USG Proceedings Journal

Meeting Abstracts | Open Access

Proceedings of the 3rd International Conference on Artificial Intelligence and Machine Learning (AIM-2026)

Day 1: Plenary Presentations

When Computers Look at Art

David G. Stork*

Stanford University, California, United States

Abstract

Recent triumphs and future opportunities for computer-assisted connoisseurship of fine art paintings and drawings. Our cultural patrimony of fine art paintings and drawings comprises some of the most important, memorable, and consequential images ever created, and presents numerous problems in art history and the interpretation of "authored" stylized images. While sophisticated imaging (by numerous methods) has long been a mainstay in museum curation and conservation, it is only in the past few years that true image analysis—powered by computer vision, machine learning, computer graphics, and artificial intelligence (AI)—have been applied to fine art images. Fine art paintings differ in numerous ways from the traditional photographs, videos, and medical images that have commanded the attention of most experts up to now: such paintings vary extensively in style, content, non-realistic conventions, and especially intended meaning. Rigorous computer methods have outperformed even seasoned connoisseurs on several tasks in the image understanding of art and have provided new insights and settled deep disputes in art history. Additionally, the classes of problems in art analysis, particularly those centered on inferring meaning from images, are forcing computer experts to develop new algorithms and concepts in AI. This talk, profusely illustrated with fine art images and computer analyses, argues for the new discipline of computer-assisted connoisseurship, a merger of humanist and scientific approaches to image understanding. Such work will continue to be embraced by art scholars and address new grand challenges in AI.

Generative Artificial Intelligence in Education: Issues and Experiences from San Francisco State University

Dragutin Petkovic*

San Francisco State University, California, United States

Abstract

Issues and experiences from teaching Computer Science (CS) at San Francisco State University (SFSU) generative artificial intelligence (GenAI) emerged as the fastest growing commercial technology with strong and disruptive impact on all aspects of society: from business, medicine, creative arts, education and science to everyday things we do. GenAI also has its challenges and limitations. In this talk we will overview some general issues and limitations we see with GenAI, then we will summarize GenAI impact on education from student and faculty perspective (using data from published surveys and our own experience), and finally we will present some recommendations on teaching with GenAI based on our teaching experiences at SFSU CS department and SFSU graduate certificate in ethical AI.

Key Note Presentations

Artificial Intelligence: Tool or Intelligent Entity?

Craig Kaplan*

iQ Company, California, United States

Abstract

Many of us have assumed that artificial intelligence (AI) is just the latest tool. But recently, leaders ranging from Jensen Huang to Dr. Geoffrey Hinton have suggested that AI may be more than just a tool. AI agents are acting with increasing autonomy and behaving more like workers who use tools rather than tools themselves. How we classify AI determines how we design it, regulate it, and ultimately whether it remains safe. Current narrow AI systems fit reasonably within a tool definition. The problem arises when we apply that same thinking to the most advanced forms of AI, including advanced agentic AI, artificial general intelligence (AGI), and superintelligence. AGI is defined as AI with average human ability across cognitive domains. Once a system reaches that threshold, it will set its own goals, learn from experience, and operate across open-ended domains. When a system can outperform humans and eventually rewrite its own cognition, the tool framing not only fails but also becomes dangerous. When we view AI as a tool, we consider safety to be a matter of applying technological fixes: guardrails, output filtering, and reinforcement learning patches. These methods do not scale with a system that may exceed human monitoring capacity. They also sidestep the central problem: an intelligent entity needs values, not just guardrails. If AGI is an intelligent entity, we must embed human values at the system level from the beginning, through design rather than fixes. The design of the systems of tomorrow begins with the classification we adopt today.

From Software Test Automation to Agentic Artificial Intelligence Test Automation

Jerry Zeyu Gao*

San Jose State University, California, United States

Abstract

In this talk, Dr. Jerry Gao discusses the major problems, challenges, and needs in test automation in today's artificial intelligence (AI) world and industry. Based on his over 25-year software test automation research and R&D experience, he first examines and reviews diverse software test automation methods, achievements, and technologies in the past. Then, he shares the current applications of AI and machine learning models in software test automation in both industry and research community. Moreover, Dr. Gao shares his group's 6-year AI test automation research work and teaching experience for modern AI-powered applications and discusses the major challenges and needs in test automation for modern AI-powered mobile apps and intelligent applications. Furthermore, Dr. Gao presents his vision and views about agentic AI test automation for future. Including its major concepts, key features, tasks, road map, and platform. In addition, Dr. Gao presents a comparative view among software test automation, AI-based test automation, and agentic AI test automation.

Featured Presentations

Bridging Data Science and Data Engineering: An Industry-aligned Curriculum Approach

Mahesh Balkrishna Chaudhari*

University of San Francisco, California, United States

Abstract

Artificial intelligence (AI) has been at the forefront of major tech companies, transforming industries through intelligent automation, predictive analytics, and generative systems. Yet the promise of AI is only as strong as the data infrastructure that powers it. As models mature and deployment scales across multicloud environments, industry has been shifting its focus towards building scalable, autonomous, secure data infrastructure. With this shift comes a critical talent gap. Demand for skilled data engineers who can architect and operate modern data infrastructure has never been higher. The data engineering curriculum at the University of San Francisco is designed from the ground up with this reality in mind. Spanning courses in distributed computing, distributed data systems, data stream processing, and the data Lakehouse architecture, the program equips students not just with technical skills but with the systems level thinking that production AI environments demand. This presentation offers a practitioner educator's perspective on bridging this gap between data science and data engineering, the two disciplines that are increasingly inseparable in real world AI deployments. Attendees will gain insight into how the curriculum models industry workflows using tools such as Apache Spark, Airflow, Kafka, among others, within cloud native platforms such as Amazon Web Services, Google Cloud Platform, and Databricks; how practical projects emulate building the end-to-end AI ready data pipelines; and how student faculty collaborative research has produced peer reviewed contributions at Institute of Electrical and Electronics Engineers and Association for Computing Machinery conferences. The talk also focuses on reasonable data practices—governance, lineage, and security—recognizing that trustworthy AI begins with trustworthy data.

From Proof of Concept to Production: Operationalizing Artificial Intelligence in Enterprise Systems

Darren Kimura*

AISquared, California, United States

Abstract

Most artificial intelligence (AI) initiatives stall at the proof-of-concept stage. In this session, Darren Kimura explores why model performance is rarely the real barrier and explains what it takes to securely deploy AI inside complex enterprises and federal environments. Drawing on real-world experience, he outlines the architectural, governance, and leadership disciplines required to move AI from experimentation to measurable operational impact.

Global Artificial Intelligence Governance in the United States, European Union, and China

Anshu Arora*

University of the District of Columbia, Washington, United States

Abstract

Over the past decade, social robotics has evolved significantly through the integration of explainable artificial intelligence (XAI) and interpretable AI (IAI) to advance user experience and customer experience (CX). This research presents a decade-long review (2015 - 2025) of frameworks in social robotics and human-robot interaction that emphasize transparency, trust, and sustained engagement. Early advancements focused on equipping robots with internal drives and motivations to produce self-explanations understandable to users. As the field matured, scholars shifted toward user-centered designs, autonomous social behaviors, and socially adaptive explanations. By the early 2020s, research increasingly explored how XAI and IAI enhanced long-term engagement, emotional connection, and service quality in CX contexts. By 2025, social robotics research had progressed toward integrated frameworks that blend technical explainability with ethical, social, and experiential dimensions. This review synthesizes a decade of findings demonstrating how XAI and IAI have shaped trustworthy, transparent, and human-aligned social robots, ultimately transforming CX through more meaningful and interpretable interactions.

Evaluating a ReAct-based Agent for Agricultural Planning

Karanbir Singh*

Salesforce, United States

Abstract

Planning is a core aspect of intelligence, enabling both humans and artificial intelligence (AI) to address complex challenges. This talk explores the capabilities of AI-driven agents to reason and plan, using the intersection of AI and agricultural planning as a practical testbed. Focusing on the gardening planner, a ReAct-based agent, the study integrates advanced tools such as retrieval-augmented generation, weather forecast application programming interface, and dynamic web search to generate personalized, context-aware crop plans. While the agent demonstrates potential, achieving effective plans only 50% of the time, the findings reveal critical limitations in its reasoning and planning processes. This talk offers insights into the challenges of deploying AI in dynamic, real-world scenarios and its implications for advancing agent-based intelligence.

