

ABSTRACT BOOK

3rd International Conference on
**ARTIFICIAL INTELLIGENCE
AND MACHINE LEARNING
(AIM-2026)**

APRIL 27-29, 2026
SAN FRANCISCO, CA

DoubleTree by Hilton
San Francisco Airport,
835 Airport Blvd Burlingame, CA

Plenary Presentations

When computers look at art

David G. Stork

Stanford University, CA

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 comprise some of the most important, memorable, and consequential images ever created, and present 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—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 artificial intelligence.

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 addresses new grand challenges in artificial intelligence.

Biography:

David G. Stork, PhD holds Adjunct faculty positions in one departments and two programs at Stanford University and is a graduate in Physics from MIT and the University of Maryland; he also studied Art History at Wellesley College. He has held faculty positions in Physics, Mathematics, Computer Science, Statistics, Electrical Engineering, Neuroscience, Psychology, Computational Mathematical Engineering, Symbolic Systems, and Art and Art History variously at Wellesley and Swarthmore Colleges, Clark, Boston, and Stanford Universities, and the Technical University of

Vienna. He is a Fellow of seven international societies and has published nine books/proceedings volumes, 220+ scholarly articles, and 64 US patents. His *Pixels & paintings: Foundations of computer-assisted connoisseurship* (Wiley) appeared this year and he is completing *Principled art authentication: A probabilistic foundation for representing and reasoning under uncertainty*.

GenAI in Education: Issues and Experiences from SFSU

Dragutin Petkovic

San Francisco State University, CA

Abstract:

GenAI in Education: Issues and Experiences from teaching Computer Science at SFSU Generative AI (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.

Biography:

Prof. D. Petkovic obtained his Ph.D. at UC Irvine, in the area of biomedical image processing. He spent over 15 years at IBM Almaden Research Center as a scientist and in various management roles. His contributions ranged from use of computer vision for inspection, to multimedia and content management systems. He is the founder of IBM's well-known QBIC (query by image content) project, which significantly influenced the content-based retrieval field. Prof. Petkovic received numerous IBM awards for his work and became an IEEE Fellow in 1998 and IEEE LIFE Fellow in 2018 for leadership in content-based retrieval area. Prof. Petkovic also had various technical management roles in Silicon Valley startups. Research and teaching interests of Prof. Petkovic include AI and Machine Learning with emphasis on Explainability, Ethics and Trustworthiness, GenAI in education, and the design and development of easy to use systems. He is also Associate Editor of IEEE Transactions on Technology and Society. Currently, Dr. Petkovic is the Professor Emeritus at the SFSU Department of Computer Science.

Key Note Presentations

AI: Tool or Intelligent Entity?

Craig Kaplan

iQ Company, CA

Abstract:

Many of us have assumed that 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.

Biography:

Dr. Craig A. Kaplan is the CEO of iQ Company and the founder of SuperIntelligence.com, specializing in AGI and SuperIntelligence. He previously founded PredictWallStreet, where he invented intelligent systems used by hedge funds. He holds numerous patents on AI-related technologies, and his work focuses on applying collective intelligence and quantitative modeling to design systems that learn, adapt, and make decisions in complex environments. His publications include a book, scientific journal articles, and white papers on SuperIntelligence. Kaplan earned his PhD from Carnegie Mellon, where he co-authored research with Nobel Laureate Dr. Herbert A. Simon, a pioneer in artificial intelligence and cognitive science.

From Software Test Automation to Agentic AI Test Automation

Jerry Zeyu Gao

San Jose State University, CA

Abstract:

In this talk, Dr. Jerry Gao discusses the major problems, challenges, and needs in test automation in today's 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.

Biography:

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 27 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 three 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 382K+. His current research areas include machine learning, smart cities, AI test automation, AI clouds and mobile clouds. In 2010, Jerry Gao has been recognized by University of Texas at Arlington as a distinguished Alumna 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, Journal of Smart Cities, Energies, and Agriculture. Dr. Gao has been included and listed in Marquis American Who's Who 2020-2022. In last 10 years, 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, SEKE06-2016, IEEE SOSE2011-2013, ICYCS'05, TQACBS2005-2006, WMCS2004-2010, IEEE EMOBS07-08, TEST'07, and EECC2006. Now, Dr. Gao has served in editorial boards for different archived journals. Recently, Dr. Gao has been invited by many international conferences as a keynote speaker to address topics in Ai Test Automation, Smart Cities, UAV AI cloud, Green Energy Ai Cloud, Smart Agricultures, and Intelligent Wildfire Platforms. 2025.

Featured Presentations

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

Mahesh Balkrishna Chaudhari

University of San Francisco, CA

Abstract:

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 systemslevel thinking that production AI environments demand. This presentation offers a practitionereducator's perspective on bridging this gap between data science and dataengineering, the two disciplines that are increasingly inseparable in realworld AI deployments. Attendees will gain insight into how the curriculum models industry workflows using tools such as Apache Spark, Airflow, Kafka, among others, within cloudbnative platforms such as AWS, GCP, and Databricks; how practical projects emulate building the end-to-end AI ready data pipelines; and how studentfaculty collaborative research has produced peerreviewed contributions at IEEE and ACM conferences. The talk also focuses on responble data practices—governance, lineage, and security—recognizing that trustworthy AI begins with trustworthy data.

Biography:

Mahesh Chaudhari is an assistant professor in the Master of Science in Data Science and Artificial Intelligence (MSDSAI) program at the University of San Francisco (USF). He is also the director of the Data Engineering Program. He specializes in databases, cloud computing, and big data. At USF, Dr. Chaudhari teaches courses on distributed systems, data lakehouse, stream processing, and data architecture, with emphasis on incorporating the latest data infrastructure technologies from industry into academics. Prior to joining USF, Dr. Chaudhari spent 11 years in the industry, where he successfully

managed and led teams in constructing autonomous data infrastructure in a multi-cloud environment.

From Proof of Concept to Production: Operationalizing AI in Enterprise Systems

Darren Kimura

AI Squared, CA

Abstract:

Most 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 enterprise 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.

