




IEEE 2026 CIC/COGMI/TPS/RISC JOINT CONFERENCES

CONFERENCE PROGRAM*

**San Jose, CA, USA
Nov. 4-6, 2026**

**VERSION
October 6**

Useful Resources

Conference Websites

- IEEE CIC 2026: <https://cic.ieee-cs.org>
- IEEE CogMI 2026: <https://cogmi.ieee-cs.org>
- IEEE RISC 2026: <https://risc.ieee-cs.org>
- IEEE TPS 2026: <https://tps.ieee-cs.org/>

Overview Conference Day 1: November 4, 2026

7:15 AM - 8:30 AM	Registration (Room: Foyer) and Continental Breakfast (Atrium – hotel will have private event signs)				
8:30 AM - 8:45 AM	Welcome and Opening Remarks (Steering Committee Chair and Organizing Committee Chairs) (Room: Willow/Oak/Cypress)				
08:45 AM – 9:45 AM	Keynote 1: Edward Y. Chang , Founder and CEO, QuadriumAI, USA (Room: Willow/Oak/Cypress) Title: Beyond Intelligence: Engineering AGI That Can Reason, Recover, and Reconsider (Chair: Huan Liu, Arizona State University, USA)				
9:45 AM – 10:00 AM	Break Room: Atrium				
10:00 AM – 11:30 AM	TPS / CIC Panel 1 (Room: Willow/Oak/Cypress) Title: Toward Safe and Trustworthy Digital Future: Challenges and Road Ahead Panelists: Daniela Oliveira (NSF), Heiko Ludwig (IBM), Indrajit Ray (Colorado State University), Jaideep Vaidya (Rutgers University) (Moderator: James Joshi, University of Pittsburgh, USA)				
11:30 AM – 12:00 PM	Fireside Chat with IBM Executive Suja Viswesnan , Vice President, Security Products, IBM, USA (Room: Willow/Oak/Cypress)				
12:00 PM – 01:00 PM	Lunch Break (provided by conference) Room: Atrium				
01:00 PM – 3:00 PM	TPS Research Session 1: Security and Privacy in Smart IoT and CPS (Room: Cypress) Session Chair: Pushpjeet Shrivastava (Bank of America)	CogMI Research Session 1: Cognitive Benchmarking and Evaluation of Foundation Models (Room: Oak) Session Chair: Reza Zafarani (Syracuse University)	Invited Session 1: Physical AI, Cyber-Physical Intelligence and Smart Systems (Room: Willow) Session Chair: Hongxin Hu (University at Buffalo, SUNY)	RISC Session 1: Post-Quantum Cryptography & Secure Space Communications (Room: Cedar/Pine) Session Chair: Dr. Efrén López-Morales (New Mexico State University)	CIC Research Session 1: AI for Cybersecurity (Room: Boardroom) Session Chair: Wenqi Wei (Fordham University)
3:00 PM – 3:15 PM	Break				
03:15 PM – 5:15 PM	TPS Research Session 2: Secure and Trustworthy Agentic AI (Room: Cypress) Session Chair: Badhan Chandra Das (Florida International University)	CogMI Research Session 2: Explainability, Safety, and Security (Room: Oak) Session Chair: Indrajit Ray (Colorado State University)	Invited Session 2: Agentic AI, Autonomous Web Agents and Distributed Intelligence (Room: Willow) Session Chair: Wenqi Wei (Fordham University)	RISC Session 2: Spacecraft & Cyber-Physical Systems Security (Room: Cedar/Pine) Session Chair: Mai Abdelhakim (University of Pittsburgh)	CIC Research Session 2: Securing and Advancing AI for Data Systems (Room: Boardroom) Session Chair: Yanzhao Wu (Florida International University)
06:00 PM – 08/09:00 PM	Networking/Reception (provided by conference): Room: Atrium				

