Sunday, November 1 - Early Registration (5:00pm - 7:00pm)

Monday, November 2 - Day 1

	ROYAL PALM 1	ROYAL PALM 2	ROYAL PALM 3	ROYAL PALM 4
7:00	**Registration (till 5:00pm)**			
7:30	**Continental Breakfast**			
8:00	**Welcoming and Opening Remarks by Dr. Taghi Khoshgoftaar, Florida Atlantic University, Conference Chair** **Welcoming Remarks by Dr. Ken Dawson-Scully, Florida Atlantic University, Provost and Executive Vice President for Academic Affairs and Dean Stella Batalama, College of Engineering and Computer Science at FAU** **Room : Grand OASIS Ballroom**			
8:15	**Keynote Speaker : Gary S. Connor** **Independent Aerospace Consultant (Brig Gen, USAF, retired)** **Room : Grand OASIS Ballroom**			
9:15	**Coffee Break**			
	Session 11A (full papers) **Chair: Andy Sinclair** **Room: Royal Palm 1**	**Session 11B (full papers)** **Chair: Sonia Moshfeghi** **Room: Royal Palm 2**	**Session 11C (full papers)** **Chair: Sareh Taebi** **Room: Royal Palm 3**	**Session 11D (short papers)** **Chair: Hamid Usefi** **Room: Royal Palm 4**
9:30	**[161] DREG: A Layer-Wise Jacobian Regularization as a General-Purpose Penalty** *Martnishn Rowan*	**[810] Multi-Class Mental Health Detection on Reddit: A Zero-Shot Comparison of Reasoning and Non-Reasoning Large Language Models** *Saquer Jamil*	**[804] When Quantum Entanglement Meets Transformer Attention: A Controlled Study of Quantum Attention Heads** *Zangoti Hussein*	**[756] CServe: A vLLM-Native Control Plane for Shared Bare-Metal GPU Clusters** *Agarwal Nitin*
9:40				**[123] PromptTab: Text-Conditioned Diffusion for Tabular Data Generation** *Shinners Timothy*
9:45	**[172] Wave Function Backpropagation with Explicit Temporal-Interval Dynamics** *Yu Byunggu*	**[825] A Comparative Evaluation of Transformers, Zero-Shot LLMs, and LoRA-Adapted Open-Source Models for Suicidal Ideation Detection** *Abusaqer Mahmoud*	**[155] Q-HyDRA: Quantum HyperDimensional Representation for Efficient Machine Learning** *Masum Abu Kaisar Mohammad*	**[178] Radial Basis Function Networks as Projection Heads in Self-Supervised Learning** *Schliebitz Andreas*
10:00	**[557] Regime-Aware Evaluation of Sequence Models under Partial Observability** *Sotiropoulos Dionisios*	**[423] Knowledge Tracing Representations for Interpretable Learning-State Trajectory Analysis** *Vu Minh Hoang*	**[540] Do Hybrid Quantum Components Improve Rare-Event Failure Ranking? A Controlled Evaluation with Mechanism and Budget Controls** *Malluvajhula Sri Bhargav*	**[424] FedMLCom: Multiple Local Computations for Communication-Efficient Federated Learning** *Palenzuela Karlo*
10:15	**[584] KESTREL: Coordinate- Wise Secant Steps with Post-Hoc Failure Benching for Deterministic Neural Optimization** *Fujimoto Takumi*	**[836] Toward Course-Aligned Feedback in a CS1 Grader Using Fine-Tuned Small Language Models** *Manrique Ruben*	**[673] QPredSGC: Hybrid Quantum Predicate Learning for Long-Tailed Scene Graph Generation** *Innan Nouhaila*	**[698] ProtoRetain: Rare-Class Prototype Retention for Federated AIoT Learning under Partial Participation** *Saha Prianka Rani*
10:30	**[636] Geodesic Step Budgeting for Neural Cellular Automata** *Dang Ian*	**[139] Factorized Multi-Agent MCTS with Teammate Modeling for Netted Radar Deception** *Zeng Weizhi*	**[798] BNKE: Kernel-Sensitive Shot Allocation for Projected Quantum Kernels Under Finite Measurement Budgets** *Islam MD Majedul*	**[238] SafeCouple: A Verification-Guided Repair Framework for Neural Network Controllers under Input-Output Coupled Constraints** *Mo Zihao*
10:45	**Coffee Break**			**[522] EXPLAIN-Bench: Benchmarking Explanation Robustness and Stability for Multimodal Emotion Recognition** *Juba Agoun*
	Session 12A (full papers) **Chair: Kehan Gao** **Room: Royal Palm 1**	**Session 12B (full papers)** **Chair: Sareh Taebi** **Room: Royal Palm 2**	**Session 12C (full papers)** **Chair: Michail Alexiou** **Room: Royal Palm 3**	**Session 12D (short papers)** **Chair: Hamid Usefi** **Room: Royal Palm 4**
11:00	**[778] DARE-FS: A Selector-Agnostic Operator for Detecting and Repairing Class-Specific Feature-Selection Harm** *Shorna Jannatul Ferdous*	**[175] ArtifactAuditBench V2: A Multimodal Benchmark for Auditing Scientific Experiment Artifacts** *Li Hongchen*	**[543] Learning Explainable Image–Text Alignment Assessment with Vision Language Models** *Gerhardt Christoph*	**[719] Advanced Methods for Weakspot Identification in Regression Models: From Standard Detectors to Structured Ensembles** *Schroven Leonard*
11:10				**[724] Leveraging Optimal Decision Trees with Logical Analysis of Data** *Ouled Dlala Imen*
11:15	**[32] Entropy-Punctured Bloom Filters for Memory-Efficient Machine Learning** *Cartmell John*	**[336] From Evaluator Geometry to Readout Control: Low-Dimensional Decision Axes in LLMs** *Grundetjern Morten*	**[593] A Behavioral Taxonomy of Attention Heads That Transfers Across Transformer Language Models** *Rathnasekara Heranga*	**[357] Multi-Objective Structured Pruning of LLMs for Latency and Model Size Optimization** *Ali Muhammad Junaid*
11:30	**[570] Hashrope: A Persistent Data Structure for Mechanisms of LLM** *Syed Muntaser*	**[350] VAST-NLG: Separating Policy Selection from Response Generation in Multi-Turn LLM Safety Evaluation** *Partow Soraya*	**[643] Validating a Composite Explanation Drift Index Across Attribution Methods and Medical-Imaging Domains** *Hammond Scott*	**[457] LANet: A Lightweight Network Architecture for Real-Time AI-ISP** *Chen Guo*
11:45	**[641] WSA-Aug: Enabling Efficient Data Augmentation for Wafer-Scale Accelerators** *Liu Dai*	**[499] Benchmark Leadership Is Not General Capability: A Validity-Centered Evaluation Framework Applied to GPT-5 and Gemini** *Alsaid Manar*	**[327] MedSentiBERT: Distilling a Frontier-LLM Sentiment Labelling Policy into a Compact Clinical Classifier** *Kalinzi Joseph*	**[632] The Effect of Quantization on Self-Certainty for Scalable LLMs' Reasoning Performance in Resource-Constrained Settings** *Reglo Murilo*
12:00	**[466] DiscoAct: Securing Autonomous Multimodal LLM Agents Against Cross-Modal Trust Flattening** *Dogdu Erdogan*	**[702] Configuration, Not Conscience: A Large-Scale Empirical Study of LLM System Prompts** *Argyropoulos Vasileios*	**[509] A Dual-Path Large Language Model Framework for Reliable Zero-Shot Feature Extraction** *Qian Yiliang*	**[791] XNano: Using XAI to Explain How Preprocessing Affects Nanoparticle Image Classification** *Li Yaqia*
12:15	**[710] Deployable W4A4 LLM Inference on Consumer Blackwell: Per-Token PE-QAT and Native INT4 mma.sync Kernels for SM121** *Mishra Shresth*	**[737] SpecDriftBench: A Matched Benchmark for Distinguishing Tool-Specification Drift from Agent Errors and Transient Failures in Tool-Using LLM Agents** *Siyuan Yang*	**[294] SemVerBench: Benchmarking LLM Comprehension of Version-Constraint Resolution Semantics** *Liu Zeming*	**[838] When Does Explanation Regularisation Improve Faithfulness? A Multi-Dataset Audit in Federated Fall Detection** *Roy Kaushik*
12:20				**[445] Name2Pkg: Lightweight One-Class Android Malware Screening via Name-Package Correspondence Modeling** *Sung Changyeop*
12:30	**[689] Physics-Aligned Modality Fusion for Next-Day Wildfire Spread Prediction** *Zhang Xin*	**[552] Cost-Aware Abstention for LLM-Based IT Ticket Classification** *Gupta Sachin*	**[805] Capacity-Pretraining Tradeoffs in YOLOv8-Based Dental Pathology Detection in Panoramic Radiographs** *Baby Kriza*	**[322] Update Slope Is Bounded by Prior Calibration: A Structural Bound and Its Contingent Empirical Signature** *Cardei Mihaela*
