms, and the Random Forest classifier achieved the highest performance, with an Accuracy of 0.9694, a Precision of 0.9749, a Recall of 0.9628, and an F1 of 0.9688. Further feature correlation and network centrality analyses showed the most influential attributes and their contribution to characterizing scam websites. Additionally, we built a multi-agent LLM system to identify and analyze linguistic and contextual indicators of scams in website content. The system contains four specialized agents (Market Researcher, Language Analyst, Threat Analyst, and Evaluator). The agents work together to identify urgency, vagueness, grammatical inconsistencies, and deceptive pressure tactics in web content. We tested the LLM-based agent system on our custom dataset of real and synthetic website text samples. The team of agents achieved high detection performance with an Accuracy of 0.9931, a Precision of 0.9867, a Recall of 1.0, and an F1 score of 0.9933. Both ML models and LLM-driven agents were integrated into a web-based application named TrustNet. TrustNet is a platform for detecting scams and phishing, and for explainable risk assessment. The system provides an interactive dashboard with risk-level visualization, linguistic explanation panels, and real-time recommendations to users.

Keywords: Phishing detection, scam websites, machine learning, large language models (LLMs), multi-agent systems, hybrid intelligence, online deception, cybersecurity, GPT-4o, feature analysis, linguistic analysis

Abstract of Paper Accepted in ICAIC-2026

338

Causal-Contrastive Graph Neural Networks for Robust, Explainable Multi-Modal Intrusion Detection

Sangharsh Agarwal
Antioch, California – 94531, USA
sangharshcs@gmail.com

ABSTRACT

Robustness and explainability are major challenges for machine learning-based intrusion detection in dynamic cyber-physical environments. Current models often overfit to dataset-specific traits, which results in poor generalization when concepts change. This work presents Causal-Contrastive Graph Neural Networks (CCGNN), a new framework that aims to learn consistent, meaningful representations of malicious behavior. CCGNN uses a method called differentiable causal masking to remove misleading correlations. It also uses graph-based relational modeling to capture the complex interactions between hosts and network flows. In addition, a temporal contrastive goal ensures that representations remain consistent over time. This approach reduces the impact of changing traffic patterns. Extensive tests on a multi-modal dataset that combines network telemetry and host events show that CCGNN greatly exceeds current leading models in detection accuracy and calibration. This model also shows strong resilience to changes in data distribution and adversarial attacks while significantly reducing alert fatigue for operators. This approach offers a reliable path to effective, clear, and ready-to-use intrusion detection systems.

Keywords: Intrusion Detection, Graph Neural Networks, Causal Learning, Contrastive Learning, Robustness.

Abstract of Paper Accepted in ICAIC-2026

339

Confidential and Attack-Resilient Edge LLM Serving for Multilingual Chatbots

SATYA KARTEEK GUDIPATI
Peritus Inc
Naveen Anand Mishra
Cigna
Akbar Mohammed, Sambasiva Rao Akkisetti
American Express
Gopikanth Ankam
IT America Inc
SATEESHKUMAR PONUGOTI
Publicis Groupe
sskmaestro@gmail.com, naveenmishra5@gmail.com,
akbarjntuus@gmail.com, asr.akkisetti@gmail.com,
gopikanth.ankam@gmail.com

ABSTRACT

Large Language Model (LLM) chatbots increasingly run at the network edge to reduce latency and cost, but edge deployments expand the attack surface and amplify privacy risk. We present Aegis-Edge, a confidential and attack-resilient serving stack for multilingual LLM inference on untrusted edge nodes. Aegis-Edge combines (i) remote attestation to verify the serving runtime before model and key release, (ii) confidential inference using trusted execution environments with a sealed key–value cache and tenant-scoped encryption, (iii) policy-as-code guardrails for locale-aware PII redaction and tool-use restrictions, and (iv) a lightweight adversarial input filter targeting multilingual prompt-injection and cache-poisoning attempts. We formalize security invariants over the request–response path and show how they compose across edge, gateway, and KMS. Our implementation supports common edge accelerators and multilingual tokenizers. We empirically evaluate security (attack success rate, leakage proxies, integrity violations), privacy (PII detection/retention bounds), and systems overhead (added latency, throughput, energy) under red-team workloads spanning diverse languages and scripts. Results demonstrate that robust privacy and integrity guarantees can be enforced at the edge while maintaining application-level service objectives. We discuss limitations (e.g., side-channels) and pathways for incremental deployment.

Keywords: security and privacy; confidential inference; remote attestation; edge computing; multilingual NLP; policy enforcement.

Abstract of Paper Accepted in ICAIC-2026

342

Enhancing Trust in AI: Addressing Vulnerabilities and Ensuring Privacy with ML Security Protocols Across the Lifecycle

Ramkinker Singh Carnegie Mellon University Pittsburgh, USA
Om Narayan New York University New York, USA
Praveen Baskar Google LLC Chicago, USA
ramkinks@alumni.cmu.edu, on371@nyu.edu,
praveenbaskar@google.com

ABSTRACT

The rise of machine learning systems in hostile and sensitive environments has only raised their security and privacy concerns. This work gives a deep analysis of the dynamic threat landscape: adversarial attacks, data poisoning attacks, evasion attacks, and model extraction attacks, as well as diverse capabilities and intentions of adversaries. This paper is an attempt to review attack surfaces that are core to machine learning pipelines. The paper reviews a bunch of security protocols like adversarial training, differential privacy, secure multiparty computation, input pre-processing strategies, ensemble methods, etc. It speaks about the interaction between robustness and privacy and focuses on domain-specific considerations in finance, healthcare, cybersecurity, smart cities, and new domains. The metrics of quantification and processes of determining vulnerability and robustness of models have been reviewed and importance of both quantitative and qualitative analyses has been emphasized. Comparative analyses emphasize diversified strengths and weaknesses of different defense strategies that have been created to counter specific types of attacks and their effect on model performance and fairness. Design principles such threat modeling and secure-by-design strategies and adaptability to dynamic and changing threats are defined to drive the design of robust machine learning systems.

Keywords: Adversarial Attacks, Adversarial Training, Data Poisoning, Differential Privacy (DP), Federated Learning, Machine Learning Security, Model Extraction, Robustness, Secure Multi-Party Computation (SMPC), Threat Landscape.

Abstract of Paper Accepted in ICAIC-2026

343

Interval-Based Estimation of Generalization Accuracy in Supervised Learning via Confusion Matrix Resampling and Bayesian Inference

Amir Liron
Texas State University, USA
amir_liron@txstate.edu

ABSTRACT

Accurate model evaluation in supervised classification requires not only reporting performance metrics but also accounting for the uncertainty inherent in sampling and training. Traditional point estimates often mask this variability, leading to overconfident or misleading conclusions. This paper proposes a comparative framework for performance estimation using four methods: single-run theoretical intervals, repeated resampling (Monte Carlo Cross Validation), bootstrap resampling of confusion matrices, and a Bayesian approach based on Dirichlet-multinomial modeling of the confusion matrix. We assess their ability to characterize effectiveness, efficiency, and bias across models of varying complexity. Experimental results on a multiclass classification task demonstrate that interval-based methods—especially Bayesian and simulation-based techniques—offer superior reliability and insight, even in low-data regimes. Notably, several methods were able to approximate a model's full generalization behavior using less than one-fifth of the available data, underscoring their potential for efficient and scalable evaluation in constrained environments.

Keywords: Model evaluation, uncertainty quantification, Bayesian inference, confusion matrix, Monte Carlo sampling, bootstrap resampling, classification performance, interval estimation, reproducibility, low-data learning, reliability

Abstract of Paper Accepted in ICAIC-2026

344

Lanyard Policy Tracker: A Secure, Privacy-Aware Student Compliance System for K–12 Environments

Curtis Faughnan, Xiantian Zhou, Aobo Jin, Hardik Gohel, Qixin Deng

ABSTRACT

Security and efficiency requirements in K–12 environments have attracted considerable attention. To address these needs, we propose Lanyard Policy Tracker, a system designed to satisfy these requirements. The Lanyard Policy Tracker was developed as a secure, desktop-based application designed to monitor and enforce student compliance with school lanyard policies while prioritizing data protection and system integrity. Implemented in Python using the Tkinter framework, the system integrates with Google Sheets through encrypted API communication, providing centralized storage and real-time synchronization across multiple authenticated devices. All data transactions occur through a secured Google Cloud service account, preventing unauthorized access and ensuring auditability. The application features layered validation and error-handling mechanisms to safeguard against duplicate or corrupted entries, while local caching maintains operational continuity in case of network interruptions. A modular design separates the user interface, data synchronization, and administrative control layers, reducing the attack surface and enhancing maintainability. Password-protected administrative tools restrict access to configuration and reset functions, ensuring that sensitive operations are performed only by authorized users. Deployed successfully within a middle school network environment, the system demonstrated high reliability, rapid synchronization, and resilience to connectivity and concurrency challenges.

Keywords: Lanyard Policy Tracker, Secure, Privacy-Aware, Student Compliance System

Abstract of Paper Accepted in ICAIC-2026

345

AI-Driven Prediction of Flexural Properties in Kevlar/Carbon/Glass Fiber Hybrid Composites Using Random Forest Regression

Sathish Pandurangan, Syed Mudassir, John C. (Chris) Becking, Prasanna Ranjith Christodoss, Mithun V Kulkarni, Ahmed Shammas

ABSTRACT

The research introduces a methodology to predict the flexural behavior of a laminate before conducting any experimental tests. To derive such a model, an artificial intelligence (AI)-based learning process was adopted. More specifically, a Random Forest regression model was trained to identify the nonlinear behavior associated between the laminate stacking sequence, the fiber properties, and its flexural response, based on the flexural test dataset. To practically implement this AI–predictive approach, an interactive, real-time interface was set up via a Python package, Streamlit, to serve as a tool that can be used to visualize flexural strength and deflection under different stacking sequences. In this way, both laminate material selection and experimental design can be more efficiently informed. In addition to visualizing flexural performance metrics, this approach also allows the identification of the most important features based on the model's feature-importance functionality. As part of this study, mechanical testing was also conducted on various hybrid laminates ([K/G/C/C/G/K], [C/G/K/K/G/C], and [G/C/K/K/C/G]) and non-hybrid laminates ($[K]_6$, $[G]_6$, and $[C]_6$). The flexural strength and flexural modulus of all the laminate configurations were determined to confirm the predicted model and to study the influence of stacking sequence on the mechanical performance. The results showed that the flexural strength of the [C/G/K/K/G/C] laminate (3.07 MPa) was the highest among all the laminates, and the contribution of hybridization is confirmed. The AI–experimental framework described in this study paves the way for future efforts to accelerate composite optimization and minimize material waste and provides a reliable AI-assisted design tool for high-performance fiber-reinforced laminates.

Keywords: Carbon/Glass/Kevlar hybrid composites; Flexural; FRP laminates; Stacking, Machine Learning, Streamlit, Random Forest.

Abstract of Paper Accepted in ICAIC-2026

346

Optimizing Customer Engagement through IoT Data Integration in CRM Ecosystems

Sathish Kumar Velayudam, Independent Researcher, USA
sathish.velayudam@ieee.org

ABSTRACT

This research aims to illustrate what actually takes place with the data that is collected from IoT devices as it flows into Customer Relationship Management (CRM) systems, showing how this integration transforms the way businesses reach and engage customers. Traditional CRM systems lean too heavily on historical transaction data and static demographic attributes, limiting their ability to interpret real-time behavioral context. Integrating live streams – such as usage traces, performance variations, and device-level error logs – closes this latency gap by enabling immediate, context-aware business actions. Using a synthetically generated dataset of 455 smart-home user transactions, this study demonstrates how dense, high-velocity IoT data can be used to reveal behavioral patterns that remain invisible in conventional CRM environments. Although the technical pipeline is straightforward - data preprocessing using Python; the structured SQL queries, and Tableau visualizations mapping device activity to customer engagement – the data itself is rich, containing signals of genuine user behavior. Real-time data enables both informed decision-making as well as hyper-personalization. Empirical results show higher retention, lower churn and faster support resolution when CRM systems respond the moment an IoT signal is generated – whether due to malfunction, usage anomalies, or performance degradation. The objective extends beyond automation for its own sake; the focus is on optimal timing—engaging customers precisely when intervention is most valuable. The key finding demonstrates that when IoT and CRM operate synergistically, they create a feedback loop capable of generating value at increasing speed. Data velocity translates into enhanced trust, relevance, and ultimately customer loyalty. This research illustrates how customer engagement scales effectively in continuously connected environments.

Keywords: IoT integration, CRM ecosystems, customer engagement, predictive analytics, data fusion.

Abstract of Paper Accepted in ICAIC-2026

347

Cybersecurity in Blockchain-based fintech platforms: Social Perceptions and Adoption Intention

ABSTRACT

This study examines how cybersecurity related perceptions shape Gen Z intention to use blockchain based fintech platforms. Drawing on UTAUT and extended valance framework, a structural model is tested in which digital literacy, social influence, perceived security, risk tolerance, perceived benefit and trust in technology explain behavioural intention. Survey data from 200 Malaysian university students are analysed using partial least squares structural equation modelling. The results show that digital literacy, social influence, perceived security, risk tolerance and trust in technology significantly and positively predict intention, with trust emerging as the strongest driver. Perceived benefit does not directly affect intention, but it strongly strengthens trust, and trust mediates the perceived benefit to intention relationship. The findings highlight the need to build digital skills, communicate safeguards clearly and reinforce ecosystem level trust.

Keywords: digital literacy, perceived security, risk tolerance, trust in technology, social influence, perceived benefit, blockchain-based fintech platforms

Abstract of Paper Accepted in ICAIC-2026

348

Security Challenges and Attack Vectors in Modern Steganography

Sheikh Thanbir Alam, R. Badlishah Ahmad, Md Maruf Hassan, Ku Nurul Fazira Ku Azir, Amit Banwari Gupta

ABSTRACT

Steganography, the ancient art of concealing information within innocuous carriers, has evolved into a sophisticated discipline at the intersection of cybersecurity, digital media, and data privacy. This comprehensive review explores the progression of steganographic techniques from historical ciphers to modern digital embedding methods, emphasizing its critical role as an invisible shield for data protection in the digital age. The paper categorizes and examines the primary types of steganography such as text, image, audio, video, and network-based highlighting their methodologies, strengths, and applications. A detailed literature review presents key advancements, current trends, and influential research contributions in the field. Furthermore, the paper evaluates quality metrics used to assess steganographic performance, including imperceptibility, capacity, and robustness. Security challenges and attack vectors such as steganalysis and active adversarial threats are discussed to underscore the ongoing arms race between steganographic methods and detection techniques. Recognizing the importance of empirical research, the review also catalogs widely used datasets tailored to various steganographic domains. The paper concludes by identifying current limitations and proposing future research directions aimed at enhancing security, adaptability, and real-world deployment of steganographic systems. This review serves as a foundational reference for researchers, practitioners,
and students seeking to understand and advance the field of steganography.

Keywords: Steganography, Steganalysis, Stego Attacks, Attack Metircs

Abstract of Paper Accepted in ICAIC-2026

349

Enclave-Driven Tokenization: Reducing PCI DSS Scope in Cloud-Native Checkout Systems

Rajgopal Devabhaktuni, Gopalakrishnan Venkatasubbu
Independent Researcher
Atlanta, USA

devabhaktuni.rajgopal@gmail.com, gopalakrishnan.venkatasubbu1@gmail.com

ABSTRACT

The shift of ecommerce platforms to cloud-native architectures (that are based on microservices and containers) has greatly made PCI DSS compliance more complex. The compliance scope of these systems is often extremely wide as their components are generally distributed and the scope includes almost all parts of such system, they have high audit costs and present more security risks. In this paper, we introduce a new design that uses confidential computing with secure enclaves to realize data tokenization. In this model, sensitive cardholder data is handled only within a hardware-enforced secure environment removing the exposure of raw PAN values out to the cloud native application. This strategy allows you to lock PCI DSS scope down into that enclave, and in essence removes compliance responsibilities from the application scale. A synthetic dataset of 410 transaction instances was used to replicate a live checkout setting for an online shopping platform. The system was implemented as a Python microservices application orchestrated using Kubernetes and executed utilizing AWS Nitro Enclaves as the confidential computing technology. It shows how the scope of coverage is cut down from the whole application stack to a minimal service. This approach is a practical, high-assurance path toward simplifying PCI DSS compliance in complex, modern cloud environments.

Keywords: Enclave-Based Tokenization, PCI DSS Scope Reduction, Cloud-Native Security, Confidential Computing, Secure Checkout Systems.

Abstract of Paper Accepted in ICAIC-2026

350

Conditional Password Generation Using a FiLM-Enhanced WGAN: A Controlled Comparison Against Standard GAN Baselines

Pranav Shetty, Pavan P, Srinivas P M, Pruthvi C V, Vishal Pujar
Sahyadri College of Engineering and Management, India
studytimemail24@gmail.com

ABSTRACT

Feature-wise Linear Modulation (FiLM) offers an efficient mechanism for conditioning deep networks; however, its impact on password-distribution modeling has not been systematically evaluated. This paper investigates a FiLM-enhanced conditional Wasserstein GAN for synthetic password generation and analyzes how FiLM influences conditional fidelity, structural consistency, and diversity. To contextualize its behavior, we conduct controlled comparisons against three widely used baselines: PassGAN, PaC-GAN, and a concatenation-based WGAN–CGAN, under a unified preprocessing pipeline and shared training configuration. The results show that FiLM improves conditional accuracy and better preserves local structural patterns compared to standard concatenation. At the same time, baseline models highlight known trade-offs such as PaC-GAN's sensitivity to dominant modes and the stabilizing effect of Wasserstein training. All experiments follow strict ethical and privacy safeguards, and only aggregate statistics are reported.

Keywords: Password modeling, generative adversarial net- works, FiLM, PassGAN, PaC-GAN, WGAN-GP, cybersecurity.

Abstract of Paper Accepted in ICAIC-2026

351

The Evolution of Agentic AI in Cybersecurity: From Single LLM Reasoners to Multi-Agent Systems and Autonomous Pipelines

Vaishali Vinay,
Microsoft, USA
vaishali.papneja@microsoft.com

ABSTRACT

Cybersecurity has become one of the earliest adopters of agentic AI, as security operations centers increasingly rely on multi-step reasoning, tool-driven analysis, and rapid decision-making under pressure. While individual large language models can summarize alerts or interpret unstructured reports, they fall short in real SOC environments that require grounded data access, reproducibility, and accountable workflows. In response, the field has seen a rapid architectural evolution from single-model helpers toward toolaugmented agents, distributed multi-agent systems, schemabound tool ecosystems, and early explorations of semiautonomous investigative pipelines. This survey presents a five-generation taxonomy of agentic AI in cybersecurity. It traces how capabilities and risks change as systems advance from text-only LLM reasoners to multi-agent collaboration frameworks and constrained-autonomy pipelines. We compare these generations across core dimensions - reasoning depth, tool use, memory, reproducibility, and safety. In addition, we also synthesize emerging benchmarks used to evaluate cyber-oriented agents. Finally, we outline the unresolved challenges that accompany this evolution, such as response validation, tool-use correctness, multi-agent coordination, long-horizon reasoning, and safeguards for highimpact actions. Collectively, this work provides a structured perspective on how agentic AI is taking shape within cybersecurity and what is required to ensure its safe and reliable deployment.

Keywords: agentic AI, cybersecurity automation, large language models, multi-agent systems, benchmarking of AI agents, AI safety and verification, security operations center (SOC)

Abstract of Paper Accepted in ICAIC-2026

354

Enhancing Intrusion Detection with Image-Based CNN and CTGAN Synthetic Oversampling

Leo Martinez III, Avdesh Mishra, Md Habibur Rahman, Mais Nijim, Ayush Goyal, David Hicks

Texas A and M University-Kingsville, USA
avdesh.mishra@tamuk.edu

ABSTRACT

Network intrusion detection has become a critical component in maintaining secure systems, but effective classification of malicious traffic remains challenging due to heavily imbalanced datasets. Datasets widely used for intrusion detection research suffer from severe class imbalance, which impacts the performance of machine learning models in past research. In this article, a novel approach is proposed to tackle these challenges by utilizing a Convolutional Neural Network (CNN) for 5-class classification of network intrusion types. Unlike prior methods that predominantly rely on tabular data, this approach converts the NSL-KDD dataset into image representations to leverage the spatial learning capabilities of CNNs. Furthermore, he heavy class imbalance issue is handled by generating synthetic data using Conditional Tabular Generative Adversarial Networks (CTGAN), which allows for more equitable learning across all classes. By combining CNN-based image classification with CTGAN-generated synthetic data, the model aims to improve classification performance and provide a robust framework for network intrusion detection. The proposed model is benchmarked against contemporary state-of-the-art methods on real-world data and achieves average gains of 6.61%, 5.36%, 7.74%, and 12.42% in accuracy, precision, recall, and F1-score, respectively.

Keywords: Intrusion detection, Convolutional neural networks (CNN), Generative adversarial networks (GAN), Deep learning, Synthetic Data Augmentation, Cybersecurity

Abstract of Paper Accepted in ICAIC-2026

355

Non-Intrusive Machine Learning-Based Anomaly Detection for Heterogeneous Embedded Platforms

Pranav Gangwani, Alexander Perez-Pons, Himanshu Upadhyay
Florida International University, USA

ABSTRACT

As the number of embedded devices in commercial markets increase, it becomes increasingly important to secure these devices from malicious actors. In this paper, we demonstrate various ways of non-intrusively extracting data from embedded devices through debugging interfaces and using this data with machine learning and deep learning models to detect anomalies in embedded devices. The results exhibit that deep learning models have great potential in embedded cybersecurity, with Convolutional Neural Network and Long Short-Term Memory network models scoring over 95% accuracy on the testing set of our data. This study not only contributes to the ongoing cybersecurity discussion for embedded devices but also creates a foundation for further advancements in the field.

Keywords: Anomaly detection, CNN, Deep learning, Embedded systems, Internet of things, LSTM, Machine learning

Abstract of Paper Accepted in ICAIC-2026

358

Leveraging Artificial Intelligence to Predict Unfunded Loans in Peer-to-Peer Lending Platforms

Mousumi Munmun, Queen Booker
Metro State University, USA
mousumi.munmun@metrostate.edu,
queen.booker@metrostate.edu

ABSTRACT

This study explores the use of Artificial Intelligence (AI) methodologies to predict and characterize loans that are unlikely to receive investor funding within peer-to-peer (P2P) lending markets. Specifically, it evaluates a Heterogeneous Ensemble (HEE) model that integrates three supervised machine learning algorithms: Logistic Regression (LR), Random Forest (RF), and Support Vector Machine (SVM). While AI and machine learning have been widely applied to borrower credit assessments, limited research has focused on predicting the second stage of P2P loan decision-making: whether approved loans will receive funding from investors. Addressing this gap, the HEE model leverages the complementary strengths of its component algorithms: LR for interpretability, RF for handling missing data and overfitting, and SVM for high-dimensional classification. Results show that the HEE model significantly outperforms individual classifiers in predicting unfunded loans, with improvements validated through statistical hypothesis testing. This research demonstrates the value of ensemble-based AI approaches in advancing intelligent financial decision-support systems.

Keywords: Artificial Intelligence, Machine Learning, Heterogeneous Ensemble Model, Peer-to-Peer Lending

Abstract of Paper Accepted in ICAIC-2026

359

AI-Optimized VLSI Architecture for Energy Efficient and Sustainable IoT Systems

Shujaatali Badami, Anshul Sharma, Bhaskar Reddy, Shailesh Kadam, Biky Chouhan, Brijeena Rana

ABSTRACT

The artificial intelligence (AI) and Internet of things (IoT) boom have caused the necessity to implement energy-efficient hardware platforms capable of delivering smart computations in real-time at the network border. Traditional VLSI design tooling can generally not be used to implement sustainable IoT designs, as they cannot realize high performance and low power consumption at the same time. It proposes an AI-Optimized VLSI Architecture, a design approach that uses machine learning-based design-space exploration and provides machine learning-based adaptive power management and energy harvesting schemes to achieve significant performance-per-watt advantages. The framework presented applies reinforcement learning, genetic algorithms and Bayesian optimization to optimize the parameters of synthesis and layout intelligently, maximizing power delay trade-offs. According to the simulations of Cadence and Synopsys design tools, the company has saved power by 43.3 percent, delay by 29.7 percent and energy by 52 percent compared to conventional VLSI systems. Further, the architecture incorporates dynamic voltage and frequency scaling (DVFS), clock gating, and AI-based power gating to achieve leakage and computation energy minimization that prolong device life in energy-constrained internet of things. This paper empirically demonstrates that allowing AI-directed optimization, the sustainability, scalability, and flexibility of next-generation clean-energy-based electronic systems can dramatically increase their viability.

Keywords: AI-optimized VLSI, low-power design, energy-efficient hardware, IoT systems, clean energy electronics, reinforcement learning, design-space exploration, dynamic power management, edge computing, sustainable semiconductor design

Abstract of Paper Accepted in ICAIC-2026

360

Enhancing Safety and Performance in Intelligent Vehicular Networks Using Edge-Based Explainable AI Models

Anshul Sharma, Sravani Lingam, Gayathri Balakumar, Biky Chouhan, Milankumar Rana

ABSTRACT

The automated cars must possess real time, efficient and understandable decisions that will ensure that there is road safety and efficiency. In the case, the researchers propose to use an edge-enabled model that is utilised to include explainable artificial intelligence (XAI) architectures, which are CNN, RNN, and GNN to provide a better safety prediction rate, low latency, and throughput when driving in dense traffic. As the analysis in the performance indicates, the deployment of the edges is significantly quicker in both the latency and the throughput like the cloud based computation provided that the real time is performance is stable, always, with reference to the condition of peak traffic. The implementation of XAI models will cause minor trade-offs in raw accuracy and throughput but result in a notable distinction in the interpretability of the model, with stakeholders in a position to trust and audit model predictions. XAI may also yield better quality of collision prediction since as a safety measure the XAI is only effective within 2-5% change of the accuracy of the collision prediction. The findings are indicative that edge enabled XAI framework is an effective solution to intelligent transportation systems of the next generation because it is effective, safe, and transparent. Future research can be done in the field of multi-modal data integration and adaptive hybrid edge-cloud architecture to make it even more performance-driven and interpretable.

Keywords: Intelligent Vehicular Networks, Edge Computing, Explainable AI (XAI), CNN, RNN, GNN, Traffic Safety Prediction, Latency, Throughput.

Abstract of Paper Accepted in ICAIC-2026

361

CityCopilot X: A Real Time Explainability Panel for Retrieval Augmented Generation (RAG)

SATYA KARTEEK GUDIPATI
Peritus Inc
Naveen Anand Mishra
Cigna
Akbar Mohammed, Sambasiva Rao Akkisetti
American Express
Gopikanth Ankam
IT America Inc
SATEESHKUMAR PONUGOTI
Publicis Groupe
sskmaestro@gmail.com, naveenmishra5@gmail.com,
akbarjntuus@gmail.com, asr.akkisetti@gmail.com,
gopikanth.ankam@gmail.com

ABSTRACT

Civic information is often dispersed across portals and PDFs, while users increasingly expect grounded answers with verifiable sources. Retrieval-Augmented Generation (RAG) can provide such grounding, but most interfaces operate as black boxes, revealing little about which passages influenced an answer or how sensitive that answer is to specific evidence. This lack of transparency limits trust, hinders debugging, and complicates audit requirements in real-world deployments. CityCopilot-X is a real-time glass-box framework for RAG that exposes evidence use across the entire pipeline. The system provides (1) visual attributions for query rewriting, retrieval, and reranking; (2) token-level coverage showing how generated text aligns with cited passages; and (3) an interactive counterfactual mode that recomputes answers with selected evidence removed, quantifying changes in confidence and coverage. CityCopilot-X supports multilingual civic queries and updates responses within seconds of source edits while maintaining stable latency and cost. In a pilot study with six analysts across 24 civic tasks, the panel reduced post-hoc correction time by 31% and increased citation coverage by 18 points compared with a black-box baseline. Synthetic case studies further highlight the system's ability to reveal ranking failures, over-dominant documents, and evidence drift.