Biography:

Darren Kimura is a technology executive, entrepreneur, and investor with more than three decades of experience building and scaling companies across energy, enterprise software, cybersecurity, and artificial intelligence. He serves as Chairman and CEO of AI Squared, helping enterprises and U.S. government agencies securely deploy AI in production environments. Darren has created over \$1 billion in enterprise value, raised more than \$350 million in capital, invested in advanced energy technologies, and is developing the Claremont AI Controls Model. He is a PhD student at Claremont Graduate University focused on AI governance and leadership.

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

Anshu Arora

University of the District of Columbia, Washington DC

Abstract:

Over the past decade, social robotics has evolved significantly through the integration of Explainable Artificial Intelligence (XAI) and Interpretable Artificial Intelligence (IAI) to advance user experience (UX) and customer experience (CX). This research presents a decade-long review (2015–2025) of frameworks in social robotics and human-robot interaction (HRI) 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 customer experience through more meaningful and interpretable interactions.

Biography:

Anshu Saxena Arora, Ph.D., PMP is a Tenured Professor of Marketing, Artificial Intelligence (AI), and Social Robotics at the School of Business and Public Administration, University of the District of Columbia (UDC), Washington, DC, USA. Dr. Arora is the Founding and Current Director of AI, Social Robotics, and Behavioral Research Laboratory at UDC, funded by the National Science Foundation. She has been awarded several grants from the National Science Foundation, including the prestigious Faculty Research Initiation Award for her project titled Social Motivation Approach for Rehabilitation Through Educational Robotics (SMARTER) research where she works with social robots (NAO and Pepper robots) targeted at students diagnosed with Autism Spectrum Disorder (ASD) and other learning disorders at high schools / universities in Washington, DC. Dr. Arora is a Certified Project Management Professional (PMP) from Project Management Institute (PMI), USA, and holds a Foundation Certificate in IT Service Management from Information Systems Examination Board (ISEB), UK. Dr. Arora has published 60+ research papers in national and international journals of repute and presented 70+ papers in national and international conferences.

Evaluating a ReAct-Based Agent for Agricultural Planning

Karanbir Singh

Salesforce, USA

Abstract:

Planning is a core aspect of intelligence, enabling both humans and 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 (RAG), weather forecast APIs, 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.

Biography:

Karanbir Singh is a Senior Software Engineer at Salesforce with over seven years of expertise in AI/ML engineering, backend systems, and time series analysis. As a Senior IEEE Member and speaker, he has presented on AI, Responsible AI, and RAG systems. His research includes Kubernetes resource optimization and AI in agriculture.

Agentic AI-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 , Jon Coffman 1

Affiliation: 1 Bioprocess Technologies and Engineering, Biopharmaceutical Development, R&D

AstraZeneca, Gaithersburg, US. 2 Data Science and Modelling, Biopharmaceutical Development, R&D AstraZeneca, Gaithersburg, US.

Abstract:

Integrated Dynamic Bioprocessing (IDB) 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 UV 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.

Biography:

Tianxin received her PhD in chemical engineering from JHU in 2023, focusing on algorithm design and software development. During her doctoral studies, she developed a novel media development algorithm, which has been licensed by over 30 leading biopharmaceutical companies as a transformative tool for designing cell culture media and precipitation prevention. This innovation significantly reduced reliance on traditional experimental approaches, streamlining media design strategies across the industry. Her expertise includes machine learning, digital twin, and agentic AI. At AstraZeneca, she leverages advanced AI methodologies to drive predictive analytics and process optimization, accelerating drug development and enabling smarter, more efficient biomanufacturing solutions.

Artificial Intelligence Framework for Predicting Injury vs. Non-Injury Crashes

Roumen Vesselinov^{1*}, Elena Vesselinov¹, Mark Scarboro¹, William Teeter¹.

¹University of Maryland at Baltimore, USA.

Abstract:

In this study, we develop a complex 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 (CNN), 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. (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 further consists of three types of machine learning-based techniques: Classification and Regression Tree (CART) models, Random Forest (RF), and Extreme Gradient Boosting (XGBoost). The framework is also complemented by SHAP analysis, Principal Component Analysis and cluster analysis. The predictive modeling of injury versus non-injury crashes underscores the limitations of relying only on structured crash report data, which produced modest accuracy with 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.

Biography:

Roumen Vesselinov, PhD, is an Assistant Professor in the Biostatistics and Bioinformatics Division, Department of Epidemiology and Public Health, University of Maryland, Baltimore, and faculty at the National Study Center for Trauma and EMS. He has authored over 60 peer-reviewed publications. Dr. Vesselinov's research focuses on artificial intelligence, machine learning, and natural language processing (NLP) applications for trauma and injury analytics, particularly in crash-related outcomes. He has served as Principal Investigator and Co-Investigator on projects funded by the NIH, Department of Defense, CDC, and National Highway Traffic Safety Administration.

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

Rafael Geurgas^{1*}, Saul J Newman^{2,3,4}, Evelina T Akimova^{1,5,6}, Katherine N Thompson¹ and Robbee Wedow^{1,5,6,7}

Department of Sociology, Purdue University, USA.

Centre for Longitudinal Studies, University College, UK.

University College, University of Oxford, UK.

Institute of Population Aging, University of Oxford, UK.

Department of Statistics, Purdue University, USA.

Center on Aging and the Life Course (CALC), Purdue University, USA.

Department of Medical and Molecular Genetics, Indiana University School of Medicine, USA.

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, XGBoost, Support Vector Machine, and Neural Networks - using data from a nationally representative longitudinal study of the U.S., 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 ROC-AUC compared to the benchmark Logistic Regression model. While this is a slight performance improvement, overall, Logistic Regression performs about

as well as many of our ML 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.

Biography:

Rafael Geurgas is a postdoctoral researcher at Purdue University working in sociogenomics. Originally from Brazil, he has a background in physics, including a master's degree in ultrahigh-energy cosmic rays. He earned his Ph.D. applying machine learning and natural language processing (NLP) to large-scale textual data, studying the formation of ideological misinformation and echo chambers around COVID-19 vaccines on Twitter and their effects. His current research focuses on developing machine learning approaches that integrate genetic and environmental data to study complex human traits, particularly mental health outcomes such as depression.