Overview Conference Day 3: November 6, 2026

7:15 AM - 8:15 AM	Registration (Room: Foyer) & Continental Breakfast (Atrium – hotel will have private signs)						EEE CIC NextGen Program: Juniors at Conference. website: https://cic.ieee-cs.org/2026/next-gen-cs-juniors/
8:15 AM - 8:20 AM	Welcome and Opening Remarks (Steering Committee Chair and Organizing Committee Chairs) (Room: Willow/Oak/Cypress)						
08:20 AM – 9:20 AM	Keynote 4: Ramesh Raskar, MIT Media Lab, USA (Room: Willow/Oak/Cypress) Title: Billions of Personal Agents vs. AGI: Which Future Wins? (Chair: Danda B. Rawat, Howard University, USA)						
9:20 AM – 10:20 AM	Keynote 5: Craig Knoblock, University of Southern California, USA (Room: Willow/Oak/Cypress) Title: From Papers to Knowledge: Building Knowledge Graphs from Technical Literature (Chair: Reza Zafarani, Syracuse University, USA)						
10:20 AM – 10:30 AM	Break Room: Atrium						
10:30 AM – 12:30 PM	TPS Application Session 1: Secure Agentic AI and AI-Assisted Vulnerability Analysis and Forensics (Room: Cypress) Session Chair: Shaswata Mitra (University of Alabama)	CogMI Application Session 3: Efficient, Resilient, and Secure AI Systems (Room: Oak) Session Chair: Hoda Mehrpouyan (Boise State University)	TPS Application Session 2: AI Assurance, Auditing, Red-Teaming, and Supply-Chain Security (Room: Willow) Session Chair: Rishabh Das (Ohio University)	TPS Application Session 3: Privacy-Preserving Data Usage, Synthetic Data, and Policy Governance (Room: Juniper) Session Chair: Ram Dantu (University of North Texas)	Gov/Industry Track ion Session 1: (Room: Cedar/Pine) Session Chair: Zeba Wahab (San Jose State U)	NAND Workshop Session (Room: Maple)	
12:30 PM – 01:30 PM	Lunch Break (provided by conference) Room: Atrium						Lunch
01:30 PM – 03:30 PM	TPS Application Session 4: Trusted and Secure Systems, Networks, and Critical Infrastructures (Room: Juniper) Session Chair: Sayantan Ghosh (LinkedIn)	CogMI Application Session 4: Visual Perception and Vision-Language Models (Room: Willow/Oak/Cypress) Session Chair: KC Santosh (Oregon Institute of Technology)	Invited Session 5: AI for Societal Impact, Knowledge Integration and Information Integrity (Room: Willow) Session Chair: Danda B Rawat (Howard University)		Gov/Industry Track ion Session 2: (Room: Cedar/Pine) Session Chair: Julian Jarrett (Lutron Electronics)	NAND Workshop Session (Room: Maple)	
03:30 PM –4:00 PM	Closing remarks (Room: Willow/Oak/Cypress)						

Conference Day 1: November 4, 2026

Registration & Continental Breakfast (provided by conference)

7:15 AM - 8:30 AM

Welcome and Opening Remarks

8:30 am – 8:45 am

All Participants and Chairs

(Steering Committee Chair and Organizing Committee Chairs)

(Room: Willow/Oak/Cypress)

Keynote 1 (Room: Willow/Oak/Cypress)

08:45 AM – 9:45 AM

Keynote 1: **Edward Y. Chang**, Founder and CEO, QuadriumAI, USA

Title: Beyond Intelligence: Engineering AGI That Can Reason, Recover, and Reconsider

(Chair: Huan Liu, Arizona State University, USA)

Coffee Break (15 min)

Panel 1 (Room: Willow/Oak/Cypress)

10:00 AM – 11:30 AM

TPS / CIC Panel 1

Title: Toward Safe and Trustworthy Digital Future: Challenges and Road Ahead

Panelists: Daniela Oliveira (NSF), Heiko Ludwig (IBM), Indrajit Ray (Colorado State University), Jaideep Vaidya (Rutgers University)

(Moderator: James Joshi, University of Pittsburgh, USA)

Fireside Chat (Room: Room: Willow/Oak/Cypress)

11:30 AM – 12:00 PM

Fireside Chat with IBM Executive Suja Viswesnan, Vice President, Security Products, IBM, USA

Lunch Break (provided by conference) Room: Atrium

12:00 PM – 1:00 PM

TPS Research Session 1: Security and Privacy in Smart IoT and CPS

Time: 1:00pm – 3:00pm

Room: Cypress

Session Chair: Pushpjeet Shrivastava (Bank of America)

Web-Based Semantic Platform for IoT Privacy Policy Regulatory Coverage with Ad-Hoc Chatbot

Heinrick Janiola, Lavanya Elluri, Dejarvis Oliver, and Srujan Kotikela

Resource-Efficient Feature Selection for IoT Botnet Detection and Classification

Md. Abid Jahan, Ashikur Rahman, Abderrahmane Leshob, Rakin Rahman, and Raqeebir Rab

Gain-Dependent Entropy Characterization of a Zener Avalanche Noise TRNG via Dual NIST Standards: Toward Low-Cost Hardware Entropy for ICS Deployments

Rishabh Das

Safety Graph: When Safety Responses Become Attack Surfaces in Cyber-Physical Systems