Keywords: Retrieval-Augmented Generation (RAG), Interactive explainability, Counterfactual retrieval, Reranking attributions, Citation coverage, Civic question answering.

Abstract of Paper Accepted in ICAIC-2026

363

Reverse Entropy Spiral Deep Neural Steganography for Secure Medical Ultrasonogram Videos

Sheikh Thanbir Alam, R. Badlishah Ahmad, Md Maruf Hassan, Ku Nurul Fazira Ku Azir, Abubokor Hanip

ABSTRACT

Video steganography allows hiding secret information inside video frames so that it is hard to detect. Traditional methods often modify large blocks of pixels, which can create visible artifacts and make the hidden data vulnerable to compression, noise, or geometric changes. This work presents a reverse-spiral pixel–based deep learning framework for video steganography. Secret data is first encrypted using RSA, and video frames are shuffled to increase security. A deep neural network identifies visually safe areas in each frame where small pixel changes are unlikely to be noticed. Encrypted bits are then embedded following a reverse spiral pattern across the RGB channels, using only about 8.7% of pixels per frame. Metadata carrying keys and encryption information is stored within the video, eliminating the need for external key exchange. Experiments on a 5-second, 25 FPS ultrasound video with payloads of 10–30 KB achieved high visual quality (PSNR > 43 dB, SSIM > 0.96, MAE <0.007) and low bit errors (BER < 0.015) even under compression, noise, and geometric attacks. The results show that reverse-spiral, entropy-guided embedding with deep learning provides a secure, robust, and efficient method for hiding data in videos, suitable for medical, surveillance, and defense applications.

Keywords: multi-spiral pixel selection, deep learning, adaptive embedding, visual complexity, secure frame shuffling

Abstract of Paper Accepted in ICAIC-2026

364

Scalable Graph-Based Detection of Fraud Rings in Large-Scale Networks

Xiantian Zhou, Qixin Deng, Aobo Jin, Hardik Gohel

ABSTRACT

Graphs are widely used to model relational data in various domains, including social media, e-commerce, telecommunications, and finance. Graph analytics is one of the most popular techniques for analyzing connectivity patterns in communication networks and identifying suspicious behaviors. However, detecting fraud rings at scale is significantly challenging since the volume of graph data is growing exponentially. Moreover, many Python-based graph libraries rely on in-memory computation, which often struggles with large-scale networks. To address these limitations, we extend our previous semiring-based graph computation framework into a domain-specific solution for large-scale fraud ring detection. We formalize key graph patterns that indicate the presence of fraud rings and develop a general detection algorithm based on semiring operations. Our algebraic approach operates efficiently on a hybrid architecture that can scale beyond RAM constraints. Furthermore, it provides interpretable results, ensuring mathematical transparency to meet regulatory demands for explainability. Our approach is developed in C++, and it can be easily called in Python. An Experimental comparison with state-of-the-art Python packages shows that our approach has comparative performance for both small and large graphs.

Keywords:

Abstract of Paper Accepted in ICAIC-2026

367

Explainable Machine Learning for Malware Detection: A SHAP-Based LightGBM Framework

Abdullah Al Siam, Hazem Abu-Adaiq, Mahdee Nafis, Ferdous Anwar Anik, Sabbir Mahmud, Md Maruf Hassan

ABSTRACT

In contemporary malware detection, machine learning has proved an essential component; many models are not transparent, making them less trustworthy and rarely suitable for security use. In this paper, we propose a machine-learningbased malware classification framework trained on static features extracted from the EMBER database of Windows PE files using the LightGBM algorithm. To improve interpretability, we add SHAP (Shapley Additive exPlanations), which allows the analyst to see which features have the most significant influence on a given prediction. The results of the experiments show that our framework achieves 91.16% accuracy, 94.22% recall, and an ROC-AUC of 0.9744, compared with baseline traditional machine learning methods. Qualitative analyses demonstrate that SHAP explanations are effective at highlighting the essential binary features relevant to malware detection, thereby enhancing model comprehensibility and enabling quicker, more informed incident response. We also mention some of the limitations, e.g., the resource required for explainability on large datasets, and what we envisage working on in the future, e.g., utilizing the frameworks with threat intelligence and running on demand.

Keywords: Malware Detection, Explainable AI, Threat Intelligence, Deep Learning, Cybersecurity

Abstract of Paper Accepted in ICAIC-2026

369

ChatXplain: Interpretable Explanations for Intelligent Assistants via Modular Rationale and Saliency

SATYA KARTEEK GUDIPATI
Peritus Inc
Naveen Anand Mishra
Cigna
Akbar Mohammed, Sambasiva Rao Akkisetti
American Express
Gopikanth Ankam
IT America Inc
SATEESHKUMAR PONUGOTI
Publicis Groupe
sskmaestro@gmail.com, naveenmishra5@gmail.com,
akbarjntuus@gmail.com, asr.akkisetti@gmail.com,
gopikanth.ankam@gmail.com

ABSTRACT

Large Language Model (LLM)–based assistants increasingly support high-stakes decision workflows, yet their opaque reasoning processes limit user trust, hinder debugging, and complicate compliance. Existing explanation techniques—such as SHAP, LIME, and attention-based attribution—struggle to provide coherent, multi-turn dialogue-level interpretability while meeting production latency constraints. This paper introduces ChatXplain, a modular framework that generates real-time, interpretable explanations for LLM-driven conversational systems without modifying underlying model weights. The framework integrates four lightweight components—an intent classifier, a rationale generator, a saliency visualizer, and a dialogue tracker with auditable reasoning traces—designed to operate as an auxiliary layer atop any LLM. We provide full architectural details, input–output specifications, and reproducibility guidelines, along with quantitative comparisons against SHAP and LIME across two domains: customer service and financial advisory. Experiments on 1,800 multi-turn dialogues generated from 218 simulated user profiles, designed to mimic realistic interaction patterns, show that ChatXplain improves user trust proxies by 22%, explanation clarity by 19%, and developer debugging efficiency by 31%, while adding only 8.7% median latency overhead.

Keywords: Explainable AI, Intelligent Assistants, Dialogue Systems, Rationale Generation, Saliency Attribution, Trustworthy AI.

Abstract of Paper Accepted in ICAIC-2026

370

TransCall: A Transformer-Driven Framework for Zero-Day Malware Detection Using System Call Sequences

Abdullah Siam, Moutaz Alazab, Areej Obeidat, Nuruzzaman Faruqui, Somaya Al-Maadeed

ABSTRACT

The growing complexity of contemporary malware, along with the widespread use of polymorphism, obfuscation, and environment-aware evasion, has significantly diminished the effectiveness of conventional signature-based and static detection methods. Behavioral analysis, especially via system call sequences, provides a robust depiction of program intent; however, current machine learning and recurrent neural network (RNN) methods struggle to capture long-range dependencies and often do not generalize well to new, unseen threats. This paper presents TransCall, a streamlined Transformer-based framework for detecting zero-day malware by analyzing system call sequences. It utilizes multi-head self-attention to capture global contextual relationships within syscall streams, facilitating the practical identification of malicious behavioral patterns, even in the presence of unknown malware families. The suggested framework integrates a highly effective embedding module, a streamlined Transformer encoder, and a refined classification head designed for real-time inference in endpoint security systems. Experimental evaluations on UNM-style syscall datasets show that TransCall outperforms traditional baselines, including Random Forest, LSTM, and GRU models. In a standard train-test division, It reaches an F1-score of 0.8571 and an AUC of 0.8743. In a challenging zero-day environment where entire malware families are omitted from the training process, TransCall attains an impressive F1-score of 0.9286, highlighting its strong generalization capabilities. Moreover, visualizing attention weights improves comprehension by emphasizing crucial syscall interactions linked to potentially harmful activities. In therefore, TransCall delivers accurate, clear, and efficient detection of zero-day malware, presenting a practical and scalable solution for modern cybersecurity environments.

Keywords: Threat Intelligence, Large Language Models (LLMs), Cybersecurity, Indicator of Compromise (IOC), MultiSource Data Fusion, SOC Automation, Real-Time Security Analysis, Threat Confidence Scoring.

Abstract of Paper Accepted in ICAIC-2026

371

Real-Time Multi-Source Threat Intelligence Fusion Using Large Language Models

Abdullah Siam, Moutaz Alazab, Areej Obeidat, Nuruzzaman Faruqui, Somaya Al-Maadeed

ABSTRACT

Threat Intelligence (TI) is essential in contemporary Security Operations Centers (SOCs), but current TI platforms operate independently and often produce disjointed, inconsistent, or conflicting evaluations of Indicators of Compromise (IOCs). The chaotic organization of the information forces analysts to meticulously link insights from multiple sources, resulting in significant delays and an increased risk of misclassification. This paper presents a real-time framework powered by a Large Language Model (LLM) to integrate multi-source threat intelligence. It skillfully integrates various outputs from VirusTotal, AbuseIPDB, AlienVault OTX, GreyNoise, and MalwareBazaar into a unified, context-aware depiction. The proposed system transforms raw TI responses into a standardized format, employs semantic reasoning via an LLM to address conflicting evidence, and produces a clear Threat Confidence Index (TCI) that measures IOC severity on a scale from 0 to 100. The experimental evaluation conducted on a dataset comprising 250 IOCs reveals that the framework attains a remarkable 90.8% alignment with the majority TI consensus. It also exhibits consistent scores, with a variance of less than 1.5 across multiple iterations, while ensuring an end-to-end latency of 2.75 seconds, thereby fulfilling real-time operational demands. Analyst studies indicate a 66% decrease in investigation time when using the fused assessment compared to conventional multi-source lookup methods. The results suggest that incorporating LLM support significantly improves the consistency, interpretability, and efficiency of TI workflows. The findings highlight the potential of LLM reasoning to enhance upcoming SOC automation and prompt threat evaluation significantly.

Keywords: Threat Intelligence, Large Language Models (LLMs), Cybersecurity, Indicator of Compromise (IOC), MultiSource Data Fusion, SOC Automation, Real-Time Security Analysis, Threat Confidence Scoring.

Abstract of Paper Accepted in ICAIC-2026

372

A Systematic Security Analysis of Model Context Protocol: Vulnerabilities, Exploits, and Mitigations

Theophilus Siameh, Chun-Hung Liu, Eric Kudjoe Fiah,
Mississippi State University, USA
Abigail Akosua Addobea
Quanzhou University of Information Engineering, China

ABSTRACT

Large Language Models (LLMs) gain significantly enhanced capabilities through the Model Context Protocol (MCP), which allows them to connect directly with external tools and services. However, this expanded functionality introduces previously unexamined security vulnerabilities. Our study provides the first thorough security assessment of MCP implementations, uncovering serious weaknesses in filesystem operations, database connections, and API integrations. We conducted systematic penetration testing on 15 different MCP server implementations and successfully executed attacks including directory traversal, SQL injection, credential theft, and resource exhaustion. Our findings show that 87% of examined MCP servers have at least one critical security flaw, while 34% are vulnerable to full system takeover. We develop a detailed classification system for MCP security threats and introduce effective defensive measures that decrease successful attack rates by as much as 94%. Our research lays the groundwork for secure MCP implementation and underscores the essential importance of prioritizing security when designing LLM integration protocols.

Keywords: Model Context Protocol (MCP), Context Injection Attacks, Resource Exploitation, Tool Manipulation, Configuration Vulnerabilities

Abstract of Paper Accepted in ICAIC-2026

373

Semantic Mastery: Enhancing LLMs with Advanced Natural Language Understanding

Mohanakrishnan Hariharan
Dept of Corporate Systems Engineering,
Apple Inc. Austin, TX, USA
m_hariharan@apple.com

ABSTRACT

Software engineering processes increasingly depend on understanding vast amounts of technical documentation, code repositories, bug reports, and requirements specifications. While Large Language Models (LLMs) show promise in automating software engineering tasks, they face critical challenges in understanding domain-specific semantics, maintaining contextual consistency across codebases, and accurately reasoning about software artifacts. This paper presents Semantic Mastery, a novel framework that enhances LLMs with software engineering-focused semantic understanding capabilities through specialized knowledge graphs, codeaware training, and multi-modal analysis of software artifacts. Our approach integrates software ontologies, API documentation, and codebase relationships to create contextually-aware models capable of accurate defect prediction, automated code review, and intelligent software analysis. We introduce a hierarchical attention mechanism for multi-file code understanding and demonstrate significant improvements in software engineering tasks, including defect prediction (31% improvement in precision), automated code review (28% reduction in false positives), and requirements traceability (24% better accuracy). The framework addresses critical challenges in software development, including cross-language compatibility, analysis of legacy systems, and maintaining semantic consistency across large-scale software projects.

Keywords: Software Engineering, Large Language Models, Code Analysis, Defect Prediction, Semantic Understanding, Software Intelligence

Abstract of Paper Accepted in ICAIC-2026

374

Reinforcement Learning Integrated Agentic RAG for Software Test Cases Authoring

Mohanakrishnan Hariharan
Dept of Corporate Systems Engineering,
Apple Inc. Austin, TX, USA
m_hariharan@apple.com

ABSTRACT

This paper introduces a framework that integrates reinforcement learning (RL) with autonomous agents to enable continuous improvement in the automated process of software test cases authoring from business requirement documents within Quality Engineering (QE) workflows. Conventional systems employing Large Language Models (LLMs) generate test cases from static knowledge bases, which fundamentally limits their capacity to enhance performance over time. Our proposed Reinforcement Infused Agentic RAG (Retrieve, Augment, Generate) framework overcomes this limitation by employing AI agents that learn from QE feedback, assessments, and defect discovery outcomes to automatically improve their test case generation strategies. The system combines specialized agents with a hybrid vector-graph knowledge base that stores and retrieves software testing knowledge. Through advanced RL algorithms, specifically Proximal Policy Optimization (PPO) and Deep Q-Networks (DQN), these agents optimize their behavior based on QE-reported test effectiveness, defect detection rates, and workflow metrics. As QEs execute AI-generated test cases and provide feedback, the system learns from this expert guidance to improve future iterations. Experimental validation on enterprise Apple projects yielded substantive improvements: a 2.4% increase in test generation accuracy (from 94.8% to 97.2%), and a 10.8% improvement in defect detection rates. The framework establishes a continuous knowledge refinement loop driven by QE expertise, resulting in progressively superior test case quality that enhances, rather than replaces, human testing capabilities

Keywords: Reinforcement Learning, Agentic Systems, Software Testing, Retrieval-Augmented Generation, Multi-Agent Systems, Continuous Learning, RAG

Abstract of Paper Accepted in ICAIC-2026

375

ENHANCING BREAST CANCER DIAGNOSIS USING LIGHTWEIGHT DEEP LEARNING MODELS

Jordan Mozebo, Xiaohong Yuan, Madhuri Siddula, Olusola Odeyomi
North Carolina A&T State University, USA jtmozebo@aggies.ncat.edu

ABSTRACT

Breast cancer remains one of the leading causes of mortality among women worldwide, underscoring the need for early and accurate detection methods. This study investigates the use of deep learning for breast cancer detection, with a focus on lightweight models to improve accuracy. Although there are various imaging techniques exist, this research specifically uses histopathological images from the BreakHis and IDC dataset; and a MRI dataset. Advanced deep learning models, including MobileNetV3-Small, BCDNet, ShuffleNetV2, and EfficientNet-B0 are evaluated. Five-fold cross-validation was applied during training to ensure reliable and unbiased performance estimation, reduce overfitting risk, and provide stable performance. Experimental results show that MobileNetV3-Small and EfficientNet-B0 consistently achieved the highest testing and validatiion accuracy across all three datasets (BreakHis, IDC, and MRI), when evaluated with 5-fold cross-validation. On the BreakHis dataset, MobileNetV3-Small achieved a testing and validation accuracy of 99.62% and 98.62%, respectively. On IDC dataset, MobileNetV3-Small achieved 94.26% testing accuracy and 93.06% validation accuracy. On MRI dataset, MobileNetV3-Small achieved 99.53% testing accuracy and 99.23% validation accuracy. Our findings highlight the promise of lightweight deep learning models in histopathological and MRI image analysis, offering reliable solutions for early breast cancer detection while reducing computational costs. This study emphasizes the value of streamlined architectures in improving the accuracy of medical diagnosis, thus contributing to better patient outcomes.

Keywords: Breast Cancer Diagnosis, Lightweight Deep Learning Models, Histopathological Images, MRI, Computer-Aided Diagnosis

Abstract of Paper Accepted in ICAIC-2026

376

A Comprehensive Assessment Tool for Prompt Injection Attacks in Large Language Models (LLMs)

Ailene Dao, Richard VanGorder, Mohamed Gebril, Gael Abboud
George Mason University, USA, **mgebril@GMU.EDU**
Sherif Abdelhamid, VMI, abdelhamidse@vmi.edu

ABSTRACT

A prompt injection attack is a type of cyberattack that targets Large Language Models (LLMs), where attackers embed malicious instructions into a LLM with the intent of manipulating the model to the attacker's desire. Due to the relevance, growth, and demand of LLMs within small to mediumsized businesses, it is essential that they have access to a tool that can comprehensively assess model vulnerabilities. Based on an established framework for Prompt Injection attacks, this paper introduces a tool that can simulate types of injection attacks across different widely used LLM models and versions, test different defenses against these attacks, and measure the performance and how vulnerable the LLM's responses are. Additionally, this paper incorporates cost analysis, CPU utilization metrics, and a prompt hardening methodology for LLM personas. The results of these experiments provide deeper insights into testing and a comparative analysis against commonly used LLMs today.

Keywords: Artificial intelligence, Cybersecurity, Large language models, Prompt injection attacks

Abstract of Paper Accepted in ICAIC-2026

377

Multi-Layered Defense Proxy for Home Assistant utilizing Large Language Models

Faran Yazdani, Stephen Antezana, Mohamed Gebril
George Mason University, USA, mgebril@GMU.EDU
Sherif Abdelhamid, VMI, abdelhamidse@vmi.edu

ABSTRACT

The ever-increasing integration of LLMs in IoT environments necessitates the implementation of robust security measures. Because such systems have access to physical devices, attacks against them can have real-world consequences. Therefore, to address the growing need for security in this realm, we propose a proxy that leverages LLMs to detect prompt injection attempts in a home IoT setting. In this paper, we will discuss various methods of performing prompt injections on LLMs that have been given access to control physical IoT devices in a home. We will discuss the relationship between our work and related efforts aimed at securing LLMs, and provide an overview of the contributions this work makes to the field. We will discuss the benefits of our approach and briefly outline potential extensions to our work, both in the realm of IoT devices and in other contexts where LLMs are used.

Keywords: Home assistant, large language model, prompt injection

Abstract of Paper Accepted in ICAIC-2026

381

Reliable Extraction of Cyber-Physical System Threat Knowledge with Multi-Run LLMs

Lara Habashy, Illia Poplawski, Francois Rheaume
Mission Critical Cybersecurity
DRDC – Valcartier Research Centre
Quebec, Canada
lara.habashy@drdc-rddc.gc.ca

ABSTRACT

Large language models (LLMs) enable automated extraction of cyber-physical system (CPS) threat knowledge from technical documents, but variability, layout complexity, and hallucinations limit their reliability. In this paper, we propose a schema-guided, provenance-aware structured CPS threat extraction pipeline that combines multimodal PDF parsing, zero-shot LLM extraction, and strict source grounding. Extractions are aggregated across multiple runs and verified through exact string matching against the source document to reduce hallucination risk and support auditability. We evaluate the framework on CPS-focused aviation cybersecurity documents, comparing Docling, an open-source layout-aware parser, with OpenAI's proprietary multimodal parser. We manually annotate systems, threats, and system–threat linkages to establish ground-truth labels for evaluation. When aggregating 25 independent runs, the OpenAI parser achieves up to 97% System F1, 95% Threat F1, and 86% Linkage F1, whereas Docling remains below 75%, 30%, and 67%, respectively. Multi-run aggregation substantially improves recall, with most gains realized within 10 to 15 runs. The results demonstrate that LLM-assisted CPS threat extraction is feasible for safety-critical contexts when outputs are grounded, aggre

Keywords: Information Extraction, Structured Data Extraction, Named Entity Recognition, Relation Extraction, Large Language Models, Threat Extraction, Cyber-Physical Systems

Abstract of Paper Accepted in ICAIC-2026

382

LightGBM-Based Multi-Class Botnet Detection Framework for IoT Networks

Vidhubala J, Kayalvizhi R
SRM IST, India
vj0154@srmist.edu.in, kayalvir@srmist.edu.in

ABSTRACT

The rapid growth of Internet of Things (IoT) had made the network heterogen more complicated and eous kind of environment, opened up a way to create large scale cyberattacks, especially botnet-driven Distributed Denial of Service (DDoS) attack. Traditional Intrusion Detection Systems (IDS), predominantly based on signature or rule-matching, are insufficient for detecting novel and evolving attack patterns due to their high dependency on predefined signatures and limited adaptability to multi-class traffic. As a result there is an increasing demand for smart, scalable, and data-driven detection techniques which are capable of operating efficiently in large volume of IoT networks. This paper presents a resilient machine learning based multi-class botnet detection framework built around the Light Gradient Boosting Machine (LightGBM) model. Although LightGBM is the primary method proposed, its performance is benchmarked against three widely used algorithms like XGBoost, Random Forest (RF), and K-Nearest Neighbors (KNN) to provide comprehensive comparison. Using the network traffic benchmark dataset, the study follows a structured workflow that includes data cleaning, categorical encoding, anomaly management, feature transformation, and balanced model training. The result shows that LightGBM model perform well than the other classifiers in prediction and evaluation. Its leaf-wise tree growth strategy and histogram-based split techniques allows it to capture complex, non-linear traffic patterns more efficiently. These techniques make the model suited for real-time botnet detection in resource-constrained IoT environments.The experimental result proves the efficiency of the proposed framework, which offers a scalable solution for the multi-class botnet detection. This study shows potential of machine learning could be a key part in next-generation IoT security systems, and capable of mitigating emerging cyber threats with high accuracy and low computational cost.

Keywords: LightGBM, Botnet Detection, Multi-Class Classification, Cybersecurity, XGBoost, KNN, Random Forest

Abstract of Paper Accepted in ICAIC-2026

385

Graphene: Leveraging Transformers with Control Flow Modalities for Malware Detection

Andrew Wheeler, Kshitiz Aryal, Maanak Gupta
Tennessee Technological University, USA
amwheeler43@tntech.edu

ABSTRACT

The proliferation of malicious programs for the Windows operating system has highlighted the need for detection frameworks that protect these vital systems. Traditional malware detection approaches struggle to generalize across the diverse and evolving behavioral patterns that are demonstrated by modern malware. In order to gain an upper hand against malware developers, researchers have experimented with varying representations of programs that can be used in conjunction with machine learning models to automatically detect and mitigate malware threats. This work presents Graphene, a graph-based malware detection framework that leverages a RoBERTa-based classification model for malware detection. Graphene encodes program behavior through graph representations that are subsequently linearized into function-call sequences, enabling the use of advanced language-model architectures for malware detection. We evaluate Graphene on a dataset of approximately 150,000 Windows Portable Executable (PE) files and benchmark its performance against two conventional machine learning models: a deep neural network (DNN) and a recurrent neural network (RNN). Experimental results demonstrate that Graphene significantly outperforms these baselines, achieving 94% accuracy with a false positive rate below 4%.

Keywords: Control Flow Graphs, Graph Algorithms, Malware Detection, RoBERTa, Transformers, Windows PE

Abstract of Paper Accepted in ICAIC-2026

386

A Survey on Machine Learning Applications for Operating System Fingerprinting

Siri Siqveland, Alexandra Newcomb, Omar Ochoa
Embry-Riddle Aeronautical University, USA
SIQVELAS@my.erau.edu

ABSTRACT

In the modern age of computers and interconnected networks, cybersecurity and cyber-attackers are evolving in tandem to exploit each other's vulnerabilities. One technique used by both parties is Operating System Fingerprinting (OSF): with the knowledge of what Operating System a target system is running, innate vulnerabilities can be identified and patched or exploited. Historically, OSF utilizes two main methods: passive and active—the former trades accuracy with undetectability while the latter is generally more detectable but more accurate. However, recent work has combined OSF with Machine Learning (ML) to improve accurate identification. The work presented here is a survey for the applications of ML on OSF for both active and passive methods and discusses how ML methods compare to the traditional non-ML methods. Various ML techniques are discussed in terms of their applications and accuracy, e.g., K Nearest Neighbor, Decision Trees, and Support Vector Machines. New tools have also been developed that apply ML to OSF, and the accuracy and methods of these tools are also detailed. The data compiled in this survey are then used to determine gaps in the research and possible direction for future work.

Keywords: Operating System, Machine Learning, fingerprinting, cybersecurity

Abstract of Paper Accepted in ICAIC-2026

389

Multi-Layered Security Framework for Financial AI Solutions – PISA

Aruun Kumar, George Belsian, Santhosh Srinivasan, Giridhar Sankararaman

ABSTRACT

Generative Artificial Intelligence (GenAI) is rapidly transforming organizations, enabling intelligent customer interactions, automated analysis, and agentic decision workflows. At the same time, GenAI introduces a new class of probabilistic and semantic vulnerabilities like prompt manipulation, hallucinations, sensitive-information leakage, and unauthorized agentic actions, that traditional cybersecurity controls are not designed to mitigate. These risks are particularly acute in financial services, a sector widely recognized as part of national critical infrastructure, where security failures can trigger immediate financial loss, regulatory non-compliance, and systemic instability. This paper introduces PISA, a four-layer security framework designed specifically to protect GenAI applications in financial institutions. PISA integrates prompt-layer defense, information-boundary controls, semantic verification, and agent-authorization safeguards into a unified defense-in-depth model. We present a reference architecture implementing PISA, along with an adversarial evaluation across high-risk financial use cases. The results demonstrate substantial reductions in prompt-injection success, hallucination rates, and unauthorized action attempts compared to baseline GenAI systems. This research provides financial institutions with a practical and sector-aligned methodology for deploying secure, trustworthy, and regulation-aware GenAI systems.

Keywords: Security Framework, Generative AI, Artificial Intelligence, Financial Services

Abstract of Paper Accepted in ICAIC-2026

395

Improving Temporal Ordering with External Data

Rahul Cherekar
Chewy, USA
rahul.cherekar@gmail.com

ABSTRACT

This paper explores the efficacy of leveraging external data sources to improve the temporal ordering of historical events. Through machine reading techniques and classification models, we investigate the extraction of contextual information from articles to enhance event ordering accuracy. Our experiments reveal significant improvements, with up to 30% enhancement in accuracy compared to using event titles alone. Leveraging Wikipedia as a primary external source, we demonstrate promising results, paving the way for further advancements in this domain.

Keywords: Temporal Ordering, External Data, Data Sources, Historical Events

Abstract of Paper Accepted in ICAIC-2026

404

Mid-Generation Jailbreaks in Open-Source LLMs Using a Pause-and-Edit Attack

Mark Spanier, Aman Singh, Komal Subhash More, Edward French, Samyam Aryal

ABSTRACT

Large Language Models (LLMs) often rely on safety-aligned system prompts to prevent harmful responses, yet most defenses assume that safety must hold only at the start of generation. We examine a mid-generation jailbreak, called pause-and-edit, that interrupts a refusal, removes the partially generated text, and resumes generation with a cooperative continuation prefix. Because LLMs decode autoregressively, this recontextualization leads the model to reconsider the request and often disregard its earlier safety decision. We test the attack on 390 harmful prompts across 13 categories using three open-source models: Mistral-7B-Instruct-v0.2, Qwen3-14B, and CodeLlama-7B-Instruct. All three models overturned nearly all initial refusals, producing harmful outputs at rates of 99.7%, 85.67%, and 62.56%, respectively. Limited tests on Gemma-27B-IT and Phi-14B-Reasoning show that robustness is influenced more by training and reasoning behavior than by model size. These results indicate that refusal at the start of generation does not guarantee safety. Defenses that operate throughout the decoding process are needed to address this mid-generation vulnerability.