A Quantum-Ensembled QLSTM-QCNN Framework for Wind Power Forecasting and Grid Stability

Nithya Venkatesan^{1*}, Sohorab Hossain¹, Rameez Kurshi², Eldad Avital³

*1Vellore Institute of Technology, Chennai, India; 2University of Hull, UK; 3Queen Mary University of London, London, UK
Email: nithya.v@vit.ac.in, +917708706730*

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 (QLSTM-QCNN) Forecasting Framework, which integrates classical deep learning models—Long Short-Term Memory (LSTM) and Convolutional Neural Networks (CNN)—with quantum computing elements via variational quantum circuits. The ensemble leverages quantum parallelism and entanglement to enhance model generalization and robustness. We evaluate the proposed framework on the NREL 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 an R^2 score of 0.9312, RMSE of 0.020106, MAE of 0.011823, and an inference latency of 0.01 seconds, highlighting its suitability for real-time control and integration within smart grid infrastructures.

Biography:

NITHYA VENKATESAN received the B.Sc. and M.Sc. degrees in Industrial Electronics from Bharathidasan University, Tiruchirappalli, Tamil Nadu, India, in 1992 and 1994, respectively. She received the Ph.D. degree in Industrial Electronics Process Control from Bharathidasan University, Tiruchirappalli, Tamil Nadu, India, in 2009. She is currently a professor in the School of Electrical Engineering at Vellore Institute of Technology, Chennai, Tamil Nadu. She has published around 50+ papers in different conferences and journals. Her research interests include nonlinear control, intelligent controllers, optimization techniques, fractional-order controllers and artificial intelligence with applications to robotics, different nonlinear systems, wind and ocean engineering systems.

AI in Advanced Manufacturing: Educating the Next Generation Technician Workforce

Presenter* and Co-author names: Karen Wosczyzna-Birch^{1*}

Affiliation (Institution/organization name, Country): 1National Center for Next Generation Manufacturing, USA

Abstract:

This session explores how artificial intelligence 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 Artificial Intelligence Consortium, participants will learn about current and emerging uses of artificial intelligence in advanced manufacturing companies and what future manufacturing technicians need to know. The advanced manufacturing industry is using artificial intelligence 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.

Biography:

Dr. Karen Wosczyzna-Birch is Executive Director and PI of the National Science Foundation funded National Center for Next Generation Manufacturing and Professor Emeritus in Applied Technology, Connecticut State Community College. She has received over \$40M in NSF funding. Since 1995, she served as CT College of Technology Director and created nationally recognized pathway programs in engineering and technology. Karen has received numerous awards including the CT Women of Innovation Award, the New England Board of Higher Education Award, and two ITEEA recognitions. She serves on various boards including those for the Connecticut Technical Education and Career System and MIT's TechAMP.

The Real Constraints of Building Agent First Search

Presenter* and Co-author names: Jay Chia

Affiliation (Eventual, USA)

Abstract:

Search for 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.

Biography:

Jay Chia is one of the Co-Founders of Eventual, the company behind Daft, an open-sourced project. Prior to Eventual, he was a software engineer working on ML infrastructure for computational biology at Freenome and self-driving cars at Lyft, building large-scale data and computing platforms for diverse industries. Jay hails from the sunny island nation of Singapore, and a fun fact is that he used to command a platoon of tanks in the Singapore military.

Poster Presentations

Empirical Evaluation of Deterministic Validation Layers for Large Language Model

Gary Langford

3x3 Institute, USA

Abstract:

Large Language Models (LLMs) are increasingly deployed in domains where correctness, consistency, and traceability are critical, yet their outputs remain 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.

Biography:

Gary Langford is an independent AI researcher and systems engineer focused on improving the reliability and trustworthiness of large-scale AI models. His work develops and empirically evaluates deterministic validation and repair frameworks that

operate alongside probabilistic systems to improve correctness, consistency, and stability without retraining or modifying the underlying models. He evaluates these methods across multiple foundation model platforms and task domains, including reasoning, compliance, and decision support. His background spans applied mathematics, physics, and complex systems engineering, with prior work in aerospace and scientific computing. Interests include practical approaches to AI governance and safe deployment across real-world applications.

Technological Reciprocity: Modeling the AI Productivity Dividend in Automated Economies

Reza Abdolee

California State University Channel Islands

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 \$110M 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*, Renu Balyan

State University of New York at Old Westbury, USA

Abstract:

There is an unprecedented incidence of diagnosed anxiety and depression among U.S. 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 (NLP) 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 architectures (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 (ML) 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.

Biography:

Anika Bokka is a senior Computer & Information Sciences student at SUNY Old Westbury. Previously, she was an artificial intelligence & machine learning fellow at Cornell Tech's Break Through Tech program, during which she worked on computer vision in autonomous vehicle systems.

Dr. Renu Balyan is a professor in the Computer & Information Sciences department at SUNY Old Westbury. She teaches graduate and undergraduate courses. Her research interests lie in the integration of Machine Learning, Deep Learning, and NLP for Health & Education Data, Health Literacy, and Mental Health.

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

Seyedezahra Paylakhi¹, Michel G. Nivard^{2,3*}, Robbee Wedow^{1*}

1Department of Sociology, Purdue University, West Lafayette, IN, USA

2Medical Research Council Integrative Epidemiology Unit, University of Bristol, Bristol, UK

3Population Health Sciences, Bristol Medical School, University of Bristol, Bristol, UK

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 priors within existing GWAS, TWAS, 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.

Biography:

Seyedezahra Paylakhi is a postdoctoral researcher at Purdue University working at the intersection of statistical genetics, machine learning, and functional genomics. Her research focuses on integrating multi-omic data and developing computational methods to better understand how genetic variation influences molecular and cellular phenotypes. She is particularly interested in combining deep learning approaches, including protein and DNA language models, with population-scale genetic analyses to improve the interpretation of complex trait genetics.