12:45	**Lunch** **Room : Grand OASIS Ballroom**			**[649] Physical-Direction Gates for Trustworthy Wildfire-Risk Models: Trading In-Distribution Accuracy for Monotone Domain Consistency** *Cohen Ido*
	13A (full papers) **Chair: Sareh Taebi** **Room: Royal Palm 1**	**13B (full papers)** **Chair: Huanjing Wang** **Room: Royal Palm 2**	**13C (full papers)** **Chair: Kehan Gao** **Room: Royal Palm 3**	**13D (short papers)** **Chair: Hamid Usefi** **Room: Royal Palm 4**
14:00	**[817] JOLT: Jointly Optimized Layerwise Termination for Budgeted Early-Exit Inference** *Shouborno S. A. I.*	**[414] CoTRA: Contribution-Tracked Risk Auditing of LLM Coding-Agent Trajectories** *LI HAO*	**[666] VTracker: A Tool to Identify Recommendation Traps and Measure AI Bias Intensity** *Agarwal Nitin*	**[612] From CVE to CWE: Multi-Label Weakness Prediction with Hierarchy-Aware Losses** *Cuccarini Marco*
14:10				**[613] Attention- and Depth-Aware Generator Priors for NAS-Based Gradient Inversion** *Gao Yali*
14:15	**[947] BRISK-KV: Reuse-Aware Binary Retrieval for Efficient Long-Context LLM Inference** *Xue Weiru*	**[467] CausalEvo-Agent: Stability-constrained Causal Self-evolution for LLM-based Agent** *Zhu Tiansheng*	**[229] Bridge Alerts for Institutional Portfolio Similarity Networks: An Explainable Monitoring Framework** *Shen Yinru*	**[668] Synthetic Data-Driven Segmentation of Sparse Linear Structures in Rock CT Images** *Usefi Hamid*
14:30	**[34] LWCAL: Loss-Weighted Calibration for Tabular Classifiers with Noisy Calibration Labels** *Liu Zeming*	**[189] TerraProbe: A Layered-Oracle Framework for Detecting Deceptive Fixes in LLM-Assisted Terraform Security Repair** *Nebolisa Chimdumebi*	**[797] When Is Another Retrieval Round Worth the Cost? Cost-Aware Depth Control for Iterative RAG** *Yang Zhekai*	**[895] Beyond Grounding: Bias-Aware Selective LLM Question Answering over Biased Records** *Esfar E Alam A M*
14:45	**[202] Super-Resolution of Solar Magnetograms via Adaptive Stratified Ensemble Learning with Uncertainty Estimation** *Norouzi kandalan Sina*	**[190] OrderBench: Benchmarking Cleanup-on-Exception in LLM-Generated Code** *Thayyil Suhaan*	**[848] Ulysses-MTRAG: A Synthetic Multi-Turn Benchmark for Retrieval-Augmented Generation over Brazilian Legal Text** *Dutra Guilherme*	**[680] HQTN-SER: Speech Emotion Recognition with Hybrid Quantum Tensor Networks** *Innan Nouhaila*
15:00	**[671] CodecBackdoor: Compression Residuals as Backdoor Carriers for Deepfake Detection** *Yuan Xiaohui*	**[267] Continuous Retrieval-Grounded Reward Design for Secure Code Generation on Small Language Models (CARGO)** *Siwali Suryansh*	**[860] When Does Retrieval Still Matter? Operating Regimes for Full-Document Reasoning: Evidence from ContractNLI** *Mutchi Sai Teja*	**[506] Detecting Student Misalignment through the Analysis of Interactions with an Educational Chatbot** *Segarra Montesinos, María Teresa*
15:15	**[321] Two-Sided Conformal Evidence Sets for Retrieval-Augmented Generation** *Cardei Mihaela*	**[831] Structured Optimization-Based Dynamic Window Approach with Predictive Relative-Motion Risk** *Hasan Md Zahid*	**[563] Label-Semantic-Guided Token Refinement for Speech Emotion Recognition** *Chen Haozhe*	**[569] CPQG: Generating Complete Educational Items with Reasoning-Path Fine-Tuning and Constraint-Guided Generation** *Wu Zhenyu*
15:30	**Coffee Break**			**[715] Latency Is Surprisal: An Exact Spiking Primitive for Local, Self-Training Prediction** *Zandian Jazi Nikan*
	14A (full papers) **Chair: Gonzalo Vivian** **Room: Royal Palm 1**	**14B: (full papers)** **Chair: Michail Alexiou** **Room: Royal Palm 2**	**14C : (full papers)** **Chair: Minghan Wei** **Room: Royal Palm 3**	**14D (short papers)** **Chair: Andy Sinclair** **Room: Royal Palm 4**
15:45	**[493] Combining Foundation Model Confidence and Monocular Depth for Training-Free Out-of-Distribution Segmentation** *Varghese Serin*	**[352] LLM-AFTQ: An LLM-Assisted Framework for Evaluating Test Quality of Code Generation Benchmarks** *Li Zikai*	**[806] SCOPE: Decoupling Condition Scaffolds from Diffusion Targets for Probabilistic Traffic Imputation** *Li Xinjin*	**[620] Multimodal AI for Expert-Supervised Graduate Thesis Quality Assurance: Evidence from a Decision-Support Study** *Liu Shenling*
15:55				**[734] FoCoKT: Knowledge Tracing with Bounded Forgetting and Long-Term Consolidation** *Cai Jiaming*
16:00	**[500] RAGScope: A Leakage-Controlled, Cost-Aware Evidence-Gating Protocol for RAG Hallucination Triage** *Liu Zeming*	**[464] Specification-Guided Multiagent AI for Context-Governed Software Life Cycle Automation and Repository-Aware Software Generation** *Alves de Lima, Vitor Mesaque*	**[280] Adaptive Low Light Preprocessing for End-to-End Robotic Perception and Grasp Planning on Edge Devices** *Rahman Md Mostiur*	**[93] NormBench: A Two-Stage Benchmark for Measuring Compliance and Regulatory Behavior in Local LLM Societies** *Manik, Md Motaleb Hossen*
16:15				**[748] Surrogate-Assisted Quantum-Inspired Cellular Genetic Algorithms for Accelerating Mobility Management Optimization** *Hamitouche Monsef*
16:15	**[606] KCP: Kernel-Smoothed Conformal Prediction for Risk-Conditional Coverage Balance in Biometric Verification** *Li Weibang*	**[773] DAISY: Fault-Guided Assertion Inference for Dafny** *Martins Ruben*	**[431] Energy-Aware Soft Actor-Critic-Based Cooperative Multi-Robot Coverage Path Planning** *Gao Tiying*	**[676] A General Vision--Language Framework for Cellular Toxicity Concentration Grading from Phase-Contrast Microscopy** *Song Yuqi*
16:35				**[925] VIROPLEX: A GNN-Driven Assessment of the Pandemic Potential of Zoonotic Bat Viruses** *Agnihothram Arjun*
16:30	**[762] Assessing Dataset Shift in Early Attorney Involvement Prediction from Adjuster Notes** *Vivian Gonzalo*	**[897] Towards Understanding Graph Neural Networks for Functionality Recognition in Binary Programs** *Alexiou Michail*	**[508] Parameter-Efficient Continuous Trajectory Prediction for Vision-Language Navigation** *Jarraya Imen*	**[928] Quality-Aware Neural Pupil Tracking and Signal Reconstruction for Dynamic Pupillometry** *Violeau Pierre*
16:45	**[933] Selective Next Location Prediction with a Lightweight Post-Hoc Confidence Metamodel** *Jary Calvin*	**[251] Recoverability-Aware Retrieval Operator Auditing: From Candidate Visibility to Ranking Accessibility** *Zhang Guoqing*	**[685] MARCC: Scalable Multi-robot Coverage with Connectivity Maintenance Via Deep Reinforcement Learning** *Corrieri Jacomo*	**[151] Explainable Solar Flare Prediction Using Deep Learning and Large Language Models** *Wang Zhenduo*
16:55				
17:00	**[603] Detecting LLM-Assisted Vietnamese Writing via Keystrokes under Behavioral Manipulation** *Dong Thanh*	**[283] Dependency-Guided Graph-based Retrieval-Augmented Generation** *Li Xinyu*	**[356] MIFLOW: Dynamic Recent-Behavior Priors for Reference-Calibrated Online Flow-Policy Reinforcement Learning** *Wang Chenhao*	**[800] On Evaluating Quantum Kernel Robustness for Low-Resource Cross-Corpus Audio Deepfake Detection** *Al Amin Lisan*
17:05				
17:15	**[241] Per-Layer-Group Bottlenecking for Cross-Domain Audio Deepfake Detection** *Abri Faranak*	**[807] Signal2Symbol: Neuro-Symbolic Temporal Reasoning for Explainable Physiological Time-Series Anomaly Detection** *Benabderrahmane Sidahmed*	**[781] ChunkVLA-AM: Parallel Action Chunking for Vision-Language-Action Robot Control in Additive Manufacturing** *Liu Zhugang*	**[827] Supporting Student Decision-Making in Personalized Academic Curriculum Balancing using an Evolutionary Algorithm** *Godoy Barrera Pedro*
17:30	**Break**			
17:45	**Reception** **Room : Grand OASIS Ballroom**			
19:00	**End of the day**			