Keywords: LLM jailbreaks, mid-generation LLM jailbreaks, pause-and-edit continuation, safety alignment, autoregressive decoding, model vulnerability evaluation, context manipulation, token-level safety failures, refusal override

Abstract of Paper Accepted in ICAIC-2026

405

Jailbreaking Large Language Models: Techniques, Trends, Defenses, and Open Challenges

Vaishali Vinay,
Microsoft, USA

vaishali.papneja@microsoft.com

ABSTRACT

Large language models (LLMs) are rapidly becoming embedded in everyday applications, security workflows, and decision-support systems. Along with their capabilities, we have also seen a rise in techniques aimed at bypassing their guardrails. This paper presents a consolidated analysis of recent jailbreaks attacks and groups them into three high-order categories based on their underlying mechanism. Unlike prior surveys, this work introduces a mechanism-grounded taxonomy and an explicit attack-to-defense mapping that unifies cognitive, representational, and system-level jailbreak pathways. We highlight the emerging trends that make modern jailbreaks more effective and discuss where current defenses tend to fail in practice. Building on these insights, we outline practical, engineering-focused best practices to lower jailbreak risk in deployed systems. The paper closes with key open challenges—from benchmark gaps to reproducibility issues and agentic safety—that must be addressed to build more robust and dependable LLM safety architectures.

Keywords: jailbreak, LLM, adversarial prompting, agentic systems, defense mechanisms, prompt injection, LLM vulnerabilities

Abstract of Paper Accepted in ICAIC-2026

422

MBIST++: An Adaptive March Algorithm Generator for Memory Test Coverage Enhancement in Post-Silicon Validation

Deepika Bhatia,
NVIDIA, United States
reachdeepikabhatia@gmail.com

ABSTRACT

As embedded memories continue to dominate System-on-Chip (SoC) designs, their testing and validation have become critical bottlenecks for post-silicon reliability. Conventional Memory Built-In Self Test (MBIST) architectures typically rely on static March algorithms which, although effective for classical fault models, struggle to expose complex or emergent defects introduced by advanced technology nodes and layout-dependent anomalies. This paper proposes MBIST++, a dynamic and hardware-efficient framework that generates adaptive March test sequences guided by runtime fault profiling. In contrast to traditional approaches, MBIST++ embeds an on-chip evolutionary engine that synthesizes high-coverage test patterns by learning from post-silicon fault signatures in real time. The architecture couples a microcoded execution engine with a feedback-driven generator and fault classification module, enabling responsive and targeted validation of embedded memory blocks. Experimental results on industrial-grade memory macros show that MBIST++ achieves up to 96.2\% fault coverage, outperforming March SS and machine-learning-enhanced BIST methods, while preserving competitive test latency and incurring only modest area overhead. These findings suggest that MBIST++ offers a scalable path toward post-silicon resilience and adaptive memory test automation for emerging semiconductor platforms.

Keywords: MBIST, March Algorithms, Memory Testing, Post-Silicon Validation, Fault Coverage, SoC Verification

Abstract of Paper Accepted in ICAIC-2026

423

DFT AI: Machine Learning--Guided Test Point Insertion for Pre-Silicon Debug and Post-Silicon Diagnosis in Complex VLSI Systems

Deepika Bhatia,
NVIDIA, United States
reachdeepikabhatia@gmail.com

ABSTRACT

The rising complexity of modern System-on-Chip (SoC) designs has made it increasingly important to deploy intelligent Design-for-Testability (DFT) strategies that support both pre-silicon debug and post-silicon fault diagnosis. Conventional test point insertion (TPI) techniques are typically driven by static heuristics or ATPG-based analyses, and often struggle to scale in the presence of diverse microarchitectures, tight timing budgets, and strict area constraints. This paper introduces DFT AI, a machine learning--guided framework that automates TPI using data-driven reasoning and structural learning. The approach combines gradient-boosted decision trees with graph neural networks (GNNs) to estimate signal importance from structural, dynamic, and contextual features extracted from synthesized RTL designs. A constraint-aware selector ensures that only high-impact, timing-safe test points are chosen. We evaluate DFT AI on ITC'99 and OpenCores benchmark suites and observe up to 28\% improvement in observability gain and a 27\% reduction in diagnostic ambiguity over state-of-the-art methods, with modest area penalty and significantly reduced runtime. The framework generalizes well across heterogeneous designs and provides interpretability via SHAP-based signal analysis, offering a scalable and trustworthy solution for lifecycle-wide DFT automation.

Keywords:

Abstract of Paper Accepted in ICAIC-2026

438

A Lightweight Deep Learning Based Classification Models for Non-Human Identity Threat Detection

Shiva Kumara
T-Mobile, USA
reachkumaras@gmail.com

ABSTRACT

The rise of non-human entities such as bots and automated scripts that interact with online platforms is a significant challenge to cybersecurity. This research offers a lightweight and efficient based system to identify the source of non-human identity threats from a publicly available Initially, Twitter Human-Bot's dataset underwent a detailed preprocessing pipeline after which feature inconsistencies were addressed and class imbalance was reduced by handling missing values, label encoding, data scaling, reshaping, normalization, and SMOTE-based class balancing. Subsequently, the processed data were divided into training and testing sets. The hybrid DL models, CNN-BiLSTM and GRU, advanced to capture the temporal behavior and extract the relevant features efficiently. The findings in the table show that CNN-BiLSTM and GRU models achieved the highest accuracy of 81.22% and 80.94% respectively, compared to the traditional methods, such as Adaboost, accuracy is 73.5%, DNN accuracy is 70%, CAE accuracy is 74.69%, and Logistic Regression accuracy is 77.64%. The high accuracy and reliability of the proposed method in detecting non-human objects is a development in cybersecurity because it offers a credible detection system that can effectively address identity-based threats in dynamic cyber environments.

Keywords: Threat Detection, Non-Human Identity, Twitter Human-bot, Machine Learning Deep Learning, Convolutional Autoencoder.

Abstract of Paper Accepted in ICAIC-2026

446

Hierarchical Attention Distillation for Real-Time Cyber Threat Detection and Mitigation in Large-Scale Networks

Ramkinker singh, CMU, singhramkinker@gmail.com
Om Narayan, New York University, on371@nyu.edu
Praveen Baskar, Google LLC, praveenbaskar@google.com
Sabitha Muppuri, Independent researcher, sabitha594@gmail.com

ABSTRACT

We present a Hierarchical Attention Distillation (HAD) framework that tackles real-time detection and mitigation of cyber threats, where traditional intrusion detection systems often struggle with traffic diversity and regional threat patterns. This integration combines local specialized models, a global meta-model, and a contextual attention mechanism to facilitate hierarchical knowledge fusion and adaptive threat response. Local models, implemented as Temporal Convolutional Networks, process regional traffic data in low latency, while a transformation-based meta-model transmits cross-region threats. This is achieved through bidirectional knowledge distillation. In addition, a context- aware attention mechanism dynamically focuses on high-risk areas by weighting local contributions based on network conditions. The proposed "system interfaces seamlessly with existing security infrastructure, such as SIEM systems and firewalls, allowing for automated actions." Experimental results show that HAD has superior detection accuracy and response speed compared to centralized or static architectures. This modularity promotes scalability and compatibility with heterogeneous network environments, making it a practical solution for contemporary cybersecurity issues.

Keywords: Cybersecurity, Intrusion Detection Systems, Hierarchical Learning, Knowledge Distillation, Attention Mechanisms, Federated Learning, Temporal Convolutional Networks, Transformers, Real-Time Threat Detection, LargeScale Networks

Abstract of Paper Accepted in ICAIC-2026

447

TabNet-IDS: A TabNet-Driven Tabular Deep Learning Framework for Intrusion Detection Systems

Tarek Mahmud
Texas A&M University-Kingsville, USA
Tarek.Mahmud@tamuk.edu

ABSTRACT

The rapid growth of networked systems has increased exposure to sophisticated cyberattacks, demanding intrusion detection methods that can learn directly from complex traffic data. This paper proposes a TabNet-driven intrusion detection framework built on the UNSW-NB15 benchmark dataset. The approach first performs data cleaning, label encoding of categorical attributes, and standardization of numerical features, followed by a binary labeling scheme that groups all attack types into a single intrusion class against normal traffic. The TabNet classifier is trained using 75% of the labeled data, while 25% is reserved for validation, and its generalization is assessed on an independent unseen test set. Experimental results show strong and balanced performance, achieving a precision of 92.68%, recall of 94.23%, F1-score of 93.45%, accuracy of 92.72%, and a Matthews Correlation Coefficient of 85.28%. The ROC and Precision–Recall curves further confirm high separability between normal and intrusive flows, with areas under the curve of 0.975 and 0.977, respectively. A comparison with several representative state-of-the-art models, including GRU-based, decision tree, ensemble, and geometric-analysis methods, demonstrates that the proposed TabNet framework provides the most favorable combination of evaluation metrics, indicating its suitability for accurate and reliable network intrusion detection.

Keywords: Intrusion Detection, UNSW-NB15, Cybersecurity, Network security, TabNet, Deep Learning

Abstract of Paper Accepted in ICAIC-2026

457

DeepNetDetect: A Deep Learning-Based Approach for Early Anomaly Detection in Network Traffic

Henry Cyril
T-Mobile, USA
henry.cyril.tech@gmail.com

ABSTRACT

Network traffic anomaly detection is a current topic in network security. Based on unsupervised learning, this paper constructs a Model for Detecting Irregularities in Network Data to solve the problems of high dimensions of abnormal traffic. This paper provides a fully connected DNN of multi-class intrusion detection. The CICIDS2017 dataset underwent thorough preprocessing, such as feature selection, and SMOTE was used, with the minority attack classes being balanced. DNN model was proposed, which included four dense layers with ReLU activations and a single sigmoid output layer, and was trained on 75: 25 data split. The experimental results are outstanding with an accuracy of 99.70, perfect recall, 99.89 precision and 99.80 F1-score. DNN model performs better than both the traditional machine learning models used, Decision Trees 80.84% accuracy, Naive Bayes (97.48% accuracy), and Auto Encoder + Logistic Regression (98.06% accuracy) and DLmodels such as CNN (96.50% accuracy) and LSTM (98.20% accuracy). The model is highly generalized, converges quicker, and has few false positives, which depict its suitability in practical conditions. The present research offers a strong, extensible and high accuracy framework to the contemporary network security system.

Keywords: Cybersecurity, Network security, Intrusion detection systems (IDS), Anomaly detection, Zero-day attacks, Traffic analysis, Machine learning,

Abstract of Paper Accepted in ICAIC-2026

459

FraudGNN: Self-Supervised Graph Neural Anomaly Detection for Real-Time Financial Fraud with Adversarial Robustness and Explainable Reasoning

Srikumar Nayak
Sr Member IEEE
Incedo Inc, USA
Srikumar.nayak2025@gmail.com

ABSTRACT

Financial fraud detection systems must operate at a very large scale, under extreme class imbalance, and with strict constraints on false alarms, where small changes in attacker behavior can quickly reduce model reliability. Recent research has shown that graph learning and self-supervised objectives can improve anomaly detection by capturing relations between transactions and shared identifiers, but many practical pipelines still struggle to jointly achieve strong recall at low false-positive rates, stable performance under feature manipulation, and efficient inference for real-time screening. In this paper, we propose \textbf{FraudGNN}, a relation-aware graph neural network that converts transaction records into a transaction-entity heterogeneous graph and learns fraud scores using a combined objective that includes (i) weighted supervised learning for imbalanced labels, (ii) contrastive self-supervised learning across augmented graph views to improve label efficiency, and (iii) adversarial training to improve robustness to small feature perturbations. Experiments on the IEEE-CIS Fraud Detection benchmark show that FraudGNN outperforms strong tabular and graph baselines, achieving an AUC-ROC of 0.915 and an AUC-PR of 0.862, with Recall@1\%FPR reaching 0.607 while also improving calibration (Brier score 0.024). Robustness evaluation further shows that FraudGNN degrades more slowly under increasing perturbation strength, maintaining AUC-PR of 0.832 at $\epsilon=1.0$. Ablation results confirm that self-supervision, adversarial training, and heterogeneous relation modeling each contribute to the final performance. Overall, FraudGNN provides a practical, accurate, and robust fraud detection framework that aligns with real-world deployment requirements.

Keywords: fraud detection, graph neural networks, self supervised learning, adversarial robustness, federated learning, anomaly detection

Abstract of Paper Accepted in ICAIC-2026

461

Adaptive Loyalty Systems: A Reinforcement Learning Framework for Dynamic and Context-Aware Benefit Allocation

Tarun Kalwani, Balakumaran Sugumar
Independent Researcher
Atlanta, GA USA
tarun.kalwani17@gmail.com, sugumar.balakumaran@gmail.com

ABSTRACT

The traditional loyalty programs rely on static reward structures incapable of engaging customers at an individual level, thereby leading to standstill retention rate and inefficient budget allocations. This paper proposes a reinforcement learning framework for an Adaptive Loyalty System that optimizes benefit allocation in a dynamic fashion. Unlike rule-based systems depending on historical segmentation, in our system, the agent learns optimal reward strategies by continuously interacting with customer behaviour environments. The goal shall be to maximize Customer Lifetime Value while reducing incentives cost as much as possible. We used a dataset of 446 distinct customer instances, including transaction history, browsing behaviour, and response to previous offers. The system was implemented using Python; for neural network approximation, its library TensorFlow was used, while Pandas served for data manipulation. It perceives the customer's current state, depicted in dimensions of recency, frequency, and depth of engagement, and chooses an action ranging from monetary discounts to experiential rewards. Results have shown that the reinforcement learning model has turned out to perform considerably better than traditional static methods of allocation in huge increases of redemption rates and overall customer sentiment. Treating loyalty management as a problem of sequential decision making, businesses can shift away from discounting reactively toward proactive, context-sensitive relationship building. The present work provides an overview of the proposed framework, demonstrating architecture, training process, and performance metrics while explaining how this can revolutionize customer retention strategies across the digital commerce industry.

Keywords: Reinforcement learning; Dynamic loyalty; Context awareness; Personalization; Benefit allocation.

Abstract of Paper Accepted in ICAIC-2026

469

Supply Chain with Sixth Sense Agentic AI

Venkatesh Prabu Parthasarathy, Pepperdine University USA

venkateshprabu2003@gmail.com
Madhusudan Sharma Vadigicherla, Purdue University USA
reachmadhusv@gmail.com

ABSTRACT

The complexity and instability of today's supply chains are forces that are going to require decision-making frameworks that are proactive, adaptive, and intelligent. The paper being presented here is about the introduction of the Sixth Sense Agentic AI, a new architecture whose purpose is to steer the traditional planning of supply chains into a nonstop loop of intelligence that is always anticipating. The system consists of four layers, namely Perception, Cognitive, Agentic, and Human–AI Collaboration, which support the ingestion of data in real-time, predictive analytics, the execution of operations by autonomous entities, and supervision by the responsible party. Uses of the system are, among others, demand sensing, multi-echelon inventory optimization, scenario planning, and collaborative Sales & Operations Planning (S&OP). The trial of the framework was done on 24 months of mixed data: historical and real-time data was used together with POS, weather, and supplier performance data, plus social sentiment information. The operational results of this study were refreshing: accuracy of forecasts went up by 35%, cost incurred from inventory outweighed by 18%, time for planning cycles cut by 50%, time to respond to scenarios made better by 89%, stock-outs down by 61%, and 13% more deliveries being on time than in the past. Feedback from users confirms the quality of decision-making has improved, that collaboration across different functions has increased, and that planners are more confident, but they also have stressed the importance of having governance and supervision in place. The study says that Sixth Sense Agentic AI can bring about a supply chain ecosystem that is proactive, autonomous, and resilient, thus making it possible to gain both operational and strategic advantages.

Keywords: Sixth Sense Agentic AI, Autonomous Supply Chain, Predictive Analytics, Scenario Planning, Inventory Optimization, Human-AI Collaboration, Real-Time DecisionMaking, Supply Chain Resilience.

Abstract of Paper Accepted in ICAIC-2026

470

A Time-Series and Machine Learning Framework for Forecasting GDP and Unemployment Using Global Economic Indicators (2020–2024)

Vidya Rajasekar, Baby Munirathinam,
Malla Reddy Deemed To Be University, India
vidyarajesh23@gmail.com, babyrathinam@gmail.com
Rajesh Vayyala, Harrisburg, NC, vayyalarajesh@gmail.com
Krithika Rajendran, Arlington, Texas, rkrithika1993@gmail.com
Arul N, St.Peter's Engineering College, arulthala82@gmail.com
Safiya Begam G, B.S. Abdur Rahman Crescent Institute of Science and Technology, safiya.begam@gmail.com

ABSTRACT

This paper is based on a hybrid forecasting framework, a combination of classical time-series models (ARIMA), multivariate econometric models (VAR), machine learning models (Random Forest, Gradient Boosting), and deep learning models (LSTM) to predict Gross Domestic Product (GDP) and unemployment over 2020-2024, using global economic indicators. According to Exploratory Data Analysis (EDA), positive trends are observed in GDP growth, and the decrease in unemployment rates suggests that the world has recovered from the pandemic. The correlation analysis claims the relationship between GDP and unemployment is moderate and negative. The simulations demonstrate that ensemble-based machine learning approaches are better than standard time-series models, and VAR can be used to predict the dependency between GDP and unemployment. The resulting insights illustrate the usefulness of hybrid forecasting models, especially in short-horizon macroeconomic forecasting.

Keywords: Prediction model, Hybrid framework, Machine learning, Multivariate analysis, ARIMA, Deep learning, Unemployment rate

Abstract of Paper Accepted in ICAIC-2026

472

Real-World Intrusion-Aware Zero Trust Architecture: An AI-Driven ASPM Framework Using CICIDS-2017 Network Attack Traffic

Sreenivasulu Gajula, Subhash Bondhala, Madhuri Margam

Independent Researcher, United States
sreenivasgajulausa@gmail.com, subhashbondhala@gmail.com,
Madhurimargam2@gmail.com

ABSTRACT

The rise of enterprise networks and the general popularity of advanced cyber threats being targeted at organizations, specifically zero-day attacks that are able to circumvent traditional signature-based security systems. This paper presents an Intrusion-Aware Zero-Trust Modelling Framework with Application Security Posture Management (ASPM) to provide a suitable model for application security, using data from the CICIDS-2017 dataset. The framework achieves very accurate traffic classification using ANN, Random Forest, and XGBoost models. In addition to regular intrusion detection, the ANN model builds and provides a low-risk score from available input data, allowing organizations to make zero-trust access decisions and inform ASPM policies related to attack awareness. The experimental findings indicate that ANN outperforms other models, achieving 98.56% accuracy, 96.01% F1-score, and 0.998 AUC, whereas XGBoost offers a good trade-off between accuracy and real-time performance. The suggested models were also compared with baseline models, such as Logistic Regression, AdaBoost, and CNN-LSTM, which demonstrated worse detection performance. The proposed framework is thus a scalable, adaptive, and practical means of real-time intrusion detection and the enforcement of zero-trust policies in modern networks.

Keywords: Intrusion Detection System, Zero-Trust, Application Security Posture Management, Artificial Neural Network, CICIDS-2017, Zero-Day Attack.

Abstract of Paper Accepted in ICAIC-2026

473

Iterative Verification and Batch Processing for Enhancing Accuracy and Confidence Computation in LLM-Based Phishing Email Detection

Aisvarya Adeseye, Jouni Isoaho
University of Turku, Finland

ABSTRACT

Large Language Models (LLMs) are preferred over traditional machine learning because they analyze unstructured email text without requiring extensive feature engineering, large labeled datasets, or specialized domain expertise. However, outputs can be unstable or hallucinated. Hence, iterative verification prompting is required to refine and validate responses, especially for smaller local models. Confidence scoring is important for decision-making, yet direct confidence estimation (DCE) by LLMs is unreliable; therefore, a structured confidence metric computation is required to improve accuracy and trustworthiness. The proposed framework was evaluated on a dataset with 2,000 emails using GPT-5.1 and LLaMA models (8B, 3B, and 1B). Classification performance reached 99.9% for GPT-5, 97.3% for LLaMA-8B, 97.15% for LLaMA-3B, and 96.95% for LLaMA-1B, representing an accuracy gain of approximately 32% to 78% compared to baseline. Batch processing contributed more to accuracy improvements than verification prompting alone, while combining both yielded the strongest results. Additionally, the WFBCC method consistently outperforms DCE across all confidence levels, with the largest gains observed in smaller LLMs.

Keywords: Phishing Email Detection; Large Language Models (LLMs); Iterative Verification; Weighted Factor-Based Confidence Computation; Hallucination Mitigation

Abstract of Paper Accepted in ICAIC-2026

474

Novel Textual Data Cleaning Techniques for Cybersecurity Recommendation Extraction and Prioritization using Local LLMs

Aisvarya Adeseye, Jouni Isoaho, Seppo Virtanen, Mohammad Tahir
University of Turku, Finland

ABSTRACT

It is important to understand different perspectives on how people perceive security risks associated with the use of digital services to improve user protection and system security. Interviews help elicit rich information to capture user perspectives. Traditional qualitative analysis of interviews is slow and time-consuming. However, Large Language Models (LLMs) offers a faster alternative; cybersecurity analysts can quickly gain helpful insights with little expertise in qualitative studies. Also, while expert interviews contain clear technical terms, non-experts rely on simple, non-technical language that makes responses unclear, noisy, and difficult to interpret for both human analysts and LLMs. To improve the extraction and prioritization of cybersecurity recommendations for qualitative transcripts, this study proposes nine novel textual data cleaning techniques with roots in digital signal processing (DSP). This systematic cleaning pipeline reduces textual noise, structures interview data, and improves ambiguous language for a more consistent and accurate data analysis. The impact of the proposed cleaning pipeline was evaluated on an interview dataset of 82 (28 cybersecurity experts and 54 non-experts from diverse organizational sectors) using both software-assisted (NVivo) manual analysis and local LLM-based analysis with LLaMA v3.1 (8B) for theme extraction, recommendation extraction, and impact-based prioritization. The pipeline performance was measured using the F1-score, Precision, False Positive Rate (FPR), Spearman's correlation (ρ), and Rank Hallucination Rate (RHR). The results showed improved accuracy with significantly lower hallucinations for both evaluation methods, with the strongest improvements observed in the LLM output.

Keywords: Qualitative Data Cleaning, Hallucination Reduction, Cybersecurity Recommendation Analysis, Large Language Models (LLMs), Digital Signal Processing.

Abstract of Paper Accepted in ICAIC-2026

480

A Cloud-Native Framework for the Petabyte-Scale Purge of Regulated Data: Achieving Compliance and Performance in the Financial Domain

Santosh Kumar Kotakonda
Independent Researcher, USA
santoshkumarkotakonda84@gmail.com

ABSTRACT

Financial record management at global organizations falls under very stringent, absolutely unyielding retention mandates (for example, SEC Rule 17a-4). And thus arises the problem of compliant data deletion—wiping petabytes of data once its retention period has expired—from legacy systems that were never built with such a purpose in mind. Manual processes are often invoked thereby introducing compliance risk and massive performance bottlenecks. This paper proposes Regulated Data Purge Framework (RDPF), a revolutionary cloud-native multi-stage architecture on AWS meant to fill in for such old systems.The RDPF orchestrates a high-throughput polyglot purge pipeline between Amazon Aurora, DocumentDB (MongoDB), and Amazon S3. The main contribution can be described as a four-stage pipeline that enables compliant purging without any performance loss: identification, extraction, purging and non-repudiable certification. RDPF scalability benchmarking has been evidenced by an ability to identify 4 billion records for destruction within less than 4 hours at 1 billion records per hour. Fully automated regulatory compliance is built into the system through generation and immutable storage of the Attestation of Dispose (AoP) in AWS.

Keywords: RegTech, Cloud-Native Architecture, Petabyte-Scale Computing

Abstract of Paper Accepted in ICAIC-2026

482

Benchmarking Container Orchestration Platforms: A Comparative Analysis of Kubernetes and AWS ECS for Stateful Microservices

Santosh Kumar Kotakonda
Independent Researcher, USA
santoshkumarkotakonda84@gmail.com

ABSTRACT

Choosing the right container orchestration platform becomes one of the most important architectural decisions of deploying high-throughput microservices, especially those running stateful workloads typical in such industries as financial services. This paper compares empirically the two options available from Amazon, Elastic Container Service (ECS) and Elastic Kubernetes Service (EKS), when both are run with the serverless AWS Fargate compute engine. It undertakes a four-dimensional differentiation based on operational complexity, performance metrics, security architecture, and Total Cost of Ownership attributable purely to differences at the orchestration layer. It deployed a simulated stateful Spring Boot microservice environment. ECS on Fargate has superior operational simplicity and beats all in Task startup latency, which means it is the best for rapid scaling bursts. EKS does expose the right hooks to maintain finer granular controls over resource scheduling, and thus adopt highly advanced security models such as service mesh architectures that are tantamount to high compliance environments. The TCO analysis did show that the dominating economic factor is not the marginal EKS control plane fee but rather substantially higher operational labor costs due to EKS complexity. The study finds an optimal platform selection strategic tradeoff—ECS to prioritize operational velocity and TCO simplicity against EKS with superior long-run architectural control, multi-cloud optionality, and advanced security maturity.

Keywords: Container Orchestration, Microservices, Kubernetes, AWS ECS, Fargate, Stateful Applications, TCO, Performance Benchmarking, Spring Boot

Abstract of Paper Accepted in ICAIC-2026

483

Intelligent Repair Bots for Power BI: An Agentic Automation Approach

Mohith Reddy Patlolla
Cyma Systems Inc, USA, mohithrpatlolla@gmail.com

ABSTRACT

This paper introduces an architecture that integrates Large Language Model (LLM)-powered autonomous agents within the Power Automate cloud platform to continuously diagnose and attempt remediation of failures in Power BI datasets. Termed here as "Agentic Repair Bots," these agents would perform continuous monitoring over data pipelines for connection timeouts, authentication errors, and schema drift, the most common problems in any pipeline. Contextually aware pre-remediation scripts are triggered by the agents to take immediate corrective actions like refreshing credentials or mapping new schema columns, thus bringing downtime to a minimum while significantly reducing manual intervention not only from analysts but also from all users involved. The research then details how agentic workflow orchestration architecture works within a mobile-cloud environment and evaluates Mean Time to Resolution (MTTR) reduction efficacy.

Keywords: Agentic AI, Power Automate, Power BI, Large Language Models (LLMs), Data Pipelines, Mean Time to Resolution (MTTR), Low-Code/No-Code (LCNC)

Abstract of Paper Accepted in ICAIC-2026

484

Integrating Deep Learning with Power BI Admin APIs for Intelligent Data Governance

Mohith Reddy Patlolla
Cyma Systems Inc, USA, mohithrpatlolla@gmail.com

ABSTRACT

A paper is proposed that will detail a new automated, context-aware data sensitivity classification framework inside Microsoft Power BI that should go some way towards alleviating the primary data governance pains associated with self-service and mobile BI. At the core of this framework sits a Named Entity Recognition hybrid deep learning model designed to both recognize identifiers and apply richer transformer-based capabilities for sensitivity inference in tabular contexts. This model interfaces with the Power BI Admin APIs and also works with Power Automate to provide true real-time end-to-end governance including automatic RLS application plus alerting—designed to adapt based on continuous learning as it sees more examples of new data types or regulation. Most importantly for adoption, it leverages XAI capabilities to not only be transparent regarding human-readable explanations behind its decisions but also toward building trust proactively among fragmented landscapes.

Keywords: Deep Learning, Data Classification, Power BI, Data Governance, Pervasive Computing, Mobile BI, Explainable AI, Hybrid Model

Abstract of Paper Accepted in ICAIC-2026

485

Agentic Metadata Cataloging for Power BI via LLM Scanners

Mohith Reddy Patlolla
Cyma Systems Inc, USA, mohithrpatlolla@gmail.com

ABSTRACT

The paper introduces a scalable ensemble of Agentic Metadata Scanning Framework, Large Language Model-based agents capable of autonomously scanning enterprise Power BI tenants. The AMSF goes beyond simple technical inventory by leveraging the LLMs in semantic inference to translate complex DAX and technical identifiers into business definitions that humans can read, implied data quality rules, and detailed, end-to-end data lineage. The framework is architected with true scalability in mind across cloud API rate limitations and extreme LLM inference costs leveraged through async processing and tiering model strategy. The metadata output finally lands into dynamic mobile-optimized lineage visualization-app-ready enhancing discoverability building digital trust required toward pervasiveness computing environments. Human-in-the-Loop (HITL) governance is hence introduced to keep such AMSF setups faithful while promoting collaborative community-driven metadata curation.