Agentic Artificial Intelligence-powered Digital Twins for Integrated Dynamic Bioprocessing: Real-time Multi-scale Modeling and Automated Error Detection in Biomanufacturing

Tianxin Zhang[1*], Ashna Dhingra[1], Xiaoli Duan[2], Michael Mollet[1] and Jon Coffman[1]

[1]Bioprocess Technologies and Engineering, Biopharmaceutical Development, R&D Unit, AstraZeneca, Gaithersburg, United States

[2]Data Science and Modelling, Biopharmaceutical Development, R&D Unit, AstraZeneca, Gaithersburg, United States

Abstract

Integrated dynamic bioprocessing remains a critical challenge in modern biologics manufacturing, where continuous monitoring and proactive control are essential for consistent product quality and operational efficiency. To address these needs, we present a digital twin system that leverages agentic artificial intelligence (AI) combined with multi-scale machine learning (ML) modeling to deliver real-time prediction and error detection across upstream and downstream bioprocessing stages. Our architecture continuously integrates live process data from bioreactor sensors and at-line analytical measurements, creating a synchronized virtual profile that mirrors the physical process. Multi-scale ML models are applied to essential process variables—such as glucose uptake, total titer, and ultraviolet absorbance—fitted across different bioreactor volumes and filtration steps. Agentic AI automates analysis and actively detects anomalies and deviations in equipment and process performance, triggering corrective actions and maintenance recommendations without operator intervention. Pilot deployment on AstraZeneca's front-end platform demonstrates that the integrated digital twin can proactively identify potential failures, reduce unplanned downtime, and optimize scheduling for maintenance and process interventions. Real-time model outputs allow operators and engineers to investigate root causes swiftly, improving decision-making and enhancing bioprocess safety and resiliency. By combining agentic AI-driven error detection with multi-scale ML modeling, our digital twin approach establishes a new paradigm for autonomous and adaptive biomanufacturing. It sets the foundation for scalable, smart production environments capable of continuous self-optimization through actionable, data-driven insights.

Artificial Intelligence Framework for Predicting Injury vs Non-injury Crashes

Roumen Vesselinov*, Elena Vesselinov, Mark Scarboro and William Teeter

University of Maryland, Baltimore, United States

Abstract

In this study, we develop a complex artificial intelligence (AI) framework for incorporating latent topics from crash narratives and latent features from crash diagrams to predict injury vs non-injury crashes, based on over 43,000 police reports in Maryland, 2019 - 2023. Crash diagrams remain an underutilized resource in predicting injuries and we offer an innovative approach in this respect. The research goals are three: (1) Perform AI analysis of the crash diagrams using convolutional neural networks, extract latent features from the images, and reduce their dimensionality; we also develop a new classification based on this analysis. (2) Discover new factors for injury vs non-injury crashes using extracted latent topics from crash narratives, based on natural language processing. And (3) build predictive models for injury vs non-injury crashes based on crash reports data, including the extracted crash latent topics and features. Our AI framework consists of three types of machine learning-based techniques: Classification and regression tree models, random forest, and extreme gradient boosting (XGBoost). The framework is also complemented by Shapley additive explanations (SHAP) analysis, principal component analysis, and cluster analysis. The predictive modeling of injury vs non-injury crashes underscores the limitations of relying only on structured crash report

data, which produced modest accuracy with area under the curve (AUC) values near 0.65. In contrast, models that incorporated narrative-extracted features achieved significantly higher performance, with random forest and XGBoost models reaching AUC values of 0.84. SHAP analyses confirmed the importance of narrative-derived topics, hit-and-run status, collision type, and roadway division as the most influential predictors of injury outcomes.

What Machine Learning Teaches Us About Depression Prediction Across the Life Course: An Exploratory Comparison of Predictive Models

Rafael Geurgas[1*], Saul J. Newman[2-4], Evelina T. Akimova[1,5,6], Katherine N. Thompson[1] and Robbee Wedow[1,5-7]

[1]*Department of Sociology, Purdue University, United States*

[2]*Centre for Longitudinal Studies, University College, United Kingdom*

[3]*University College, University of Oxford, United Kingdom*

[4]*Institute of Population Aging, University of Oxford, United Kingdom*

[5]*Department of Statistics, Purdue University, United States*

[6]*Center on Aging and the Life Course, Purdue University, United States*

[7]*Department of Medical and Molecular Genetics, Indiana University School of Medicine, United States*

Abstract

Identifying individuals at risk for depression early is important for preventing long-term mental health issues. However, the variability in depression severity, duration, and triggers complicates predictions. This study explores whether machine learning models can outperform traditional methods, like logistic regression, in predicting self-reported depressive symptoms and clinical depression during adolescence and adulthood. We applied five machine learning models with varying complexity levels - logistic regression, decision tree, extreme gradient boosting (XGBoost), support vector machine, and neural networks - using data from a nationally representative longitudinal study of the US, which tracked participants for 20 years. The models were trained with early-life predictors (ages 12 - 18) from wave I, including environmental factors (family, school, health) and genetic predispositions (polygenic scores) from wave IV. Models were evaluated on their ability to predict depressive symptoms and clinical diagnoses in both adolescence and adulthood. After evaluating the performance of all five models, XGBoost emerged as the most effective, with a 0.02 increase in area under the receiver operating characteristic curve compared to the benchmark logistic regression model. While this is a slight improvement, overall, logistic regression performs about as well as many of our machine learning models. Early-life data showed strong predictive value for depressive symptoms and clinical diagnoses in adolescence and adulthood, highlighting adolescence as a critical period. Polygenic scores do not add predictive power when combined with environmental data. Feature importance analyses identified self-perception and physical health as key predictors of depressive symptoms, while trauma and life-changing events were more influential for clinical depression.

A Quantum-ensembled Quantum Long Short-term Memory-quantum Convolutional Neural Network Framework for Wind Power Forecasting and Grid Stability

Nithya Venkatesan[1*], Sohorab Hossain[1], Rameez Kurshi[2] and Eldad Avital[3]

[1]*Vellore Institute of Technology, Chennai, India*

[2]*University of Hull, United Kingdom*

[3]*Queen Mary University of London, London, United Kingdom*

Abstract

The increasing reliance on renewable energy sources, particularly wind power, necessitates accurate and reliable forecasting systems to ensure grid stability and efficient energy management. However, the inherent variability and intermittency of wind introduce significant forecasting challenges, including high error rates and instability in voltage and frequency regulation. Traditional machine learning approaches often underperform when dealing with small-scale or highly uncertain datasets. To address these limitations, we propose a quantum-ensembled quantum long short-term memory-quantum convolutional neural network (CNN) forecasting framework, which integrates classical deep learning models-long short-term memory and CNN—with quantum computing elements via variational quantum circuits. The ensemble leverages quantum parallelism and entanglement to enhance model generalization and robustness. We evaluated the proposed framework on the National Renewable Energy

Laboratory wind integration national dataset and the hourly energy demand and weather dataset, demonstrating superior performance in terms of forecasting accuracy, noise resilience, and computational efficiency. The model achieves a coefficient of determination score of 0.9312, root mean square error of 0.020106, mean absolute error of 0.011823, and an inference latency of 0.01 s, highlighting its suitability for real time control and integration within smart grid infrastructures.

Artificial Intelligence in Advanced Manufacturing: Educating the Next Generation Technician Workforce

Karen Wosczyna-Birch*

National Center for Next Generation Manufacturing, United States

Abstract

This session explores how artificial intelligence (AI) is transforming manufacturing and reshaping technician education. Drawing on results from the National Science Foundation (NSF) funded National Center for Next Generation Manufacturing's national industry survey, their national summit for educators in manufacturing-related programs, and national partnerships including the NSF funded National Applied AI Consortium, participants will learn about current and emerging uses of AI in advanced manufacturing companies and what future manufacturing technicians need to know. The advanced manufacturing industry is using AI to enhance precision, productivity, and safety from company operations down to how individual products are made. Participants will learn about strategies and partnerships to bring these innovations to their campuses.