Tuesday, November 3 - Day 2

7:30 — Registration (till 5:00pm) and Continental Breakfast

8:00 — Daily Remarks
Room : Grand OASIS Ballroom

8:15 — Keynote Speaker : Dr. Ling Liu
Georgia Institute of Technology
Room : Grand OASIS Ballroom

9:15 — Coffee Break

ROYAL PALM 1

Session 21A (full papers)
Chair: Miltos Alamaniotis
Room: Royal Palm 1

9:30
- [239] Event-Native Symbolic-Temporal Spike Encoding Framework for Heterogeneous Cyber Streams — *Diez Dalton*

9:45
- [256] Towards Precise Forgery Localization via Dynamic Boundary-aware Progressive Feature Fusion — *Li Qingyan*

10:00
- [289] LLM-Triage for Real-Time Smishing Detection in Financial IoT — *Jurkuch Abit*

10:15
- [348] An Interpretable AI Decision-Support Framework for N-Level Hardware Trojan Testing Under Strategic and Behavioral Uncertainty — *Nan Satyaki*

10:30
- [936] HARP-LoRA: Hierarchy-Conditioned Rank Projection for Low-Rank Adaptation in Scientific Table Understanding — *Sholehrasa Hossein*

10:45 — Coffee Break

Session 22A (full papers)
Chair: Miltos Alamaniotis
Room: Royal Palm 1

11:00
- [349] SABRE: Stream-Adaptive Balanced REplay for Class Incremental EDR Alert Classification under Class Imbalance — *Marasco Isabella*

11:15
- [379] Categorical Features Don't Lie: Detecting Poisoning in Gradient-Boosted Tree Ensembles — *Colaco Valency Oscar*

11:30
- [406] Artifact-Aware Spectral Amplitude Randomization: Zero-Overhead Regularization for Lightweight, Generalizable Forgery Detection — *Zhang Xuecen*

11:45
- [411] Prompt Injection Detection for Email Agents Through Attack Chain Modeling — *Yin Yunting*

12:00
- [669] Agentic AI for Distributed Cognitive Surveillance (DCS) of Gamma-Ray Spectra in Nuclear Security — *Alamaniotis Miltiadis*

12:15
- [475] Secured Agentic Multimodal Orchestration for Zero- Touch 6G Network Containment — *Yadav Sahil*

12:30 — Lunch
Room : Grand OASIS Ballroom

23A (full papers)
Chair: Sareh Taebi
Room: Royal Palm 1

13:45
- [523] StegaQR: Learned Private Data Embedding in Standard-Decodable QR Codes — *Alghamdi Mohammed*

14:00
- [578] Boosting Transferable Adversarial Attacks on Vision Transformers via Spatial-Frequency Guided Patch Transformation and Deviation-aware Momentum Token Embedding — *Cao Bo*

14:15
- [648] Parameter-Efficient Open-Set Facial Deepfake Detection with Frozen DINOv3 Features — *Gomes Torres Barbalho Costa*

14:30
- [661] Backdoors Leave Structural Traces: FedMAST for Backdoor Detection and Containment in Federated Learning — *Subramanian Srinivasan*

14:45
- [235] *Learning to Navigate in Stochastic and Uncertain Environments Ensemble Learning with Uncertainty Estimation* — *Choton Jahid Chowdhury*

15:00
- [690] Unveiling the Unseen: Exploring Membership Inference Leakage via the Lens of Explainability — *Tabassoum Nafiza*

15:15 — Coffee Break

24A (full papers)
Chair: Huanjing Wang
Room: Royal Palm 1

15:30
- [938] Calibrated Membership Inference Auditing of NLP Text Classifiers — *Abusaqer Muhammad*

15:45
- [296] CausalFairAgent: A Causal-Structure-Learning Framework for Trustworthy Evaluation and Automatic Bias Repair in Educational Agents — *GUO Meixia*

16:00
- [316] Hidden Dependencies: Intersectional Bias in Text-to-Image Diffusion Models — *Liebrand Nina*

16:15
- [337] When Is In-Game AI Welcomed or Rejected? A Steam Review-Based Analysis of User-Perceived AI Characteristics and Game Recommendation Behavior — *Joung Yeonwoo*

16:30
- [458] Emotionformer: Transformer with Gaussian Range Encoding for Emotion Recognition Based on Short Keystroke Sequence — *Nabila Mansouri*

16:45
- [465] Auditing Socioeconomic Concentration in Automobile Repossession Prediction — *Sinclair Andy*

17:00
- [582] When Correlation Isn't Enough: Using Causal Methods for Discriminatory Bias Detection — *Padela Joyce*

17:15
- [597] AutoCoRe-Select: Verifiable LLM-Guided Potential-Based Reward Shaping for Cooperative Multi-Agent Reinforcement Learning — *Bai Jinyin*

17:45 — Break

18:00 — Panel : AI Applications in Government: Including the Virtues of Western Open-Weight Models
Moderator : Van Hipp
Panelists : Brion Bennett, Hamlet Yousef, John Alden Haynes, Mike Lu, Daniel Ratner
Room : Royal Palms (Combined Rooms)

19:00 — Dinner and Awards
Dr. Gregg B. Fields, Vice President for Research at FAU, Dinner Speaker
Room : Grand OASIS Ballroom

21:00 — End of the day

ROYAL PALM 2

Session 21B (full papers)
Chair: Andy Sinclair
Room: Royal Palm 2

9:30
- [921] SpaRL: Reinforcing Spatial Layout in Language Models for Structured Diagram Generation — *Steinigen Daniel*

9:45
- [922] Dynamic Retrieval-Augmented LLM-based Multimodal Sarcasm Detection — *Longze Fan*

10:00
- [378] Improving Constraint Models with LLM Agents — *Szeider Stefan*

10:15
- [166] Surviving Sensor Failure with Modality-Robust Depth Completion — *Aiersilan Aizierjiang*

10:30
- [253] Towards Brain-Guided Person Search: A Pilot Study on EEG--Visual Alignment — *Jiang Fanzhi*

Session 22B (full papers)
Chair: Andy Sinclair
Room: Royal Palm 2

11:00
- [338] MonoPeakMatch: Projection-Guided Native-Peak Matching for Monocular 3D Car Detection — *Sun Yongqing*

11:15
- [492] Unsupervised Point Cloud Semantic Segmentation Based on Point Shuffle Attention — *Qin Chiquan*