Keywords: Agentic AI; Metadata Management; Power BI; Large Language Models (LLMs); DAX; Data Lineage; Pervasive Computing; Mobile Computing; Data Governance.

Abstract of Paper Accepted in ICAIC-2026

488

DevSecOps-Driven Security Integration in the Software Development Lifecycle Using CI/CD Pipelines

Henry Cyril
T-Mobile, USA
henry.cyril.tech@gmail.com
Shiva Kumara
Independent Researcher
University of Washington
skum33915@gmail.com

ABSTRACT

Continuous Integration and Continuous Deployment (CI/CD) pipelines are essential in today's lightning-fast software development environments for dependable application delivery. However, the growing complexity and speed of modern development practices often lead to overlooked security flaws, leaving applications vulnerable to cyber threats. As a logical extension of DevOps, DevSecOps addresses this issue by integrating security testing and compliance checks directly into the CI/CD pipeline, ensuring that security is treated as a shared responsibility throughout the development process. Conventional post-development security models do not support the fast release cycle requirements, leading to inconsistent vulnerability coverage, high false positives, and slow remediation. To tackle these, this paper proposes a DevSecOps-oriented, CI/CD-built security automation system, which entails vulnerability discovery, reporting, and measures-based assessment, directly integrated into the development pipeline. The system injections are automated, inject insecure commits, run multi-tool static and dependency scanning, and finally check the results against a ground-truth dataset of labeled vulnerabilities. Experimental results display Recall of 0.75, Precision 0.38, F1 score 0.50 with the level of rule detection differing widely (precision with a range of 0.0 to 1.0). Moreover, the commit-level analysis shows gradual improvement in detections as fixes are spread across the workflow. The paper illustrates a data-driven, repeatable approach to assessing DevSecOps toolchains and the need for proper rule engineering, containerized pipelines, and automated reporting to scale security.

Keywords: CI/CD Pipelines, DevSecOps, Vulnerability Detection, Security Automation, Semgrep, SDLC Security.

Abstract of Paper Accepted in ICAIC-2026

490

An Explainable Machine Learning Framework for Predicting Software Defects in Large-Scale Software Systems

Srikanth Kavuri

Independent Researcher, USA

srikanthkavuri.research@gmail.com

ABSTRACT

Software-intensive systems are rapidly growing in size and complexity, and early defect prediction is very important in enhancing reliability and minimizing the costs of maintenance. Traditional defect prediction models focus on performance but offer limited transparency, thereby hampering their application. This research introduces a transparent ML model for large-scale software defect prediction using the JM1 dataset. The framework also includes preprocessing systematic data, addressing class imbalance using SMOTE, and scaling features to improve learning efficiency. An Extreme Gradient Boosting (XGBoost) model is used to model complex, non-linear associations among software metrics. By combining local and global explanations of a prediction, explainable AI systems such as LIME and SHAP address the interpretability gap in black-box models. Experimental testing shows that the suggested method achieves 96.65% recall, 96.65% accuracy, and 96.65% F1-score. The findings indicate that the framework offers high predictability and transparency, supporting informed decision-making in software quality assurance.

Keywords: Large-Scale Software System, Software Defect Prediction, Machine Learning (ML), Artificial Intelligence (AI), Explainable AI (XAI).

Abstract of Paper Accepted in ICAIC-2026

491

Comparative Insights: A Multigroup Analysis of Privacy Management Across Youths, Parents, Educators, and AI Professionals in AI Applications

Molly Campbell, Yulia Bobkova, Ajay Shrestha,
Vancouver Island University, Canada Ajay.Shrestha@viu.ca

ABSTRACT

The integration of Artificial Intelligence (AI) into digital ecosystems has significantly reshaped privacy dynamics, particularly for young digital citizens (ages 16–19). However, the complexities of data management have raised concerns among various stakeholders, necessitating an understanding of how different groups perceive privacy. This study aimed to explore the interactions between key constructs such as Education and Awareness, Data Ownership and Control, and Perceived Risks and Benefits among young digital citizens, parents, educators, and AI professionals, and analyze variations in privacy management perspectives across these groups using multigroup partial least square structural equation modeling (PLS-SEM) on survey data from 424 participants. The findings indicate that greater awareness leads to higher risk sensitivity in young digital citizens compared to AI professionals, while parents and educators weigh risks and benefits more heavily in decisions about sharing youths' data. The importance of data ownership and control was common across all groups. Trust played a stronger mediation role between perceived control and data sharing for AI professionals than for young users. Measurement invariance was only partially achieved, demonstrating fundamental differences in construct perception. These findings highlight the need for stakeholder-driven privacy frameworks that incorporate evolving transparency strategies, user-centric controls, and targeted education to foster ethically aligned AI systems.

Keywords:

Abstract of Paper Accepted in ICAIC-2026

499

Privacy-Preserving Federated Learning for Multi-Tenant CRM Systems

Nidhi Sharma
Independent Researcher, United States
nidhi.sharmatechlead@gmail.com

ABSTRACT

Modern CRM systems face competing demands: ultra-personalization versus stringent privacy requirements. Multi-tenant SaaS architectures containing data from multiple corporate customers on shared infrastructure create high risk of unintended data exposure and regulatory breach. This work introduces and analyzes a privacy-preserving federated learning framework for multi-tenant environments. By separating model training from data storage, a global machine learning model is refined using insights from multiple tenants without raw customer data leaving each tenant's local enclave. The synthetic dataset includes 457 instances containing complex B2B customer interactions, churn indicators, and engagement measures. Implementation uses Python with TensorFlow Federated for decentralized training and Scikit-learn for preprocessing. Results demonstrate that the federated approach achieves predictive performance comparable to centralized training while maintaining strong privacy guarantees through differential privacy mechanisms including gradient clipping and noise injection. The federated model achieves higher accuracy than isolated models, demonstrating a substantial collaborative advantage for resource-constrained tenants. The computational overhead is practical at 4.5 MB per communication round, with convergence in 50 rounds. This work validates that high-performance predictive analytics are achievable in multi-tenant CRM without compromising confidentiality, enabling regulatory-compliant AI deployment in competitive business ecosystems.

Keywords: federated learning, multi-tenant systems, customer relationship management, differential privacy, privacypreserving machine learning

Abstract of Paper Accepted in ICAIC-2026

500

AI-Driven experimentation for user experience: An Architectural Blueprint for Self-Optimizing iOS Experiences at Enterprise Scale

Snehal Mehta, Wal Mart Associates Inc, USA
snehal.mehta203@gmail.com

ABSTRACT

Today we see a shift in the mobile user experience (UX) space from what was mainly static and which had been put together by hand to very dynamic, self-optimising systems. What we also have is that traditional experiment models, which for the most part were A/B testing, are breaking down into what we term "stagnant metrics" that is to say that we are not seeing value out of incremental design changes, which in turn does not translate to long-term growth or engagement. This research presents a comprehensive architectural plan for enterprise-level iOS apps that use on-device artificial intelligence (AI) and reinforcement learning (RL) to enable real-time UI adaptation. We look at what Apple's native hardware, like the Neural Engine, has to offer in terms of high-concurrency analytic streams, and we see how our put-forth framework moves away from a one-size-fits-all approach to a very granular, personal interaction model. Also, we identify key performance issues within the UIKit and SwiftUI rendering pipeline, such as main thread blocking and high-level abstraction, which we, in turn, present solutions for in terms of what it takes to do complex visual changes within the 16.7ms frame budget we have for 60 FPS performance. At the core of this plan is a closed-loop feedback system that uses on-device foundation models to determine user intent and, in turn, change interface elements in real time.

Keywords: Self-Adaptive User Interfaces; Reinforcement Learning; iOS Architecture; On-Device AI; Core ML; Enterprise Mobile Systems; User Experience Optimization; Mobile Analytics; UIKit Performance; Foundation Models.

Abstract of Paper Accepted in ICAIC-2026

504

Machine Learning-Enabled Classification of Diabetes Using Clinical via Lifestyle Health Data

Sanjoy Mukherjee, Cognizant, USA sanjoymukherjee302@gmail.com

ABSTRACT

Diabetes is a recurrent ailment that has enduring effects on the health of a person that is characterized by an increase in blood sugar levels. Because of outliers or missing data in diabetes datasets, as well as the limited number of labeled data, it is hard to accurately predict the existence of diabetes. Machine learning (ML) methods are common among studies in diabetes to analyze data and predict the onset of the disease. This paper is aimed at making predictions on diabetes on the basis of clinical and lifestyle health data of Pima Indians Diabetes Dataset (PIDD). It employed a complete chain of data preprocessing, including managing missing data, outliers, normalization, label encoding, feature selection based on the Principal Component Analysis (PCA), and balancing data with SMOTE. a number of machine learning (ML) classifiers were tested on the PIDD. AdaBoost, Random Forest (RF), and ResNet18 comparison models attained 94.67, 80, and 80.86 in terms of accuracy, respectively. The ANN + XGBoost model proposed performed better as it had an accuracy (ACC) of 97.46, precision (PRE) of 96.44, recall (REC) of 95.42, and an F1-score (F1) of 96.46. These findings indicate the success of using the combination of artificial neural networks and gradient boosting in enhancing predictive power and strength to detect diabetes early on.

Keywords: Diabetes Diagnosis, PIMA Indian diabetes dataset, Machine learning, Hybrid model (ANN+XGBoost), SMOTE, Health Care.

Abstract of Paper Accepted in ICAIC-2026

506

An Empirical Evaluation of Deep Neural Networks as Hash-Like Mappings Under Digital Signature Threat Models

Juan Couder, Lynn Vonderhaar, Omar Ochoa
Embry-Riddle Aeronautical University, USA
ortizcoj@my.erau.edu, vonderhl@my.erau.edu, ochoao@erau.edu

ABSTRACT

Hashing is a crucial step in the Digital Signature (DS) generation and verification process. Traditional DSs rely on mathematically-defined hashing algorithms and public-private key encryption that offer strong formal guarantees. Our previous work introduces ML-256, a deep hashing model trained on random input-output pairings to produce a fixed-length hash-like output. That work demonstrates the feasibility of integrating Machine Learning (ML) based hashing into DS generation and highlights potential benefits in adaptability and maintainability. In this paper, we evaluate ML-256 through a more rigorous empirical evaluation of its behavior as a hash-like mapping algorithm. Rather than claiming formal cryptographic guarantees, we assess ML-256 through empirical tests of determinism, collision behavior under finite sampling, and resistance inversion using ML-based reconstruction attacks. This work additionally analyzes the computational trade-offs introduced by an ML approach by comparing throughput, DS generation and verification latency, and memory consumption against traditional DS pipelines. Results show that while ML-256 is slower and more resource expensive than traditional cryptographic functions, its outputs are difficult to invert using the tested reconstruction methods and exhibit no collisions within the evaluated regimes. Our findings demonstrate that robustness against ML-based attacks does not imply suitability as a cryptographic hash for DS. However, these results hint that ML-256 could be experimentally robust and easily configurable hash-line function for DS.

Keywords: Machine Learning, Digital Signatures, Cybersecurity, Cryptography, Hashing, Deep Learning

Abstract of Paper Accepted in ICAIC-2026

510

A Self-Adaptive Red Teaming Framework for Vulnerability Profiling with Dynamic Attack Graphs and Retrieval Augumented Generation

Jothsna Praveena Pendyala, Clark University, USA
jothsnapraveena1421@gmail.com
Sundeep Bobba, East Texas A&M, USA sundeepbobba@gmail.com
Ali Azghar Hussain Syed Abbas, Independent Researcher, USA
saaahussain@ieee.org

ABSTRACT

The increased complexity of enterprise networks, which has been associated with microservice architectures, transient cloud infrastructure, and short software development cycles, has exceeded the capacity of conventional vulnerability management tools. Often static scanners miss threat due to banner obfuscation, while Large Language Models suffer from hallucination and an inability to reason about newly discovered vulnerabilities in red teaming. To overcome these shortcomings, this paper introduces Graph Retrieval Augmented Generation agent, a self-adaptable multi-agent red teaming architecture combining Retrieval-Augmented Generation with dynamic attack graph formation. The semantic retrieval and time relevance of our system are based on the CISA Known Exploited Vulnerabilities (KEV) catalog and the NVD 2025 Common Vulnerabilities and Exposures (CVE) feed as a hybrid vector database, which are highly precise, thereby guaranteeing accurate semantic retrieval and time relevance. Our proposed agent detects CVEs with 100 percent accuracy and retrieval latency of sub-20ms upon controlled experiments on 15 differentiated targets including obfuscated and novel services, significantly outperforming regex-based scanners and standalone standoff Large Language Models. The paper gives an explicable, scalable, and hallucination-resistant architecture of autonomous vulnerability discovery in practical security settings.

Keywords: Large Language Models, Retrieval Augmented Generation, Graph RAG, Common Vulnerabilities and Exposures, Autonomous Agents, Red Teaming, Cybersecurity

Abstract of Paper Accepted in ICAIC-2026

514

Artificial Intelligence (AI)-Driven Network Traffic Anomaly Detection for IT Infrastructure Security

Dileep Jain, Sharad Jain

ABSTRACT

Anomalies are common in the monitoring of network systems. It is critical to recognize such anomalies in network behavior when it materialize as network hazards to be mitigated, service outages to be avoided, and security issues to be addressed. To determine some suspicious network traffic on CICIDS-2017 dataset, this study develops a CNN-BiLSTM model of hybrid convolutional neural networks (NN). The proposed architecture will contain systematic preprocessing procedures to guarantee quality input information, such as data cleaning, one-hot encoding, feature selection, minmax normalization, and data balancing. The BiLSTM is used to extract the temporal relations and the CNN portion is used to reconstruct the qualities of the network data in space, enabling the realistic modeling of sequential behaviors that are interrelated to cyberattacks. Experimental findings indicate that the suggested model reached 96.58% accuracy (ACC), 97% precision (PRE), 95.64% recall (REC), and F1-score (F1) of 96.09% which are better than the traditional models including Random Forest, Logistic Regression, LGBM, Autoencoder, CNN-LR, XGB and SVM. The comparison shows that the model is better at detecting various types of network anomalies, which make it an incredibly powerful and scalable next-generation intrusion detection system.

Keywords: IOT, Network Traffic, Cybersecurity, Machine Learning, Anomaly Detection, CICIDS-2017 Dataset.

Abstract of Paper Accepted in ICAIC-2026

517

Age-Differentiated Pathways to Privacy Protection in Smart Voice Assistants: A Multigroup PLS-SEM Study of Youth

Yulia Bobkova, Molly Campbell, Trevor De Clark, Ajay Shrestha
Vancouver Island University, Canada
Ajay.Shrestha@viu.ca

ABSTRACT

This paper investigates the way youth interactions with smart voice assistants (SVAs) are influenced by their stage of development, household dynamics, and level of control they have over privacy settings and data collection. Using survey data from 412 Canadian participants aged 16-24, we conduct a multigroup Partial Least Squares Structural Equation Modeling (PLS-SEM) analysis to compare younger youth (16-18; N=245) and older youth (19-24; N=167) across five privacy related constructs: perceived privacy risk, perceived privacy benefits, algorithmic transparency and trust, privacy self-efficacy, and privacy-protective behavior. The results show age-based differences in perceptions, where older youth report greater concern and stronger protective behavior, while younger youth express higher trust and perceived benefits. The fundamental pathway between privacy self-efficacy and protective behavior remains consistent across age groups. Notably, the pathway from algorithmic transparency and trust to privacy self-efficacy is significantly stronger among older youth ($\beta = 0.567$) than younger youth (16-18; $\beta = 0.356$; $p = 0.024$), indicating that developmental stage is associated with stronger links between trust and self-efficacy. These findings point to the importance of age-appropriate privacy design, emphasizing guided and household-aware controls for younger users, and greater autonomy for older youth.

Keywords: Privacy, Smart Devices, Multigroup Analysis, Smart Voice Assistants, PLS-SEM, User control, Youth, Age

Abstract of Paper Accepted in ICAIC-2026

519

Early Detection and Prediction of Depression Based on Data-Driven Machine Learning Techniques in Mental Healthcare

Sanjoy Mukherjee, Cognizant, USA sanjoymukherjee302@gmail.com

ABSTRACT

Depression is a significant public health concern, as it is among the causes of the burden of disease in the world. Both genetic and environmental factors define whether one is at risk of becoming depressed or not. Although genetic influences cannot be changed, it is vital to find out possible reversible environmental factors and make an attempt to restrict the manifestation of depression. As a timely intervention and effective mental care, it is important to identify cases of depression as early as possible. In this work, the authors provide a machine learning (ML) predictive model of depression severity as applied to the DASS (Depression Anxiety Stress Scales) dataset. The LSTM (Long Short-Term Memory) model and XGBoost (Extreme Gradient Boosting) were developed and tested with better results on capturing sequential patterns and crucial emotional features. A lot of pre-processing of data, feature selection, and class balancing methods, such as SMOTE (Synthetic Minority Over-sampling) and random oversampling, were used to increase the reliability of models. The LSTM model got 99.73% in accuracy, precision, recall and F1-score, whereas XGBoost got 99.48% as compared to the baseline models, namely BERT (Bidirectional Encoder Representations from Transformers), Gradient Boosting (GB), Random Forest (RF), Logistic Regression (LR) and Naïve Bayes (NB). The findings support the effectiveness of the suggested methodology in effective, consistent, and almost perfect depression classification, which contributes to its possible use in the field of practice as a mental health assessment and early intervention.

Keywords: Healthcare Domain, Depression, Mental Disorder, Machine Learning, Deep Learning, Dass Dataset.

Abstract of Paper Accepted in ICAIC-2026

521

AI-Based Cybersecurity in Healthcare: A Data-Driven, Governance-Aware Framework for Secure Clinical Systems

Rajani Kant, Bishwajeet Pandey, Ramachander Rao Thallada, Palak Shrivastava

ABSTRACT

Artificial Intelligence (AI) is reshaping the healthcare cybersecurity practices by offering capabilities to detect proactively, flexibly, and with data, threats in clinical, administrative, and medical Internet of Things (IoT) infrastructures. Healthcare is among the most desired fields because of the high cost of electronic health records (EHRs), use of outdated infrastructure, fast deployment of cloud computing, and high-regulatory standards. The current paper will provide an elaborate AI-based cybersecurity framework tailored to healthcare settings in particular. Based on real-world breach intelligence data provided by the U.S. Department of Health and Human Services (HHS), the Verizon Data Breach Investigations Report (DBIR), and the IBM X-Force Threat Intelligence Index, we study incidences of attacks and assess AI-enhanced intrusion detections, anomaly detections, and automated response systems. The suggested framework combines machine learning, deep learning, governance controls, and explainable AI to achieve the compliance of the regulations and the trust of the operations. The experimental analysis provides evidence of better detection accuracy of over 95 percent, decreased incident response time, and improved privacy protection. The article adds a new, governance conscious AI cybersecurity architecture applicable to enterprise healthcare system and in accordance with the emerging regulatory, ethical and operational requirements.Recent researches indicate that artificial intelligence used on real-world healthcare data can have a considerable positive impact on clinical decision-making, risk stratification, and operational efficiency.

Keywords: Artificial Intelligence, Healthcare Cybersecurity,Electronic Health Records, Intrusion Detection Systems, Explainable AI, Data Privacy, Critical Infrastructure Security, Healthcare IoT Protection, Smart Cities, Industrial IoT, Autonomous Cyber-Response

Abstract of Paper Accepted in ICAIC-2026

524

Future-Proofing Cloud Infrastructure: AI/ML-Driven Automation for Predictive Cloud Operations

Vishwa Lakhnakiya
Tata Consultancy Services Inc, USA
vishwalakhnakiya1628@gmail.com

ABSTRACT

The increasing complexity and rapid dynamism of the modern cloud and Mobile Edge Computing (MEC) landscape calls for a real fundamental shift from reactive incident response to proactive predictive operations. In such an environment, traditional CI/CD practices lag in accommodating the scale and velocity required in self-governance, leading to extended downtime accompanied by increased operational expenditure. The CAIRO paper proposes an AI/ML-powered architecture that will preempt infrastructure failure through dynamic resource optimization as well as enhancing overall system reliability. It does so by leveraging cutting-edge Machine Learning (ML) models in the example of Transformer based log detectors and hybrid Neural Networks incorporated directly into CI/CD autonomously extended as CI/CD/CM/CC pipelines. Testing shows the current prototype results in a large decrease in operational lag, about 40% reduction in mean time to repair (MTTR), raising resource utilization up to 85.7% from previous baselines with an over 20% gain related to intelligent resource forecasting. The use of CAIRO Literal heightened operational resilience accompanied by an explicit competitive edge for enterprises within dynamically transforming low latency environments.

Keywords: AIOps, Continuous Correction (CC), Predictive Operations, Machine Learning, Cloud Resilience, CI/CD, Mean Time to Repair (MTTR).

Abstract of Paper Accepted in ICAIC-2026

525

Automating Multi-Cloud Deployments at Scale Using an Advanced GitOps Framework

Vishwa Lakhnakiya
Tata Consultancy Services Inc, USA
vishwalakhnakiya1628@gmail.com

ABSTRACT

Enterprise Kubernetes has led 71% of these organizations to run complex hybrid or multi-cloud infrastructures, while being in production use by 96% of organizations. Managing this complexity across different providers (AWS, Azure, GCP) brings about crucial challenges with respect to configuration drift and maintaining compliant operations in regulated environments. This paper introduces a scalable GitOps architecture based on ArgoCD deployment integrated with proactive Policy-as-Code (PaC) governance and Git-native automated resilience. This architecture exploits version-controlled Git repositories as the source of compliant truth for enabling continuous reconciliation and atomic recovery of the system.

Standard DORA metrics empirically validated this method. The framework realized very substantial quantifiable improvements in delivery performance, by driving to the tune of 45% reduction in Change Failure Rate and up to 85% reduction in Mean Time to Recovery within a large-scale enterprise environment. It also empirically validates that such advanced GitOps frameworks are prerequisite components toward enabling consistent, auditable, and resilient infrastructure management at enterprise scale across distributed cloud-native ecosystems extending out all the way to the edge-cloud continuum.

Keywords: GitOps, Multi-Cloud Orchestration, ArgoCD, Policy-as-Code (PaC), DORA Metrics, Mean Time to Recovery (MTTR), Cloud-Native Governance, Edge Computing.

Abstract of Paper Accepted in ICAIC-2026

528

Machine Learning–Based Fault Prediction in Large- Scale Distributed Systems

Amit Kumar Padhy, University of Illinois Urbana-Champaign, USA
a.padhy@ieee.org

Tejas Patel, Vinay Soni, Sandeep Shivam, Gajendra Babu Thokala, Bharadwaj Vulugundam
Independent Researcher, IEEE, USA

ABSTRACT

Cloud computing has piqued the curiosity of both academics and regular people. Global cloud services are offered by tech giant Google. The framework detailed in this paper is a ML-based fault prediction system that is proposed in the form of an ensemble, which is applied to Google 2019 Cluster Sample dataset and provides a depiction of the actual workload execution and resource utilization patterns on large-scale distributed systems. The given methodology follows a systematic flow which includes data preprocessing feature engineering, labeling of the categorical variables, and tackling the issue of class imbalance. XGBoost, LightGBM and a Voting Classifier are sophisticated ensemble models that are checked and assessed with the help of the standard performance matrix. The experimental results prove that the Voting Classifier is superior to the individual models with a high accuracy rate of 97.97, high precision and recall as it identifies all the cases of faults. The proposed ensemble framework is superior to the existing ML and DL methods, which are further demonstrated through comparative analysis. The findings affirm that feature optimization combined with ensemble learning is a powerful and scalable fault prediction system in distributed systems.

Keywords: Fault Diagnosis, Distributed System, Machine Learning, Failure Prediction, Voting Classifier.

Abstract of Paper Accepted in ICAIC-2026

529

FraudSentinel: Federated Multi-Agent Reinforcement Learning for Privacy-Preserving Cross-Marketplace Fraud Detection in Distributed E-Commerce Ecosystems

ABSTRACT

E-commerce fraud costs the global economy $48 billion annually, with sophisticated fraud rings operating across multiple marketplaces. While individual platforms deploy fraud detection systems, fraudsters exploit the lack of cross-platform intelligence sharing. Traditional centralized fraud databases violate privacy regulations and create single points of failure. We present FraudSentinel, a federated multi-agent reinforcement learning framework that enables privacy-preserving fraud pattern sharing across distributed e-commerce marketplaces without exposing sensitive customer data. We evaluate FraudSentinel on the IEEE-CIS fraud detection dataset (590K transactions) extended with synthetic multi-marketplace scenarios across five verticals. Our framework achieves 96.8% fraud detection accuracy with 0.31% false positive rate—a 9.1% relative improvement over isolated systems—while maintaining <8ms detection latency and providing formal differential privacy guarantees (ε=1.31). Ablation studies demonstrate that federated learning contributes 17.9% of the accuracy gain, while multi-agent coordination adds 24.1%. Security analysis confirms zero customer data leakage across marketplace boundaries. FraudSentinel provides a practical blueprint for collaborative fraud fighting while respecting privacy regulations like GDPR and CCPA.

Keywords: Federated Learning, Multi-Agent Reinforcement Learning, E-Commerce Security, Fraud Detection, Privacy-Preserving Machine Learning, Distributed Systems, Cybersecurity

Abstract of Paper Accepted in ICAIC-2026

530

AI-Driven Data Architecture: Building Intelligent Analytics Platforms with Azure and Python

Suresh Pairu Subramanyam

Avanade, USA sureshpairusubramanyam@gmail.com

ABSTRACT

An Energy Company under transformation at the digital core, led by the urgent requirement to move away from standard post-mortem Business Intelligence. In this respect, this paper proposed an AI-driven data architecture as a platform for intelligent analytics in oil and gas producers. The framework leverages Azure's end-to-end pipeline management scalability and the open flexibility of the Python data science ecosystem to deliver both predictive and prescriptive analytics. Included is a detailed plan for a multi-layered Medallion architecture for data ingestion, cleansing, and consumption, addressing some fundamental issues related to Data fragmentation, quality, and trust. It demonstrates the utility of a framework through an analysis review of those applications that have high impacts toward positive results when applied with AI/ML — these being Predictive Maintenance, Operations Optimization, and Reservoir Characterization. It introduces strategic mitigation techniques for data drift and other critical issues, leveraging transfer learning as well as MLOps. This paper shall serve as a practical blueprint that makes AI real, and sets the groundwork for a data-driven culture, making the long-term dependability and sustainability of intelligent systems in energy achievable.

Keywords: Artificial Intelligence, Machine Learning, Data Architecture, Hydrocarbon Exploration, Azure, Python, Predictive Analytics, Data Drif

Abstract of Paper Accepted in ICAIC-2026

531

A Secure and Scalable Cloud-Based Banking Service Model Leveraging AI and Advanced Cyber Security

Sreenivasulu Gajula, Principal Full-Stack Engineer, USA,
sreenivasgajulausa@gmail.com
Madhuri Margam, Director, Software Engineer, USA
madhurimargam2@gmail.com

ABSTRACT

Security, scalability, and efficiency in cloud-based banking systems have become a critical issue especially with the fast-paced implementation of digital banking. This paper presents a secure and scalable cloud-based banking service model that would combine Artificial Intelligence (AI) and leading cybersecurity mechanisms. The model uses AI algorithms to detect fraud in real-time, to predict analytics and offer tailored banking services, and to use multi-layered cybersecurity strategies such as intrusion detection, encryption, and anomaly detection based on behavior to secure sensitive financial information. The proposed system, through simulation and evaluation demonstrates greater ability to scale up to support growing transaction loads and greater ability to resist advanced cyber-attacks. The most important insights are that the combination of AI and the use of powerful security measures will greatly decrease fraudulent operations and system downtimes and enhance customer satisfaction. The research also determines the real-life challenges connected with cloud resource management, interpretability of AI models, and regulatory compliance and presents the possible way forward on further optimization.