Rakesh Podder, Philip D. Hopkins, Sarath Sreedharan, Indrakshi Ray, and Indrajit Ray

REVIVE: A Multi-Modal Framework for Vandalism Detection and Recovery in Autonomous Vehicles

Abdullah Tariq Choudhry and Tapadhir Das

VISTA: Visual Interpretation of Security and Trust Assessment for Satellite Constellations

Quincy Bayer, Karthik Raja, Samuel Mergendahl, Mai Abdelhakim, and Robert Cunningham

CogMI Research Session 1: Cognitive Benchmarking and Evaluation of Foundation Models

Time: 1:00pm – 3:00pm

Room: Oak

Session Chair: Reza Zafarani (Syracuse University)

Measurement Invariance Between Human and Machine Respondents on Cognitive Assessments

Mohammad Hamza Siddiqui

Multi-Judge Consensus Scoring for IRT-Based Adaptive LLM Evaluation

Omkar Podey, Philip Heller, Wendy Lee, and Amith Kamath Belman

Agreeing on the Answer but Not the Reason: Measurement Limits of LLM Reasoning Failure**Taxonomies**

Madhu Sudhan Rao Beesetty

The Silicon Meridian: If LLMs Can Speak Like Us, Can They Decide Like Us?

Vibha Santhanakrishnan

Interventional Semantic Interface Gradients: Evaluating Interface-Conditioned Robustness in Language Models

Dominic Okonkwo, Adetayo Okunoye, and Ismailcem Arpinar

Diagnosing the Synthetic-to-Real Gap in Small VLMs for Deployable Flood Monitoring

Kunal Chevuri, Angel Peredo, and Sergei Chuprov

Invited Session 1: Physical AI, Cyber-Physical Intelligence and Smart Systems

Time: 1:00 PM – 3:00 PM

Room:

Session Chair: Hongxin Hu (University at Buffalo, SUNY)

Physical AI-Based Human Performance Assessment Providing Explainable Personalized Feedback

Dimitrios Georgakopoulos

Vision: Physical Cognition as a Continuous Belief for Cyber-Physical Systems and Internet of Things

Tarek Abdelzaher

Cities That Sense, Reason, Plan and Act: Trustworthy Hybrid AI for Human-Centred Mixed-Autonomy Urban Mobility

Jerry Jialie Shen

Offload Shaping is Split Computing Done Right

Mahadev Satyanarayanan

Uncertainty-Aware Causal Discovery from Sparse Multi-Modal Image Datasets: When Does Reweighting Help?

Adrienne Raglin

RISC Session 1: Post-Quantum Cryptography & Secure Space Communications

Time: 1:00pm – 3:00pm

Room: Cedar/Pine

Session Chair: Dr. Efrén López-Morales (New Mexico State University)

RAPTOR: Risk-Aware Post-Quantum Trust Orchestration for Resource-Constrained Space and Critical Networks *(application)*

Sylvia Llosa (University of Colorado Boulder), Yueqi Chen (University of Colorado Boulder), Sangtae Ha (University of Colorado Boulder), and Elliot Eichen (University of Colorado Boulder)

Double Ratchet Key Evolution for CCSDS Telecommand and Telemetry *(research)*

Erik Wadell Ledin (RISE Research Institutes of Sweden), Alfonso Iacovazzi (RISE Research Institutes of Sweden), and Shahid Raza (RISE Research Institutes of Sweden)

ZT-FedLEO: Zero-Trust Federated Anomaly Detection for High-Risk Path Isolation in LEO Satellite Networks *(research)*

Abdiasis Ahmed Mohamed (National Taiwan University of Technology) and Wei Chung Teng (National Taiwan University of Technology)

SpaceNASim: Contact-Aware Penetration-Testing Simulation for Satellite-Ground Systems *(research)*

Tyler Cody (University of Virginia) and Peter A. Beling (University of Virginia)

A Cryptographic Agility Framework for Satellite Constellations Using Post-Quantum Cryptography *(research)*

Shilpi Mittal (University of North Texas), Ankit Gupta (University of North Texas), and Gahangir Hossain (University of North Texas)

Side Channel Resilience of Post Quantum Cryptography Algorithms in Radiation Prone Space Environments *(research)*

Ankit Gupta (University of North Texas)

CIC Research Session 1: AI for Cybersecurity

Time: 1:00pm – 3:00pm

Room: Boardroom)

Session Chair:Wenqi Wei (Fordham University)

FALCON: Transforming Cyber Threat Intelligence into Deployable IDS Rules with Self-Reflection