End of Day-1

Keynote Presentations

Beyond Probability: The Certainty of AI

Seth Dorbin

Chief Executive Officer, Qantm AI, CT

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 WLDA 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."

Highlights: Strategy: Directly led the development of an AI strategy with quantifiable metrics tied directly to organizational KPI at 2 Fortune 250 companies and at many of the Global Fortune 500 companies using a repeatable methodology honed over a decade of application. Execution: Developed scalable and robust process for the execution of an AI strategy that is tied directly to the business goals, leading to meaningful KPI for the projects resulting in measurable cost savings and new revenues.

Innovation: Combining his scientific background with business acumen, he drives innovation at the intersection of AI, data analytics, and business transformation while advocating for responsible and inclusive AI. This is especially important with the recent advances in generative AI, in which weekly step changes occur. Eminence: He is a sought-after speaker headlining technology conferences, industry conferences, and private equity and VC events. He has appeared multiple times in tier-one popular and technology press on the BBC, PBS, NPR, Forbes, The Atlantic, and others.

Recognition: His pioneering work has garnered multiple prestigious accolades such as 2024 AI for Good Champion by DataIQ and 2024 Women Leaders In Data & AI (WLDA) AMBASSADOR OF THE YEAR, AIConics Innovator of the Year, Top 10 Most Influential Chief AI Officers 2023, and Top 50 Tech Entrepreneurs 2023, signifying global recognition for his leadership and commitment to ethical AI practices.

Featured Presentations

Machine Learning

Machine Learning for Physical Layer Communications: Let There Be Lightness

Benjamin Belzer^{1*}, Krishnamoorthy Sivakumar¹, Ahmed Aboutaleb², Amirhossein Sayyafan³, and Ali Kanaani¹

1Washington State University School of EECS, USA; 2Marvell Technology, Inc., USA; 3Western Digital Corporation, USA

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-than-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 (CNNs) 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 SVD-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.

Biography:

Benjamin Belzer received a B.A. in Physics from U.C. San Diego in 1982, and a Ph.D. in Electrical Engineering from UCLA in 1996. In 1996, he joined the faculty of the School of EECS at Washington State University, Pullman, WA, where he is currently an Associate Professor. His research interests are in physical layer digital communications, including machine learning (ML) for equalization, channel estimation and non-orthogonal multiple access (NOMA) on wireless channels, and ML for two- and three-dimensional magnetic recording. He has co-authored over 100 refereed conference and journal publications that have collectively received over 4200 citations, per Google Scholar.

Risk-averse Reinforcement Learning in Portfolio Optimization

Shimin Zhang 1 *, Kandethody M. Ramachandran 2

Department of Mathematics and Statistics/University of South Florida, USA

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 Measure (ERM) is used to estimate the risks involved with uncertain outcomes. In this study, we propose the DDPGERM algorithm, which incorporates ERM into the DRL algorithm, Deep Deterministic Policy Gradient (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 (EM). 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.

Biography:

Shimin Zhang is a Ph.D. candidate in the Department of Mathematics and Statistics at the University of South Florida. Her research focuses on applying machine learning methods for risk-aware decision making, with a special focus on deep reinforcement learning for portfolio optimization under uncertainty. Her work integrates risk measures with deep reinforcement learning algorithms to address financial decision-making problems.

Annotation-Efficient Object Detection Under Domain Shift Using Hybrid Active and Semi-Supervised Learning

In Seon Kim*, Ali Moghimi

Computer Science Department, University of California, Davis, USA

Biological and Agricultural Engineering Department, University of California, Davis, USA

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 (RF-DETR with a DINOv2 backbone) 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.

Biography:

In Seon Kim is a Master's student in Computer Science at the University of California, Davis, expected to graduate in June 2026. Her research interests include computer vision and machine learning with a focus on annotation-efficient learning strategies and domain adaptation methods for large-scale visual data. She is currently a member of the UC Davis Robotics Challenge Team, where she leads the AI team in segmenting objects from 3D-reconstructed scenes. Her work explores the application of machine learning techniques to real-world problems involving visual data and photogrammetry.

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

Matthew Waismann¹, Somansh Budhwar², Armin Schwartzman^{2*}

¹Amazon Inc., USA; ²University of California, San Diego, USA

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 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. The 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 $0.2\% \pm 0.01\%$ per year.

Biography:

Armin Schwartzman received his bachelor's and master's degrees in electrical engineering from the Technion - Israel Institute of Technology and the California Institute of Technology; and his PhD in Statistics from Stanford University. He was an R&D engineer at Rockwell Semiconductor and Biosense Webster, and has held faculty positions in Biostatistics at Harvard University and Statistics at North Carolina State University. He is now Professor at the University of California, San Diego, with a joint position between the Halicioglu Data Science Institute and Biostatistics. Professor Schwartzman's research encompasses theoretical and practical aspects of statistical signal and image analysis in a variety of scientific applications. These include spatio-temporal and high-dimensional data analysis, geometric statistics and smooth Gaussian random fields, with applications in neuroscience, genetics and the environment.

Architecting and scaling a unified, role-aware agentic AI chatbot for field operations in a large telecom enterprise

Dipjyoti Das^{1*}

¹Ascendion Inc, USA

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 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, REST-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 JWT-based authentication, LLM-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 APIs. A Tech Assist agent leverages retrieval-augmented generation (RAG) over large-scale document corpora, while supervisor agents provide BI-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 APIs, the architecture supports horizontal scaling. Observability is implemented using self-hosted Langfuse through hierarchical tracing of agent and LLM 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.

Biography:

Dipjyoti Das is a Generative AI Architect with over 13 years of experience designing and deploying enterprise-scale AI and ML systems across startups and Fortune 150 companies. He specializes in building production-grade GenAI and Agentic AI platforms that deliver real-world impact. His expertise spans LLM systems, RAG pipelines, traditional Machine Learning algorithms, AIOps, and scalable model serving on AWS and Azure Kubernetes environments. Currently, he leads GenAI initiatives for Fortune 100 clients, including deploying an Agentic AI chatbot platform serving over 80,000 users and reducing operational support calls by 60%. He has previously spoken at Ai4 and REWORK AI conferences.