11:30
- [519] Balancing Global and Local: A Lightweight Dual-Branch Network for Low-Light Image Enhancement — *Wang Zhongyuan*

11:45
- [542] ClearSplat: Reliability-Guided Degradation-Decoupled Gaussian Splatting for Adverse-Weather 3D Vision — *Yao Tenghui*

12:00
- [581] DecTD: Decoupled Transformer Decoding for End-to-End Container Code Spotting — *Yang Jingyi*

12:15
- [615] SVGT: Scene Visual Geometry Transmorfer — *Taoui Tarik*

23B (full papers)
Chair:Xiaohong Yuan
Room: Royal Palm 2

13:45
- [651] GridVCN: VLM-Generated Four-View Grids for Object Classification — *Dadzie Frank*

14:00
- [851] Learning Violence Dynamics in World Coordinates: Metric 3D Human Motion for Violence Detection — *Yuan Xiaohong*

14:15
- [274] Skeleton-Aware Medical Image Segmentation via Sinkhorn Optimal Transport — *Cheng Hong*

14:30
- [496] Stochastic k-Simplex Embedding in High-Dimensional Cancer Classification — *Curtis Ethan*

14:45
- [562] FAG-MAE: Feature-Aware Graph Masked Autoencoding for Protein Surface Representation Learning — *Kamel Madi*

15:00
- [639] OncoTwin: Pathway Aware Digital Twins for Cross-Domain Survival Prediction and Virtual Treatment Simulation — *Islam, Dewan*

24B: (full papers)
Chair: Yifan Yu
Room: Royal Palm 2

15:30
- [694] Double-Decoder U-Net with Geometric Priors for Bone Segmentation — *Frikel Youssef*

15:45
- [735] Cross-Attentive Multi-Task Learning for Diabetic Retinopathy Grading and Lesion Segmentation — *Zhou Xiang*

16:00
- [763] PACE: Parameter-Matched Deep Learning for Continuous ECoG Decoding and Cross-Patient Transfer — *Yu Yifan*

16:15
- [828] Multi-Representation Facial Feature Fusion for Cross-Domain Infant Distress Recognition — *Shah Mohammad Akbar Bin*

16:30
- [855] Retinal Digital Twin via Interpretable AI for Forecasting Diabetic Retinopathy Progression — *Agarwal Vedaant*

16:45
- [892] How much does transfer learning help personalized glucose forecasting, and what limits it? — *Fulgencio da Cunha Menezes Beatriz*

17:00
- [902] ARMED: Adaptive Retrieval-Augmented Medication Recommendation — *Chung Kyungkeon*

17:15
- [956] LinguaX-PAAD: Pure Acoustic Hierarchical Attention for Cross-Lingual Alzheimer's Disease Detection — *Liu Xian-Xian*

ROYAL PALM 3

Session 21C (short papers)
Chair: Sonia Moshfeghi
Room: Royal Palm 3

9:30
- [343] Trustworthy AI Requires Trustworthy Benchmarks: Version Drift and Syntactic Sensitivity in Self-Referential LLM Code Hallucination — *Imran Hasan Md*

9:40
- [390] Evaluating LLM Safety: Can Models Interpret Emerging Slang — *Albujasim Zainab*

9:50
- [405] Do Multimodal LLMs Understand Misleading Visualizations? Evaluating Visualization Literacy and Reasoning Errors in Chart Interpretation — *Tisha Sadia Nasrin*

10:00
- [427] Won by a Single Question: Item-Level Attribution and the Robustness of LLM Leaderboard Decisions — *Mehta Madhur*

10:10
- [746] Poltard: Governed Semantic Execution Trees for Framework-Independent AI Agent Control and Replay — *Sfaghi Marius*

10:20
- [882] LIMBO: Lifelong Inference-Time Memory and Budget Optimization for LLM Agents — *Sharma Siddharth*

10:30
- [853] Hybrid Quantum-Classical Deep Learning for Pneumonia Detection from Chest X-Ray Images — *Salma Fahmina Nur*

Session 22C (short papers)
Chair: Sonia Moshfeghi
Room: Royal Palm 3

11:00
- [307] Priority-Aware Grid Sensing for Scalable Heterogeneous Multi-Robot Coordination — *Shawon Rehab Uddin*

11:10
- [952] Internet of Building Blocks: A Graph-Grounded Multi-Agent Orchestration Framework for Bi-Objective (Cost, CO_2) Construction Optimization on Earth and Tri-Objective (Cost, Mass, Lifespan) Optimization in Outer Space — *Heidari Negin*

11:20
- [707] Co-Evolving Policy and Reward for Code Generation: A Mechanistic Study of Single-Adapter Drift — *Tomuro Noriko*

11:30
- [811] RefTestModel: Catalog-Guided Detection of Test-Specific Refactorings with Foundation Models — *Gheyi Rohit*

11:40
- [551] From Text to Visual Images: A Scalable Multimodal LLM for Math Question Answering over Heterogeneous Data — *Ng Yiu-Kai*

11:50
- [573] Factual Inconsistency in Chart Captioning: Error Analysis and Two-Stage Fine-Tuning for MLLMs — *Kim Jong-Kyou*

12:00
- [608] Creating Hybrid Interactive Narrative Systems with LLM: From Formal Description to Evaluation — *Clerc Jules*

12:10
- [757] Improving LLM-Based OCR Through Document Segmentation and Semantic Prompting — *Jiang Britny , Miller Chreston*

23C (short papers)
Chair: Miltos Alamaniotis
Room: Royal Palm 3

13:45
- [759] Quality-Feedback Cognitive Memory RAG for Long-Text Question Answering — *Fan Wenyan*

13:55
- [872] Continue, Stop, or Confirm? Cost-Aware Intent-Bounded Response Control for Large Language Models — *Yu Chen-Yeou*

14:05
- [257] Multi-Prior Guided Image Editing for Architectural Sketch-to-Photorealistic Generation — *Zhao JiaYu*

14:15
- [691] FaunaCodec: ROI-Aware Video Compression and Reconstruction for Wildlife Monitoring — *Sakevych Mykhailo*

14:25
- [753] Budget-Aware Learned Measurement Allocation for Low-Shot Shadow Distillation — *Stenzel Gerhard*

14:35
- [133] EEFML: An Explainable Multi-Domain Stacked Ensemble for Chest X-Ray Classification under Small-Data Constraints — *Almasi Naghash Omid*

14:45
- [277] Multi-Label Classification in Brazilian Gastroscopy: Clinical Co-occurrence, Shortcuts, and Explainability — *Ferreira Santos Willgnner*

14:55
- [544] Localizing Discriminative Retention-Time Regions in LC–MS MS1 Signals via Contrastive Representation Learning — *Phillips Joshua*

15:05
- [812] Quantum FM2I: An Empirical Study of a Quantum-Inspired Gramian Encoding for Forecasting by Image Inpainting — *Noufel Saad*

24C (full papers)
Chair: Kehan Gao
Room: Royal Palm 3

15:30
- [654] Transferability of SHAP Feature Rankings Across Classifiers in Credit Card Fraud Detection — *Liang Qianxin*

15:45
- [796] Explaining Knowledge Utilization in Heterogeneous GNNs for 3D Semantic Scene Graphs — *Stelzner Tim*

16:00
- [871] Causal Head Selectors Agree with Each Other but Not with Attention Heuristics — *Dewage Kasun*

16:15
- [910] Auditing Accuracy, Explanation Fidelity, and Robustness for Trustworthy UAV Power Prediction — *Le Thu*

16:30
- [156] Exact Privacy-Preserving Similarity Search via Two-Server Additive Secret Sharing — *Alsulami Ameerah*

16:45
- [203] Structural Admissibility Certificates for Secure State Transitions in Tool-Using LLM Agents — *Osmani Aomar*

17:00
- [596] Validator-Gated Local SLMs for Offline Safety-Policy Translation — *Dadi Sriram Chowdary*

17:15
- [631] The Cost of Caution: Trading Over-Inference Against Information Loss in an LLM Prescription-Intent Classifier — *Garcia Lusk Maria*

17:30
- [130] Asynchronous Byzantine-Robust Multi-Agent Policy Gradient with Staleness-Aware Centered Clipping — *Cao Junjie*