The Real Constraints of Building Agent First Search

Jay Chia*

Eventual, San Francisco, United States

Abstract

Search for artificial intelligence (AI) agents is becoming one of the most important challenges in modern computing. Traditional search systems were created for humans who naturally adapt to the information they explore. Humans remember context, recognize structure, and refine their approach from one query to the next. AI agents cannot do this. Each query begins from a blank state without intuition or memory, which creates failure modes that do not appear in human-centered search. This talk introduces the core ideas behind search that is built for agents rather than people. It explains why unstructured and constantly changing documents make retrieval significantly harder, and why simple semantic matching or keyword matching is rarely enough. The session explores the role of metadata, hybrid retrieval, chunking, summarization, and global execution, and why these components matter when agents must operate at high speed and at large scale. Attendees will learn a practical framework for thinking about agent-oriented search, along with the design questions that arise when building systems for fragmented knowledge, continuous updates, and large volumes of slightly varied queries. The goal is to give the audience a clear view of how the next generation of search systems will need to behave in order to support reliable AI reasoning.

Poster Presentations

Empirical Evaluation of Deterministic Validation Layers for Large Language Model

Gary Langford*

3x3 Institute, United States

Abstract

Large language models (LLMs) are increasingly deployed in domains where correctness, consistency, and traceability are critical, yet their output remains stochastic and prone to logical and structural errors. We present a deterministic validation and repair framework that operates as a post-generation control layer for LLMs, enforcing task-specific constraints without retraining or modifying the underlying models. We evaluate this framework empirically across five widely used models - Claude Sonnet 3.5, GPT-4.5, GPT-5.2, Grok (free tier), and Gemini (free tier) on a curated benchmark of logical reasoning, structured compliance, and planning tasks. Results show consistent reductions in logical violation rates, improved output stability across repeated runs, and high repair success rates across all models. These findings indicate that deterministic validation layers can substantially enhance the reliability and trustworthiness of LLM-based systems and provide a practical complement to existing alignment and safety techniques.

Technological Reciprocity: Modeling the Artificial Intelligence Productivity Dividend in Automated Economies

Reza Abdolee*

California State University Channel Islands, California, United States

Abstract

Artificial intelligence (AI) is reshaping global productivity by automating administrative, analytical, and industrial tasks once dependent on human labor. While automation boosts efficiency and corporate profitability, its economic gains remain concentrated among technology firms and capital owners. This paper introduces the concept of an AI productivity dividend (AIPD), a mechanism through which governments capture a portion of AI-driven profits and redistribute them to society via direct income, infrastructure investment, and essential-goods subsidies. Using a quantitative manufacturing case study, we demonstrate how modest taxation of automation-based profits (20 - 30%) can enhance equality, stabilize aggregate demand, and lower living costs. The AIPD reframes the relationship between automation and social welfare, offering a policy framework where machines generate efficiency and humans retain purpose. This approach advances a balanced, human-centered economy in the era of intelligent automation. Simulation results indicate that a 25% automation tax could generate $110 million in public revenue from a mid-sized manufacturing firm while preserving over 70% of post-AI profits.

Leveraging Natural Language Processing to Identify and Predict Mental Health Issues Among College Students Through Social Media Analysis

Anika Bokka* and Renu Balyan

State University of New York, Old Westbury, United States

Abstract

There is an unprecedented incidence of diagnosed anxiety and depression among United State college students, yet many encounter systemic barriers to care. This research addresses the escalating mental health crisis among students by engineering a proactive intervention framework powered by natural language processing and large language models (LLMs). The study's primary objective is to move beyond traditional reactive support models by analyzing discourse and sentiment from social media platforms such as Reddit, Instagram, and X to identify the latent environmental and systemic root causes of student distress. We are evaluating the capabilities of LLM architecture (such as GPT-4o) as autonomous research assistants capable of high-level feature classification and labeling. By developing a granular taxonomy, we have successfully identified 10 major stress categories and 40 specific sub-triggers, ranging from academic rigor to financial instability. This research focuses on the automated quantification of stress levels, transforming qualitative student narratives into actionable, numerical data points. The integration of these LLM-derived insights into machine learning predictive models enables the university to flag high-risk cases within internal databases with unprecedented precision. By synthesizing linguistic markers of psychological disorders with advanced algorithmic labeling, this project establishes a data-driven pipeline for mental health advocacy. Ultimately, this work

provides institutions and their leaders with the tools to transition from a self-reporting dependency to a preventive ecosystem, ensuring that support reaches students before they reach a point of crisis.

Beyond Additive Models: Integrating Protein Language Models with Polygenic and Rare Variant Frameworks

Seyedezahra Paylakhi[1], Michel G. Nivard[2,3*] and Robbee Wedow[1*]

[1]*Department of Sociology, Purdue University, West Lafayette, United States*

[2]*Medical Research Council Integrative Epidemiology Unit, University of Bristol, Bristol, United Kingdom*

[3]*Population Health Sciences, Bristol Medical School, University of Bristol, Bristol, United Kingdom*

Abstract

Additive models have been central to modern human genetics, enabling genome-wide association studies (GWAS), polygenic risk scores (PRS), and rare-variant burden tests to scale across hundreds of thousands of individuals. By treating genetic variants as independent contributors to phenotype, these frameworks achieve statistical power and computational feasibility, but at the cost of biological context: haplotype phase, sequence continuity, and molecular interactions are largely discarded. As a result, additive models capture population-level associations but provide limited insight into how genetic variation perturbs molecular function. In parallel, recent advances in genomic deep learning, particularly protein, DNA, and RNA language models—offer a complementary perspective. These models operate directly on contiguous, phased sequences and learn biochemical and structural constraints underlying molecular function, enabling context-dependent and non-additive variant effect prediction. However, they are typically trained on canonical or multispecies sequences and do not explicitly model population-level genetic diversity. Rather than replacing additive models, we focus on near-term strategies in which sequence-based variant scores are incorporated as biologically informed annotations or prior within existing GWAS, transcriptome-wide association studies, PRS, and rare-variant frameworks. We outline concrete integration points enabled by phased genomic data, allowing population-scale association pipelines to interface with protein and nucleotide language models using currently available tools and datasets. By reframing additivity as a population-level abstraction rather than a biological constraint, this work emphasizes feasible strategies for combining statistical power with molecular realism. Integrating phased genomic variation with sequence-aware deep learning models provides a practical path toward more interpretable and biologically grounded genetic risk modeling.

Day-2: Keynote Presentations

Beyond Probability: The Certainty of Artificial Intelligence

Seth Dorbin*

Qantm AI, Connecticut, United States

Abstract

Dr. Seth Dobrin is a pioneering artificial intelligence (AI) and data science leader renowned for his transformative work across industries. At Monsanto, he revolutionized agriculture by integrating AI to enhance farming through precision techniques. As IBM's first Global Chief AI Officer, he led the data science elite team, driving ethical AI development and diversity initiatives. Dobrin founded Qantm AI, championing a human-focused approach to ethical AI in business. His work guides fortune 500 companies and governments worldwide in developing responsible AI strategies. He's also a general partner at 1infinity ventures and Silicon Sands Venture Studio. A strong advocate for diversity and against technological colonialism, Dobrin supports organizations like World Leaders in Data & AI and Wonder Women Tech. His multidisciplinary expertise, combining molecular and cellular biology with industry experience, sets AI leadership benchmarks. Dobrin is shaping the future of AI to be responsible, ethical, and inclusive, as explored in his book "AI iQ for a human-focused future: strategy, talent, and culture.

Featured Presentations: Machine Learning

Machine Learning for Physical Layer Communications: Let There be Lightness

Benjamin Belzer[1*], Krishnamoorthy Sivakumar[1], Ahmed Aboutaleb[2], Amirhossein Sayyafan[3] and Ali Kanaani[1]

[1]*Washington State University School of EECS, United States*

[2]*Marvell Technology, Inc., United States*

[3]*Western Digital Corporation, United States*

Abstract

Due to the need to minimize latency, power consumption and silicon area, machine learning (ML) for physical layer communications must deliver the required performance metrics (e.g., maximum data rates or information storage densities, minimum bit-error-rates, or minimum channel parameter estimation errors) with minimum computational complexity; i.e., ML solutions must be "light weight." Fortunately, decades of research in communication theory have yielded techniques and algorithms that are optimal (or near optimal) in many important scenarios. Thus, a system designer should utilize these legacy algorithms where possible and employ ML only where the legacy algorithms are sub-optimal and/or prohibitively complex. This talk provides examples of such ML-enhanced designs in three emerging technologies: two- and three-dimensional magnetic recording (TDMR, 3DMR) for hard disk drives, time-offset sub-faster-then-Nyquist non-orthogonal multiple access (TO-sFTN-NOMA) for 6G wireless communications, and channel estimation and equalization for doubly dispersive (DD) wireless communication channels encountered in 5G, 6G, and underwater communications. For TDMR and 3DMR, use of legacy methods for intersymbol interference equalization together with convolutional neural networks for non-linear media noise estimation increases data storage density with manageable complexity. For TO-sFTN-NOMA, an autoencoder-based design is presented that improves performance and lowers complexity compared to previously proposed singular value decomposition-based solutions. For DD channel estimation, we propose to combine frequency-domain Karhunen-Loeve-based statistical signal processing for DD channel estimation with reinforcement learning for tracking the time-varying channel statistics. In each case, proposals towards future work to enhance the performance and/or lower the computational complexity of the considered ML-based systems are provided.