Unbinding Prometheus: How Generative AI is Liberating Computational Thinking from Coding

Douglas C Schmidt

William & Mary, VA

Abstract:

In my talk "Unbinding Prometheus: How Generative AI is Liberating Computational Thinking from Coding," I describe how generative 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.

Biography:

Dr. Douglas C. Schmidt is the inaugural Dean of William & Mary's School of Computing, Data Sciences & Physics. Prior to his role with William & Mary, Dr. Schmidt served as the Director of Operational Test and Evaluation (DOT&E), a President-appointed Senate-confirmed position that advises the Secretary of Defense. He is a highly cited author with over 51,000 citations to date and has published more than 700 papers and co-authored more than 10 books. His recent research focuses on prompt engineering techniques and patterns that enhance the accuracy and expressiveness of large language models and generative augmented intelligence platforms.

Steering LLMs with Light-Weight Auxiliary Policy Models using Reinforcement Learning

Haipeng Chen

Affiliation: William & Mary, USA

Abstract:

Reinforcement Learning (RL) has become a cornerstone in advancing the capability and alignment of large language models (LLMs). Existing works, such as Reinforcement Learning from Human Feedback (RLHF) and Reinforcement Learning with Verifiable Reward (RLVR) 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.

Biography:

Haipeng Chen is an assistant professor and graduate program director of the Data Science Department at William & Mary. He directs the Data-Driven Decision Intelligence (D3i) Lab and serves as the vice chair of the AIBRC at CT&L. Previously, he was a postdoc at Harvard University and Dartmouth College. He obtained his Ph.D. from Nanyang Technological University and B.S. from the University of Science and Technology of China. His research interest lies in reinforcement learning, generative AI, and applications in health, environment, and physics. He has received the best paper nomination at AAMAS-2021, Innovation Demonstration Award runner-up at IJCAI-2019, and Champion of the 2017 Microsoft Malmo Collaborative AI Challenge. He publishes regularly in premier AI conferences such as ICLR, NeurIPS, AAAI, IJCAI, KDD, and journals.

Enhancing NEC Prediction in Preterm Infants with TVAE-Augmented Data Modeling

Presenter* and Co-author names {Shilpa Muralidhar1*, Mohan Pammi2}

1, 2Division of Neonatology, Department of Pediatrics, Baylor College of Medicine & Texas Children's Hospital, USA

Abstract:

Preterm birth impacts 1 in 10 newborns globally, with over 500,000 cases annually in the U.S., increasing risks of severe complications like necrotizing enterocolitis (NEC), a life-threatening intestinal condition. Diagnosing NEC is challenging due to symptom overlap and low prevalence (6.3%), complicating predictive modeling. This study introduces an innovative approach using Tabular Variational Autoencoders (TVAE), a GAN-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, IRB-approved), imputing missing data via Multiple Imputation by Chained Equations (MICE). To counter the 6.3% NEC prevalence, SMOTE-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 (ROC 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.

Biography:

Dr. Shilpa Muralidhar a data scientist at Texas Children's Hospital, and Dr. Mohan Pammi, a neonatologist, are advancing data-driven neonatal care solutions. Dr. Muralidhar, with a PhD in Biomedical Sciences, designs machine learning models, including Tabular Variational Autoencoders, to tackle data imbalance in rare disease prediction. Dr. Pammi, an MD and PhD specializing in preterm infant care, provides critical clinical insights for early necrotizing enterocolitis (NEC) diagnosis. Collaborating with the Vermont Oxford Network, our ongoing research aims to enhance predictive accuracy for neonatal health, with forthcoming publications in medical informatics.

AI as Co-Designer: Enhancing Conceptualization with Generative AI in Interior Design Education

Qun Zuo*

Department of Fashion, Interior Design, and Merchandising, Central Michigan University, Mount Pleasant, Michigan, 48859, USA.

Abstract:

Generative 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 CAD-based outputs. The analysis combined qualitative observations with computational image processing techniques: (1) semantic similarity using CLIP 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.

Biography:

Qun Zuo is an Associate Professor in the Department of Fashion, Interior Design, and Merchandising at Central Michigan University. With a background in architecture and multimedia design, her research focuses on interactive virtual environments, spatial visual attention, and inclusive design pedagogy. She also studies indoor environmental quality for the health and well-being of occupants across age groups.

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

Liubov Kuposova

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.

Biography:

Liubov Koposova is an art historian specializing in Buddhist art and cultural heritage. Her research focuses on issues of attribution, iconographic analysis, and provenance, with a particular interest in the methodological challenges posed by complex cultural objects. She has worked extensively with a large corpus of Buddhist artworks, addressing questions of stylistic diversity, historical transformation, and cross-cultural exchange. Recently, she has extended her work to the intersection of art history and Data Science, exploring how computational approaches can contribute to problems of attribution, authentication, and cultural heritage preservation. Her current research highlights the potential of interdisciplinary collaboration between AI and the humanities. I am from the Russian Federation, relocated to Serbia.

End of Day-2

Featured Presentations

Statistical learning with machine generated data

Presenter* and Co-author names {For Ex: (Yuchen Wu1*, Kangjie Zhou2, Weijie Su3)}

Affiliation: 1. Cornell University, USA 2. Columbia University, USA 3. University of Pennsylvania, USA

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.

Biography:

Yuchen Wu is an Assistant Professor in the School of Operations Research and Information Engineering at Cornell University. Prior to Cornell, she was a postdoctoral researcher in the Department of Statistics and Data Science at the Wharton School, University of Pennsylvania. She received her Ph.D. in 2023 from the Department of Statistics at Stanford University, advised by Professor Andrea Montanari. Her research focuses on establishing rigorous foundations for statistical and machine learning methods, while also developing new algorithms guided by theoretical insights.

Human care and AI regulation

Carlos Montemayor*

**San Francisco State University, Department of Philosophy, USA.*

Abstract:

Recent discussions on 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, Geoffrey Hinton has recently proposed that one way to guarantee this care for humanity is for AI to have maternal instincts. In this talk I 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.