ROYAL PALM 4

Session 21D (short papers)
Chair: Michail Alexiou
Room: Royal Palm 4

9:30
- [196] True/False at Scale: An AI-Assisted Tool for Building Educational T/F Question Datasets — *Dutulescu Andreea*

9:40
- [504] TRaCe: A Translate–Rephrase–Check Pipeline for Enhancing Low-Resource Programming Language Support in LLMs — *Dutulescu Andreea*

9:50
- [688] MMEE-GCP: A Stage-wise Multimodal Event Extraction Method for Chinese Military News — *Hu Weijian*

10:00
- [273] EchoNav: Embodied Audio-Visual Memory and Language-Guided Reasoning for Emergency Robot Navigation — *Cheng Wen-Chung*

10:10
- [834] EmotiQ: A Quantum Residual Specialist Framework for Text Emotion Classification — *Yin Yunting*

10:20
- [888] Forget to Protect: A Machine Unlearning Framework for Responsible Artificial Intelligence — *Asadinia Marjan*

10:30
- [252] VERA: Authority-Preserving Edge Revocation for Federated AI-Agent Workflows — *Yu Haoran*

Session 22D (full papers)
Chair:Xiaohong Yuan
Room: Royal Palm 4

11:00
- [134] Beyond the Snapshot: A Longitudinal Study of AI Agent Integration in Open-Source Software Development — *Ahuja Vinod*

11:15
- [683] LiteCMT-Net: Lightweight Structural-Contextual Modeling for Crack and Pothole Segmentation — *Hu Weijian*

11:30
- [743] TSF-Former: Temporal-Statistical Token Fusion for Road Distress Recognition with Multi-Sensor Vehicle-Motion Signals — *Hu Weijian*

11:45
- [594] Cost-Aware Post-Hoc Deferral Under Calibration and Shift: An Environmental AI Case Study — *Yu Haoran*

12:00
- [693] Beyond Final-Answer Accuracy: Paired-Protocol Evaluation of LLM Graph Reasoning with Multiple Optima — *Yuan Xiaohong*

12:15
- [188] Equivalent-State Q-Learning: Non-Local Credit Assignment for Mixed-Quality Robotic Demonstrations — *Acs Matthew*

Wednesday, November 4 - Day 3

7:30	Registration (till 5:00pm) and Continental Breakfast				
8:00	Daily Remarks Room : Grand OASIS Ballroom				
8:15	Keynote Speaker : Dr. Jerry Gao San Jose State University Room : Grand OASIS Ballroom				
9:15	Coffee Break				
	ROYAL PALM 1		**ROYAL PALM 2**		**ROYAL PALM 3**
	Session 31A (full papers) Chair: Kehan Gao Room: Royal Palm 1		Session 31B (full papers) Chair: Huanjing Wang Room: Royal Palm 2		Session 31C (short papers) Chair: Sonia Moshfeghi Room: Royal Palm 3
9:30	[453] DREAM: Delay-aware Reconstruction and Alignment Model for Tobacco Drying Outlet Moisture Forecasting *Li Xiao*	9:30	[121] Optimized Track Completion in Railway Stations Using Constraint Programming *Löffler Sven*	9:30	[173] SWIFT: Specification-Driven Flow for Smart-Home Web of Things Applications *Euripides Petrakis*
				9:40	[816] GRAM-HAR: Grammar-Aware Learning of Latent Motion Primitives for Complex Occupational Activity Recognition *Bouchekif Djalal*
9:45	[652] A review of AI-enabled integrated data management in IoT-Fog-Cloud Environments *Kobti Ziad*	9:45	[377] Components of information flow underlying heuristic decisions in constraint satisfaction search *Wallace Richard*	9:50	[142] ERDNet-TE: A Spatiotemporal Graph Neural Network for Enterprise Risk Prediction with Transfer Entropy and Dynamic Industry Features *Zhang Weichen*
				10:00	[534] ARCADE: Adaptive Reasoned Composition for Anomaly Detection in Graphs *Benhalima Ahmed Youcef*
10:00	[742] CPTR-TabM: Physics-Embedded Cross-Packet Token Restoration for IPMSM Inductance Identification *Zhu Chenyi*	10:00	[554] One Heuristic Across Trade-offs: Evolving Preference-Conditioned Heuristic Programs with Large Language Models *Zhong YiHao*	10:10	[319] Real-World Field-Service Vehicle Routing through Geospatial Clustering and Combinatorial Optimization *Domingues João*
10:15	[775] Explainable Deep Learning for Martian Terrain Classification: A Comparative Study of CNN and Transformer Models *Gao Kehan*	10:15	[579] Accelerating ZDD-Based Solvers for Independent Set Reconfiguration via Path Decomposition *Ito Takehiro*	10:20	[802] SparseTemporalPIE: Cross-Attention over Sparse Temporal Context for Pedestrian Crossing Intention Prediction *Moradi Hamidreza*
10:30	[803] Sensor-Permutation Equivariant Mamba for Indoor Temperature Forecasting *Wang Yi*	10:30		10:30	[334] Localized Spatiotemporal Trajectory Prediction for Heterogeneous Multi-Agent Traffic *Jang Jinwoo*
10:45	Coffee Break				
	Session 32A (full papers) Chair: Andy Sinclair Room: Royal Palm 1		Session 32B (full papers) Chair: Sonia Moshfeghi Room: Royal Palm 2		Session 32C (short papers) Chair: Minghan Wei Room: Royal Palm 3
11:00	[819] When Features Stay Behind: Stress-Testing Feature-Mismatched Soybean Yield Prediction Across Environments *Ghanaatian Amirhossein*	11:00	[712] Accelerating Discrete Diffusion Models for Combinatorial Optimization via Momentum Inference *Azizi Bessam*	11:00	[903] T-SMART: Mechanism-Level Attribution for Tool-Augmented Time-Series Question Answering *Delgado Ivan*
				11:10	[545] Autonomous Manipulation System for Seaweed Harvesting in Tank Aquaculture *Wei Minghan*
11:15	[125] HDForecast: Hyperdimensional Inherent Representations for Time Series Forecasting *Morris Justin*	11:15	[755] Local Rules, Global Rectangles: An Efficient SAT Encoding for Shakashaka *Berton Carl*	11:20	[738] EI-MCTS: Dynamic Information Map-Guided Decentralized Monte Carlo Tree Search for Multi-UAV Wildfire Search *Cai Xinlong*
				11:30	[751] Language-Guided Feature Anticipation for Goal-Conditioned Trajectory Prediction *Holzmann Dennis*
11:30	[222] Exact Incremental Maintenance of Delaunay-Graph Clustering for Dynamic Spatial Data *Purushothama Sushanth*	11:30	[840] Solving Constraint Optimization Problems Using Learning-Based Techniques *Ouali Abdelkader*	11:40	[682] When Does Closed-Loop Curriculum Adaptation Help GRPO? An Analysis on Countdown *Pant Yuv Raj*
11:45	[362] AusAML: A Synthetic Australian Bank Transaction Benchmark for AI-Driven Money Laundering Detection *Kute Dattatray*	11:45	[898] Distributed Conflict-Based Search for Multi-Domain Virtual Network Embedding *Zheng Yi*	11:50	[729] Context-Aware Exploration in Reinforcement Learning Using Large Language Models *G V, Vijay Raj*
12:00	[387] In-Context Learning for Foundation Time-series Models with Prior-Data Fitted Networks *Vu Long*	12:00	[300] The Operating Room as a Knowledge Graph: Symbolic Reasoning for Surgical Flow Optimization *Ben Othman Sarah*	12:00	[731] Partial-Solution-Guided Heatmap Refinement for Neural TSP Solvers *Liu Zeyu*
12:15	[625] RMG-TTA: Residual-Memory-Guided Test-Time Calibration for Non-Stationary Day-Ahead Load Forecasting *Tan Jingjing*	12:15	[345] A Doxastic Logic for Boolean Networks *Rosenblueth David A*	12:10	[769] Symmetry Detection and Breaking in Cryptographic XNF Instances *Blomme Anthony*
				12:20	[318] From Authentication Log Anomalies to Short-Term Attack-Risk Forecasting: An AI Framework for Proactive Cyber Defense *Moghadam Armin*
12:30	Lunch Room : Grand OASIS Ballroom				
	33A (full papers) Chair: Sonia Moshfeghi Room: Royal Palm 1		33B (full papers) Chair: Andy Sinclair Room: Royal Palm 2		33C (short papers) Chair: Kehan Gao Room: Royal Palm 3
13:45	[626] TRIO: Tri-Axial Decoupled Alignment of Temporal, Spatial, and Cross-View Representations for Multivariate Time-Series Forecasting *Du Wenjun*	13:45	[924] ReGSS: Regional Graph State-Space Modeling for EEG Abnormality Detection *Yang Baoyao*	13:45	[813] On the Use of Random-Walk Planning *Chrpa Lukas*
				13:55	[886] Unsupervised Graph Neural Networks with Multi-Head Latent Attention for the Traveling Salesman Problem *Liu Shenyang*
14:00	[665] RAMAT: A Regime-Aware Memory-Augmented Transformer for Unsupervised Anomaly Detection in Financial Time Series *Jiang Zhuocheng*	14:00	[419] KDA: A Cross-Modal Knowledge-Driven Alignment Framework for Unstructured Engineering Texts and Parametric Spatial Models *Peng Wenjie*	14:05	[899] Solving versus Verifying: Catching Contradictions in Tax Reasoning Systems *Chudziak Jaroslaw*
				14:15	[351] ToMyRiS: Tsetlin Machines for Interpretable Embedding-based Recommender Systems *Borgersen Karl Audun*
14:15	[777] Target-Queried Channel Attention: Measuring and Exploiting Cross-Channel Context in Multivariate Time-Series Forecasting *Rahman Quazi*	14:15	[896] ARGUS: Text-to-SPARQL Natural Language Interfaces with Embedding-Guided Grounding and Self-Healing Generation *Kona Srividya*	14:25	[658] Consistent Aggregation of Qualitative Preferences with Maximal Intent Preservation *Qi Zihao*
14:30	[862] When 10,000 Windows Are Not 10,000 Tests: Auditing Statistical Confidence in Sliding-Window Time-Series Classification *Shi Xinze*	14:30	[901] Combining Relevant Expert Knowledge Using Mixtures of Feedback Fuzzy Cognitive Maps *Adigun, Olaoluwa*	14:35	[547] A Neuro-Symbolic Verification Loop for 5G-AKA: From LLM-Extracted Requirements to Solver-Checked Repairs *Dalai Binay*
14:45	[346] Uncertainty-Aware Vehicle Routing Under Correlated Travel-Time Variability: A GPS-Derived Experimental Study *Jang Jinwoo*	14:45	[904] Semantic Objective Compression for AI-Guided Engineering Design *Spence Cameron*	14:45	[589] Does a Quantum Kernel Need Its Entanglement? A Reusable Test, with a ZZ Case Study across Clinical Data *Hidila Zineb*
15:00		15:00	[931] POLARIS: Patterns over Local Cohorts for Auditable Recommendation in Strict Cold Start *Gendarme Steve*	14:55	[701] Towards Mitigating Fabricated Consensus: The Active Provenance Gate for Multi-Agent Debate Synthesis *Chudziak Jaroslaw*
15:15	Coffee Break				
	34A (full papers) Chair: Huanjing Wang Room: Royal Palm 1		34B (full papers) Chair: Sonia Moshfeghi Room: Royal Palm 2		34C (full papers) Chair: Kehan Gao Room: Royal Palm 3
15:30	[360] β-Light: A Spatiotemporal Beta Policy for Traffic Signal Phase and Duration Control *Yang Baoyao*	15:30	[494] OMoE-Net: Orthogonally Disentangled Mixture-of-Experts Network for Image Restoration under Low-Level Degradations *Tao Haojie*	15:30	[389] Static Analysis of Logic Programs through Boolean Networks *Benhamou Belaid*
15:45	[469] CARE:Context-aware Feature Restoration for Weakly Supervised Instance Segmentation under Adverse Weather *Yang Baoyao*	15:45	[537] Route Fine, Explain Compactly: A Soft-Tree-Gated Mixture-of-Experts Surrogate *Folino Francesco*	15:45	[555] Edge-Optimized Drone Acoustic Detection: A Comparative Study of Machine Learning, Deep Learning, and Domain Adaptation *Kona Srividya*
16:00	[674] UMSCL-POI: User-Aware Multimodal Fusion, Spectral Enhancement, and Multi-View Contrastive Learning for Next POI Recommendation *Zhang Yan*	16:00	[700] DMEN: A Dynamic Weighting Multi-Expert Synergistic Network for Class-Imbalanced Facial Expression Recognition *Li Haitao*	16:00	
16:15	Break				
16:30	Closing Session Room : Grand OASIS Ballroom				
17:00	End of the program				