Keywords: Cloud Banking, Artificial Intelligence, Cybersecurity, Fraud Detection, Scalable Systems, Data Privacy, Secure Banking Services.

Abstract of Paper Accepted in ICAIC-2026

532

DevOps and CI/CD Maturity in Large-Scale Organizations: A SonarQube and Jenkins Approach

Suresh Pairu Subramanyam

Avanade, USA sureshpairusubramanyam@gmail.com

ABSTRACT

This paper explains how the combination of DevOps with Continuous Integration, Continuous Deployment/Delivery (CI/CD), and the Identity and Governance Administration (IGA) components creates an actionable pathway to maturity elevation in large-scale organizations. It is based on the premise that standardized, scaled, and optimized software delivery can be supported by strong pillars formed through the judicious use of Jenkins and SonarQube. The framework facilitates demonstrating how Jenkins serves as the primary orchestration engine of a CI/CD pipeline automating build, test, and deployment phases together with SonarQube which will then provide for continuous code quality and security analysis. Therefore, besides serving its purpose as an academic endeavor toward integrated policy-enforcing workflow architecture based on architectural blueprints wherein specifically within this workflow architecture SonarQube Quality Gates stop code not meeting set standards right in the pipeline. It takes that idea further into a complete DevSecOps approach by illustrating how one can leverage automation to identify and map resource owners in the cloud and CI/CD pipeline so as to enforce periodic access reviews making sure all users have the appropriate level of access at any given moment. This paper also references previous works on DevOps maturity models in addition to conducting empirical case studies with top industry practitioners as a means of synthesizing the required place of automation, cultural change, and continuous governance for successful applications of the DevOps framework.

Keywords: DevOps, CI/CD, Jenkins, SonarQube, DevSecOps, Static Analysis (SAST), Quality Gates, Identity and Governance Administration (IGA), Cloud-Native, Maturity Model, Governance, Continuous Delivery, Continuous Integration, CI as Code

Abstract of Paper Accepted in ICAIC-2026

533

TrustGraph: Federated Graph Neural Networks for Cross-Platform Trust and Fraud Propagation Analysis

Tejas Pravinbhai Patel, Arun Kumar, Madhushree Kumari, Rajesh Purushothaman, Rakesh Keshava, Milan Parikh

ABSTRACT

Centralized fraud detection systems in e-commerce ecosystems face significant limitations due to stringent data privacy regulations, platform heterogeneity, and the inherently distributed nature of sophisticated fraud rings operating across multiple marketplaces. Existing approaches predominantly rely on isolated platform-specific models or centrally aggregated data, fundamentally limiting their ability to capture cross-platform trust relationships and fraud propagation dynamics that characterize modern coordinated fraud campaigns. Through extensive experiments on the YelpChi benchmark dataset under realistic non-IID data distributions with K=10 federated clients, we demonstrate that TrustGraph achieves 0.93 AUC, approaching within 0.02 of centralized performance while significantly outperforming local-only models by 0.06 AUC and federated non-graph baselines by 0.08 AUC. Statistical significance tests across five random seeds confirm the robustness of our results, establishing TrustGraph as an effective solution for privacy-preserving fraud detection in distributed e-commerce ecosystems.

Keywords: Federated Learning, Graph Neural Networks, Fraud Detection, Trust Modeling, Privacy-Preserving AI, Cybersecurity, E-Commerce

Abstract of Paper Accepted in ICAIC-2026

534

AgentSCO: A Multi-Layer Agentic Framework for Security Operations Automation

Joyjit Roy, KForce Inc, USA, joyjit.roy.tech@gmail.com
Samaresh Kumar Singh, HP Inc USA, ssam3003@gmail.com

ABSTRACT

Security Operations Centers (SOCs) increasingly encounter difficulties in correlating heterogeneous alerts, interpreting multistage attack progressions, and selecting safe and effective response actions. This study introduces AgentSOC, a multi-layered agentic AI framework that enhances SOC automation by integrating perception, anticipatory reasoning, and riskbased action planning. The proposed architecture consolidates several layers of abstraction to provide a single operational loop to support normalizing alerts, enriching context, generating hypotheses, validating structural feasibility, and executing policycompliant responses. Conceptually evaluated within a large enterprise environment, AgentSOC improves triage consistency, anticipates attackers' intentions, and provides recommended containment options that are both operationally feasible and wellbalanced between security efficacy and operational impact. The results suggest that hybrid agentic reasoning has the potential to serve as a foundation for developing adaptive, safer SOC automation in large enterprises. Additionally, a minimal ProofOf-Concept (POC) demonstration using LANL authentication data demonstrated the feasibility of the proposed architecture.

Keywords: Agentic AI, Security Operations Center, Threat Detection, MITRE

Abstract of Paper Accepted in ICAIC-2026

537

The Architect of Advantage: How Robust Data Curation and Edge-Case Analysis Provides a Disproportionate Market Edge

Rupam Priya,
Manager - Marketing Analytics, USA, rupampriya.001@gmail.com

ABSTRACT

This is a technologically advanced era and most organizations have access to a lot of raw data, and therefore, there is a great deal of parity among them with regard to their ability to compete. This paper illustrates one key strategic flaw in losing valuable data signals due to the use of BI statistical filtering of complexity based on abnormal behavior (edge cases). Data curation represents a process that validates, adds value to, and provides strategic interpretation to complex proprietary data anomalies - which is a true source of sustainable competitive advantage. Therefore, this paper introduces a two-stage ML framework with adjustment to transform edge cases into high fidelity signals from simple statistical noise. In terms of application and impact, within a high velocity pervasive gaming environment, data curation produced an increase in top-line operating revenue of 14.5% during a time of year when tourist activity was at its lowest. From a technical standpoint, the proposed framework increases predictive model accuracy by a factor of 5. The findings of this research provide a direct quantitative measure of the relationship between superior data quality capability and market outperformance, and thus prove that data curation methodology is an essential component of the architectural foundation of developing a reliable interactive and mobile computing system.

Keywords: Competitive Strategy; Data Curation; Edge Case Analysis; Predictive Analytics; Pervasive Computing; Resource-Based View (RBV); Player Lifetime Value (pLTV).

Abstract of Paper Accepted in ICAIC-2026

540

Quantum-Resilient Secure Framework for Agentic LLM Workflows in E-Commerce and FinTech Systems

Phaneendra Yerra, Bank Of America, USA, phaneendravkyerra@outlook.com

Naga Subrahmanyam Cherukapalle, Catalyst Brands, USA
nagacherukupalle@outlook.com

ABSTRACT

In e-commerce and FinTech platforms, agentic large language model (LLM) systems are becoming more common to automate the decision-making process, examine frauds, and communicate with customers. The systems are autonomous and work with the help of planning, using tools and constant feedback. This enhances efficiency; however, new security and governance risks are introduced. Meanwhile, with the fast development of quantum computing, classical cryptography leading to the practical penetration of AI systems and electronic transactions is under threat. This poses a two-fold problem to new business ventures. The given paper suggests and discusses a quantum-resistant secure architecture of agentic LLM processes. It is a framework that includes post-quantum cryptography, zero-trust access control, and risk governance at the level of agents, which are dynamic. The proposed system is compared with a baseline agentic system and these comparisons are done based on a quantitative experimental methodology on the basis of classical and quantum-inspired threat models. Measures to be undertaken through the use of large-scale simulations are metrics like attack success rate, probability of risk detection, system latency and completeness of the audit trace. It reveals the outcomes that the suggested framework brings about the reduction of successful attacks greatly, enhanced auditability, and stable performance, even in the case of simulated quantum threats. Though the framework incurs a moderate implementation cost, the security and control services are by far better than the cost. These conclusions prove the fact that quantum-resistant and properly governed agentic LLM systems not only are practical but also required to exist in the future of e-commerce and FinTech worlds.

Keywords: Agentic Large Language Models, Quantum Threat Modeling, Cryptographic Auditability, Monte Carlo Simulation, Enterprise AI Security

Abstract of Paper Accepted in ICAIC-2026

541

An Autonomous Governance Framework for Generative AI: Real-Time PII Redaction and Compliance in LLM-Driven Data Pipelines

Vatsal Mavani
CVS Health Inc, USA

vatsalkishorbhai.mavani@gmail.com

ABSTRACT

As Large Language Models (LLMs) are rapidly being implemented into high-stakes applications in order to provide immediate decision-making capabilities, a substantial gap in governance exists; this is because LLM's probabilistic nature cannot be governed under existing deterministic frameworks that govern Personally Identifiable Information (PII) or Protected Health Information (PHI), such as the General Data Protection Regulation (GDPR) and/or the Health Insurance Portability and Accountability Act (HIPAA); therefore, an Autonomous AI Governance Framework (AAGF) has been developed to proactively inject real-time PII and PHI redactions at the high-throughput input/output interfaces of all LLM data pipelines in the form of a Context-Aware Output Sanitation Layer (OSL). The scalability and data residency compliance of the AAGF is based upon the development of multi-cloud serverless architectures utilizing Google Cloud Platform and Amazon Web Services in order to ensure compliance. As a result, the AAGF is able to serve as a determinate compliance gate that intervenes between an LLM's output and its end-user(s) in order to sanitize LLMs' output prior to delivery in order to mitigate potential data leakage risks. Furthermore, we have quantitatively evaluated both the detection efficacy of the combined use of structured metadata extraction from LLM function calls versus traditional techniques, as well as the "compliance tax" (i.e., latency overhead) associated with implementing the real-time governance functions of the AAGF, which validates our framework for implementation in any interactive and pervasive computing application where an unrepudiably auditable log is required for regulatory purposes.

Keywords: Generative AI Governance, PII Redaction, LLM Compliance, HIPAA, GDPR, Serverless Architecture, Output Sanitation, Real-Time Systems.

Abstract of Paper Accepted in ICAIC-2026

542

Quantifying Risk Reduction: AI-Driven Model for Automating Access Certification and Segregation of Duties (SOD) Controls

Sunnykumar Kamani
SoftSages Technology, USA

ABSTRACT

The traditional Access certification and SOD controls are an integral part of compliance yet they hardly ever possess quantitative metrics to reflect financial ROI. The QRR Model and its basic metric, ARI — Access Risk Index, attempt to initiate a conversation towards bridging the existing huge gap between qualitative compliance and measurable security value that is currently associated with these areas. The model uses AI and ML which will convert from manual access reviews campaigns to fully automated continuous governance control. Access Risk Index assigns risk score 0 or 1 to users by calculating potential risk and threat with entitlement criticality, user role, activities in system, amount of SOD violation and behavioral anomalies. This paper describes an API based automation framework which drives human intervention based on the result evolutions. This framework improves overall enterprise security posture and based on six-month comparison it has shown a significant drop in average Access Risk Index (ARI) to 89% and 96% fall in Toxic Combination Density which could have created high security risk. Based on results it proves that this automation framework is effective in detecting and remediating security risk.

Keywords: Identity Artificial Intelligence (AI), Identity and Access Management (IAM), Machine Learning (ML), Segregation of Duties (SOD), Quantifying Risk Reduction (QRR)

Abstract of Paper Accepted in ICAIC-2026

543

Codebase Aware Generative Agents for the SDLC: Automating Documentation, Dependency Analysis and Test Generation

Vatsal Mavani
CVS Health Inc, USA

vatsalkishorbhai.mavani@gmail.com

ABSTRACT

Autonomous Large Language Model agents are increasingly being used for complex software development tasks. One main challenge they face is a lack of sufficient codebase context, particularly in large, ever-evolving systems. Addressing this gap, researchers present CA-SAF (Codebase-Aware Self-Adapting Framework), describing the multi-agent system designed to eliminate architectural blindness with active, precise codebase information. Central to this contribution is ZSDM (Zero-Shot Dependency Mapping), an AST-based hybrid approach bound with LLM structural extraction constrained to generate all entities and relationships in the codebase as a queryable Knowledge Graph. ZSDM works on highly consolidated structured summaries of codes that achieve approximately 80% compression through hierarchical summarization levels suitable for contextual efficiency.We tested CA-SAF on the best Foundation Models—GPT-4o, LLaMA 3.1, and Claude 3 Opus—the context window of Claude 3 Opus (200,000 tokens) beats them by a slight margin in accuracy for complex non-local mapping tasks over their 128,000-token capacity. This was implemented in an operational six-month longitudinal study that measured a significant decrease in Cycle Time and Lead Time for maintenance tasks as a result of using CA-SAF. The framework also leveraged its codebase awareness to successfully mitigate the documented trade-off between increased velocity and high-impact technical debt introduction by controlling the density of BLOCKER/CRITICAL bugs injected with high-reasoning LLM code generation.

Keywords: Generative AI, Autonomous Agents, Software Development Life Cycle (SDLC), Knowledge Graph, Retrieval-Augmented Generation (RAG), Zero-Shot Dependency Mapping, Technical Debt.

Abstract of Paper Accepted in ICAIC-2026

546

Machine Learning Integration in Loan Decision Systems: Enhancing Kafka Based Workflows with Predictive Analytics

Jigar Solanki, INCEDO INC, USA, jigarmahendrabhaisolanki@gmail.com

ABSTRACT

Financial technology (FinTech) requires low-latency, high-throughput systems to make decisions for mission-critical applications such as dynamic credit scoring and fraud detection. While underlying streaming infrastructure for Event-Driven Architectures (EDAs) that utilizes Apache Kafka, is capable of handling throughputs of (up to 1.2 million events per second) at a low base latency, (e.g., 5 ms at the 99th percentile). A severe architectural challenge is raised by more accurate Machine Learning (ML) models compared with traditional statistical methods. Such models are typically trained in polyglot environments—(e.g., in Python), require remote inference via REST or gRPC microservices. A very large network and serialization overhead are thus introduced, leading to bloated end-to-end transaction latencies that may often exceed the allowable threshold of 50 ms. The polyglot bottleneck may be partially addressed by deploying models natively inside the Kafka Streams Java Virtual Machine (JVM), as discussed as part of a hybrid architecture herein. We use the Open Neural Network Exchange (ONNX) Runtime Java API, which allows safe inference execution within the process. This method lowers communication between processes and network delays, changing model run time from milliseconds to microseconds, hence making sure that total credit decision delay always meets tough financial Service Level Agreements(SLAs). The setup makes use of Kafka Streams Processor API so as to get detailed control and performance checking, thus showing a very strong, fast plus ready for tomorrow's highly automated lending.

Keywords: Apache Kafka; Kafka Streams; Real-Time Credit Scoring; Machine Learning Inference; Low Latency; JVM-Native Deployment; ONNX Runtime

Abstract of Paper Accepted in ICAIC-2026

548

Evaluating Modern Software Design Trends for Efficient and Maintainable Application Development

Ashish Pokhrel
Vivint, USA
ashishbdm90@gmail.com

ABSTRACT

Contemporary software development relies on the adequate choice of architectural paradigms that may guarantee performance, scalability, and maintainability in diverse environments. This paper is a comparison of four modern software design paradigms (Functional, Object-Oriented (OOP), Microservices, and React-based) involving a controlled setting and utilizing Node.js. The metrics were collected using automated benchmarking tools (e.g., Lighthouse and Puppeteer) to compare Lines of Code (LOC), Performance Score, First Contentful Paint (FCP), and Throughput quantitatively, in terms of measurability, trade-offs, performance, and maintainability. The React-based architecture delivered the highest run-time performance with FCP of 810 MS and was measurably throughput at 230 requests/second which would be responsive to existing web systems in general use. The Functional architecture offered the highest code readability (137 LOC), and OOP offered the greatest modularity and maintainability. Microservices are the most scalable, however, during load testing, they consumed the largest system resources (CPU at 70%; memory at 66%). This comparative method applies the principles of design science to define an executable strategy that guides developers and researchers in making informed architectural decisions based on data. This endeavor is helping to be able to maintain high-performance, high-quality, and maintainable software-engineering practice in current application development.

Keywords: Design Science, Software Architecture, Maintainability, Scalability, Benchmarking, Microservices, React Framework, Performance Evaluation.

Abstract of Paper Accepted in ICAIC-2026

549

dataRLsec: Safety, Security, and Reliability With Robust Offline Reinforcement Learning for DPAs

Shriram KS, Naresh Kshetri

Rochester Institute of Technology, USA
kshetrinaresh@gmail.com

ABSTRACT

Data poisoning attacks (DPAs) are becoming popular as artificial intelligence (AI) algorithms, machine learning (ML) algorithms, and deep learning (DL) algorithms in this artificial intelligence (AI) era. Hackers and penetration testers are excessively injecting malicious contents in the training data (and in testing data too) that leads to false results that are very hard to inspect and predict. We have analyzed several recent technologies used (from deep reinforcement learning to federated learning) for the DPAs and their safety, security, & countermeasures. The problem setup along with the problem estimation is shown in the MuJoCo environment with performance of HalfCheetah before the dataset is poisoned and after the dataset is poisoned. We have analyzed several risks associated with the DPAs and falsification in medical data from popular poisoning data attacks to some popular data defenses. We have proposed robust offline reinforcement learning (Offline RL) for the safety and reliability with weighted hash verification along with density-ratio weighted behavioral cloning (DWBC) algorithm. The four stages of the proposed algorithm (as the Stage 0, the Stage 1, the Stage 2, and the Stage 3) are described with respect to offline RL, safety, and security for DPAs. The conclusion and future scope are provided with the intent to combine DWBC with other data defense strategies to counter and protect future contamination cyberattacks.

Keywords: cyberattacks, data poisoning attacks, hash verification, reinforcement learning, safety, security

Abstract of Paper Accepted in ICAIC-2026

550

Zero-Shot Tokenizer Transfer for Targeted Attacks on Retrieval-Augmented Generation

Mark Spanier, Edward French, Samyam Aryal, Aman Singh, Komal More, Joe Hammond
Dakota State University, USA
mark.spanier@dsu.edu

ABSTRACT

Retrieval-Augmented Generation (RAG) is widely used to improve the reliability of Large Language Models (LLMs) by grounding responses in external documents. However, RAG systems that ingest content from internet sources remain vulnerable to prompt injection attacks embedded within retrieved documents. We introduce ZBEAST, an adaptation of Beam Search-Based Adversarial Attacks (BEAST), designed to generate malicious documents specifically for RAG settings.

We evaluate ZBEAST using an embedding-based heuristic built on BAAI-BGE-base-v1.5 and compare it against the state-of-the-art BEAST approach, which relies on an instruction-tuned LLM during search. ZBEAST replaces the generator's tokenizer with that of the embedding model via Zero-Shot Tokenizer Transfer (ZeTT), enabling more effective adversarial document optimization. Our results show that ZBEAST scales reliably with beam size, while the baseline approach exhibits diminishing returns. We further demonstrate that ZBEAST-generated documents evade common filtering mechanisms used in RAG pipelines. These findings expose a practical and underexplored vulnerability in RAG systems and suggest that existing defenses may be insufficient against optimized adversarial document generation.

Keywords: Retrieval-Augmented Generation (RAG), Large Language Models (LLMs), Prompt injection attacks, Adversarial document generation, Beam search, LLM security

Abstract of Paper Accepted in ICAIC-2026

552

Predictive Green FinOps: Joint Optimization of Cost, Carbon, and Reliability in AI-Intensive Clouds

Venkata Thej Deep Jakkaraju, Jayaprakasan V

ABSTRACT

The exponential proliferation of Artificial Intelligence (AI) workloads, particularly Large Language Models (LLMs) and generative AI systems, has precipitated a critical operational trilemma in modern cloud computing: the simultaneous minimization of financial expenditure and carbon emissions while maintaining strict Service Level Objectives (SLOs). Traditional Financial Operations (FinOps) frameworks, which prioritize unit economic efficiency, often inadvertently exacerbate carbon footprints by incentivizing the usage of low-cost but carbon-intensive cloud regions. Conversely, emerging Green Operations (GreenOps) initiatives can incur prohibitive costs or introduce unacceptable latency penalties, creating friction with business objectives. This paper introduces Predictive Green FinOps, a comprehensive, formal methodology for the joint optimization of cost, carbon, and reliability in AI-intensive cloud environments. By integrating advanced predictive analytics for spot instance interruption risks and grid-level carbon intensity forecasting, the framework enables dynamic, policy-driven workload placement and temporal shifting. A rigorous simulation using 2024 data across ten major cloud regions demonstrates that this approach can achieve a 28% reduction in carbon emissions (tCO_2eq) and a 19% reduction in training costs compared to baseline schedulers, while maintaining reliability scores above 99.9%. Furthermore, the research addresses the "hidden costs" of reliability failures, providing a mathematical basis for evaluating the trade-off between checkpointing overhead and spot market volatility. These findings provide a robust foundation for enterprise-grade, sustainable AI infrastructure strategies in the 2025–2035 era, ensuring compliance with emerging regulations such as the EU Corporate Sustainability Reporting Directive (CSRD) and SEC climate disclosure rules.

Keywords: Predictive FinOps, GreenOps, carbon-aware scheduling, AI workload placement, multi-objective optimization, sustainable cloud computing, spot instance reliability, cost-carbon trade-off, EU CSRD, SEC climate disclosure.

Abstract of Paper Accepted in ICAIC-2026

554

Cybersecurity Architecture for Cloud Database Infrastructure- A Case Study of Oracle Autonomous Database on Microsoft Azure

Chaitanya Kulkarni, Oracle America Inc, USA, ckulkarni@ieee.org

ABSTRACT

Enterprise cloud adoption continues to accelerate, yet securing database infrastructure across public cloud platforms remains a persistent challenge for organizations handling sensitive data. This paper presents a comprehensive security framework for Oracle Autonomous Database deployed on Microsoft Azure, addressing the complete lifecycle of cloud database security from network isolation through compliance automation. Our five-layer architecture integrates Oracle's database security capabilities with Azure's native security services, creating a defense-in-depth approach validated through production deployments. The framework encompasses network security using private endpoints and Virtual Network isolation, federated identity management via Azure Entra ID, customer-managed encryption with Azure Key Vault integration, database-level controls through Oracle Database Vault and Data Safe, and unified compliance monitoring across 25+ regulatory frameworks. Production implementations demonstrate 99.99% uptime while maintaining automated security patching, sub-5-minute threat detection response, and zero-downtime key rotation. We provide detailed implementation patterns for Azure Key Vault transparent data encryption, OAuth2 authentication flows, and DevSecOps automation templates. While focused on Azure deployment, the architectural principles establish foundational patterns applicable to broader multicloud database strategies. This work contributes practical security guidance for enterprise architects deploying mission-critical database workloads in public cloud environments.

Keywords: Cloud Database Security, Oracle Autonomous Database, Azure Security, Encryption Key Management, Federated Identity, Zero-Trust Architecture, Database Vault, Compliance Automation

Abstract of Paper Accepted in ICAIC-2026

555

HopeBot: An AI-Powered Mental Health Chatbot Built to Support College Students

Prasanthi Sreekumari
University of Louisiana at Monroe, USA
sreekumari@ulm.edu

ABSTRACT

Mental health strongly affects college students' well-being and academic performance. As social pressure and campus life become more stressful, mental health problems are increasing, and there are still not many effective ways to assess and support students. In this paper, we present HopeBot, an AI-powered mental health chatbot that gives students a private space to share their feelings and receive supportive guidance. HopeBot uses Amazon Lex to understand user messages, Amazon Bedrock to generate responses, and Amazon S3 to store mental health resources. The main aim is to make emotional support, easy to access, especially during difficult times. Based on system testing and user feedback, HopeBot gives accurate, helpful, and emotionally sensitive responses, with 91.4% accuracy, an average response time of 850 ms, and a user satisfaction score of 4.6 out of 5. These results demonstrates that HopeBot can be a useful tool for supporting student mental health.

Keywords: Conversational AI, Mental Health Support, Chatbot Systems, Natural Language Processing (NLP), HumanComputer Interaction (HCI)

Abstract of Paper Accepted in ICAIC-2026

556

Breaking the Monolith: A Data-First Strategy for Enterprise Microservices Migration with Measured Improvements in Scalability and Resilience

Venkaiah Chirumavilla, DIGITAL SCRIPTS INC, USA

ABSTRACT

The modernization of large-scale enterprise applications in the .NET ecosystem has traditionally had the drag effect from monolithic data layers. Even when there's agility enabled at the application tier through refactoring to microservices, without an equally structured way to decompose the database, distributed monoliths often result and perpetuate legacy system coupling across a network landscape that's even more complex. This paper proposes a detailed migration approach that puts decoupling and transforming the persistence layer up front before any architectural evolution can take place. Drawing upon lessons learned from two benchmark modernization initiatives—TriZetto Facets migration from Sybase ASE to Microsoft SQL Server and Einstein 360 platform microservices transformation—this paper will be defining all those technical and strategic patterns critical for high-volume data migration. The finding optimizes SSIS and SSMA for achieving sub-second query response time with near-zero downtime possible. Methodologically rigorous measurements prove that after migration, time to extract claims reduced by more than 50% ($p < 0.001$) and time for recovery, RTO, was also cut in half. Only a data-centric approach can fundamentally enable scalability to the highest level together with operational resilience in the new modern cloud-native environment.

Keywords: Microservices Architecture; .NET Core Modernization; Database Decomposition; Sybase ASE; Microsoft SQL Server; Zero-Downtime Migration; ETL Optimization

Abstract of Paper Accepted in ICAIC-2026

561

Strategic Unification: How Blazor is Redefining Enterprise Web & AI Stack

Venkaiah Chirumavilla, DIGITAL SCRIPTS INC, USA

ABSTRACT

The enterprise web development landscape is currently moving away from traditional polyglot stacks and adopting a unified, high-performance environment. This paper will appraise the future sustainability and strategic benefits of Microsoft's Blazor framework for enterprise.NET investments through comparing architectural patterns, TCO, and actual implementations like Texas Air Notification System (TANS) and Property Damage Assessment (PDA) platforms. In the long run, Blazor has more staying power than competing JavaScript-based Single Page Application (SPA) frameworks such as React or Angular. It also examines developer productivity enabled by C# skill unification on par with performance using WebAssembly (WASM), backed by stability assurance through annual Long-Term Support (LTS) roadmaps from Microsoft. If already inside the Microsoft ecosystem, it minimizes cognitive load while providing technical debt support.

Keywords: Blazor,.NET 10, WebAssembly, Enterprise Software Sustainability, Total Cost of Ownership, Single Page Application, C# Skill Consolidation, Software Architecture.

Abstract of Paper Accepted in ICAIC-2026

562

Software Engineering Challenges in the Deployment of Generative AI Models at Scale

Venkata nagendra Satyam, Devisharan Mishra,
Amazon Web Services
Braja Gopal Mahapatra
Mastercard Technologies Inc

ABSTRACT

One of the most significant advancements in the field of generative AI, AI software engineering has recently gained incredible progress. Deep generative modeling is a fast-evolving field in recent years. Generative AI models have proved to be capable of doing great things with natural language generation, although their use on a large scale poses serious software engineering problems, such as latency, reproducibility, complexity of integration, and scalability. This paper creates a systematic deployment pipeline based on the Kaggle Fitness Exercises dataset to solve these problems. Preprocessing entailed data cleaning, merging instructions and non-data normalization and tokenization to generate standardized inputs. A GPT-2 model was fine-tuned and embedded in a FastAPI application and containerized using Docker so it can move around and run in any cloud environment to facilitate real-time API interaction. The experimental outcomes were effective low-training (1.58) and validation loss (0.14), powerful predictive (3.16) and stable text generation patterns. Nevertheless, the deployment showed inference latency (2.55 s) and low throughput (0.39 requests/sec), which demonstrates the necessity to continue its optimization. A comparative analysis with reinforcement learning and edge-cloud methods has shown that the proposed pipeline has better reproducibility, deployment ease and scalability. The paper highlights the importance of effective software engineering practice to surmount technical challenges to the massive implementation of generative AI systems.

Keywords: Generative AI, software engineering, model deployment, docker, FastAPI, scalability, reproducibility.