Risk-averse Reinforcement Learning in Portfolio Optimization

Shimin Zhang[*] and Kandethody M. Ramachandran

Department of Mathematics and Statistics, University of South Florida, United States

Abstract

Portfolio optimization is a major problem in investment. The traditional approach to portfolio optimization is based on minimizing risk or maximizing returns. Deep reinforcement learning (DRL) has been used to improve portfolio optimization strategies in recent years. Entropic risk measures (ERM) are used to estimate the risks involved with uncertain outcomes. In this study, we propose the deep deterministic policy gradient ERM (DDPGERM) algorithm, which incorporates ERM into the DRL algorithm, DDPG. This incorporation enables risk-sensitive updates to both the actor and critic networks within DDPG. We further propose the DDPGEERM algorithm by incorporating ERM into DDPG augmented with episodic memory. The proposed algorithms are tested on low-and high-volatility stock portfolio datasets. The experimental results indicate that the proposed DDPGERM algorithm outperforms conventional DRL algorithms and that it substantially outperforms the DDPG algorithm, thereby demonstrating the effectiveness of the ERM approach. This study contributes to the growing body of research on risk-averse reinforcement learning and offers practical guidance for portfolio optimization under uncertainty.

Annotation-efficient Object Detection Under Domain Shift Using Hybrid Active and Semi-supervised Learning

In Seon Kim[1*] and Ali Moghimi[2]

[1]*Department of Computer Science, University of California, Davis, United States*

[2]*Department of Biological and Agricultural Engineering, University of California, Davis, United States*

Abstract

Deep learning has enabled large-scale object detection in large-scale remote sensing and satellite imagery. However, its effectiveness remains constrained by the high cost of manual annotation and limited cross-domain generalization. Models trained on one visual domain often fail to transfer reliably to new environments without substantial labeled data, creating a major bottleneck for scalable deployment. This study investigates annotation-efficient learning strategies for tree detection in urban areas under domain shift using multimodal imagery from satellite and Google street view. We introduce a framework that integrates transfer learning, semi-supervised learning, and active learning to reduce annotation requirements while maintaining robust detection performance. A transformer-based detector is first pretrained on a visually distinct source-domain dataset and then adapted to a target domain with limited labeled samples. To further minimize manual labeling, we evaluate three learning strategies: semi-supervised learning (auto-label annotation), active learning (human-in-the-loop annotation), and a hybrid strategy combining both approaches. Experimental results show that purely semi-supervised training suffers from confirmation bias that gradually reinforces model errors, while active learning improves performance by directing human annotation toward uncertain samples. The proposed hybrid strategy achieves the strongest performance by combining scalable pseudo-label propagation with targeted human supervision, improving the F1-score from 0.80 to 0.90 while substantially reducing annotation requirements compared to conventional fully supervised pipelines. These results highlight the importance of integrating transfer learning with hybrid annotation-efficient training strategies to overcome domain adaptation challenges and enable scalable deep learning systems under limited labeled data conditions.

A Large-scale Analysis of Mountain Glacier Shrinkage Using Convolutional Neural Networks

Matthew Waismann[1], Somansh Budhwar[2] and Armin Schwartzman[2*]

[1]*Amazon Inc., United States*

[2]*University of California, San Diego, United States*

Abstract

Mountain glaciers are undergoing widespread shrinkage. While large glacier inventories exist, these datasets are temporally sparse, and traditional methods of collecting glacier data, such as ground-based measurements or manual analysis of satellite imagery, are not scalable for continuous global monitoring. Recent advancements in computer vision techniques now allow us to programmatically analyze extensive archives of satellite images. In this work, we develop a framework for globally estimating the area of individual mountain glaciers from Landsat images over time. After data pre-processing and filtering, the main engine is a U-net deep convolutional neural network (CNN) with a modified resnet encoder. CNN is trained on 7174 glacier contours and their corresponding images, each at only one time point, and then applied to time series of satellite images of 1083 glaciers, comprising a total of 505835 images, to obtain time series of estimated surface area for those glaciers over the 40-year-period 1984 - 2023. Our global analysis confirms a consistent trend of glacial surface area shrinkage in all five major glacierized regions: Asia, the Caucasus, Europe, North America and South America. We estimate an overall glacier surface area decline of 0.2% ± 0.01% per year.

Architecting and Scaling a Unified, Role-aware Agentic Artificial Intelligence Chatbot for Field Operations in a Large Telecom Enterprise

Dipjyoti Das[*]

Ascendion Inc, United States

Abstract

Field technicians in large telecom enterprises often rely on multiple disconnected systems for job context, troubleshooting, and operational insights, leading to longer resolution cycles and inconsistent support experiences. This work presents the architecture and production deployment of a unified, role-aware agentic artificial intelligence (AI) chatbot, along with key design and

operational insights to address these challenges across field technicians, supervisors, construction and maintenance engineers. We propose a stateless, representational state transfer-based agentic AI chatbot architecture embedded within an existing tech mobile application that serves over 80,000 technicians. At its core, a LangGraph-based orchestration layer enables multi-agent coordination, dynamically routing requests based on user role, intent, and job context. The pipeline integrates JSON web token-based authentication, large language model-driven intent classification, and role-aware agent selection. The system comprises specialized agents, including a job insights agent that aggregates and summarizes enterprise data from downstream application programming interface. A tech assist agent leverages retrieval-augmented generation (RAG) over large-scale document corpora, while supervisor agents provide business intelligence-driven operational visibility. The RAG pipeline incorporates intelligent chunking, embedding generation, metadata tagging, and re-ranking to ensure grounded and traceable responses. Deployed on AWS Kubernetes pods with fully stateless application programming interface, the architecture supports horizontal scaling. Observability is implemented using self-hosted Langfuse through hierarchical tracing of agent and large language model interactions. Production-scale deployment results indicate reduced system switching, improved troubleshooting consistency, and efficiency gains of approximately 30%, with configuration accuracy exceeding 90%. This study provides practical design patterns and architectural insights for building scalable, role-aware agentic AI systems in enterprise environments, focusing on orchestration, grounding, security, observability, and deployment at scale.

Unbinding Prometheus: How Generative Artificial Intelligence is Liberating Computational Thinking from Coding

Douglas C. Schmidt*

William & Mary, Virginia, United States

Abstract

In my talk, I describe how generative artificial intelligence (AI) is transforming the nature of computing, enabling a shift from traditional coding to natural language-driven problem solving. Drawing from the myth of Prometheus, I present generative AI as a tool that "steals fire" for humanity—democratizing computational thinking for non-programmers across disciplines. My talk provides four key themes: (1) AI as a transformational force reshaping industry, (2) the commoditization of computational thinking through natural language interfaces, (3) the emergence of prompt engineering as the new literacy, and (4) the need to adapt educational systems, workforce development, and regulatory frameworks accordingly. I illustrate through examples how AI tools allow users to perform complex tasks—such as data analysis, system modeling, and workflow automation—without traditional programming. This evolution mirrors the history of the automobile, where the barrier to driving fell dramatically over time. I also describe "prompt engineering in the small" (day-to-day AI use) and "in the large" (team-based, version-controlled prompt workflows), comparing the latter to software engineering's evolution toward collaborative development. I stress the value of "prompt patterns" to guide interaction with AI models. I conclude by discussing the implications of this shift: potential job displacement for junior developers, the rise of "vibe coding" for rapid prototyping, and the urgency of equipping students and professionals with AI fluency, metacognitive skills, and ethical awareness. Computational thinking is now accessible to all—but must be wielded with wisdom and rigor.