Keynote Speaker: Gary S. Connor

Title: Engineering the Decision: What Comes After Intelligence?

Abstract:

Artificial intelligence is rapidly increasing our capacity to generate analysis, hypotheses, predictions, recommendations, and courses of action. But more intelligence does not necessarily produce better decisions. As models, agents, data, and human expertise proliferate, an emerging challenge is not simply how to produce intelligence, but how to evaluate, challenge, reconcile, and translate competing intelligence into consequential action under uncertainty and time.

This keynote examines the evolution from traditional human decision-making to quantitative methods and increasingly abundant machine intelligence. It introduces the **Decision Gap**: the growing challenge of converting intelligence into timely, adaptive decisions as information accelerates, conditions change, and actions themselves alter the environment in which subsequent decisions must be made.

The talk introduces **Quantitative Decision Science** as an emerging field of inquiry centered on decision quality and explores the need for **Decision Architecture** capable of integrating human judgment, quantitative methods, and machine intelligence. The central proposition is simple: as we engineer increasingly intelligent systems, the next great engineering challenge may be the **decision itself**.

Bio:

Gary Connor brings over 31 years of experience in program management and senior executive leadership in the U.S. Air Force. Gary's last executive assignment was the Deputy Program Manager for the $8 billion per year, Ballistic Missile Defense System, Missile Defense Agency (MDA), Huntsville, Alabama. Since retiring from the Air Force in late 2009, Gary serves as an Aerospace Industry consultant, providing strategy and implementation advice to a diverse array of Federal, Defense and commercial technology endeavors. A recognized and well-respected "systems thinker", Gary offers specific, in-depth expertise in Large-scale Program Management, Enterprise Management, Strategic Planning, Logistics Transformation (such as performance-based logistics) and Complex Systems Integration. Since December 2019 he serves as Chairman of the Board of Directors for nLogic in Huntsville, and he's an advisor to several small technology companies in the area.

Gary's senior executive management experience focused on leading multi-billion dollar enterprises and thousands of people to develop and field fighter, missile defense, command & control, and intelligence/surveillance/reconnaissance systems globally for the US, NATO and our Allies. Gary's a seasoned leader and experienced manager with superior communications, organizational, execution, international relations and people skills. His other notable senior leadership positions include: Program Director, Ground-based Midcourse Defense, MDA; Deputy Chief of Staff for Communications, Multi-National Force, Iraq; Commander, Command & Control (C2) Intelligence Surveillance & Reconnaissance (ISR) and Battle Management Systems Wings, Hanscom AFB, Massachusetts, and; Director Reconnaissance Systems Program Office, Wright Patterson AFB, Ohio.

Gary earned his Bachelor of Science in Electrical Engineering and Master of Science in Computer Engineering from the University of Lowell, Lowell, Massachusetts. He also earned a master's degree in National Resource Strategy from the Industrial College of the Armed Forces, Ft McNair, Washington DC. Gary is a Life Member of both the Air Force Association (AFA) and the National Defense Industrial Association, and he is active in AFA events at the local and national levels.

Keynote Speaker: Dr. Ling Liu

Title: Efficiency and Robustness Challenges of Agentic AI Systems and Tools

Abstract:

Agentic AI, as the next evolution of Generative AI, represents a new breed of AI agent systems. Agentic AI tools aim to increase GenAI utility and productivity beyond multi-modal perception and reasoning by working autonomously with goal decomposition and multi-step agentic workflow execution, capable of completing tasks autonomously with minimal human supervision, and integrating with other software systems. Efficiency of Agentic AI requires continuous learning, effective goal decomposition, and autonomous scaling to accelerate multi-turn operations and optimize for future decisions. Robustness of Agentic AI requires trustworthy governance and evaluation to prevent policy violations and cascading effects of operational errors or drifts. This keynote will give an overview of technical challenges for Agentic AI efficiency and robustness through research projects currently being developed at my research lab.