Abstract of Paper Accepted in ICAIC-2026

566

Accelerating AI Maturity: A Framework for Reducing Development Lifecycles and Increasing Predictive Accuracy

Gopi Krishna Pamula
Mitchell Martin Inc, USA
gopipamula@gmail.com

ABSTRACT

A Model Manager Framework is proposed as a structured approach toward streamlining the total machine learning model lifecycle that will directly attack challenges at the heart of AI maturity within enterprise environments. Beyond simply codifying and automating best MLOps practices, including advanced versioning and continuous evaluation integrated with AI governance, it will elevate workflow from a mere patchwork of manual operations into an extensible scientific process. The current study applies these claims in practice by presenting evidence on how implementation reduced average model development cycle time by 30% and increased predictive accuracy by 23%. It thus sets up a strategic discussion on the interplay of technical efficiency, ethical governance, and business value—stressing strong dimensions that move implementation of such a framework away from being conceived as merely a plumbing exercise into constituting strategically sound business practice to enable responsible, scalable, high-performance AI.

Keywords: Model Lifecycle Management (MLOps), AI Governance, Model Versioning, Reproducibility, Predictive Accuracy.

Abstract of Paper Accepted in ICAIC-2026

567

Scalable Data Governance Through Engineering-Driven Quality and Consistency Controls

Gopi Krishna Pamula
Mitchell Martin Inc, USA
gopipamula@gmail.com

ABSTRACT

Conventional, centrally-controlled data governance is mostly seen as a bureaucratic compliance exercise and thus suffers from minimal user adoption and high failure rates, predicted to impact up to 80% of organizations that try to scale digitally by 2025. Legacy DG thereby becomes highly inappropriate due to the high velocity and complexity of modern ecosystems—particularly those propelled by ubiquitous, mobile computing. This paper will describe Data Governance as a Product (DGaaP). It is fundamentally designed from the service delivery aspect, rather than on control objectives—focusing not on mandates but on delivering value to users. DGaaP Governing derives empathetic product management (PM) practices which include user research, Agile development between the cycles of the Plan-Do-Study-Act (PDSA), continuous roadmap reprioritization support the governance policies that would maximize value return to internal users—data scientists and analysts. Its technical architecture is implemented as a Data Catalog that provides transparent interfaces and implements Policy as Code (PaC) through Continuous Integration/Continuous Deployment providing automatic real-time enforcement of consistency. A theoretical case study application validated this approach in which iterative, user-focused governing led to $\mathbf{90\%}$ reduction of data errors measured on accuracy and consistency dimensions. It can scale up progressively, giving room for highly effective DQ plus compliance by dissolving organizational friction into trust and replacing barriers with enablers via technology.

Keywords: Data Governance, Product Management, Agile, Data Quality, Policy as Code, Data Catalog, Pervasive Computing.

Abstract of Paper Accepted in ICAIC-2026

568

Anomaly Detection in Financial Payment Transactions Using Efficient Data-Driven Machine Learning Techniques

Sandeep Shivam, Tavant, USA, sandeep.shivam@ieee.org,
Tejas Patel, Amazon, USA, tejas.patel@ieee.org
Amit Padhy, University of Illinois Urbana-Champaign, a.padhy@ieee.org
Chaitanya Kulkarni, Independent Researcher, IEEE, ckulkarni@ieee.org
Chandrashekhar Medicherla, Independent Researcher, IEEE,
chandrashekhar.medicherla@ieee.org
Vinay Soni, Independent Researcher, IEEE, USA, Vinay.soni@ieee.org

ABSTRACT

Sophisticated fraud cases due to the growth of Digital Payment Systems occurred within an environment where thousands of new forms of Digital Payments were developed every month; thus, there must be an increasing demand for effective reasons for detecting Financial Transaction anomalies. This research will focus on developing a framework using ML to detect fraudulent payment behaviors within a very imbalanced Credit Card dataset. The proposed approach will incorporate several key processes of data mining including Exploratory Data Analysis, Duplicates Removal, Robust Scaling (through Robust Scaler), SMOTE (Synthetic Minority Over Sampling Technique), and Three Models (XGBoost, Random Forest and Hybrid XGBoost with Isolation Forest anomaly scoring feature). From the results of this experiment, it shows that the proposed Hybrid XGBoost model had the highest Accuracy (99.95%), Precision (97.22%) and F1-Score (83.83%) of all models tested, thus outperforming all base models. The originality of the study lies in the integration of Unsupervised Anomaly Scores with a Supervised Classifier, which has been used to produce a system that is more sensitive to the detection of slight fraudulent patterns and at the same time reduces the rate of false positives. This validates the effectiveness and scalability of the proposed model to practical Financial Fraud Detection in the Financial Industry.

Keywords: Machine Learning, digital payment systems, financial fraud detection, robust scaler, hybrid model.

Abstract of Paper Accepted in ICAIC-2026

569

Securing the LLM Supply Chain: Analyzing Threats and Mitigation Strategies

Atish Kumar Dash
ADVANCE Solutions Corp., USA

atish.dash.7@gmail.com

ABSTRACT

Over the past several years, with the rise of ChatGPT, there has been a paradigm shift in the way large language models (LLMs) have accelerated the growth of technology, captured public imagination, and propelled the promise of attaining artificial general intelligence (AGI). However, such advancements have also simultaneously brought with them techniques in which such models can be altered and exploited by malicious actors for harm. With both large organizations, from finance to healthcare, and nonprofits and government organizations adopting such technologies at an unprecedented rate, it becomes imperative for them to secure their LLM pipelines end-to-end. This study delves into how organizations can take control of the situation and secure their LLM pipelines. It requires them to focus on the entire LLM lifecycle from data acquisition and model training to deployment and post-production monitoring. The text covers methods for protecting the LLM pipelines through techniques rooted in zero trust principles and effective data governance to make them more robust against any sort of adversarial attacks. It also focuses on deploying strategies such as secure-by-design and model verification to ensure the integrity of the LLM supply chain. The text also covers ways in which organizations can protect their AI investments through implementing proper access controls and adopting secure MLOps practices. Audit and compliance oversight of LLM models have also been thoroughly analyzed. The document further discusses research on the resultant outcomes of such security breaches, and the future implications specifically from a societal impact perspective.

Keywords: LLM Pipeline Security, Generative AI, Zero-Trust Principles, Secure MLOps, Adversarial Attacks

Abstract of Paper Accepted in ICAIC-2026

570

AI-Assisted Zero Trust Architecture for Continuous Risk Assessment of Programmable Logic Controllers in Food Processing Infrastructure

Sakshyam Ghimire
Minnesota State University, Mankato, USA sakshyam.ghimire@mnsu.edu

ABSTRACT

Food processing infrastructure relies heavily on programmable logic controllers (PLCs) to manage physical processes that support safety and continuous operation. As these environments become increasingly interconnected through the integration of information technology and operational technology, cyber risks continue to grow. Traditional perimeter-based security approaches and static Zero Trust implementations often struggle to address continuously changing controller behavior in operational environments. In addition, many PLCs lack built-in security capabilities that provide the device-level visibility needed for ongoing trust evaluation. This paper presents a conceptual architecture that combines artificial intelligence (AI) with Zero Trust principles to support continuous monitoring and risk assessment of PLCs in food processing environments. The proposed approach uses passive behavioral observation and network-level monitoring to track controller activity over time and enable adaptive trust decisions based on observed risk. By improving visibility without requiring changes to controller hardware or control logic, the architecture helps preserve safety and availability. Future work will focus on validating the proposed architecture in simulated operational technology environments and assessing its impact on detection accuracy, safety, and system performance.

Keywords: Programmable Logic Controllers (PLC), Zero Trust Architecture (ZTA), Artificial Intelligence (AI), Operational Technology (OT), Industrial Control Systems (ICS), Critical Infrastructure Security

Abstract of Paper Accepted in ICAIC-2026

571

AI-Guided Adaptive Compression for Secure and Efficient Web Resource Delivery

Kateryna Babii, Eleks Inc., USA, katrusyabb@gmail.com>
Oleh Polishchuk, Independent researcher, USA oleg.polischuks@gmail.com

ABSTRACT

This paper presents an AI-guided adaptive compression framework that dynamically selects transport-level compression strategies for web resource delivery. We instantiate the decision layer with a lightweight tree-based classifier (Random Forest) that predicts the compression choice per resource from runtime features such as Content-Type, payload size bucket, observed RTT, and concurrency context. The policy minimizes a cost function that balances expected transfer-time reduction against CPU overhead. A controlled experimental setup models representative web workloads composed of heterogeneous static assets delivered under parallel and sequential loading strategies.Results show that adaptive compression outperforms static baselines across the evaluated workload and network profiles. While individual asset compression yields moderate gains, cumulative effects across concurrent loading achieve aggregate transfer reductions exceeding ten percent, leading to measurable improvements in First Contentful Paint and Largest Contentful Paint. Beyond performance, adaptive compression reduces network exposure duration and can lower the effectiveness of timing-based traffic analysis by shortening and smoothing observable transfer windows. These findings position adaptive compression as a practical, under-explored control surface for AI-enabled, security-aware web infrastructure.

Keywords: AI-guided optimization, adaptive compression, web performance, cybersecurity resilience, content delivery networks, First Contentful Paint, Largest Contentful Paint

Abstract of Paper Accepted in ICAIC-2026

572

AI-Based Cross-Layer Vulnerability Management for Cloud-Native Systems

Oleh Polishchuk
Sofo-Group, USA
oleg.polischuks@gmail.com

ABSTRACT

Cloud-native systems introduce complex and dynamic attack surfaces driven by microservice architectures, identity-centric access models, and continuously evolving runtime behavior. Traditional vulnerability management approaches rely on static severity scoring and siloed security signals, resulting in excessive alert noise and limited insight into real-world exploitability. This paper presents a novel AI-driven cross-layer vulnerability management framework that models exploitation risk as an evolving behavioral process rather than a static property of individual vulnerabilities. The proposed approach represents cloud-native environments as temporal cross-layer graphs and applies self-supervised sequence learning, graph-based risk propagation, and anomaly detection to identify multi-stage exploitation paths spanning application, identity, runtime, and endpoint layers. Unlike conventional prioritization techniques, the framework learns how risk propagates across layers over time and adapts to previously unseen attack strategies without reliance on labeled exploit data. A simulated Kubernetes-based case study demonstrates substantial alert noise reduction while preserving exploitation-relevant signals, outperforming severity-based and rule-based correlation approaches.

Keywords: AI-based Vulnerability Management, Cybersecurity, Temporal Graph Learning, Cloud-Native Systems, Self Supervised Anomaly Detection

Abstract of Paper Accepted in ICAIC-2026

573

The Adaptive AI SOC Agent – Moving Beyond Linear Playbooks with Cognitive Reasoning

Valentin Chichurov, Adelia Ibragimova, Anton Tararykov

ABSTRACT

This paper presents an architecture for an autonomous Tier 1 security analyst designed to operate across heterogeneous SOC environments. Unlike static automation, the proposed framework utilizes a schema-agnostic abstraction layer to normalize diverse security logs into a unified investigation context. The system employs a Solver-Critic cognitive loop, where distinct agents propose and verify investigation traces against raw evidence to mitigate hallucination. This enables the autonomous triage of high-volume alerts without hardcoded integrations.

Keywords: Security Operations Center (SOC), Large Language Models (LLM), Autonomous Agents, Security Automation, Incident Triage, Solver, Critic Architecture, Hallucination Mitigation, Security Information and Event Management (SIEM).

Abstract of Paper Accepted in ICAIC-2026

574

BRUJO: Baseline-calibrated Risk from Unified Joint Observations for MITRE ATT\&CK-based Time Series Forecasting in Security Operations Centers

Leonardo Octavio Perez-De La Torre, Andrew Wilson, Marco Perez-Cisneros, Gabriel Sanchez-Perez, Aldo Hernandez-Suarez

ABSTRACT

Security Operations Center (SOC) teams face incident volumes that exceed analyst capacity. Time series analysis in cybersecurity frequently relies on aggregated counts and opaque scores, which limits customer-level action and weakens alignment with MITRE ATT&CK tactics. This work presents BRUJO, a baseline-calibrated risk profiling method designed for shift-level SOC operations. BRUJO builds hourly incident series per customer and tactic, forecasts one week ahead conditioned on the 168 hours of the week, and separates sustained and intermittent regimes. Continuous series are forecast with exponential smoothing models and intermittent series with Croston-type estimators. A calibration layer clips forecasts using hour-of-week statistics to produce stable and plausible baselines. Forecasts and realised incidents are then combined into hourly risk vectors and shift-aligned health indices along four dimensions: incident load, temporal stability, tactic dispersion, and forecast alignment. Evaluation on Microsoft Sentinel telemetry shows an average root mean squared error of 0.28 incidents per hour, an average mean absolute error of 0.18 incidents per hour, and an average determination coefficient close to 0.80 across the main tactics. Over the evaluated week, 66.1% of hours fall in excellent health, 22.6% in moderate health, and 11.3% in poor health. BRUJO provides interpretable signals for triage, capacity planning, and rule tuning by indicating which customers, tactics, and shift windows drive deterioration.

Keywords: cyber risk profiling, time series forecasting, MITRE ATT&CK, Security Operations Center, incident prioritisation, ETS, Croston, health index

Abstract of Paper Accepted in ICAIC-2026

578

AI-Driven Frontend Cybersecurity for Real-Time Phishing and Threat Detection in ReactJS and Swift Applications

Mykola Savenko, Docusign, mykola.savenko21@gmail.com

Ilia Sedoshkin, Wehead, i@wehead.com

ABSTRACT

The rapid adoption of modern client-side frameworks such as ReactJS for web applications and Swift for iOS has transformed the frontend layer into a critical cybersecurity boundary. Phishing interfaces, malicious scripts, and deceptive UI overlays increasingly operate within the client runtime, often bypassing backend-centric security controls. This paper proposes a unified AI-driven frontend cybersecurity framework that embeds lightweight machine learning models directly into ReactJS and Swift applications using TensorFlow.js and CoreML. The framework leverages publicly available phishing and malicious webpage datasets for model training and constructs a multi-modal frontend feature representation encompassing URL structure, DOM and UI composition, JavaScript execution patterns, and client-side network behavior. Experimental evaluation across web and mobile environments demonstrates that the proposed lightweight transformer-based model achieves up to 97.2% detection accuracy with minimal runtime overhead (<4% CPU and ~20 ms latency), enabling practical real-time deployment. The results highlight the effectiveness of frontend-integrated AI security as a proactive defense mechanism that detects phishing and anomalous behaviors at the moment of user interaction.

Keywords: Frontend Security, Phishing Detection, ReactJS Security, Swift iOS Security, Artificial Intelligence, Machine Learning, Client-Side Cybersecurity

Abstract of Paper Accepted in ICAIC-2026

579

Integrating Customer Data Platforms (CDPs) with Experimentation Tools: A Framework for Optimizing Customer Journey Revenue

Gopi Krishna Pamula
Mitchell Martin Inc, USA
gopipamula@gmail.com

ABSTRACT

Granular real-time personalization cannot be implemented to fine-tune complex customer journeys within highly pervasive digital environments due to data silos and some architectural flaws pertaining to client-side A/B testing methodologies. The paper, therefore, introduces a novel closed-loop architectural framework for both the technical and strategic integration of an enterprise Customer Data Platform (CDP), namely Adobe Real-Time CDP (RT-CDP) that runs on the Adobe Experience Platform (AEP), with a server-side experimentation engine. In particular, the framework describes critical implementation challenges surrounding low-latency audience activation by taking explicit advantage of the Edge Network architecture. It allows decisioning on the server to leverage unified, dynamic customer profile attributes (that is, affinity scores, loyalty status) for a highly targeted experiment assignment and delivery of personalization. It formalizes the inbound feedback loop by systematically ingesting attribution data and experiment results back into the CDP for continuously enriching profiles and further targeting model refinements. This architecture has been validated empirically under very high traffic that produced a 19% uplift in conversion rate from optimized customer journeys pivoting on personalized product discovery-exacting measurement validation of technical feasibility as well as revenue impact from high-fidelity, data-informed server-side experimentation.

Keywords: Customer Data Platform (CDP), A/B Testing, Adobe Experience Platform, Real-Time Personalization, Server-Side Experimentation, MarTech Stack.

Abstract of Paper Accepted in ICAIC-2026

580

Hybrid Transformer and XGBoost Model for Federated IoT Intrusion Detection

Madhu Siddharth Suthagar, Velu Natarajan, Harish Namasivayam Muthuswamy

ABSTRACT

The rapid expansion of Internet of Things (IoT) devices has increased both the scale and complexity of network traffic, making intrusion detection a critical but challenging task. Although Federated Learning (FL) provides privacy-preserving distributed training, existing FL-based intrusion detection systems typically rely on either deep learning or tabular models alone, leading to suboptimal performance on heterogeneous IoT network flows. Furthermore, high class imbalance, non-IID client distributions, and excessive communication overhead limit their real-world applicability. This paper proposes a hybrid FL architecture integrating a Transformer backbone enhanced with Low-Rank Adaptation (LoRA) adapters, an XGBoost tabular specialist, and a learnable fusion head. The system incorporates stratified data partitioning, class-weighted cross-entropy, and balanced oversampling to address dataset imbalance. Only LoRA adapter updates are exchanged with the server, significantly reducing communication costs while preserving privacy. Experiments on the CICIoT2023 dataset demonstrate that the fusion model achieves improved performance on minority attack classes and exhibits faster convergence compared to single-model FL baselines.

Keywords: Federated Learning, LoRA adapters, IoT intrusion detection, hybrid models, XGBoost, Transformers, anomaly detection, class imbalance.

Abstract of Paper Accepted in ICAIC-2026

581

Architecting Agentic AI Systems with Multimodal Reasoning for Scalable Visual Pattern Recognition

Linga Reddy Alva, Hari Shanker Reddy Resu
IT Spin Inc, USA alvalingareddy@gmail.com

ABSTRACT

Modern progress in agentic and multimodal AI, including ReAct, HuggingGPT, and MMReAct, show that large language models can coordinate vision tools by using planner executor loops. Nevertheless, all these frameworks are of ad hoc nature: they do not include a principled model of cost-conscious decision making, formal memory-verification, and reproducible architectures of largescale visual reasoning. In a bid to fill these gaps, we propose Agentic Multimodal Pattern Recognition (AMPR) a formal reasoning and planning system that combines hierarchical decomposition, probabilistic self-checking and dynamic costconscious inference with a common optimization problem. In contrast to earlier models, AMPR is a clear graphical reasoning as a constrained optimization problem to trade-off accuracy, latency and cost, and integrates episodic and semantic memory to promote instead of a single step of reasoning. We submit theoretical background as well as empirical performance across benchmarks of classification, detection, segmentation, and visual question answering. Findings demonstrate that AMPR has better accuracyefficiency tradeoffs, and is better behaved to distribution shifts with demonstrable reasoning consistency guarantees. AMPR defines a new standard of scalable, interpretable and resourceefficient visual intelligence by integrating formal algorithmic contributions and system-level validation.

Keywords: Agentic AI, Multimodal Reasoning, Visual Pattern Recognition, Cost-Aware Inference, Scalable AI Systems, Explainable Artificial Intelligence (XAI)

Abstract of Paper Accepted in ICAIC-2026

582

The Medallion Architecture in Practice: A Framework for Building Scalable and Governed Data Lakehouses on Microsoft Fabric

Vamshi Krishna Pamula
Artek Solutions Inc, USA
vamshikpamula@gmail.com

ABSTRACT

The present day we see an extensive amount of data produced by what is basically the universal use of digital systems and mobile computing which in turn put out large volumes of very varied data. What we are seeing is that companies today are putting into play and at the same time are trying to manage structured, semi structured and unstructured data flows in near real time also at the same time they are looking at issues of scalability, reliability and strong governance. Also what we are seeing is that the legacy Enterprise Data Warehouse (DWH) structures which are very much the norm at present are put to the test by these requirements. This paper reports on a practical full scale implementation of the Medallion Architecture which includes the Bronze, Silver and Gold layers we did in the Microsoft Fabric Lakehouse unified environment which we used Delta Lake as the base storage format. We are putting forth here an improvement over the past DWH based designs which did not do a great job structurally and also in terms of cost.

By we mean we broke the storage from the compute and we went with open and very efficient data formats. What we found is that we did see great improvements in data quality, in processing flexibility and in performance also we were able to put in place the AI and ML work flows within the same architecture.

Keywords: Data Lakehouse, Medallion Architecture, Microsoft Fabric, Unity Catalog, Data Governance, Attribute-Based Access Control, Pervasive Computing

Abstract of Paper Accepted in ICAIC-2026

583

RAGSec: Retrieval-Augmented Generation for Cybersecurity Threat Intelligence in Enterprise Networks

Shravya Bussari

HCLTech Inc, USA
shravyabussari@gmail.com

Gayathri Balakumar, Prateek Punj

ABSTRACT

This work presents RAGSec, a retrieval-augmented threat intelligence framework tailored to enterprise environments. RAGSec introduces (i) sensitivity-scoped ingestion for heterogeneous sources, (ii) retrieval-confidence thresholds mapped to incident severity, (iii) abstention and evidence-citation templates aligned to SOC governance, and (iv) a SOC Productivity Score (SPS) to measure impact beyond factual accuracy [6]. Evaluation on public reports and synthetic enterprise incidents compares RAGSec with LLM-only and search-assisted baselines across factuality, relevance, completeness, and analyst workload [6]. Experiments show RAGSec improves factual accuracy by up to 32%, decreases unsupported statements by 41% relative to LLM-only baselines [7], and reduces incident triage time by 27%. These results demonstrate that domain-governed RAG pipelines can enhance analyst decision-making while maintaining operational control and auditability [5]. The paper concludes with limitations, ethical considerations, and deployment guidance for adopting retrieval-enhanced GenAI within live SOC environments

Keywords: Cybersecurity, Threat Intelligence, RetrievalAugmented Generation, Large Language Models, Enterprise Networks, SOC Automation.

Abstract of Paper Accepted in ICAIC-2026

584

Prompting for LLM Security and RAG: A Survey from Zero-Shot to Automatic Prompt Optimization (APO) and Prompt-Injection Defenses

Ritesh Ruparel
CSG, USA
ritesh.ruparel@yahoo.com
Aravinda Jatavallabha, Long Health, aravindaraman04@gmail.com
Srinivasa Maringanti, T-Mobile, maringanti.srinivasa@outlook.com

ABSTRACT

Large Language Models (LLMs) are increasingly embedded in security-sensitive workflows such as incident triage, code review, threat hunting, and retrievalaugmented assistants. In these settings, prompting is not only a performance tool but also a security control surface: LLMs may consume untrusted content (tickets, logs, web pages, retrieved documents) that can adversarially manipulate outputs via prompt injection. This paper presents a structured, practitioner-oriented survey that unifies (i) core prompting methods (zero-shot, few-shot), (ii) reasoning-structured prompting (chains/trees/graphs of thought), (iii) robustness methods (self-consistency), (iv) efficiency workflows (skeleton-of-thought), and (v) automatic prompt optimization (APO) and learned prompting (directional stimulus prompting). We additionally synthesize cybersecurity-specific research on prompt injection, indirect prompt injection benchmarks, retrieval-augmented generation (RAG) poisoning, coordinated Prompt–RAG attacks, and interface-level defenses (structured queries). We provide a taxonomy, summary tables, and a decision guide for selecting prompting strategies under accuracy–cost–security constraints.

Keywords: Prompt engineering, prompt injection, retrieval-augmented generation (RAG), LLM security, automatic prompt optimization, survey

Abstract of Paper Accepted in ICAIC-2026

585

Integrating AI and Cloud to Advance Scalable, Secure, and Automated Information Management in Enterprises

Shiva Krishna Kodithyala, Bread Financial, USA
reachkodithyala@gmail.com
Hari Shanker Reddy Resu, IT Spin Inc, USA

ABSTRACT

The hyper-growth of data and the increased need of automated decision-making have made businesses turn to creating scalable, secure, and efficient ways of managing information. To overcome these issues in enterprise information management, this paper suggests and discusses a layered structure that is based on the combination of Artificial Intelligence and cloud technologies. The framework is based on four fundamental layers, namely AI, Cloud, Integration, Governance, and Data Management. The AI layer allows intelligent automation using machine learning, natural language processing, and robotic process automation, and the cloud layer provides elastic infrastructure to sustain high volumes of activity. The Integration & Governance layer applies secure data flows, compliance, and auditability, and the Data Management layer is used to be in control of the structured and unstructured information. The framework is designed in such a way that it can be adopted gradually according to the size of the organization and its level of maturity. Its practical usefulness was shown by a pilot deployment in a mid-size retail bank that saw the number of customer service response times drop by 40 percent, and the workload of the call center was reduced by 25 percent. These findings confirm the modular adoption, scaling, and compliance integration that the framework has the potential to offer, and it can be defined as a methodological contribution and a strategic map on which businesses in the ever-more data-driven world should be guided.

Keywords: Artificial Intelligence (AI), Cloud Computing, Enterprise Information Management (EIM), Scalability, Automation, Data Governance, Data Management

Abstract of Paper Accepted in ICAIC-2026

586

AI-Driven Multi-Cloud Strategies for Financial Services: Ensuring High Availability with AWS and Kubernetes

Jigar Solanki, INCEDO INC, USA
jigarmahendrabhaisolanki@gmail.com

ABSTRACT

The financial services sector faces unprecedented regulatory pressure regarding operational resilience, driving the adoption of multi-cloud strategies to mitigate single-vendor concentration risk and ensure continuous service availability. This paper analyzes an AI-driven architecture for mission-critical financial applications, focusing on an Active-Active deployment pattern utilizing Amazon Web Services (AWS) Elastic Kubernetes Service (EKS) and Kubernetes-native resilience mechanisms. Artificial intelligence and machine learning techniques are integrated to enable predictive workload scaling, anomaly detection, and intelligent traffic steering across cloud environments. The methodology combines AI-assisted global traffic management using AWS Route 53 to achieve sub-minute Recovery Time Objectives (RTOs) with intelligent data persistence and replication strategies necessary for near-zero Recovery Point Objectives (RPOs). Furthermore, the paper examines the role of AI-enhanced observability platforms, centralized Kubernetes management, and cloud-agnostic data replication technologies in achieving proactive fault mitigation and seamless cross-cloud state synchronization. Finally, the paper highlights the importance of AI-augmented Chaos Engineering for empirical validation against pervasive network failures, establishing trade-offs between throughput resilience in centralized cloud environments and response stability in latency-sensitive, cloud-edge financial deployments.

Keywords: AI-Driven Multi-Cloud, Kubernetes, AWS EKS, Machine Learning, Disaster Recovery (DR), RPO, RTO, FinTech, Operational Resilience, Chaos Engineering, Pervasive Computing.

Abstract of Paper Accepted in ICAIC-2026

587

GXMalDetect: A Hybrid GA–XGBoost Architecture for Malware Detection Using Static Image Features

Abdul Kareem Uddin Mohd, Md Habibur Rahman, Md Wahidur Rahman, Avdesh Mishra, Tarek Mahmud, Maleq Khan

ABSTRACT

Malware is growing in both quantity and complexity. Techniques such as polymorphism, packing, and obfuscation weaken traditional signature-based detection, motivating the development of efficient learning-based methods. This paper presents a hybrid malware classification architecture that combines Genetic Algorithm (GA)-based feature selection with an Extreme Gradient Boosting (XGBoost) classifier to achieve high accuracy with low computational overhead and improved interpretability. Using the Malimg benchmark (9,339 grayscale malware images across 25 families), we first extract a compact set of handcrafted static image features capturing statistical characteristics, texture/GLCM relationships, and morphological (shape) properties. The dataset is split using stratified 80:20 sampling to preserve class distributions, and GA optimization is conducted only on the training set to prevent data leakage. GA evaluates candidate feature subsets via cross-validated performance with a sparsity-aware fitness objective, converging to nine highly informative features spanning statistical (variance, energy, entropy, skewness, kurtosis), texture/GLCM (dissimilarity, homogeneity, correlation), and shape (compactness) descriptors. Trained on the selected subset, XGBoost achieves 99.30% test accuracy with 99.32% precision, 99.30% recall, 99.30% F1-score, and an MCC of 0.9918, while ROC and precision–recall analysis indicate near-perfect separability across classes. Overall, the proposed approach demonstrates that carefully selected handcrafted features combined with gradient boosting can match or surpass heavier deep-learning pipelines, offering a practical, fast, and interpretable solution for real-world malware classification.