Steering Large Language Models with Light-weight Auxiliary Policy Models using Reinforcement Learning

Haipeng Chen*

William & Mary, Virginia, United States

Abstract

Reinforcement learning (RL) has become a cornerstone in advancing the capability and alignment of large language models (LLMs). Existing works, such as RL from human feedback and RL with verifiable reward focus mostly on fine-tuning the underlying LLMs, which introduce inductive bias toward the LLMs and are costly in both training and inference. This talk highlights our recent works on RL for training separate, light-weight, auxiliary policy models (APMs) that interact with or guide LLMs. These APMs learn to evaluate, interpret, or modulate LLM outputs, providing adaptive feedback loops that extend model performance in downstream LLM tasks.

Enhancing Necrotizing Enterocolitis Prediction in Preterm Infants with Tabular Variational Autoencoders-augmented Data Modeling

Shilpa Muralidhar* **and Mohan Pamm**

Division of Neonatology, Department of Pediatrics, Baylor College of Medicine, Texas Children's Hospital, United States

Abstract

Preterm birth impacts 1 in 10 newborns globally, with over 500,000 cases annually in the United States, increasing risks of severe complications like necrotizing enterocolitis (NEC), a life-threatening intestinal condition. Diagnosing NEC is challenging due to symptom overlaps and low prevalence (6.3%), complicating predictive modeling. This study introduces an innovative approach using tabular variational autoencoders (TVAE), a generative adversarial networks-inspired deep generative model, to address data scarcity and class imbalance in neonatal disease prediction. We analyzed a retrospective cohort of ~3,500 preterm infants from Texas Children's Hospital (Vermont Oxford Network, Institutional Review Board-approved), imputing missing data via multiple imputation by chained equations. To counter the 6.3% NEC prevalence, synthetic minority over-sampling technique-rebalanced data trained TVAE to generate synthetic samples, with similarity validated by cosine (0.167) and Wasserstein (0.155) distances. A realism-aware dataset was created by filtering synthetic data to reduce false positives, enhancing clinical applicability. We evaluated logistic regression, XGBoost, CatBoost, and ensemble models (LogReg + CatBoost, LogReg + CatBoost + LightBGM) on this dataset. Logistic regression outperformed others (area under the receiver operating characteristic AUC: 0.789, precision: 0.212, recall: 0.581, F1: 0.311), though all models benefited from improved data quality. This novel TVAE-based approach demonstrates significant potential for augmenting rare clinical data, enhancing predictive accuracy in imbalanced neonatal disease models, and supporting early NEC diagnosis, aligning with the conference theme of advancing medical informatics through innovative data solutions.

Artificial Intelligence as Co-designer: Enhancing Conceptualization with Generative Artificial Intelligence in Interior Design Education

Qun Zuo*

Department of Fashion, Interior Design, and Merchandising, Central Michigan University, Michigan, United States

Abstract

Generative artificial intelligence (AI) is transforming interior design by streamlining workflows and enabling creative exploration through iterative image generation. In design education, these tools present both pedagogical opportunities and challenges. This study investigates the impact of generative AI on students' design thinking and creative autonomy. Twelve junior interior design students participated in a semester-long project focused on the adaptive reuse of a vacant shopping mall into a community educational center. During the early conceptualization phase, they used multiple text-to-image AI tools to generate visual stimuli that informed materiality, space planning, and thematic expression. To evaluate AI's role as a co-creative partner, a comparative analysis was conducted between AI-generated visuals and students' final computer-aided design-based outputs. The analysis combined qualitative observations with computational image processing techniques: (1) semantic similarity using contrastive language–image pre-training embeddings to measure conceptual alignment; and (2) visual saliency prediction using DeepGaze to assess overlaps in visual features. Student surveys further captured perceptions of AI's influence across key stages of the design process. Preliminary findings indicate that generative AI enhanced conceptual clarity, supported design decision-making, and encouraged broader spatial experimentation. Students also developed greater technical fluency, critical reflection, and creative autonomy. Identified challenges included aesthetic inconsistency across AI platforms, prompt sensitivity, and difficulty translating complex AI outputs into feasible design solutions. This study proposes a pedagogical framework for integrating generative AI into design education, grounded in human-computer interaction and computational analysis. Future research will explore platform variability, effective prompt engineering strategies, and methods for aligning speculative AI-generated imagery with practical design constraints.

The Future of Art Attribution: Data Science, Artificial Intelligence, and the Complexities of Cultural Heritage

Liubov Koposova*

Independent Scholar, Russian Federation

Abstract

The convergence of artificial intelligence (AI), data science, and art history is opening a transformative research frontier. Cultural heritage objects, especially works of art, represent some of the most challenging data imaginable: they embody not only visual forms but also layers of symbolic meaning, historical evolution, and regional variation. Precisely because of this complexity, art and cultural heritage provide an unparalleled testing ground for advancing AI methodologies beyond conventional pattern recognition. This paper frames the study of art attribution and authentication as a critical arena where humanistic expertise and computational innovation must intersect. Conventional methods often suffer from subjectivity and incomplete provenance, while AI and machine learning offer the potential to enhance accuracy, transparency, and scalability. Yet the success of such systems depends on more than technical solutions: it requires curating reliable datasets, incorporating nuanced iconographic interpretation, and embedding historical context. Buddhist art exemplifies these challenges and opportunities. Its extraordinary diversity of iconography and the assimilation of local traditions across centuries create a profound "semantic gap" for standard AI models. Addressing this gap is not a matter of simply applying existing tools; it calls for new interdisciplinary frameworks where art historical expertise actively shapes the design and interpretation of computational methods. Rather than treating cultural heritage as a peripheral application, we propose it as a domain capable of driving innovation in AI itself. By confronting the most intricate data, AI research and art history together can redefine both fields—building a future of more trustworthy, ethical, and insightful approaches to cultural heritage.

Day-3: Featured Presentations

Statistical Learning with Machine Generated Data

Yuchen Wu[1*], Kangjie Zhou[2] and Weijie Su[3]

[1]*Cornell University, United States*

[2]*Columbia University, United States*

[3]*University of Pennsylvania, United States*

Abstract

This work explores the emerging paradigm of statistical learning with machine-generated data, where synthetic samples produced by generative models' supplement or replace traditional real-world datasets. We investigate the statistical principles that govern learning in this setting, focusing on how generative data influences estimation accuracy.

Human Care and Artificial Intelligence Regulation

Carlos Montemayor[*]

Department of Philosophy, San Francisco State University, United States

Abstract

Recent discussions on artificial intelligence (AI) regulation and safety have highlighted the importance of developing an AI that is not only roughly aligned with human values, but one that actually cares about humanity. For instance, Geoffry Hinton has recently proposed that one way to guarantee this care for humanity is for AI to have maternal instincts. In this talk, I will address two issues with respect such proposals. First, I explain why without the biologically grounded empathic needs humans, and at least mammals have, there is no realistic way to create AI with maternal instincts or any other care instincts. Second, I argue that the best way of understanding these needs is in terms of their temporal dimension and their symmetry with respect to other agents with similar needs. This temporal and social account contrasts with the two leading views about how to conceive of biological constraints on intelligence, based on metabolism and our biological substrate respectively.

Artificial Intelligence-powered Profiling of Human Trafficking Buyers

Shreyas Ramani[*], Dhanashree Pokale, Alexandra Joukova, Vikas Singh, Nnenne Onyioha-Clayton, Spencer A. Murray, Catherine Scott and Kendra Taylor

Informs Pro Bono Analytics, United States

Abstract

This research paper explores the use of artificial intelligence (AI), particularly large language models (LLMs), to identify and analyze buyers involved in human trafficking. Unlike traditional studies that focus primarily on victims or traffickers, this work shifts attention to understanding buyer behavior, motivations, and characteristics. By applying advanced natural language processing techniques, the study analyzes existing research to uncover patterns and profiles associated with trafficking buyers. The methodology combines LLM-driven analysis with persona modeling to develop detailed buyer personas. Using one-shot and few-shot prompt engineering approaches, the researchers extract insights from multiple sources, enabling the identification of behavioral trends, motivations, and fears among buyers. The study ultimately presents more than eight distinct buyer personas, which can be used for academic research, public awareness campaigns, and policy development. A key contribution of this work is its novel application of LLMs for social impact, moving beyond traditional machine learning, surveys, or basic natural language processing methods. The findings demonstrate how AI-powered profiling can support law enforcement, policymakers, and researchers in proactively disrupting trafficking networks and designing targeted interventions. Additionally, the research emphasizes its broader societal impact. By focusing on prevention and awareness, the identified personas can be shared through public platforms to educate communities and discourage demand. Conducted as part of an INFORMS pro bono analytics volunteer initiative in collaboration with a Statewide Human Trafficking Task Force, the study highlights how data science and AI can contribute to policy development and efforts aimed at combating human trafficking.