Bio:

Ling Liu is a Professor in the School of Computer Science at Georgia Institute of Technology. She directs the research programs in the Distributed Data Intensive Systems Lab (DiSL), examining various aspects of Internet-scale big data powered artificial intelligence (AI) systems, algorithms and analytics, including performance, reliability, privacy, security and trust. Her research in the ML systems area is mainly centered on efficient AI systems and Algorithms, as well as trustworthy AI through developing AI security and AI privacy guardrails. Prof. Liu is an elected IEEE Fellow, a recipient of IEEE Computer Society Technical Achievement Award (2012), and a recipient of the best paper award from numerous top venues, including IEEE ICDCS, WWW, ACM/IEEE CCGrid, IEEE Cloud, IEEE ICWS. Prof. Liu served as the editor in chief of IEEE Transactions on Services Computing (2013-2016), the

editor in chief of ACM Transactions on Internet Technology (2019-2024), and currently the editor in chief of IEEE Transactions on Big Data (since 2025). Prof. Liu is a frequent keynote speaker in top-tier venues in Big Data, AI and ML systems and applications, Cloud and Edge Computing, Privacy, Security and Trust. Her current research is primarily supported by USA National Science Foundation, CISCO and IBM.

Keynote Speaker: Dr. Jerry Gao

Title: Agentic AI Test Automation for Smart Apps. - Intelligence, Agents, and Service Platform

Abstract:

According to recent market research, global AI-enabled testing market is projected to expand significantly to between $4.64 billion and $35.96 billion by the early 2030s. With the fast advance of AI and machine learning technologies, more and more AI-powered intelligent applications are developed to be ready for deployment. Now many people start to pay close attention to the quality of the developed AI-powered intelligent products in safety, security, accuracy, and reliability to make sure that we are building trustworthy and human friendly AI systems. In this talk, Dr. Gao first provides a comprehensive review on AI testing and automation for intelligent applications, including AI test issues, challenges, needs, test modeling, test generation, and automatic validation. Then, he discusses model-based AI app test modeling, test design and AI-based data augmentation, LLM-based test generation, as well as automatic test validation. Moreover, Dr. Gao shares his vision and insights about the roadmap of agentic AI testing and automation including agentic testing services, diverse AI-based agents, infrastructures, and the required service support platform. Finally, Dr. Gao will share an agentic AI testing platform to support AI app test automation.

Bio:

Jerry Zeyu Gao is a professor at the Department of Computer Engineering at San Jose State University. Now, he is the director of SJSU research center on Smart Technology, Computing, and Complex Systems. He had over 28 years of academic research and teaching experience and over 10 years of industry working and management experience on software engineering and IT development applications. He has published four technical books and over hundreds (>350) publications in IEEE/ACM journals, magazines, international conferences and workshops. His Google Scholar citation is over 10K+ and his ResearchGate reads is over 398K+. His current research areas include machine learning, smart cities, AI test automation

and agent-based AI testing, AI clouds and mobile clouds. In 2010, Jerry Gao has been recognized by University of Texas at Arlington as a distinguished alumnus for College of Engineering at its 50th anniversary. In 2011, he was award as a KSI Fellow in SEKE2011. In 2013, Dr. Gao has received the College of Engineering Faculty Award for Excellence in scholarship, Dr. Gao served as an editorial board member and an associate editor for numerous international journals in software engineering, machine learning, and smart cities, such as IEEE Software, Smart Cities, Energies, and Agriculture. Dr. Gao has been included and listed in Marquis American Who's Who 2020-2022.

Since 2010, Dr. Gao has played as one of leaders in organizing many international conferences and workshops as a conference co-chair, program co-chair, and workshop co-chair. Now, he is the steering board chair for IEEE CISOSE congress and IEEE Future Technology Summit, and the steering board member for IEEE Smart World Congress. IEEE CISOSE congress includes IEEE BigDataService, IEEEAITest, IEEE Intelligent Mobile Cloud, IEEE SOSE, and IEEE DAPPS. In addition, he served as a chair for many other conferences, such as IEEE CAI2025, IEEE AITest2025, IEEE BigDataService*, IEEE MobileCloud*, IEEE Smart World Congress 2017, Smart City Innovation 2017.

Since 2019, Dr. Gao established IEEE AITest as a co-founder, and started his pioneer research on AI Test Automation. Up to now, his group has published over 20 archived technical papers on AI testing and automation. Recently, he has delivered two IEEE tutorials on AI testing automation for different topics, and delivered several keynote presentations on Ai testing and automation. In addition, Dr. Gao has been invited by numerous international conferences to give his keynote speeches on other topics in Smart Cities, mobile AI cloud intelligence and services, UAV intelligence and applications, Green Energy Ai Cloud, Smart Agricultures, and Intelligent Wildfire Platforms.

Panel

AI Applications in Government: Including the Virtues of Western Open-Weight Models

Moderator: Van Hipp

Van Hipp is Chairman of American Defense International, Inc. He's a former Chairman of the South Carolina Republican Party, as well as a former member of the Presidential Electoral College. In 1990, he was sworn in as Deputy Assistant Secretary of the U.S. Army (Reserve Forces and Mobilization). In this capacity, he served as the Army Secretariat's "point man" for the successful mobilization, and then demobilization, of the Army's reserve forces for Operation Desert Shield/Storm.

Following the "Tailhook Scandal," Hipp was named as the Principal Deputy General Counsel of the U.S. Navy. Hipp served in this capacity until January 1993.

Hipp is a veteran of the U.S. Army and served on active duty in both Operation Desert Storm and Operation Restore Democracy. He is a recipient of the Bronze Star Medal. He continues to speak on defense issues at public forums across the country and writes a regular column for Newsmax on International Policy. His book, "The New Terrorism: How to Fight It and Defeat It," was published in February 2015. One hundred percent of all proceeds go to scholarships for the children of fallen Guardsmen and Guardsmen wounded in action in the War on Terror. All scholarships are being administered by the National Guard Educational Foundation (NGEF).

Since the September 11th attacks, Hipp has appeared on the Fox News Channel and Fox Business Channel well over 600 times as an expert commentator on the War on Terror and

has been a guest on most of the network's major news programming. In addition, he has appeared on MSNBC, CNN, the London-based Sky News Channel, Newsmax TV, and the CBS Evening News.

He currently serves as a member of the Board of Visitors of Charleston Southern University, as a member of the National Capital Salvation Army Advisory Board, and as a member of the Board of Directors of the Palmetto Promise Institute. In addition, Van Hipp was the recipient of the Salvation Army National Capital Area Command's 2015 "Compassionate Citizen Award."

On November 3, 2016 he presented a lecture entitled "The Greatest Challenge Of Our Next President: Keeping America Safe During These Times" at the Gerald R. Ford Presidential Museum in Grand Rapids, MI. He also served as moderator for two National Security Town Halls for President Donald Trump during the 2016 presidential election cycle. Recently, he has been elected to the Board of Trustees for the American Battlefield Trust.

Hipp received his bachelor's degree in Economics from Wofford College and earned his Juris Doctor from the University of South Carolina School of Law. In 2014, Hipp was the recipient of Wofford College's Distinguished Service Award. In addition, he was the recipient of the Queen Elizabeth II September 11th Garden Leadership Award for National Security in May 2018.

Van Hipp is married to the former Jane Grote of Nashville, Tennessee, and they have three children: Trey, Sarah Camille, and Jackson. Van and his family split their time between Alexandria, VA and Georgetown, SC.

Panel

AI Applications in Government: Including the Virtues of Western Open-Weight Models

Panelist: Brion Bennett

Brion Bennett is a healthcare technology executive and U.S. Naval Academy graduate with more than 20 years of leadership in information technology and medical device sales, business development, strategic partnerships, and federal contracting. He is Chief Executive Officer and Co-Founder of code626, Inc., in Fort Lauderdale, Florida, where he leads go-to-market for Cast, a first-to-market compiled programming language and runtime whose goal is to raise the standard of critical backend infrastructure in support of the Agentic Era.