Biography:

Carlos Montemayor (Law degree from UNAM, MA in Philosophy from the New School for Social Research, PhD in Philosophy from Rutgers University) is Professor of philosophy at San Francisco State University, California, USA. He works on consciousness, attention, time perception, knowledge, and memory. Montemayor is the author of *Minding Time: A Philosophical and Theoretical Approach to the Psychology of Time* (Brill, The Netherlands, 2013), and of multiple papers and co-authored books on attention and consciousness. His most recent book is *The Prospect of a Humanitarian Artificial Intelligence: Agency and Value Alignment*.

AI-Powered Profiling of Human Trafficking Buyers

Shreyas Ramani (Presenter) and Dhanashree Pokale, Alexandra Joukova, Vikas Singh, Nnenne Onyioha-Clayton, Spencer A. Murray, Catherine Scott, Dr. Kendra Taylor

Affiliation (Informs Pro Bono Analytics, USA)

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 NLP 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

The Honorable Xiąże John McClellan Marshall

Senior Judge, Fourteenth Judicial District of Texas

Honorary Professor of the University, UMCS, Lublin, Poland

Member, International Academy of Astronautics

"Everything we hear is an opinion, not a fact. Everything we see is a perspective, not the truth." ~ Marcus Aurelius

Abstract:

If Europe imagined the Enlightenment, it was 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 an 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 the Enlightenment to a world view characterized by a subjective interpretation of "reality".

Biography:

Judge Marshall holds degrees from Virginia Military Institute, Vanderbilt University, and Southern Methodist University Dedman School of Law, as well as an honorary Doctor of Laws from the Mexican Academy of International Law. Judge Marshall is an internationally recognized prize-winning author whose works include "The Terminator Missed a Chip!: Cyberethics", International Astronautical Congress, Oslo 1995; MEDEL Conference, Kraków, Poland, 2018 - "The Constitution and the Independent Judiciary"; "The Modern Memory Hole", Athenaeum Review, Fall 2019 (University of Texas at Dallas); "Cyberethics In The 21st Century: The Reign Of The Machines", Mensa Bulletin, January 2020; "Machine-Machine To Major Tom: Cyberethics", IAA/UT Conference on Space Traffic Management, Austin, TX, February 2020; "Technoevidence: The 'Turing Limit' 2020", Journal of AI and Society (Springer 2021), DOI 10.1007/s00146-020-01139-z. "Rewriting the Constitution: Revolution in 1787. . and Now!", Texas Bar College Member Bulletin, September 2025.

End of 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, IN 47907, USA

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 VC-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.

Biography:

Wojciech Szpankowski is the Saul Rosen Distinguished Professor of Computer Science at Purdue University, where he teaches and conducts research in the analysis of algorithms, information theory, analytic combinatorics, random structures, and machine learning for both classical and quantum data. He has held visiting positions at McGill University, INRIA, Stanford University, Hewlett-Packard Labs, the Universite de Versailles, the University of Canterbury (New Zealand), Ecole Polytechnique (France), the Newton Institute (Cambridge, UK), ETH Zurich, the University of Hawaii, Gdansk University of Technology, and Jagiellonian University (Cracow, Poland). In 2008, he launched the Interdisciplinary Institute for the Science of Information. In 2010, he became the Director of the NSF Science and Technology Center for Science of Information. Szpankowski is an IEEE Fellow, Erskine Fellow, and AIIA Fellow. He received the Humboldt Research Award (2010), the Inaugural Arden L. Bement Jr. Award (2015), and the Flajolet Lecture Prize (2020). He is a member of Academia Europaea and the National Academy of Artificial Intelligence

A Visualized Case Tool for Temporal Representation and Reasoning

Jixin Ma¹*

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 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 Artificial Intelligence 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 Artificial Intelligence; (b) overview some fundamental issues with respects 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 Artificial Intelligence.

Biography:

Dr Jixin Ma is a Full Professor of Computer Science (Artificial Intelligence) and the Director of PhD/Postgraduate Research Programme in the School of Computing and Mathematical Sciences at University of Greenwich, U.K. He has been the Director of the Centre for Computer and Computational Science and the Lead of Artificial Intelligence Research Group. His main research areas include Artificial Intelligence, Data Science, and Information Systems, with special interests in Temporal Logic, Information Security, Machine Learning and Pattern Recognition. He has been the Editor of several international journals and conference proceedings, Conference/Program Chair, and Invited Keynote Speakers of many international conferences.

Featured Presentations

Beyond RAG: Agentic LLMs as Active Investigators, Not Passive Responders

Krupa Pratap Brahmkstri*

Visa, USA

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 AI systems, capable of reasoning, decision-making, and collaborative intelligence across domains.

Biography:

Krupa B. is a Senior Consultant-level Data Scientist with over 10 years of experience delivering end-to-end machine learning solutions across industries, with a focus on cybersecurity and AI. She specializes in applying Data Science and Generative AI to complex challenges like threat detection, anomaly modeling, and security automation. Known for her analytical rigor and creative problem-solving, Krupa approaches every challenge as an opportunity for innovation. She is passionate about using AI to make a real-world impact and enjoys collaborating to push the boundaries of what data-driven technologies can achieve in securing the digital landscape.

Delta-aware edge AI caching for hybrid inference: feature–content coherence at the extreme edge

Samaresh Kumar Singh

HP Inc., USA

Abstract:

Edge inference pipelines now mix classic predictive models with LLM/RAG components, stressing bandwidth, latency, and memory at heterogeneous CPU/GPU/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 KV-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 SRAM → device DRAM → local NVMe → 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, 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 AI tasks (vision analytics, speech, retrieval-augmented QA), along with guidance on safety (PII-aware caching) and observability (cache hit SLOs, drift alarms). Attendees will leave with concrete recipes to make hybrid AI production-ready on constrained edge fleets.