Shaswata Mitra (The University of Alabama), Subash Neupane (Meharry Medical College), Martin Duclos (The University of Alabama), Sudip Mittal (The University of Alabama), Aritran Piplai (The University of Texas at El Paso), Md Rayhanur Rahman (The University of Alabama), Edward Ziegler (National Security Agency) and Shahram Rahimi (The University of Alabama)

Performance Interference Detection for Cloud-Native Applications Using Unsupervised Machine Learning Models

Eli Bakshi (California Polytechnic State University), Sumona Mukhopadhyay (California Polytechnic State University), Diwakar Krishnamurthy (University of Calgary) and Joydeep Mukherjee (California Polytechnic State University)

Graph-Aware Reinforcement Learning for Network Resilience: A Comparative Study of Reward Design

Shivam Verma (California State University Northridge), Alex Modarresi (California State University Northridge) and John Symons (University of Kansas)

An Empirical Evaluation of Large Language Models to Generate Unit Test Code for Industrial Automation Controllers

Md Humaun Kabir (Lamar University), Md Rakibul Islam (Lamar University) and Helen Lou (Lamar University)

Confidence-Gated AIOps: A Reversibility-Aware Reference Architecture for Predictive Autoscaling in Cloud-Native Systems

Jyostna Seelam (Independent Researcher), Abdul Rasheed Shaik (Independent Researcher) and Khadarvali Shaik (Independent Researcher)

Coffee Break (15 min)

TPS Research Session 2: Secure and Trustworthy Agentic AI

Time: 03:15 PM – 5:15 PM

Room: Cypress

Session Chair: Badhan Chandra Das (Florida International University)

Local Allow, Global Deny: Why Per-Action Authorization Does Not Compose in LLM Agents

Anmol Pandey, Aditya Jain, Liang Chen, and Carsten Maple

Compiling Semantic Guardrails for Autonomous Payment Agents: Deterministic, Auditable Prompt-Injection Defense at LLM-Level Detection

Yash Rajeshbhai Parikh, Janaki Ram Goteti, and Ashok Singamaneni

TrustShiftProbe: Characterizing, Benchmarking, and Defending Staged Trust Attacks on MCP Servers

Mehrdad Rostamzadeh, Sidhant Narula, Mohammad Ghasemigol, and Daniel Takabi

MPIDetector: Detecting Multi-modal Prompt Injection via LLM-based Reasoning

Yiheng Li and Jingyue Li

Privacy Policy Enforcement Guardrails for Data-Sensitive Retrieval-Augmented Generation

Osama Zafar, Alexander Nemecek, Yiqian Zhang, Wenbiao Li, Debargha Ganguly, Vikash Singh, Vipin Chaudhary, and Erman Ayday

CogMI Research Session 2: Explainability, Safety, and Security

Time: 3:15pm – 5:25pm

Room: [Oak](#)

Session Chair: Indrajit Ray (Colorado State University)

When Explanations Explain Themselves: Recursive Behavioral Drift and Behavioral Fixed-Point Training for Sparse Autoencoders

Sajib Hossain

Explainable Novel Category Discovery in Semantic Concept Space

Ifrat Ikhtear Uddin, Yang Zhou, Chaowei Zhang, and Longwei Wang

AgentGuard: Calibrated Failure Detection and Risk-Aware Routing for Multi-Agent LLM Workflows

Xiaoliang Zhang

Governance-Aware Cognitive Orchestration for Human-Supervised Cybersecurity AI Systems

Andrei Kojukhov

Trust Under Attack: Adversarial Security Analysis of GDPR-Oriented RAG Assistants

Manjil Ghimire, Lopamudra Praharaj, Md Sadun Haq, and Maanak Gupta

AXIOS: Preserving Vertical Federated Learning Model Integrity via Decentralized Training-Time Anomaly Detection

Ella Roundy, Timothy Gurevich, Dylan Gresham, and Gaby Dagher

Invited Session 2: Agentic AI, Autonomous Web Agents and Distributed Intelligence

Time: 3:15pm – 5:15pm

Room:

Session Chair: Wenqi Wei (Fordham University)

WebAgents: Unbundling Agent Creation, Hosting, and Execution for an Agent-Native Web

Ramesh Raskar

What Social Media Teaches Us About Accumulative AI Risks

Dawei Li, Bohan Jiang, Huan Liu

Psychology of AI Agents

Ashish Kundu

Integrating Heterogeneous Reliable Knowledge Sources to Support Brazilian Judicialization of Healthcare