Since August 2024 he has taken Cast to market across commercial and government channels. In 2026, code626 has booked multi-year contracts with annual recurring revenue across esports, biotechnology, and healthcare. Current agreements include a five-year exclusive partnership with World Champion Fantasy, its first private client, to serve as the backend for the PlayerX esports platform, with a per-paid-user revenue share and planned launches across the United States and Asia in 2026 and 2027. It also signed a Master Aggregator Agreement with Carahsoft, enabling Carahsoft and its reseller partners to offer Cast to federal, state, and local governments and educational institutions.

From 2016 to 2024 Bennett built ViiMed's (Hiive Health) public and private business in Washington, D.C. As Chief Growth Officer (2020–2024) he led sales, business development, partnerships, and communications, scaling the company from pre-revenue to $23 million ARR ($1 million in 2020, $4 million in 2021, $11 million in 2022, $23 million in 2023). He played a key role in the rapid configuration of the U.S. Coast Guard's Care Anywhere Virtual Health Platform, which secured a conditional Authority to Operate within ten business days at the start of the COVID-19 pandemic and was featured in an AWS case study. He negotiated the prime award for the DoD Joint Operational Medical Information Systems Healthcare Documentation System, a subcontract as the patient engagement platform for the VA Office

of Connected Care, and a subcontract to develop the Medical and Public Health Information Sharing System for the Office of Health Security. Private clients included strategic partnerships with Johns Healthcare Solutions and the Hospital for Special Surgery (HSS). He previously held P&L responsibility for ViiMed's federal business unit as Vice President of Sales and Business Development and as Executive Vice President of Federal Relationships.

At Siemens Healthcare (2005–2016) he generated more than $155 million in capital equipment bookings. As Key Account Executive he closed a five-year, $40 million exclusive multi-vendor maintenance agreement with the University of Maryland Medical System, the first of its kind for Siemens, and a five-year Strategic Alliance with MedStar Health for exclusive MRI and CT purchasing valued at more than $50 million. He converted six competitive capital imaging systems to Siemens at Walter Reed National Military Medical Center, ranked #3 of 98 account executives nationally in FY10. Earlier, as Cardiology Account Executive in Los Angeles, he ranked #5 of 39 nationally in FY08 and #2 of 38 in FY07, and established St. Joseph's Hospital in Orange, California, as a national cardiology reference site. In total, Brion was an eight-time Apex award winner, all years eligible, in his service to Siemens.

Bennett served in the U.S. Navy from 2002 to 2005 as Assistant Weapons Officer, Anti-Submarine Warfare Officer, and Visit Board Search and Seizure (VBSS) Officer aboard USS Bunker Hill (CG-52), completing two Western Pacific deployments in support of the Global War on Terrorism and Operation Enduring Freedom. He supervised 60 personnel across four enlisted ratings in the maintenance and operation of shipboard weapons systems on an Aegis-class guided-missile cruiser, served as Staff Watch Officer for the Bonhomme Richard Expeditionary Strike Group (ESG), and led the ESG's humanitarian mission to Indonesia in support of Operation Unified Assistance, coordinating delivery of more than one million pounds of relief supplies. As VBSS Boarding Officer, he led his team in the discovery of nearly three tons of illegally smuggled hashish in motor vessel Al Zahidi, placing the dhow's crew and vessel under control in what was the largest bust in CENTCOM's history in 2005. His decorations include the Navy and Marine Corps Achievement Medal and the Navy Unit Commendation Medal, National Defense Service Medal, Navy "E" Ribbon (2), Sea Service Deployment Ribbon, Global War on Terrorism Expeditionary Medal, and the Global War on Terrorism Service Medal.

He earned B.S. degrees in General Engineering and Economics from the United States Naval Academy (1998–2002), where he was 5th Battalion Executive Officer, 25th Company First Sergeant, and a member of the varsity baseball team. He currently resides in Lighthouse Point, Florida, with his family.

Panel

AI Applications in Government: Including the Virtues of Western Open-Weight Models

Panelist: Hamlet Yousef

Hamlet Yousef is the Co-Founder and Managing Partner of IronGate Capital Advisors, a venture fund investing in dual-use national security and defense technologies. Before founding IronGate in 2018, Mr. Yousef worked in the federal government's national security and diplomacy sector, supporting agencies on cybersecurity, counterterrorism, and threats from nation-state and nonconventional adversaries. His service focused on applying technology to mission-critical national security needs.

Mr. Yousef is also Managing Partner of SABiO Global Advisors, a boutique business intelligence and geopolitical advisory group. He advises private clients, family offices, and the U.S. Government on manufacturing and supply chain strategies shaped by complex geopolitical dynamics. SABiO also conducts due diligence and counterparty risk analysis for complex transactions and joint ventures.

He serves as an Economic and Legal Warfare Advisor to the Irregular Warfare Initiative, a research and policy program focused on the human, political, and informational dimensions of modern irregular and hybrid conflict. IWI is supported by the Modern War Institute at West Point and the Empirical Studies of Conflict Project at Princeton University.

Mr. Yousef is also a regular contributor to The Cipher Brief, a digital media platform focused on global security and national intelligence. He also serves as Senior National Security Advisor to Magnet Defense, which is building autonomous warships for the U.S. Government, and to Tower Strategy Group, which advises corporations, investors, and governments on engaging Washington effectively.

Before joining government, Mr. Yousef conducted business development, due diligence, and geopolitical risk analysis in the U.S., Europe, and the Middle East. He worked with private investors to assess disruptive technologies in the energy and technology sectors. Earlier in his career, he spent nearly a decade in sales and marketing roles at Stanley Black & Decker.

Mr. Yousef earned an MBA from Pepperdine University, with an emphasis in International Finance, and a BS in International Business from Florida State University.

Panel

AI Applications in Government: Including the Virtues of Western Open-Weight Models

Panelist: John Alden Haynes

John Alden Haynes leads growth and business development for Reflection's U.S. Government Federal Civilian portfolio. Reflection is an American open frontier AI lab, building open-weight models for enterprise, government, and regulated industries that need to deploy AI in environments they control. John's work strengthens American leadership in AI by advancing Reflection's mission to make open intelligence accessible to all. He also leads the company's engagement with the Department of Energy's Genesis Mission, where Reflection serves as the open model provider to the National Laboratories in support of their scientific work.

Before joining Reflection, John served a decade as a U.S. Army officer in the 1st and 25th Infantry Divisions, where he commanded reconnaissance units and deployed overseas multiple times. He is a decorated veteran of the Global War on Terror and was selected for promotion to Major before leaving active duty. He later advised Fortune 500 leaders as a management consultant at McKinsey & Company, focused on growth strategy and large-scale transformations.

His academic background includes published research on strategic leadership and communication, co-authored with Dartmouth professor Paul Argenti. John holds a B.S. in Engineering Management from the United States Military Academy at West Point and an MBA from the Tuck School of Business at Dartmouth.

Panel

AI Applications in Government: Including the Virtues of Western Open-Weight Models

Panelist: Mike Lu

Mike Lu is the co-founder and executive chairman of Aquila Health AI, bringing deep product and technology experience and the relationship capital to execute at scale. As CEO of Triller, Mike grew the AI-powered video platform from a small editing app into a global product reaching over 300 million users and surpassing a billion-dollar valuation. Across eighteen-plus years in gaming and social media, he has worked at the intersection of complex systems, high-stakes infrastructure, and the people who build and buy them, giving Aquila both engineering credibility and the network to grow.

Panel

AI Applications in Government: Including the Virtues of Western Open-Weight Models

Panelist: Daniel Ratner

Daniel Ratner is a senior staff scientist and Jefferson Professor at the Thomas Jefferson National Accelerator Facility (JLab), and is the interim deputy director for the Models Consortium and the Models pillar co-lead within the Genesis Mission. Prior to joining JLab, he earned his PhD in accelerator physics at Stanford and SLAC National Accelerator Laboratory, working on commissioning of the world's first X-ray laser, and inventing the steady-state microbunching EUV source for lithography and a new method for hadron cooling at the electron-ion collider. In 2015 he founded SLAC's program in AI for particle accelerators, and subsequently founded and led SLAC's lab-wide AI-for-science program from 2019-2026, working on domains ranging from high-energy physics to materials science to drug discovery. Before arriving at Stanford he worked at NYC's MoMA using x-ray and optical methods to support art conservation.