Keywords: Malware classification, Malimg dataset, Genetic Algorithm, Feature selection, XGBoost, Static image features.

Abstract of Paper Accepted in ICAIC-2026

589

From Science to Production: A Systematic Framework for Operationalizing and Governing Machine Learning Model

Vamshi Krishna Pamula

Artek Solutions Inc, USA
vamshikpamula@gmail.com

ABSTRACT

One of the greatest barriers remains the operationalization of machine learning models-known as the last mile-because most prototypes never make an effective transition beyond sandbox development into a stable, scalable production environment. This paper presents an MLOps framework, to be implemented structurally through tools and centered around the architectural pattern known as Model Factory, enforcing standardization, reproducibility, and governance in the lifecycle management of ML. Important prescriptions for continuous automation are articulated herein, extending traditional CI/CD through CT and CM. Among key contributions are: definition of implementation by three levels of testing-code, data, model; interpretable model drift detection integrated into the system so as to facilitate operations on non-stationary data environments; sound governance structure guided by model versioning and lineage tracking capabilities ensuring auditable accountability. The framework is intended to apply to diverse libraries-e.g., H2O or Scikit-learn-in supporting real enterprise-class systems with interfaces made standard across systems supported via centralized metadata management.

Keywords: MLOps, Continuous Integration, Model Governance, Model Factory, Model Drift, Reproducibility, CI/CD for ML.

Abstract of Paper Accepted in ICAIC-2026

595

Optimizing AI-Driven Mobile Applications on iOS: A Comparative Analysis of Core ML and TensorFlow Lite for Cross-Platform Performance

Rutul Desai,
Kunai Inc, USA,
rutuldesai193@gmail.com

ABSTRACT

This paper will address the critical challenge of real-time, energy-efficient inference at the resource-constrained edge for mobile AI applications and compare Apple's native Core ML with TensorFlow Lite, a cross-platform baseline on iOS powered by Apple Silicon (A-series and M-series). This paper introduces a benchmark presenting inference latency of several critical AI tasks, such as Image Classification, Object Detection, and NLP, on different hardware tiers consuming energy. It benchmarks hybrid execution scheduling between Apple Neural Engine (ANE) and Metal Performance Shaders (MPS), capable of delivering energy efficiency for vision tasks—for instance, battery consumption reduced by nearly 40% in the MobileNetV2 case compared to alternative TFLite delegates—against current state-of-the-art frameworks available to the public. It further shows that TFLite dominates more complex non-vision workloads, such as BERT QA. Aside from benchmarking, a developer-oriented toolkit and best practices guide will be presented in conducting hardware-aware optimization (Core ML Tools' W8A8 quantization) and a dependable model conversion pipeline (TensorFlow → Core ML). Results position Core ML as the High-Efficiency Reference Architecture for Native iOS Deployment, but also reveal Model-Specific Performance Anomalies that must be considered in architectural design when deploying cross-platform in the future.

Keywords: On-Device ML, Core ML, TensorFlow Lite, Edge AI, Neural Engine, Quantization, Mobile Performance, Energy Ef iciency, Cross-Platform Deployment.

Abstract of Paper Accepted in ICAIC-2026

596

A Secure Biometric Authentication Framework for iOS with Cross-Platform Extensions: Addressing OS-Level Vulnerabilities and Enhancing Real-Time Protection

Rutul Desai,
Kunai Inc, USA,
rutuldesai193@gmail.com

ABSTRACT

The widespread use of mobile commerce, and sensitive data will require an effective multi-layered method to authenticate users beyond the simple trade-off between ease-of-use and security. On the one hand, user-friendliness is provided through the availability of on-board mobile devices of various types of biometric sensors; on the other hand, a single check of sensor level does not protect against sophisticated mobile attacks such as those utilizing memory corruption exploits (for example, mercenary spyware) and run time application tampering. In this paper we describe the development of the Secure Biometric Framework (SBF) which includes both a Hardware Roots of Trust (Secure Enclave) and Proactive Operating System (OS) integrity mechanisms (Memory Integrity Enforcement, MIE) and Continuous Runtime Assurance (App Attest, RASP), and thus provides multiple layers of defense against sophisticated mobile attacks. The architectural design of the SBF, and its ability to protect against OS-based attack vectors, is a significant improvement over documented high rates of misuse of the traditional Android Biometric Application Programming Interface (API). For example, according to recent research, it was determined that as many as 97.15% of the tested applications had serious configuration errors which could have resulted in potential loss of sensitive user information. Additionally, the SBF also utilizes the FIDO2 standard as the cross-platform authentication mechanism, which manifests itself as Passkeys. The SBF's foundational security architecture allows for the capability to provide phishing resistant authentication across all platforms, while maintaining full regulatory compliance, and ensuring full security assurance.

Keywords: Biometric Authentication, Secure Enclave, App Attest, Local Authentication, FIDO2, Memory Integrity Enforcement (MIE), Runtime Application Self-Protection (RASP), OS-Level Vulnerabilities.

Abstract of Paper Accepted in ICAIC-2026

602

Packing Induced Bias in Deep Learning Malware Classifiers: A Systematic Experimental Study

Jeevana Swaroop Kalapala, Lan Zhang
Northern Arizona University, USA
jk2396@nau.edu, Lan.Zhang@nau.edu

ABSTRACT

Packing is a highly used malware evasion technique that compresses, encrypts, or obfuscates executable content, significantly altering the structural characteristics that static malware detectors rely on. While prior work has studied the impact of packing on traditional feature based classifiers, limited attention has been given to understanding how modern deep learning based static detectors behave under realistic packing conditions. This paper investigates the robustness and generalization capabilities of a CNN based malware detector trained on image representations of Windows Portable Executable (PE) binaries by conducting two controlled experiments. The first evaluates how increasing exposure to packed benignware affects a model's ability to correctly distinguish packed malware from packed benign binaries. Results show initial improvement in true negative rates, followed by instability as packing artifacts dominate learned representations. The second experiment analyzes cross packer generalization and demonstrates strong packer dependency, where models perform well only on packers seen during training and collapse when confronted with unseen packers. Overall, our findings demonstrate that packing significantly undermines the reliability of static deep learning based malware detectors, highlighting the need for packing aware training strategies and more resilient detection models.

Keywords: Malware Detection, Software Packing, Deep Learning

Abstract of Paper Accepted in ICAIC-2026

603

Secure Execution of Post Inference Scripts in AI Driven Vector Search Pipelines

Khrystyna Terletska, Kateryna Babii

ABSTRACT

AI powered vector search pipelines are increasingly deployed in cloud native systems to enable semantic retrieval over large scale datasets. Many production deployments combine embedding models with lightweight scripting engines such as Lua to implement post inference tasks including scoring, filtering, and ranking. This paper identifies a code injection vulnerability class in such pipelines where user controlled query parameters are concatenated into dynamically executed Lua code. We analyze realistic exploitation paths that enable unauthorized code execution and potential exfiltration of sensitive vector embeddings and inference metadata. To mitigate the issue, we propose a secure execution framework based on parameterized script invocation, strict allowlist validation, Lua runtime hardening, and least privilege data access. We evaluate the approach in a representative vector search setting and show that it prevents injection attacks while introducing negligible latency overhead. The results highlight post inference execution layers as an underexamined attack surface in AI systems and provide actionable guidance for securing script based ranking components.

Keywords: AI Security, Vector Search, Injection, Lua Sandbox, Cloud Native Systems, Secure Execution

Abstract of Paper Accepted in ICAIC-2026

604

Triangulating Digital Forensics to Detect Friendly Fraud and Abuse at Scale: A Cloud-Native, Agentic AI Framework

Sabarinathan Govindaraj
PWC, USA gsabari_89@hotmail.com

Karthik Mahalingam, Velu Natarajan

ABSTRACT

The landscape of financial fraud is experiencing accelerated proliferation and sophisticated evolution, amplified by readily available AI and loosely coordinated criminal ecosystems. Among the most difficult categories to contain are friendly fraud and policy/returns abuse, which together cost U.S. retailers roughly $101B in 2023 and $103B in 2024 even as overall return rates moderated. These patterns increasingly leverage 'refund-as-a-service' networks and exploit customer-centric policies and post-purchase protections. This paper proposes a practical, cloud-native framework that triangulates historical digital footprints against suspected abusive transactions, implemented as a system of collaborating AI agents that synthesize device, network, behavioral, logistics, and order-resolution evidence into a case file suitable for both automated controls and human dispute analysts. The approach aligns with current card-network programs (e.g., Visa Compelling Evidence 3.0) and modern identity/authentication guidance (e.g., NIST SP 800-63B), while addressing governance and privacy-by-design obligations.

Keywords: Friendly Fraud, Return Abuse, AI Agents, Cloud-Native Systems, Digital Footprint, Visa CE3.0, Graph Analytics, Fraud Detection

Abstract of Paper Accepted in ICAIC-2026

605

HRIT: A Human-Readable Framework for Phishing URL Detection using Large Language Models

Mohammad Masum, Tanya Yadav
San Jose State University, USA
mohammad.masum@sjsu.edu

ABSTRACT

Phishing URL detection remains a persistent cybersecurity challenge due to rapidly evolving adversarial tactics and the inherent limitations of rule-based and traditional machine learning approaches. Although Large Language Models offer promising capabilities in contextual reasoning and natural-language explanation, their effective use for URL-based threat detection is constrained by a semantic mismatch between numeric security features and language-model reasoning. This paper introduces HRIT (Human-Readable Indicator Transformation), a lightweight and model-agnostic framework that transforms structured URL security features into semantically meaningful, human-readable indicators optimized for LLM inference. HRIT enables LLMs to reason about phishing intent using interpretable descriptors derived from statistically validated, dataset-driven, and industry-informed thresholds, without requiring model fine-tuning or retraining. We evaluate HRIT on a balanced phishing URL dataset containing 11,430 samples and compare its performance against a strong Random Forest baseline and multiple LLMs under zero-shot prompting. Experimental results show that HRIT improves detection sensitivity, improving recall and F1-score while maintaining low inference cost and latency. These findings demonstrate that effective phishing URL detection can be achieved through semantic feature transformation rather than complex prompting or model adaptation.

Keywords: Phishing URL Detection, Large Language Models, Cybersecurity, Feature Transformation, Human-Readable Indicators, Prompt Engineering

Abstract of Paper Accepted in ICAIC-2026

607

Toward Context-Aware Alert Classification in Security Operations Centers Using LLMs

Jonathan Roy,
Université du Québec à Chicoutimi, Canada
jonathan.roy@uqac.ca

ABSTRACT

This paper evaluates the ability of a Large Language Model (LLM), Meta LLaMA 3.3-70B-Instruct-Turbo, to support alert classification in a Tier-1 SOC setting by incorporating organisational context while decoupling context management from the alert classification logic encoded in the LLM prompt. Rather than estimating population-level performance, the objective is to isolate the effect of organisational context on alert classification under controlled conditions. The experimental environment is based on the open-source SIEM Wazuh monitoring Windows and Ubuntu hosts subjected to adversary-emulation scenarios generated with CALDERA, alongside legitimate system activity. A manually labelled corpus of alerts is used to evaluate two configurations: one relying solely on raw SIEM metadata and another in which alert classification is enriched with organisational context injected directly into the model prompt as a structured JSON file. Without contextual information, the model achieves an accuracy of 85.45% (precision = 80%, recall = 96.55%, F1 = 87.5%), with false positives arising from missing organisational context. When contextual information is provided, the model produces no false positives on the evaluated corpus and yields consistent classification outcomes across repeated executions.

Keywords: SOC, SIEM, security alert classification, large language models, contextual reasoning, Meta LLaMA 3.3, Wazuh, CALDERA, false positives, cybersecurity automation.

Abstract of Paper Accepted in ICAIC-2026

608

Secure and Compliant AI/ML-Based KYC: A Cybersecurity-Aware Architecture for Regulatory Identity Verification

Varun Pandey, IEEE, USA pandey.varun.087@gmail.com

ABSTRACT

Know Your Customer (KYC) regulation requires financial institutions to verify customer identities, assess risk, and maintain ongoing due diligence. Traditional, manual KYC processes are slow, error-prone, and increasingly insufficient in the face of sophisticated fraud, synthetic identities, and tightening regulatory expectations. Recent advances in artificial intelligence and machine learning enable more accurate and scalable identity verification, but they also expand the attack surface, introducing new adversarial and cybersecurity risks alongside longstanding concerns around robustness, fairness, explainability, and legal compliance. This paper proposes a secure, multi-modal, compliance-aware KYC framework that combines OCR-free document understanding, deep learning–based face verification, liveness detection, and risk-based orchestration with explicit cybersecurity controls aligned to NIST SP 800-63, NIST SP 800-207 (Zero Trust), and the NIST Cybersecurity Framework 2.0. The framework integrates adversarial robust models, secure MLOps practices, and access control policies guided by emerging guidelines such as OWASP machine learning security risks and MITRE ATLAS. It incorporates an explainable risk engine that exposes both KYC and cybersecurity signals to human reviewers, enabling shared visibility across fraud, compliance, and security teams. An evaluation design is outlined using public and synthetic datasets, with metrics spanning accuracy, latency, fraud detection uplift, security detection coverage, and fairness across demographic groups. In a prototype synthetic evaluation (10,000 onboarding cases), the proposed framework increased straight-through processing from 62% to 78% and improved injected high-risk session detection from 61% to 89% compared to a rule-based baseline.

Keywords: Know Your Customer (KYC), Regulatory Compliance, Document Understanding, Face Verification, Liveness Detection, Adversarial Robustness, Fairness, GDPR, NIST SP 800-63.

Abstract of Paper Accepted in ICAIC-2026

609

Client Side AI Based Intent Verification for Defending Against CSRF Attacks in Modern Web Applications

Kateryna Savenko, Inspyr, USA

Kateryna Babii, Eleks, Inc., USA

ABSTRACT

Cross Site Request Forgery remains a persistent threat to modern web applications, particularly in single page application architectures where complex client side execution flows weaken traditional server side security assumptions. Existing mitigation techniques such as synchronizer tokens, SameSite cookies, and origin validation are effective under ideal configurations but frequently fail in practice due to implementation inconsistencies, legacy endpoints, and emerging attack vectors involving embedded frames and automated request triggering This paper introduces a client side AI based intent verification framework that treats user intent as a first class security signal. The proposed approach operates entirely within the browser and evaluates lightweight interaction, visibility, execution context, and request level signals to infer whether a sensitive state changing request is causally linked to an explicit user action. An on device machine learning model classifies requests based on intent likelihood, enabling real time detection of forged requests without requiring backend modifications.

Keywords: client side security, csrf defense, intent verification, web application security, machine learning

Abstract of Paper Accepted in ICAIC-2026

610	**Automated SIEM Detection Rule Translation System** Adelia Ibragimova Epam Systems, USA adelina.30stm@inbox.ru **ABSTRACT** **Keywords:**

Abstract of Paper Accepted in ICAIC-2026

611

Digital Precognition: Teaching Transformers to Map Cyber Threats Before Analysts Can

Shane Waldrop, Erdogan Dogdu, Roya Choupani, Bill Mitchell

ABSTRACT

Mapping cyber threat intelligence (CTI) reports to the MITRE ATT&CK framework is tedious but necessary work for security analysts. We present a neural reranking system that automates this classification, achieving 87.7% precision at rank 1 (P@1) on 146 test queries spanning seven APT actors (95% CI: 81.4%–92.1%, Wilson score interval). The system uses a two-stage pipeline: BM25 retrieval generates candidate techniques, then a cross-encoder fine-tuned on CTI-to-ATT&CK pairs reranks them. Compared to BM25 alone (75.4% P@1), fine-tuning yields a 12.3 percentage point improvement. A pretrained MS-MARCO model without fine-tuning scores only 57.4%, actually degrading BM25's ranking—though we tested only one pretrained baseline, limiting generalizability of this finding. Per-actor analysis shows variation from 66.7%–100%, but small per-actor sample sizes (n=7 to n=41) preclude robust conclusions. Important caveats: 52% vocabulary overlap between train and test sets, all actors appear in training, and results are from a single training run. We release the complete pipeline for reproducibility.

Keywords: Cyber Threat Intelligence, MITRE ATT&CK, Neural Information Retrieval, Cross-Encoder, Domain Adaptation

Abstract of Paper Accepted in ICAIC-2026

613

AI-Powered Economic Digest and A Composite Index for National Financial Well-being: An AI-Driven Approach to Quantifying United States Financial Health with the USFHI

Bireswar Banerjee, Rajib Maitra

ABSTRACT

This paper delivers a thorough analysis of the United States' financial health by integrating critical economic indicators such as stock market performance, Consumer Price Index (CPI), unemployment rate, federal funds rate, and housing market trends. This AI-driven solution provides a snapshot of these key areas related to the present financial and economic condition. Along with the same, it provides a unified and actionable measure in terms of an index value, which is a weighted average of the indexes pertaining to the above topics. The paper introduces the US Financial Health Index (USFHI), a composite index designed to quantify and track the nation's financial well-being, which can be useful for investors in making key decisions. It uses a RAG-based AI solution to gather required information from available financial information libraries to generate the snapshot details as a daily digest. It also performs behind-the-scenes calculation to combine the available financial indexes to generate a weighted average score or index.

Keywords: US Financial Health Index, composite index, economic indicators, normalization, AI, financial analysis.

Abstract of Paper Accepted in ICAIC-2026

614

Unified AI Framework for Real-Time Customer Churn Prediction Using Behavioral and Event-Log Anomaly Signals

Jyoti Kunal Shah, Nixalkumar Patel, Darsana Usha Devi, Mahendra Babu Iragala, Ashok Kumar Kunkala, Biky Chouhan

ABSTRACT

Predicting churnage among customers is one of the most important challenges when working with a digital service platform; however, the majority of existing methods utilize mostly behavioral and transactional attributes and do not consider the influence of an anomalous event in a system on customer disengagement. Abnormal event-log patterns, including atypical access control and authentication failures, are likely to have a severe impact on user behavior and affection and frequently to churn. The paper introduces a single AI-driven real-time customer churn prediction framework combining behavioral and event-log generated anomaly signals.

The proposed system uses unsupervised anomaly detection of system and security logs to extract behavioral deviations caused by events and models the resulting anomaly scores with behavioral features to predict churn using supervised classification. Experimentally, it has been demonstrated that the performance is significantly better than the behavior-only baselines with a churn detect probability of about 0.45 at 20 days, 0.68 at 10 days and 0.92 at 5 days before churn as compared to 0.75 with a baseline model. The structure also enhances the isolation of risks, and the median churn score is greater (0.69) as compared to the traditional methodologies (0.50). These findings indicate that event-log anomaly signals are quality early warning signs of customer churn.

Keywords: Customer Churn Prediction, event-log analytics, , Machine Learning, Security-Aware Systems, Event-Log Analytics, Real-Time Churn Prediction, Anomaly Signals, Unified AI Framework

Abstract of Paper Accepted in ICAIC-2026

616

A Dual-task Prediction Model for Starlink Maritime Performance

Richard Li, Memorial High School, USA, lirichard448@gmail.com
Md Muntasir Hossain, Xingya Liu, Ruhai Wang
Lamar Univeristy, USA
xliu@lamar.edu

ABSTRACT

While LEO satellite systems have revolutionized broadband access for maritime industries, the dual challenges of high orbital mobility and dynamic weather lead to unpredictable signal instability. This prevents the maritime industry from effectively managing bandwidth. To address this issue, this work developed a short-term, dual-task prediction model to predict future downlink throughput and latency simultaneously using a 15-minute prediction horizon. The model is trained using real-world experimental data, which includes recent throughput history, network-layer indicators, and environmental variables. By utilizing a shared backbone for simultaneous downlink throughput and latency forecasting, this dual-task approach outperforms single-task models. Furthermore, its computational efficiency makes it ideal for edge deployment and real-time, application-aware decision-making.

Keywords: LEO satellite networks, Starlink, Throughput Prediction, latency Prediction, multi-task learning, RandomForest regression, Maritime communication

Abstract of Paper Accepted in ICAIC-2026

617

Interpretable Generative AI for Predictive Project Risk and Success Analytics

Madhusudan Bangalore Nagaraja, Abhishek Jain, Renuka Yallappa

ABSTRACT

Enterprise IT projects frequently fail because early risk indicators are dispersed across structured metrics and unstructured project narratives. This paper proposes a governanceaware, multi-modal analytical framework for early prediction of project risk and success across enterprise portfolios. Using a publicly available IT project management dataset from Kaggle, the framework integrates structured project metadata with unstructured textual descriptions through a layered architecture encompassing data governance, feature representation, multimodal fusion, predictive modeling, and generative AI–based reasoning. Classical machine learning models achieve strong predictive performance under stratified cross-validation, with accuracy and F1-scores in the range of 0.98–0.99 and AUC values approaching 0.99. Multimodal fusion preserves predictive strength while enhancing explanatory depth, and uncertaintyaware calibration supports risk-sensitive decision-making. SHAPbased feature attribution aligns dominant predictors, such as project completion progress, cost, and complexity, with narrative explanations generated by large language models. Business impact analysis also demonstrated the framework's potential to reduce financial exposure through early intervention. Our proposed approach balances predictive performance with transparency, governance, and organizational trust for enterprise project analytics.

Keywords: Enterprise Project Management, Risk Analytics, Generative AI, AI Governance.

Abstract of Paper Accepted in ICAIC-2026

618

AI Driven Claims Adjudication: Optimizing Healthcare Systems with Machine Learning and Deep Neural Networks

Damodhara Reddy Palavali, Suneetha Pothireddy, Dinesh Kumar Elumalai, Madhusudan Nagaraja

ABSTRACT

Healthcare claims adjudication is an important administrative process that affects the accuracy and efficiency of reimbursements in healthcare systems and prevents fraud. Growing volumes of claims, challenging billing codes, undocumented clinical records, and sophisticated fraud schemes present weaknesses in rule-based and manual adjudication. This study introduces an AI-based claims adjudication system that combines transformer-based natural language processing (NLP), ensemble machine learning, and a two-step fraud detection pipeline. The system examines the structured attributes of claims, utilization, profile of provider behavior, and contextual embeddings of clinical text to automatically perform adjudication with great accuracy and transparency. Claim classification is performed using a weighted combination of XGBoost, Random Forest, and neural network classifiers, and fraud detection is performed using a hybrid model of unsupervised anomaly detection and supervised learning. On large-scale healthcare data, experimental analysis reveals 94.7% classification accuracy, 67 percent processing time reduction, and 37 percent better fraud detection results than rule-based systems. The interpretability, regulatory compliance, computational complexity, and deployment considerations discussed in this study make the framework a scalable platform for next-generation healthcare claims processing.

Keywords: Claims adjudication, healthcare analytics, machine learning, fraud detection, NLP, transformers, ensemble learning, XGBoost, deep learning

Abstract of Paper Accepted in ICAIC-2026

619

LLM-Assisted Codebook Development for Cybersecurity Interviews with Enhanced Accuracy and Reduced Hallucination

ABSTRACT

Beyond what numerical data captures, qualitative cybersecurity interviews reveal human behaviors, lived experiences, trust perceptions and decision-making patterns. However, today's current manual and software-assisted coding is slow, difficult to scale and subjective when distinguishing expert and non-expert perspectives. Consequent, recent development of Large Language Models (LLMs) makes them useful for qualitative analysis, but larger models remain costly despite lower hallucination, while smaller models alternatives are cheaper but less reliable. A codebook plays an essential role in structuring themes and interpreting qualitative data transparently and consistently. Therefore, this study proposes an LLM-assisted architecture to generate traceable and hierarchically structured codebooks from cybersecurity interviews. Five techniques were grouped into three areas: accuracy improvement, hallucination reduction, and reduction of context memory usage. These techniques were applied to measure performance, reliability, and coding quality from seven LLMs of various parameter sizes. The architecture produced accurate codebook that improved coding reliability by up to 75\% for non expert and 35\% for experts when compared to baseline manual extraction. Reduction of contextual memory use increased processing efficiency by over 40\%, enabling even 1B–3B models to run effectively. Hallucination dropped by 82\%, which demonstrates that trustworthy qualitative codes can be generated by small and mid-sized LLMs.

Keywords: Large Language Models (LLMs), Qualitative Data Analysis, Codebook Development, Cybersecurity, Hallucination Reduction, LLaMA, Gemma, Phi

Abstract of Paper Accepted in ICAIC-2026

620

ZT-ICAS: A Zero-Trust Integrity-Constrained Framework for Agentic Vulnerability Scanning

Jothsna Praveena Pendyala, Sumeet Jeswani

ABSTRACT

Software vulnerability discovery is currently dominated by static application security testing (SAST) and heuristic rule-based scanners. While scalable, these approaches suffer from high false-positive rates, weak exploitability assessment, and limited adaptability to evolving codebases. Recent advances in Agentic AI offer a paradigm shift, enabling adaptive scanning through autonomous tool orchestration and multi-step reasoning. However, embedding agentic reasoning into security workflows fundamentally expands the attack surface, as these systems must ingest adversarially controllable inputs such as source code, inline comments, and dependency metadata to function. This exposure allows attackers to exploit the scanner via prompt injection and reasoning manipulation. This paper presents a secure-by-design framework for agentic vulnerability scanning that explicitly treats the analysis system as a security-critical target. We propose a Zero-Trust architecture enforcing semantic input neutralization, ephemeral context sharding, and capability-bounded tool invocation to preserve analysis integrity. We empirically evaluate this framework against codebases augmented with adversarial comments and poisoned metadata. Results demonstrate that while unconstrained agents are susceptible to systematic manipulation, our constraints significantly reduce analysis compromise with manageable performance overhead.

Keywords: Agentic AI, Vulnerability Scanning, Zero Trust Architecture, Security, Adversarial attacks, Autonomous Agents

Abstract of Paper Accepted in ICAIC-2026

621

Semantic Chunking for Triple Extraction from Cyber Threat Intelligence Reports

Shane Waldrop, Sawyer Cawthon, Erdogan Dogdu, Roya Choupani, Bill Mitchell

ABSTRACT

Cyber threat intelligence (CTI) reports are gold mines of adversary tactics, malware behaviors, and indicators of compromise—but this knowledge is buried in prose that machines cannot easily query. We present a pipeline that extracts structured knowledge triples from CTI documents using semantic chunking and consensus-filtered LLM extraction. Our Max-Min chunking algorithm groups sentences by embedding similarity while a lookahead mechanism prevents the kind of topic drift that plagues naive approaches. Three prompt variants query Gemma2 (9B parameters), and a consensus filter keeps only triples that multiple prompts agree on—the logic being that if the model hallucinates differently each time, agreement suggests signal rather than noise. Validation against our MALONT-lite schema (75 entity types, 10 predicates) winnows 351 raw extractions down to 31 valid triples. A separate repair pipeline recovers 91 additional triples, though against a different schema (STIX 2.1), which creates an unfortunate incompatibility we discuss frankly. The sobering 10.2% validation rate reveals just how poorly current LLMs follow explicit schema constraints—a finding that may prove more valuable than the extraction pipeline itself.

Keywords: cyber threat intelligence, knowledge extraction, semantic chunking, large language models, knowledge graphs

Abstract of Paper Accepted in ICAIC-2026

622

Quad-Stream Deepfake Detection: Combining Spatial and Frequency Domain Analysis for Robust Video Authentication

Vaishnav Anand
The Athenian School
Ivan Rodriguez
Brown University
vaishnavanand90@gmail.com, ivan_felipe_rodriguez@brown.edu

ABSTRACT

Deepfake videos poses a great challenge in cybersecurity, it may deceive people to take actions that increase the vulnerability to fraud, misinformation, identity theft, and social engineering attacks. Current methods, for the most part, take as input RGB images and process them as a blackbox algorithm that learns to distinguish between real and fake. In this paper, we propose a quad-stream deepfake detection network that combines spatial and frequency domain analysis of both the face region and the entire video frame. Our approach first performs face detection and alignment to extract high-quality face crops from video frames, while also retaining the context of the whole frame. Each face crop and full-frame image is then transformed into both RGB spatial domain and frequency domain representations (via Fourier transform), yielding four parallel input streams. We design a neural network architecture with four specialized CNN-based sub-networks: two spatial streams (face and frame) using pretrained backbones, and two frequency streams (face and frame) using modified ResNet backbones for frequency spectrum input. Features from all streams are combined for final classification. The proposed method shows promising results as it is able to spot subtle forgery artifacts in both facial details and broader scene context, as well as anomalies revealed in frequency spectra that are imperceptible in raw pixels. We evaluate our model on popular deepfake video benchmarks including FaceForensics++ and Celeb-DF v2, achieving state-of-the-art detection accuracy. Notably, our quad-stream approach improves robustness under challenging conditions such as video compression and cross-dataset evaluation.