The Anti-enlightenment Mentality in the Law

John McClellan Marshall[1,2*]

[1]*Senior Judge, Fourteenth Judicial District of Texas, United States*

[2]*Maria Curie-Skłodowska University, Lublin, Poland*

Abstract

If Europe imagined enlightenment, it would be America that realized it. In one of his many letters to Thomas Jefferson, John Adams wrote that "the 18th century, notwithstanding all its errors and vices, has been, of all that are past, the most honorable to human nature." The period from the last quarter of the 17th century through the end of the 18th century has often been characterized as "The Enlightenment." During this period, various philosophers began the intellectual detachment from medieval scholasticism and its mystical base. The result was an emphasis on empiricism that supported what was, and still is, viewed as a healthy skepticism. The leaders of this movement included such writers as John Locke, David Hume, Voltaire, Montesquieu, and scientists such as Benjamin Franklin, collectively known as philosophes. Though their work tended to look at the natural sciences, its influence on politics and the law was at least as impactful. In modern times, in part as the result of the influence of technology on the educational process, there has been a similar disconnect from enlightenment to a world view characterized by a subjective interpretation of "reality".

Day-3: Virtual Sessions

Key Note Presentations

Towards the Statistical Limits of Learning Non-hallucinating Generative Models

Changlong Wu, Ananth Grama and Wojciech Szpankowski*

Purdue University, West Lafayette, United States

Abstract

Generative models have shown impressive capabilities in synthesizing high-quality outputs across various domains. However, a persistent challenge is the occurrence of hallucinations, where the model produces outputs that are not grounded in the underlying facts. While empirical strategies have been explored to mitigate this issue, a rigorous theoretical understanding remains elusive. In this talk, we present a theoretical framework to analyze the learnability of non-hallucinating generative models from a learning-theoretic perspective. Our results reveal that non-hallucinating learning is statistically impossible when relying solely on the training dataset, even for a hypothesis class of size two and when the entire training set is truthful. To overcome these limitations, we show that incorporating inductive biases aligned with the actual facts into the learning process is essential. We present an approach to achieve this by restricting the facts set to a concept class of finite Vapnik–Chervonenkis-dimension and demonstrate its effectiveness under various learning paradigms. Although our findings are primarily conceptual, they represent a first step towards a principled approach to addressing hallucinations in learning generative models.

A Visualized Case Tool for Temporal Representation and Reasoning

Jixin Ma*

University of Greenwich, United Kingdom

Abstract

The notion of time plays an essential role in modelling natural phenomena and human activities concerning the dynamic aspect of the real world. Virtually, most knowledge in the universe of discourse is time-dependent and suitable methodologies are needed to cope with this. In particular, temporal reference is an idea deeply integrated in human common-sense, and many artificial intelligence (AI) systems need to deal with the temporal dimension of information, the change of information over time and the knowledge about how it changes. Various areas in the domain of AI require temporal representation and reasoning, including prediction/planning, diagnosis/explanation, industrial process control, temporal database management, historical reconstruction, natural language understanding, etc. The purpose of this talk is to: (a) look at an interesting and important topic, the notion of time, in the domain of AI; (b) overview some fundamental issues with respect to temporal ontology; (c) introduce a visualized case tool, visual time. It is hoped that this talk will raise some special interest in representing and reasoning about the temporal aspects of knowledge in the community of AI.

Featured Presentations

Beyond Retrieval-augmented Generation: Agentic Large Language Models as Active Investigators, Not Passive Responders

Krupa Pratap Brahmkstri*

Data Science and Generative AI, Visa, United States

Abstract

Retrieval-augmented generation (RAG) has significantly enhanced large language models (LLMs) by grounding responses in external knowledge. However, traditional RAG pipelines remain inherently passive—limited to single-turn retrievals and static context injection. In high-stakes, evolving domains such as cybersecurity, compliance, or fraud analytics, this approach falls short. This talk introduces a next-generation approach: Agentic LLMs that operate not as passive responders, but as autonomous investigators. Using agentic frameworks like LangGraph and CrewAI, we demonstrate how LLMs can self-initiate multi-step reasoning, formulate investigative queries, iteratively refine context, and make intelligent decisions based on evolving evidence—mimicking the cognitive workflow of a human analyst. We will walk through a real-world implementation in cybersecurity incident response, where agentic LLMs correlated access logs, hypothesized lateral movement, queried historical behavior patterns, and proposed mitigation strategies—reducing triage time by 60%. The session will detail architecture (planner-executor loops, tool calling, memory graphs), evaluation metrics (accuracy, precision vs standard RAG), and design patterns to build explainable, auditable agentic systems. Attendees will leave equipped with a concrete blueprint for transforming static RAG stacks into autonomous, goal-driven artificial intelligence systems, capable of reasoning, decision-making, and collaborative intelligence across domains.

Delta-aware Edge Artificial Intelligence Caching for Hybrid Inference: Feature–content Coherence at the Extreme Edge

Samaresh Kumar Singh*

HP Inc., United States

Abstract

Edge inference pipelines now mix classic predictive models with large language models/retrieval-augmented generation (RAG) components, stressing bandwidth, latency, and memory at heterogeneous central processing unit/graphics processing unit/neural processing unit (NPU) nodes. This talk introduces a delta-aware caching framework for hybrid (predictive + generative) workloads at the edge. Rather than caching whole artifacts or responses, we cache deltas—feature shards, token key-value states, segment embeddings, operator plans, and model/runtime micro-assets—tied to content and feature coherence. A utility-driven policy (latency saved vs staleness risk vs footprint) governs admission/eviction across a multi-tier store (NPU static random-access memory → device dynamic random-access memory → local non-volatile memory express → regional cache). We combine: (i) query- and scene-similarity keys for video/vision and RAG prompts; (ii) deadline-aware refresh to prevent stampedes; (iii) model-version pinning with canary invalidation; and (iv) on-device telemetry (load, thermals, and link quality) to steer cache placement. Deployed as a pluggable module in an edge orchestrator, the approach reduces tail latency for repeated scenes/queries, lowers backhaul egress and increases effective throughput without over-provisioning accelerators. We share design patterns, scheduler hooks, cache index schemas, and reproducible benchmarks for common edge artificial intelligence (AI) tasks (vision analytics, speech, retrieval-augmented QA), along with guidance on safety (Personally identifiable information-aware catching) and observability (cache hit service-level objective, drift alarms). Attendees will leave with concrete recipes to make hybrid AI production-ready on constrained edge fleets.

Targeting the Tail: How Adaptive Deep Learning Architectures Achieve High-precision Personalization at Scale

Vinay Venkatesh*

AI/ML Lead at Google Search, Senior IEEE Member, San Francisco, California, United States

Abstract

The frontier of personalized recommendation systems is challenged by the "dilution effect," where niche interests of "power users" are overshadowed by general user behavior. This talk presents a unified research roadmap built upon the foundational transformer-based PinnerFormer architecture, expanding this state-of-the-art deep learning technology through a series of innovations. We introduce the dynamic weighted loss function, a single-model approach that automatically prioritizes rare, sparse domains (e.g., film-noir) during training, achieving up to a 50% lift in Recall@10 on niche categories. Next, we detail an efficient transfer learning framework that fine-tunes this foundation model on specialized data, achieving superior performance with a 70% reduction in computational resources. We conclude by discussing three cutting-edge methods for integrating large language models as semantic enhancers to address cold-start, improve diversity, and achieve the next generation of domain-aware personalization.