Biography:

Samaresh Kumar Singh is a Principal Engineer at HP's Technology Innovation Organization, leading architecture for edge AI/ML orchestration across heterogeneous CPU/GPU/NPU fleets. He has 21+ years building production-grade distributed systems and previously drove industrial IoT/edge platforms at Johnson Controls. His work spans hybrid AI (LLMs/RAG + predictive), low-latency microservices (C++/Go/Python), and secure cloud-edge pipelines with deep focus on observability and reliability. He currently contributes to various HP's Agentic AI based edge computation platform, and has helped deliver initiatives including HP AI Studio and HP Z Boost.

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, USA

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 (DL) 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 LLMs as semantic enhancers to address cold-start, improve diversity, and achieve the next generation of domain-aware personalization

Biography:

Vinay Venkatesh is an AI/ML Lead at Google working on Google Search in Mountain View California working on cutting edge technologies to innovate and enhance the Local Search Experience. He previously led critical projects at Google Assistance Recommendations and the Google Display Ads Recommendations.

He holds a master's degree in Electrical and Computer Engineering from Carnegie Mellon University in Pittsburgh, Pennsylvania, USA.

Bridging the AI Literacy Gap in Higher Education: The AI Literacy Competency Framework for an AI-Driven Workforce

Dr. David Frazier¹, Dr. Gwendolyn Davidson²

Hazard Community & Technical College, USA

Abstract:

Artificial intelligence is rapidly transforming 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 (ALCF), 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.

Biography:

Dr. David Frazier is an IT and higher education leader with nearly 30 years of experience as a professor, administrator, and technologist. His background spans network and database administration, web development, and programming across healthcare and education. He has also consulted on infrastructure, systems integration, and cybersecurity policy. Dr. Frazier's work centers on helping institutions use technology effectively and securely.

Dr. Gwendolyn Davidson is a higher education and technology leader with 15 years of experience in academia and a long standing career in the technology field. Her background includes computer information systems, computer security, multimedia, and educational leadership, with a focus on integrating emerging technologies into academic and workforce programs.

Transformer-based focal-pairwise attentive ranking network for logical access anti-spoofing

Sidan Yin

San Domenico School, USA

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 ASVspoof 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 (minDCF).

A Transformer-based Focal-Pairwise Attentive Ranking Network (TFPARN) 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. RawBoost waveform-level augmentation is incorporated to enhance robustness against distribution shifts and diverse spoofing manipulations. In addition, a lightweight test-time augmentation (TTA) strategy is applied during inference to improve prediction stability. Under the closed training condition, relying solely on the official ASVspoof 5 dataset without external corpora, the proposed system achieves an EER of approximately 12–13%, reaching performance comparable to approaches that utilize additional training data.

Biography:

Sidan Yin is an independent researcher specializing in computer science and computer engineering. His interests lie in speech security, deep learning, and machine learning infrastructure in Rust. His recent work focuses on Transformer-based speech anti-spoofing for AI-generated voice detection. He is also the creator of RustyML, a machine learning toolchain built in Rust.

Virtual Sessions

Keynote Presentation

How Generative AI Really Works

Neil F. Johnson^{1,2}, Dylan J. Restrepo^{1,3}, Nicholas J. Restrepo¹

1d-AI-ta Consulting, U.S.A.; 2George Washington University, U.S.A.; 3Cornell University, U.S.A.

Abstract:

Do you trust AI? The rapid deployment of generative AI systems (e.g. LLMs such as ChatGPT, Claude, 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.

Biography:

Neil Johnson (speaker) is Head of the Dynamic Online Networks Laboratory and Professor at George Washington University. He was previously Professor of Physics at Oxford University and holds a PhD from Harvard University and an MA/BA from Cambridge University. Dylan J. Restrepo is an MSc student at Cornell Tech, Cornell University. Nicholas J. Restrepo holds an MBA from HEC Paris. All 3 are members of d-AI-ta Consulting LLC.

Featured Presentations:

AI-Driven Advanced Sensing: Enhancing Reliability and Predictive Diagnostics in Complex Engineered Systems

Kunal Goyal¹

1Georgia Institute of Technology, USA

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 its crucial role in making technological systems more robust and secure.

Biography:

I am an Electrical Engineer specializing in advanced sensor technologies and AI/ML integration. I hold master's degrees in electrical engineering from Michigan Technological University and Computer Science from Georgia Institute of Technology, providing a unique interdisciplinary perspective. My work focuses on enhancing system reliability, efficiency, and safety through innovative sensing solutions and data-driven insights. As an active member of the global engineering community, I am a PC Member for ACM SIGCSE TS 2026 and regularly contribute as a peer reviewer for reputable journals. My

research interests include predictive analytics, embedded systems, and the application of AI for real-world problem-solving.

Multi-stage biomedical feature selection extraction algorithm for cancer detection

Dr Mohitkumar Pravinchandra Bhadla,

Gandhinagar University, Gandhinagar 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

Biography:

Dr. Mohitkumar Bhadla is a Principal of Gandhinagar Institute of Technology & Professor in Department of Computer Science & Engineering at Gandhinagar University, Gujarat India. He has more than 16 Years of Experience in IT Field as well as an academic Field. With the active academician, he is an active researcher also. He has published more than 21 papers in well-known journals. Moreover 3 Different Patents filled among them one patent Filled for Government of Australia for An IoT Based Sensor Network for Smart city Implementations and 2 for Government of India in Wireless Sensor Networks.

Pricing intelligence: What frontier AI economics mean for builders and buyers

Charu Maheshwari

Anthropic, USA

Abstract:

Pricing is where strategy meets science. For businesses building with or on top of 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 levers 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.

Biography:

Charu Maheshwari is on the Product Finance and Strategy team at Anthropic, where she works on financial modeling for model launches, pricing strategy for AI capabilities, and post-launch analytics. A CFA Charterholder and Wharton MBA, she previously spent several years at McKinsey and Company advising C-suite leaders across industries. She brings a rare combination of rigorous financial training and direct operational experience inside a frontier AI lab.

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

Presenter*: Ravi Mandliya¹*

¹*Discord Inc., USA*

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 ML 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.