Keywords: Deepfake detection, multi-stream network, frequency-domain analysis, video forensics, face manipulation, robust authentication.

Abstract of Paper Accepted in ICAIC-2026

623

Automated Validation and Repair of Knowledge Graph Triples for Cyber Threat Intelligence

Vitaly Andrejeus, Bill Mitchell, Sawyer Cawthon, Jesse Sullins, Erdogan Dogdu, and Roya Choupani

ABSTRACT

Cyber Threat Intelligence (CTI) knowledge graphs depend on extracting Structured Threat Information Expression (STIX) 2.1-compliant entity–relation triples from unstructuredthreat reports, but current systems oftentreat schema violations as terminal errors and discard invalid outputs. In practice, Large Language Model (LLM) extraction frequently produces near-correct triples that fail validation due to minor formatting artifacts, entity type drift across sentences, alias fragmentation, or STIX domain-range mismatches. This paper presents a multi-stage triple validation and repair framework that recovers such rejected triples while enforcing strict STIX 2.1 constraints. Starting from an invalid-triple log, the framework applies deterministic normalization (artifact stripping, type/predicate canonicalization, and domain-range enforcement), followed by probabilistic Markov smoothing to stabilize entity typing across document contexts. A unified, schema-constrained LLM repair module generates a small set of candidate repairs,
each revalidated under the same STIX rules, ensuring that no hallucinated entities or unsupported predicates enter the final output. Finally, a graph-based Skew Zero Forcing (SZF)pass then reinforces structurally consistent neighborhoods by propagating trust from high-confidence nodes and filtering in-compatible relationships. Using DNRTI as a benchmark, the results show that staged repair significantly improves the usable triple set by converting a portion of validation failures into STIX-valid triples, increasing graph connectivity and reducing information loss without relaxing ontology constraints.

Keywords: Cyber Threat Intelligence, Knowledge Graphs, Triple Extraction, Markov Smoothing, Skew Zero Forcing

Abstract of Paper Accepted in ICAIC-2026

624

Optimizing Production Variance and Yield Reporting through Cross-Modular Integration in SAP S/4HANA

Harsh Patel
UST Global Inc, USA
harshpatelv1009@gmail.com

ABSTRACT

Traditional ERP environments typically do not succeed in providing timely reconciliations between production yield and scrap captured in PP-PI, with the related financial variances posted in CO. Since data is split, manual reconciliations are required, and thus managerial intervention cannot be based on timely data; therefore, the objectives of traditional MAPs cannot be realized to their full potential. It discusses the architecture pivot that was introduced with SAP S/4HANA and defines how to leverage the Universal Journal (ACDOCA) and the Virtual Data Model (VDM), realized based on Core Data Services (CDS) views, in building an integrated link from PP-PI operational confirmation data to Financial Accounting (FI) variance calculation. The cross-modular reporting solution turns retrospective variance analysis into a proactive management process by enabling real-time yield-adjusted variance analysis. This integration brings agility and data transparency to support effective Integrated Business Planning (IBP), where the operational targets are related to measurable financial results, thereby speeding up the management decision cycle.

Keywords: SAP S/4HANA, Production Variance, Yield Reporting, Integrated Business Planning (IBP), Core Data Services (CDS Views), Management Accounting Practices, Universal Journal, PP-PI

Abstract of Paper Accepted in ICAIC-2026

626

A Hybrid Methodology for SAP Implementations: Blending ASAP with Agile Principles for Enhanced Flexibility and Stakeholder Engagement

Harsh Patel
UST Global Inc, USA
harshpatelv1009@gmail.com

ABSTRACT

SAP ERP implementations do deliver a corporate transformation foundation initiative, but substantial overruns on their costs and schedules are a standard problematic experience due to so many dimensions of their inherent complexities and chronic misalignment of fixed requirements with evolving business environments. Originally implemented, Accelerated SAP (ASAP) would provide strong governance with its sequential roadmap. Rigidly framed like any waterfall process, ASAP cannot keep up with the required pace for development in today's digital world where conditions change fast and when late user involvement raises the risk at stake. Methodological weaknesses that Hybrid ASAP-Agile is informally used to overcome will be formalized herein. Therefore, it structurally embeds what governance by prediction seems necessary from the planning phases (Project Preparation, Final) into adaptive, iterative delivery of Agile sprints (Realization). This synthesis fundamentally redefines core activities such as requirements engineering (moving to Lean Blueprint via prioritized user stories) and feedback loops (continuous Sprint Demos). Analysis shows that it increases project flexibility. It speeds up time-to-value by delivering the most important functionality first and also increases stakeholder satisfaction because of continuous involvement and increased transparency which makes the delivery of the project inherently aligned with core business value creation.

Keywords: ASAP Methodology, Agile, Hybrid Project Management, SAP Implementation, Stakeholder Management, ERP.

Abstract of Paper Accepted in ICAIC-2026

628

SARS-CoV-2 Classification Using Classical vs. Quantum Machine Learning: A Performance Comparison of SVM, QSVM, and Pegasos

Sthefanie J. G. Passo, Muntakim Haque, Vishal H. Kothavade, John J. Prevost
The University of Texas at San Antonio, USA
sthefanie.passo@utsa.edu

ABSTRACT

Quantum Machine Learning (QML) integrates principles of quantum computation with statistical learning to address classification tasks. This paper conducts a comparative study of three paradigms—classical Support Vector Classification (SVC), Quantum Support Vector Classification (QSVC) using quantum kernels, and QSVC enhanced via the Pegasos stochastic algorithm. We benchmark classification accuracy on a SARS-CoV-2 dataset, analyzing model robustness and resource requirements. Our results indicate that while classical SVC maintains superior scalability and stability for large datasets, QSVC variants show competitive accuracy across simulator backends (Central Processing Units or CPUs and Graphical Processor Units or GPUs) and may achieve quantum advantage as hardware matures. Implications for hybrid workflows and future research directions are discussed.

Keywords: Quantum machine learning, QSVC, Pegasos, support vector classification, GPU acceleration, quantum computing

Abstract of Paper Accepted in ICAIC-2026

629

Beyond Single-Hop:Link Prediction Through Multi-Hop Reasoning In Malware Knowledge Graphs

Kibrom Bahlibi, Hoang Long Nguyen, Erdogan Dogdu, Roya Choupani, and Bill Mitchell

ABSTRACT

Knowledge graphs (KGs) have become essential for representing complex relationships among entities such as malware, vulnerabilities, and attack techniques extracted from cyber threat reports. Traditional KG embedding methods primarily focus on single-hop link prediction, limiting their ability to capture longer relational chains that are critical in real-world attack scenarios. Multi-hop reasoning, on the other hand, reflects the true structure of cyber threats, where interactions often span multiple entities (for example, malware → technique → vulnerability → software → sector). Inferring such indirect connections enables early detection of potential attack surfaces, improved attribution, and richer contextual threat intelligence. In this work,
we propose a unified framework for multi-hop link prediction in malware knowledge graphs. The framework learns latent representations of entities and relations by leveraging both graph- based and language-based reasoning paradigms. Specifically, we fine-tune large language models (LLaMA-2-7B and Phi-2) for path-based reasoning using natural language representations,
and we benchmark them against graph neural network (GNN) variants such as GraphSAGE, GCN, and GAT. Experimental results show that GNNs generalize effectively on shorter paths but degrade as relational complexity increases, while LLMs maintain robust performance across longer reasoning chains due to their chain-of-thought and contextual learning capabilities. These findings highlight the complementary strengths of symbolic and neural reasoning approaches and demonstrate the potential of LLMs for dynamic malware analysis and reasoning over unseen entities

Keywords: Cyber threat intelligence, malware knowledge graph, multi-hop link prediction, graph neural networks (GNNs), large language models (LLMs), cybersecurity analytics

Abstract of Paper Accepted in ICAIC-2026

631

From UIKit to SwiftUI: A Quantitative Analysis of Migration Strategies for Large-Scale Financial Applications

Ranjith Kumar Vanaparthi

Ventois Inc, USA
ranjithkumarvanaparthi21@gmail.com

ABSTRACT

UI, which we focused on, is performance in the render cycle, development speed, and term maintainability. We saw great results, validated by empirical data showing a 32.3% reduction in peak CPU usage during data sync, which is a significant performance improvement. Also, we saw a 42.0% increase in how quickly we added features and an 18% point increase in how well the WCAG 2.1 Level AA accessibility criteria were met. We present to you a rigorous and proven path for enterprises that wish to modernise highly complex and regulated mobile platforms, and we are confident that we are mitigating operational risk.

Keywords: SwiftUI, UIKit, Software Migration, Mobile Development, Performance Benchmarking, Accessibility, Declarative Programming, Strangler Fig Pattern, Financial Technology (FinTech).

Abstract of Paper Accepted in ICAIC-2026

632

Architecting for Privacy and Performance: A Federated Learning Framework for On-Device Financial AI in Mobile Applications

Ranjith Kumar Vanaparthi

Ventois Inc, USA
ranjithkumarvanaparthi21@gmail.com

ABSTRACT

User concerns about mobile financial AI are addressed by integrating personalized services, fraud detection, and robust security. However, concerns about data centralization, the GDPR, and the global data protection streamlining make the design difficult. In this regard, this paper presents DPFed-MobileFin. This innovative cross-device Federated Learning (FL) framework specially focuses on iOS banking applications. DPFed-MobileFin combines state-of-the-art privacy-preserving methods —'Local Differential Privacy' (LDP) and 'Secure Aggregation' (SA) — with on-device ML training using Core ML and TensorFlow Lite, and an on-device Differential Privacy (DP) implementation. In addition, the system uses a Personalized FL (PFL) framework to address performance degradation caused by non-IID, heterogeneous financial data. The system architecture is designed to achieve hardware acceleration using Core ML's Neural Engine, cryptographic privacy guarantees, and a focus on mitigating computational overhead. Such a comprehensive system architecture meets data residency requirements and sets a new standard for safe and highly effective AI in cross-device financial applications and systems. The framework demonstrates its promise for privacy and usefulness by attaining a fraud detection rate of 95.1%. The low mobile overhead, low latency, and energy usage justify the mobile user experience.

Keywords: Federated Learning, Core ML, Data Privacy, On-Device AI, Financial Anomaly Detection, Mobile System Architecture, Differential Privacy, Personalized FL.

Abstract of Paper Accepted in ICAIC-2026

633

CodeGraph- Malware Detection via Control Flow Graph Embeddings and Graph Neural Networks

Chandrashekhar Medicherla, Chaitanya Kulkarni, Rajesh Purushothaman, Rakesh Keshava, Arun Kumar, Nandagopal Seshagiri
Independent Researcher, United States

ABSTRACT

Cybersecurity threats cost organizations $6 trillion annually, with traditional signature-based antivirus systems achieving only 78.5% detection accuracy and failing completely against zero-day malware variants. This paper presents CodeGraph, a novel malware detection framework that combines Control Flow Graphs (CFGs) with Graph Neural Networks (GNNs) for behavioral analysis. Unlike conventional approaches treating malware as byte sequences or images, CodeGraph captures structural program behavior through graph representations that persist across obfuscation techniques. We validate the approach using synthetic data carefully designed to model real malware characteristics, demonstrating 99.4% detection accuracy with 99.0% zero-day detection rate. This represents a 7-9 percentage point improvement over CNN (92.4%) and LSTM (90.8%) baselines, proving that graph-based behavioral analysis is superior to sequential and image-based methods. With less than 1 millisecond inference time, the system enables real-time deployment in enterprise environments.

Keywords: Malware Detection, Graph Neural Networks, Control Flow Graph, Zero-Day Threats, Behavioral Analysis, Cybersecurity

Abstract of Paper Accepted in ICAIC-2026

634

Bridging the Black Box: Explainable Anomaly Detection for Critical Infrastructure Systems

William Mitchell, John Deleon, Serena Reese, Erdogan Dogdu, and Roya Choupani

ABSTRACT

As computer systems are increasingly used in critical infrastructure systems (CIS), the need to secure and protect them has grown exponentially. Cyberattacks have become more frequent and complex, further increasing the chance of malicious actors targeting infrastructure. Deep learning models have shown promise in being effective at detecting anomalies, but are generally uninterpretable. Our CIS framework seeks to incorporate Explainable AI (XAI) and knowledge graphs to make a more interpretable model that enhances the provided context of each detected anomaly, while balancing the predictive accuracy of deep learning models with the added XAI. It includes an ETL process that ingests network and sensor logs through a scalable pipeline powered by Spark, Kafka, TimescaleDB, and Neo4j graph databases. We deploy a graph neural network autoencoder for anomaly detection and an XAI framework, such as SHAP or GNNExplainer, to indicate the most influential features, nodes, and relationships from the graph.

Keywords: Critical Infrastructure Security, Knowledge Graphs, Graph Neural Networks, Explainable AI

Abstract of Paper Accepted in ICAIC-2026

637

Explainable AI for Cloud Intrusion Detection: A User Study of SHAP and LIME in AWS GuardDuty

Zakaria Alomari, Humberto Goncalves
Department of Computer Science, New York Institute of Technology
Vancouver, British Columbia, Canada
Email: {hdasilva, zalomari}@nyit.edu

ABSTRACT

Cloud-based intrusion detection systems, such as AWS GuardDuty, provide effective threat detection but often lack transparency, which undermines analyst trust and incident response capabilities. This paper evaluates the integration of XAI techniques, specifically SHAP and LIME, into AWS GuardDuty alert triage workflow to enhance interpretability. An XGBoost model trained on the CIC-IDS2017 dataset was used to generate explanations for security alerts presented to twelve security professionals, divided into two groups: control and XAI-supported. Participants classified four attack types while providing confidence ratings and justifications for their decisions. Results show that analysts with XAI access achieved higher classification accuracy, with three participants correctly classifying all alerts, compared to none in the control group. They also demonstrated increased confidence levels and more elaborate reasoning. However, participants encountered varying difficulty interpreting SHAP and LIME visualizations, with LIME proving more immediately actionable for alert-specific analysis while SHAP's global insights presented a steeper learning curve. These findings highlight both the potential and practical challenges of deploying XAI in operational cloud security environments.

Keywords: Explainable Artificial Intelligence, Interpretability, Human-AI interaction, AWS GuardDuty, Intrusion Detection Systems, Cloud Security, Adversarial Scenarios

Abstract of Paper Accepted in ICAIC-2026

638

Toward Deployable Disinformation Defense: Benchmarking Lightweight Transformers on FakeNewsNet

Humberto Goncalves, Kossi Sam Affambi, and Zakaria Alomari*
Department of Computer Science, New York Institute of Technology
Vancouver, British Columbia, Canada
Email: {hdasilva, kaffambi, and zalomari}@nyit.edu

ABSTRACT

Disinformation poses a persistent threat to news integrity and online trust, spreading six times faster than factual information on social media platforms. While some transformer-based models achieve state-of-the-art performance in fake news detection, their high computational costs and lack of interpretability limit deployment capabilities and operational transparency. This work benchmarks three lightweight transformer models on FakeNewsNet datasets to address critical gaps in scalability, interpretability, and applicability for fake news detection. Results show that DistilBERT achieves the highest accuracy, precision, F1-score, and AUC, while ALBERT demonstrates superior recall. TinyBERT offers the fastest inference time with minimal memory footprint, making it ideal for edge deployment. Explainability analysis using SHAP and LIME reveals that all models identify similar semantic patterns as primary indicators of fake news, with model-specific attention mechanisms reflecting their respective compression strategies. This study supports the development of scalable, transparent, and high-performing disinformation detection systems suitable for real-world deployment where efficiency, interpretability, and reliability are essential.

Keywords: Lightweight Transformer Models, Natural Language Processing (NLP), Fake News Detection, Real-Time Disinformation Detection, Explainable Artificial Intelligence (XAI)

Abstract of Paper Accepted in ICAIC-2026

642

Proactive Discrepancy Detection at Scale: A Continuous Validation Framework for Distributed Data Systems

Hareendra Sura
Coupang, USA
Hareendrasura@gmail.com

ABSTRACT

The move to eventually consistent, highly distributed data architectures has brought about extreme complexities toward efforts of assuring both data integrity and digital accessibility compliance. This paper discusses an approach to high throughput continuous validation as a method of proactive discrepancy detection that can be applied during high-risk data migrations in large-scale systems. More than 100 million records per day will be validated using temporal, massively parallel validation activities comparing legacy and target environments side by side- with the added benefit of no long-lived data pipelines that bring overheads into the process. At the heart of this architecture sits putting accessibility as a first-class engineering requirement firmly within the software development lifecycle (SDLC). The proposed framework combines rule-based engines with AI-augmented visual auditing and CI/CD-integrated testing pipelines for shifting left quality assurance. The implementation in a regulated public sector enterprise results in an 80 percent reduction of post-launch data discrepancies together with a 68 percent reduction in the accessibility defect injection rate. This paper describes the design of WAI-ARIA 1.2 implementation patterns for complex single-page application (SPA) components that do not relate to any particular framework and sets up a model that can be used to quantitatively assess the impact on engineering productivity and user inclusiveness.

Keywords: Distributed Data Systems, Continuous Validation Framework, Data Consistency, Eventual Consistency, Rollout Safety, WAI-ARIA 1.2, WCAG 2.1, AI-Augmented Testing, TMHP Case Study, SDLC

Abstract of Paper Accepted in ICAIC-2026

643

Forecasting and Overcommit Pipelines for Cloud Storage: A Model for Persistent Disk Capacity Management

Hareendra Sura
Coupang, USA
Hareendrasura@gmail.com

ABSTRACT

A significant challenge for the persistent disk capacity management exercise within private cloud environments is inefficiency. Thus, this paper reviews an implemented production-grade pipeline in forecasting storage demand and safe overcommiting of capacity. It replaces static, cluster-level hard-coded overcommit percentage with a dynamic data-driven framework based on historical usage patterns storage consumption to achieve a 35% more correct resource allocation. This paper details the productization of 'historical' capacity and utilization data, coupled with real-time operational signals, processed by very advanced statistical forecasting models toward determining overcommit decisions automatically. The technical core of such a framework is hence a compliance model injected proactively into SDLC so recoveries do not hamper digital inclusion. This has been introduced as 'AI-augmented' testing in the pipeline using Axe-core and engineered implementation patterns for WAI-ARIA that were validated through an accessibility KPI quantitative impact model. A longitudinal case study of the Texas Medicaid & Healthcare Partnership, a regulated public-sector enterprise going through very large digital transformations, will prove the model accurate. Results found that predictive resource forecasting and automated infrastructure control improved not only resilience in operation but also the use of resources within mission-critical cloud environments.

Keywords: Capacity planning, predictive resource forecasting, cloud storage management, data-driven infrastructure control, infrastructure efficiency, digital accessibility, SDLC compliance.

Abstract of Paper Accepted in ICAIC-2026

645

Adaptive Honeypots using Reinforcement Learning Algorithms DQN and DDQN in Ensemble Framework: A Systematic Literature Review

ABSTRACT

The exponential growth of sophisticated cyber threats has necessitated the evolution from traditional static honeypots to intelligent, adaptive deception systems. This systematic literature review examines the integration of reinforcement learning algorithms, specifically Deep Q-Networks (DQN) and Double Deep Q-Networks (DDQN), within ensemble frameworks for adaptive honeypot systems. Following PRISMA 2020 guidelines, we analyzed 48 studies from 2015-2025, revealing significant performance improvements in adaptive honeypots employing reinforcement learning techniques. Our analysis demonstrates that DDQN-based adaptive honeypots achieve 95% detection accuracy with 7.2 minutes average engagement time, substantially outperforming traditional static honeypots (73% accuracy, 2.1 minutes engagement). Furthermore, proposed ensemble DDQN frameworks show potential for 97% detection accuracy with 4% false positive rates. Key research gaps identified include limited real-world deployments, lack of standardized evaluation frameworks, and insufficient integration of ensemble methods. This review contributes a comprehensive taxonomy of adaptive honeypot architectures, comparative analysis of reinforcement learning approaches, and identifies eight critical research directions for future work. The findings provide a foundation for developing next-generation cyber deception systems capable of autonomously adapting to evolving threat landscapes.

Keywords: Adaptive honeypots, reinforcement learning, Deep Q-Networks, Double Deep Q-Networks, ensemble learning, cybersecurity, intrusion detection, cyber deception

Abstract of Paper Accepted in ICAIC-2026

646

Securing 5G/6G Networks: Handover, Slicing, and QoS Security Issues

Hritesh Yadav, Varun Singh, Kshitij Sharma, Suyash Karmarkar, Akhilesh Keshap
Independent Researcher, USA
hritesh.yadav@ieee.org

ABSTRACT

Next-generation mobile networks rely on advanced mechanisms such as seamless handover, logical network slicing, and strict Quality of Service (QoS) enforcement to support latency-sensitive and mission-critical applications. While these mechanisms are central to the design of 5G and anticipated 6G architectures, they also introduce new and often underexplored security weaknesses. This paper analyzes security challenges associated with mobility management, slice isolation, and QoS control in 5G and emerging 6G networks. We examine how shortcomings in handover signaling, shared virtualized infrastructure, and QoS enforcement can be exploited to disrupt service continuity, violate isolation guarantees, and degrade network performance. The analysis highlights realistic attack scenarios, including denial-of-service, cross-slice interference, and control-plane manipulation, and evaluates how current standards and deployment practices address—or fail to address—these threats. Based on these findings, the paper identifies open research challenges and outlines architectural and protocol-level directions for improving resilience without undermining performance or scalability. The study underscores the necessity of integrating security considerations directly into the design of mobility, slicing, and QoS mechanisms as mobile networks evolve toward 6G.

Keywords: 5G Security; 6G Networks; Mobility Management Security; Secure Handover; Network Slicing Isolation; QoS Enforcement; Telecom Cybersecurity; Control-Plane Attacks; Denial-of-Service Mitigation

Abstract of Paper Accepted in ICAIC-2026

650

Intelligent Test Data Automation: A Python-Based Framework for Deterministic and Scalable Software Testing

Datta Snehith Dupakuntla Naga
Teladoc Health Inc, USA
dndattasnehith1989@gmail.com

ABSTRACT

As in many areas where software systems are turning to data-driven and AI-enabled elements, we see that the quality and consistency of test data are being recognised as vital to validating system performance. We see that traditional methods of test data generation are manual, tied to the environment and non-deterministic, which in turn limits their value in automated and intelligent testing systems. This work presents an AI-enabled Python-based test data automation framework that combines deterministic data generation, intelligent constraint enforcement, and neuro-symbolic data synthesis. We present a framework that combines rule-based generation and machine-assisted validation to guarantee schema compliance, referential integrity, and policy adherence in large-scale test environments. By weaning test data logic away from the infrastructure, enabling on-demand generation within CI/CD pipelines, we have a solution that supports the reliable testing of AI and data-intensive systems. We have evaluated the framework in enterprise-scale testing, which has reported improved reproducibility, scalability and automation reliability. We report a 68% reduction in defect injection rates through a shift-left approach to access and data integrity. We have also developed a new quantitative model for the return on investment of proactive accessibility engineering, which we measure using metrics such as Compliance Velocity and Defect Escape Rate. A long-term case study within the Texas Medicaid Healthcare Partnership (TMHP) system, we use as a baseline and proven model for the adoption of proactive accessibility engineering in large-scale, compliance-driven organizational settings.

Keywords: AI-Assisted Test Data Synthesis, Deterministic Automation, Python-Based Reasoning Engines, Software Development Life Cycle (SDLC), TMHP, WAI-ARIA 1.2, WCAG 2.1 Compliance, CI/CD Integration, Scalable Software Testing

Abstract of Paper Accepted in ICAIC-2026

686	**Multi-Frequency Implementation and Timing Analysis of LAES on Spartan-7 and Kintex-7 FPGAs** **Keshav Kumar, Bishwajeet Pandey, Chinnaiyan Ramasubramanian** **ABSTRACT** **Keywords:**

Abstract of Paper Accepted in ICAIC-2026

721

Resource allocation in 6G Networks:A comprehensive review of Machine Learning Approaches

Devanand Patil, Joydev Ghosh, Bishwajeet Kumar Pandey

ABSTRACT

The resource allocation in sixth generation (6G) wireless networks is important to support the stringent characteristic of the network. The characteristics include very high data rates, very low latency, massive connectivity, reliability, security, energy efficiency and scalability. The resource allocation in 6G is challenging due to heterogeneous, dynamic, ultra-low latency constraints of 6G networks. Machine learning (ML) is essential because conventional methods cannot handle the characteristics of the 6G networks. ML approaches have necessary intelligence and adaptability to fulfill the network demands. In this paper, we have studied Convolutional Neural Network (CNN), Recurrent Neural Networks (RNN),Long Short-Term Memory(LSTM), Deep Q Learning and Hybrid (RNN+CNN) machine learning schemes. It is identified that the Hybrid scheme can optimize resource allocation in 6G efficiently. Eventhough hybrid scheme is expected to be the promising approach for resource allocation, several challenges to be addressed yet to achieve the target value of key performance indicators envisioned for 6G.

Keywords: 6G, Resource Allocation, Machine learning, CNN, RNN, LSTM, Massive connectivity,

IEEE International Conference on AI in Cybersecurity (ICAIC) 2026

Houston, Texas, United States

Legal, Safety, and Participation Terms

1. Liability Disclaimer

All participants attend the conference at their own risk. IEEE ICAIC 2026, IEEE, the University of Houston, the organizing committee, volunteers, sponsors, and partners shall not be held liable for any injury, loss, damage, or incident that may occur during or in connection with the conference.

2. Code of Conduct

All attendees are expected to conduct themselves in a professional and respectful manner. Any form of harassment, discrimination, or disruptive behavior will not be tolerated and may result in removal from the conference without refund.

3. Force Majeure

The organizers reserve the right to modify, postpone, or cancel the conference in whole or in part due to circumstances beyond their reasonable control, including but not limited to natural disasters, health emergencies, or government-imposed restrictions.

4. Recording Consent

By attending the conference, participants grant permission for audio, video, and photographic recordings to be made and used for conference-related purposes, including promotional and archival use.

5. Speaker Responsibility

Speakers are solely responsible for the content of their presentations and must ensure that their material does not violate intellectual property rights or applicable laws.

6. Data Privacy

Participant information will be collected and used strictly for conference administration, communication, and operational purposes, in accordance with applicable data protection regulations.

7. Refund Policy

Registration fees are non-refundable unless explicitly stated otherwise by the conference organizers.

General Limitation of Liability

In addition to the above terms, IEEE ICAIC 2026, IEEE, the University of Houston, the organizing committee, volunteers, sponsors, and partners shall not be responsible or liable for any unfavorable, unforeseen, or adverse events, including but not limited to:

- Travel delays, cancellations, or visa-related issues
- Adverse weather conditions or natural disasters
- Health-related incidents, outbreaks, or medical emergencies
- Technical failures, including internet or system disruptions
- Loss, theft, or damage of personal belongings
- Actions or omissions of third-party service providers
- Conduct of participants, speakers, exhibitors, or vendors
- Government regulations, restrictions, or security actions
- Any other circumstances beyond the reasonable control of the organizers

By registering for or attending IEEE ICAIC 2026, participants acknowledge and agree to these terms and waive any claims against the organizers arising from such events.

NEXT CONFERENCE

11th International Conference on Green Computing and Engineering Technologies (ICGCET ®)

01-02 April 2026

Liberty Central Saigon Riverside Hotel, Ho Chi Minh City, Vietnam

https://icgcet.org/

2026 IEEE Conference on Generative AI for Secure Systems (GAISS)

28-30 October 2026

Austin Texas USA

https://gaiss.info/

https://conferences.ieee.org/conferences_events/conferences/conferencedetails/66401