Bridging the Artificial Intelligence Literacy Gap in Higher Education: The Artificial Intelligence Literacy Framework for an Artificial Intelligence-driven Workforce

David Frazier and Gwendolyn Davidson*

Hazard Community & Technical College, United States

Abstract

Artificial intelligence (AI) rapidly transforms how knowledge is created, analyzed, and applied across professional fields. AI systems now support research, writing, programming, and decision-making in disciplines ranging from business and healthcare to engineering and communication. Despite this shift, many higher education curricula continue to emphasize traditional digital literacy skills developed prior to the widespread adoption of generative AI technologies. This misalignment has created a growing AI literacy gap between workforce expectations and the competencies currently emphasized in higher education. This presentation examines how higher education institutions can address this gap by redefining digital literacy to include competencies that support effective human–AI collaboration. Drawing on emerging research in AI literacy, workforce skills reports, and instructional observations in higher education learning environments, the presentation identifies critical areas where current systems fail to prepare learners for AI-enabled knowledge work. To address this challenge, the presentation introduces the AI literacy competency framework, a conceptual model that organizes AI literacy into three dimensions: critical understanding of AI systems, responsible engagement with AI technologies, and AI-augmented knowledge work. Within these dimensions, five key competencies are identified: prompt literacy, AI output verification, ethical and responsible AI use, bias and algorithm awareness, and AI-augmented productivity. These competencies support learners' capacity to critically evaluate AI-generated information, understand the limitations of algorithmic systems, and integrate AI tools into professional workflows. By embedding AI literacy competencies across disciplines, higher education institutions can better prepare graduates to participate responsibly and effectively in an increasingly AI-driven workforce and knowledge economy.

Transformer-based Focal-pairwise Attentive Ranking Network for Logical Access Anti-spoofing

Sidan Yin*

San Domenico School, United States

Abstract

Recent advances in AI-based speech synthesis and voice conversion have significantly increased the risk of spoofed speaker authentication, posing serious threats to automatic speaker verification (ASV) systems. This work addresses the ASV spoof 5 logical access track 1 task, which focuses on stand-alone spoofing and speech deepfake detection under extreme class imbalance and ranking-oriented evaluation metrics such as equal error rate (EER) and minimum detection cost function. A transformer-based focal-pairwise attentive ranking network is proposed for this task. The model deploys a transformer encoder to capture long-range time–frequency dependencies and integrates attention pooling to highlight sparse and localized spoofing artifacts that are often overlooked by conventional global pooling methods. To better align optimization objectives with ranking-based

metrics, a joint loss combining class-imbalance-aware focal classification loss and pairwise ranking loss is adopted, encouraging improved score separation between bona fide and spoof samples. Raw boost waveform-level augmentation is incorporated to enhance robustness against distribution shifts and diverse spoofing manipulations. In addition, a lightweight test-time augmentation strategy is applied during inference to improve prediction stability. Under the closed training conditions, relying solely on the official ASV spoof 5 datasets without external corpora, the proposed system achieves an EER of approximately 12 - 13%, reaching performance comparable to approaches that utilize additional training data.

Day-4: Virtual Sessions

Keynote Presentation

How Generative Artificial Intelligence Really Works

Neil F. Johnson[1,2*], Dylan J. Restrepo[1,3] and Nicholas J. Restrepo[1]

[1]*d-AI-ta Consulting, United States*

[2]*George Washington University, United States*

[3]*Cornell University, United States*

Abstract

Do you trust artificial intelligence (AI)? The rapid deployment of generative AI systems (e.g., large language models (LLMs) such as ChatGPT, Claude, and Gemini) is outpacing our ability to explain their risks and failures. This undermines trust in AI's deployment for the employees and customers of any business, and for the public at large – particularly in the areas of health, finance and defense where lives are livelihoods are often at risk. In this talk, we present our new research that addresses this trust gap by uncovering the fundamental mechanisms behind sudden shifts in LLM behavior – when output mid-response becomes inaccurate, misleading, or harmful. We derive an exact analytical formula that predicts when such "Jekyll-and-Hyde" tipping points occur at the core of transformer-based architectures. Our work identifies the root cause as a geometric breakdown in the model's attention dynamics: when attention is spread too thinly across competing inputs, it abruptly snaps, leading to qualitative shifts in output. Crucially, this formula is derived using only secondary school mathematics, offering a rare bridge between cutting-edge AI research and public, legal, and policy discourse. It enables a transparent and accessible framework for AI auditing, interpretability, and regulation – without compromising system performance. This work has very broad implications: it empowers engineers to design more stable systems, equips regulators and businesses with scientifically grounded oversight tools, and provides the public and decision-makers with a clear, shared language for understanding the behavior and risks of modern AI. Through this approach, we demonstrate how deep technical insight can enable more explainable, governable, and ultimately trustworthy and hence profitable AI.

Featured Presentations

Artificial Intelligence-driven Advanced Sensing: Enhancing Reliability and Predictive Diagnostics in Complex Engineered Systems

Kunal Goyal*

Georgia Institute of Technology, United States

Abstract

The presentation will focus on the critical effects of combining artificial intelligence (AI) and machine learning (ML) with next-generation sensor technologies. The objective is to improve reliability and support predictive diagnostics in complex systems. The primary focus will be on how AI/ML algorithms can, when used with data from advanced sensors, provide early and smart predictions of potential issues. I will demonstrate methods for using AI-driven pattern recognition and predictive modeling with various sensor inputs, such as thermal, chemical, and pressure signatures, to predict system degradation and prevent serious accidents. The talk will highlight the real benefits of these innovative sensing systems. These include better operational resilience, improved maintenance strategies, and stronger safety rules in various critical applications. Attendees will gain a solid understanding of the latest developments in AI-driven intelligent sensing and their crucial role in making technological systems more robust and secure.

Multi-stage Biomedical Feature Selection Extraction Algorithm for Cancer Detection

Mohitkumar Pravinchandra Bhadla*

Gandhinagar University, Gujarat, India

Abstract

Cancer is a significant cause of death worldwide. Early cancer detection is greatly aided by machine learning and artificial intelligence (AI) to gene microarray data sets (microarray data). Despite this, there is a significant discrepancy between the number of gene features in the microarray data set and the number of samples. Because of this, it is crucial to identify markers for gene array data. Existing feature selection algorithms, however, generally use long-standing, are limited to single-condition feature selection and rarely take feature extraction into account. This work proposes a multi-stage algorithm for biomedical deep feature selection (MBDFS) to address this issue. In the first, three feature selection techniques are combined for thorough feature selection, and feature subsets are obtained: in the second, an unsupervised neural network is used to create the best representation of the feature subset to enhance final classification accuracy. Using a variety of metrics, including a comparison of classification results before and after feature selection and performance of alternative feature selection methods, we evaluate MBDFS's efficacy. The experiments demonstrate that although MBDFS uses fewer features, classification accuracy is either unchanged or enhanced

Pricing Intelligence: What Frontier Artificial Intelligence Economics Means for Builders and Buyers

Charu Maheshwari*

Anthropic, United States

Abstract

Pricing is where strategy meets science. For businesses building with or on top of artificial intelligence (AI), pricing decisions made inside frontier labs directly shape their costs, product roadmaps, and competitive position. And for AI companies themselves, pricing is one of the most consequential and least understood strategic levels available. This session brings an insider perspective from the frontier to practitioners navigating these decisions in real time. Drawing on direct experience in product finance and strategy at Anthropic, where pricing decisions touch everything from token economics to enterprise packaging to model launch strategy, we'll examine what makes AI pricing fundamentally different from every pricing model that came before it, and what that means for the businesses consuming, building on, and competing with frontier AI. We'll cover what actually drives AI economics, and why token-based consumption models create cost structures that behave nothing like SaaS. We'll look at how frontier labs think about the trade-offs between simplicity and sophistication in pricing design, why complex pricing can be a one-way door in a still-evolving market, and how outcome-based models are beginning to reshape what value means in AI. Attendees will leave with a practical framework for evaluating AI pricing, concrete tools for making smarter budgeting

decisions, and a clearer mental model of the strategic forces shaping how intelligence gets packaged and priced today.