Biography:

Ravi Mandliya is a Machine Learning Engineer at Discord Inc. and an IEEE member specializing in large-scale recommendation and ranking systems. His work focuses on fairness, transparency, and trust in applied ML, bridging research and production deployment. Ravi has designed scalable ML pipelines serving millions of users. He is passionate about operationalizing ethical AI principles within real-world systems and mentoring engineers to build transparent, accountable models that align with human values.

Integrating Transcendental Elements into AI-Powered Serious Games for a Holistic Enhancement of Geriatric Health and Well-Being

Hari C. Gangadharan

Academic Programs, Ocean Grove Charter School, Campbell, USA

Abstract:

The recent advancements and extended usage of AI-powered serious games in geriatric healthcare has 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 the 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.

Biography:

Hari C. Gangadharan is an interdisciplinary researcher and game developer specializing in the ethics of AI, intelligent reality, and gaming. A dual-enrollment student at Ocean Grove Charter School and Mission College, his research has

been published in IEEE Xplore and presented at prestigious conferences, including the IEEE Gaming, Entertainment, and Media Conference 2024, the IEEE 3rd International Conference on Intelligent Reality 2024, and "The Art of Storytelling: Archetypes in Focus 2025" at Oxford University. He is dedicated to implementing his findings into practical game development.

Modernizing Legacy Enterprise Data Platforms: From Traditional ETL to Cloud Lakehouse Architectures

Presenter* and Co-author names: Avik Datta1*

Affiliation (Institution name, Country):

1 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, compliant, 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 HIPAA 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 Generative AI responsibly within regulated healthcare and enterprise platforms.

Biography:

Avik Datta is a Lead Data Engineer in Sun Life Financial with over 13 years of experience solving complex data challenges in the financial & insurance sector, including P&C and healthcare. He specializes in building scalable, dependable data systems that support policy administration, claims processing, billing and member management

Avik brings hands-on expertise across legacy ETL technologies and modern cloud platforms such as Informatica Cloud, Snowflake and AWS. His work centers on practical enterprise data engineering: modernizing legacy warehouses without disruption, designing proactive data quality frameworks, and implementing governance models that meet regulatory expectations while supporting innovation. Avik is currently pursuing a part time MS in Data Science at Northwestern University, where he focuses on integrating AI and machine learning into traditional data pipelines to create measurable business value.

Shaping Scalable Data Platforms with a Human Touch

Sivadeep Katangoori

Insight Global

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 AI enablement on platforms like Google Cloud and AWS. I've led end-to-end technical transformations, from architecting secure data lakes to modernizing ML 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.

Biography:

I'm a seasoned Data Lake and Cloud Platform Specialist with deep expertise in AI, security analytics, and large-scale data systems. With a master's degree in software engineering and certifications in Google Cloud and project management, I've had the opportunity to lead strategic data initiatives at institutions like Wells Fargo and Bank of America. I thrive at the intersection of hands-on technical work and leadership translating complex challenges into efficient, scalable solutions. I'm passionate about lifelong learning, and I enjoy mentoring others and driving enterprise transformation efforts as both a technical lead and project manager.

Introduction and Evaluation of AI powered Coding Tools for Teaching Computer Programming to Software Engineering Students

Sarita Singh

Northeastern University, USA

Abstract:

Leading software companies are increasingly expecting computer science graduates to be proficient in using various AI powered coding tools to enhance productivity, accelerate debugging and improve code quality. Educators across the globe are faced with the challenges of how to teach 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.

Biography:

Sarita Singh is an Associate Teaching Professor with Khoury College of Computer Science at Northeastern University, USA. She holds a PhD in Computer Science as well as a Doctor of Education (EdD) degree. She earned her PhD in Cybersecurity as an Infosys Research fellow where her dissertation focused on designing and implementing a public key cryptosystem based on the Boolean Satisfiability Problem. Her EdD dissertation focused on use of innovative education technology to support students' learning in a computer programming class. Her area of research interests includes Internet of Things, Cybersecurity and Computer Science Education.

Simplicial nonnegative matrix factorization with applications to high-dimensional and large-scale datasets

Michail Tsagris^{1*} and Abdulaziz Alenazi²

1 Department of Economics, University of Crete, Gallos Campus, Rethimno, 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.

Biography:

Michail Tsagris is an Associate Professor at the Department of Economics of the University of Crete (UoC) and an Adjunct Professor in the Department of Mathematics and Statistics at the University of New Brunswick Saint John. Michail has published more than 70 papers in journals, conference proceedings and book chapters and has (co)developed 52 R packages. His current research interests include computational statistics, compositional and directional data analysis, applied econometrics and machine learning.

A comprehensive survey of Mixture of Experts (MoE) architectures

Dev Vyas¹ and Dr. Mohit Bhadla²

¹Georgia State University, USA; ²Gandhinagar University, India

Presenter Contact: devyas17272@gmail.com, +1 470-265-1767

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 (LLMs) 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 (MMOE). 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.

Biography:

I am an Honors Program Computer Science undergraduate at Georgia State University, expecting graduation in May 2026. My research focuses on novel AI architectures, particularly the development of Recurrent Transformer models to enhance iterative reasoning in large-scale systems. With a strong foundation in Python, PyTorch, and TensorFlow, I have experience in continual learning, reinforcement learning, and Retrieval-Augmented Generation. I hold a Machine Learning Specialization from Stanford University and DeepLearning.AI and am passionate about pushing the boundaries of efficient and scalable AI.

End of Day-4

We wish to see you at

AIM-2027

April 26-28, 2027 | Orlando, FL

Looking for additional information?

Visit us online at <https://artificialintelligence.unitedscientificgroup.org/>
for complete details on the technical program, workshop registration and hotel/ travel arrangements.

Organized by



UNITED | Scientific
Group
A non-profit organization

USG United Scientific Group

(A non-profit organization)

8105, Rasor Blvd - Suite #112, PLANO, TX 75024

Tel: +1-469-854-2280/81; Fax: +1-469-854-2278;

Email: aim@uniscigroup.org; Web: <https://artificialintelligence.unitedscientificgroup.org/>