Operationalizing Artificial Intelligence Ethics: Fairness, Transparency, and Human Alignment in Recommender Systems

Ravi Mandliya*

Discord Inc., United States

Abstract

Recommendation systems shape billions of daily digital interactions i.e., what people watch, buy, read, and believe. Yet optimizing purely for engagement can amplify bias, opacity, and echo-chambers, undermining user trust. This talk presents engineering frameworks for trustworthy recommender systems built on the principles of fairness, transparency, and accountability. We explore practical methods to embed ethical considerations throughout the machine learning pipeline from bias-resilient data sampling and fairness-aware loss functions to post-processing re-ranking, explainable user interfaces, and audit dashboards. Real-world case studies, including fairness-aware ranking in professional networking platforms, show how trustworthiness can be operationalized without sacrificing performance. Attendees will learn how to design evaluation metrics that balance accuracy with equity, implement user-facing explanation layers, and establish governance loops for ongoing bias monitoring. By reframing trust as an engineering constraint rather than a philosophical aspiration, this session bridges research and production practice to help teams build responsible, human-aligned recommendation systems.

Integrating Transcendental Elements into Artificial Intelligence-powered Serious Games for a Holistic Enhancement of Geriatric Health and Well-being

Hari C. Gangadharan*

Academic Programs, Ocean Grove Charter School, Campbell, United States

Abstract

The recent advancements and extended usage of artificial intelligence (AI)-powered serious games in geriatric healthcare have stirred a discussion surrounding the need to find a balance between technical growth to gain faster positive outcomes in patient recovery, and enhancement of holistic human well-being. As evident in the extant literature, the present focus of AI-powered serious games and their impact on the health and well-being of geriatric individuals has predominantly centered on their effects on limited goals of physical recuperation and cognitive enhancement alone. While the concept of gerotranscendence has found substantial consideration within the broader spectrum of research pertaining to geriatric health care and well-being, it has yet to be incorporated into the realm of employing AI-powered serious games as a means to accommodate and address the transcendent requirements of the elderly within healthcare settings. The present paper delineates the significance of addressing transcendent needs in the context of elderly care, drawing upon theories of Maslow's hierarchy of needs, the theories of gerotranscendence and Stebbin's philosophical discussions surrounding serious leisure activities as a theoretical framework. It strongly advocates for the inclusion of such considerations within the design and implementation of AI powered serious games in healthcare settings, with the ultimate goal of providing a more comprehensive and holistic approach to geriatric care.

Modernizing Legacy Enterprise Data Platforms: From Traditional Extract, Transform, Load to Cloud Lakehouse Architectures

Avik Datta*

Northwestern University, United States

Abstract

The adoption of generative artificial intelligence (GenAI) in healthcare offers significant opportunities to enhance clinical decision support, automate documentation, and improve personalized care. However, deploying GenAI in regulated healthcare environments introduces critical challenges related to data privacy, model transparency, regulatory compliance, and operational trust. This presentation introduces a trust-aware generative AI pipeline architecture designed to enable safe, complaint, and auditable GenAI integration within regulated clinical systems. We present a modular, cloud-native pipeline that integrates data governance controls, privacy-preserving preprocessing, model risk management, and continuous compliance monitoring across the end-to-end GenAI lifecycle. The architecture incorporates policy-aware data ingestion, secure data-in-motion and data-at-

rest mechanisms, explainability layers to support clinical interpretability, and automated audit logging aligned with regulatory requirements such as health insurance portability and accountability act and emerging AI governance guidelines. Using a synthetic healthcare dataset and simulated clinical workflows, we evaluate the pipeline across key trust metrics, including data leakage risk, inference traceability, latency overhead, and model behavior stability under regulatory constraints. The session highlights practical design patterns, governance integration strategies, and operational trade-offs involved in deploying GenAI systems in high-compliance healthcare environments. This talk provides actionable guidance for architects, data engineers, and AI practitioners seeking to operationalize GenAI responsibly within regulated healthcare and enterprise platforms.

Shaping Scalable Data Platforms with a Human Touch

Sivadeep Katangoori*

Insight Global, North Carolina, United States

Abstract

Hi, I'm Sivadeep Katangoori a Data Lake and Cloud Specialist with over a decade of experience navigating the complexities of large-scale data systems across financial, security, and enterprise sectors. My expertise lies in cloud-native architecture, big data engineering, and artificial intelligence (AI) enablement on platforms like Google Cloud and Amazon Web Services. I've led end-to-end technical transformations, from architecting secure data lakes to modernizing machine learning workflows using Vertex AI, BigQuery, and Kubernetes. Beyond the tech, I believe in people. I've successfully managed cross-functional teams, conducted client discovery calls, mentored engineers, and translated business problems into scalable technical solutions. I'm certified as a Google Professional Cloud Architect and Data Engineer and hold a PMP certification that complements my engineering mindset with structured project leadership. From deploying TLS/SSL across clusters, automating DevOps pipelines, to building disaster recovery systems and securing multi-tenant data platforms, I've worn many hats and enjoyed every challenge. What drives me is the opportunity to build solutions that not only scale with data, but scale with purpose.

Introduction and Evaluation of Artificial Intelligence Powered Coding Tools for Teaching Computer Programming to Software Engineering Students

Sarita Singh*

Northeastern University, United States

Abstract

Leading software companies are increasingly expecting computer science graduates to be proficient in using various artificial intelligence (AI) powered coding tools to enhance productivity, accelerate debugging and improve code quality. Educators across the globe are faced with the challenges of teaching students to use these coding tools effectively without compromising on software security, reliability, or ethical standards. There are few established teaching models or best practices that guide how AI tools should be integrated into programming curricula while maintaining educational rigor and teach students how to verify, test, and secure AI-assisted code. This research study focuses on providing guidelines to introducing AI powered coding tools such as GitHub Copilot into computer programming courses and evaluates its impact on students' learning outcomes, coding practices, and perceptions. A mixed method approach was adopted to understand how students felt, what they learned, and how Copilot influenced their learning behavior. The findings can help to modify existing traditional pedagogical strategies for integrating AI coding tools into software engineering curricula. The study contributes to understanding how AI can enhance programming education while preserving critical thinking, academic integrity, and prepare students to meet future career challenges.

Simplicial Nonnegative Matrix Factorization with Applications to High-dimensional and Large-scale Datasets

Michail Tsagris[1*] and Abdulaziz Alenazi[2]

[1]*Department of Economics, University of Crete, Crete, Greece*

[2]*Department of Mathematics, College of Science & Center for Scientific Research and Entrepreneurship, Northern Border University, Arar, Saudi Arabia*

Abstract

We propose simplicial nonnegative matrix factorization for simplicial, or compositional, data by imposing simplicial restrictions on the rows of the basis and the coefficient factorization matrices. We examine two widely used objective functions, the squared Frobenius norm and the KullbackLeibler divergence and propose two algorithms for either case. We further examine the cases of large-scale data and/or of high dimensional data. Experiments using real data reveal that the Frobenius norm is to be preferred and depending on the case scenario, a hybrid algorithm may be used. The resulting algorithm is quite fast and memory efficient.

A Comprehensive Survey of Mixture of Experts Architectures

Dev Vyas[1*] and Mohit Bhadla[2]

[1]*Georgia State University, United States*

[2]*Gandhinagar University, India*

Abstract

This survey provides a comprehensive review of the mixture of experts (MoE) architecture, a pivotal technique for scaling neural networks to trillions of parameters while maintaining computational tractability. We trace the evolution of MoE from its foundational concept as an ensemble of "local experts" to its modern incarnation as a sparsely activated, conditional computation framework integral to state-of-the-art large language models and other domains. The report details core methodologies, including routing algorithms (e.g., Top-k, expert choice) and load balancing strategies (e.g., auxiliary losses, expert capacity). We conduct an in-depth analysis of landmark models such as GShard, Switch Transformers, Mixtral 8x7B, and DeepSeek-MoE, and explore extensions into computer vision (V-MoE) and multi-task learning. We further discuss MoE-specific evaluation metrics, systemic challenges like training instability and inference latency, and critical ethical considerations, including environmental impact and bias amplification. Finally, we synthesize recent innovations from 2024 - 2025 and outline key future research directions, positioning MoE as a cornerstone of next-generation AI.

Received: May 04, 2026 **Accepted:** May 07, 2026 **Published:** May 08, 2026

Citation: Proceedings of the 3rd International Conference on Artificial Intelligence and Machine Learning (AIM - 2026). *USG Proc J* 2(Suppl 17): S1-S25.

Copyright: © 2026. This is an Open Access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY) (http://creativecommons.org/licenses/by/4.0/) which permits commercial use, including reproduction, adaptation, and distribution of the article provided the original author and source are credited.