Calton Pu

Staking Claims of Generalization: A Causal View

Saketh Vishnubhatla, Adrienne Raglin, Huan Liu

RISC Session 2: Spacecraft & Cyber-Physical Systems Security

Time: 3:15pm – 5:15pm

Room: Cedar/Pine

Session Chair: Mai Abdelhakim (University of Pittsburgh)

Trust Without Boundaries: An Architectural Analysis of Satellite Flight Software *(research)*

Jack Vanlyssel (University of New Mexico), Gruia-Catalin Roman (University of New Mexico), Kendra Cook (SAIC), Sazzadur Rahaman (University of Arizona), and Afsah Anwar (University of New Mexico)

Stargate: Security Analysis and Mission-Aware Firewall for the CubeSat Space Protocol *(research)*

Owen Ramsey (New Mexico State University), Carlos Gonzalez-Cortes (Universidad de Santiago de Chile), and Efrén López-Morales (New Mexico State University)

The Systems Security Assurance Discipline in the Space Domain *(research)*

James Mertz (Jet Propulsion Laboratory), Cris Hernandez (Jet Propulsion Laboratory), and Pratik Satam (University of Arizona)

Temporal Enforcement for Adaptive Cyber-Physical Containment: Critical-Infrastructure and Space-Ground Case Studies *(application)*

Hoda Mehrpouyan (Boise State University)

Orbital Grounding: A Capability Model and Resilience Evaluation for Space-Based Evidence *(application)*

Tyler Cody (University of Virginia) and Peter A. Beling (University of Virginia)

CIC Research Session 2: Securing and Advancing AI for Data Systems

Time: 3:15pm – 5:15pm

Room: Boardroom

Session Chair: Yanzhao Wu (Florida International University)

Neuro-Symbolic AI for Trustworthy Clinical Reasoning: A Verification-Stage Taxonomy and Evidence Survey

Firoj Husen Shaikh (Meharry Medical College), Subash Neupane (University of Alabama) and Himanshu Tripathi (University of Alabama)

Defending Retrieval-Augmented Intrusion Detection Against Knowledge Poisoning and Prompt Injection

Kaysarul Anas Apurba (Laurentian University), Md. Hasibul Hasan (Laurentian University), Mahedee Zaman Moon (Laurentian University), Sk. Md. Mizanur Rahman (Centennial College) and Atsuo Inomata (The University of Osaka)

TRACE: An Exposure-Aware Diagnostic Metric Suite for Corpus Poisoning in Arabic RAG Pipelines

Mitha Alshammary (The University of Texas at Dallas), Md Nahiyen Uddin (The University of Texas at Dallas) and Latifur Khan (The University of Texas at Dallas)

STRAP-FL: Adaptive Federated Learning under Concept Drift and Data Streams

Divya Bhatt (University of Oklahoma), Le Gruenwald (University of Oklahoma) and Egawati Panjei (University of Oklahoma)

Sentence Specificity Scores for Collaborative Technical Documentation

Rocker D'Antonio (Mississippi State University), Thomas Benton Townsend (Mississippi State University) and Dimitrios Michael Manias (Mississippi State University)

Goal-Adaptive Joint Action–Deferral Optimization for Capacity-Constrained Human–AI Quality Control in Automotive Remanufacturing

Jing Chen (Independent Researcher) and Liang Chen (Independent Researcher)

Networking/Reception (provided by the conference)

Room: King's Garden Foyer

06:00 PM – 09:00 PM

Conference Day 2: November 5, 2026

Registration & Continental Breakfast (provided by conference)

Registration (Room: [Foyer](#)) and Continental Breakfast ([Atrium](#) – **hotel will have private event signs**)
7:15 AM - 8:30 AM

Welcome and Opening Remarks

8:30 am – 8:45 am

All Participants and Chairs

Room - [Willow/Oak/Cypress](#)

Keynote 2 (Room: Willow/Oak/Cypress)

08:45 AM – 9:45 AM

Keynote 2: **Kevin Fu**, Professor, Northeastern University, USA

Title: Medical Device Cybersecurity: The First 178 Years

(Chair: Rob Cunningham and Mai Abdelhakim (University of Pittsburgh, USA))

Coffee Break (15 min)

TPS Research Session 3: LLM Assurance, Privacy, and Robustness

Time: 10:00AM – 12:00PM

Room: [Cypress](#)

Session Chair: Yash Rajeshbhai Parikh (University of Washington)

Auditing the Judge: Internal Signals for LLM Reasoning Assurance

Mayank Gulati and Gerhard Wunder

Confidence-Calibrated Human Deferral for Trustworthy LLM Data Governance

Muzakkiruddin Ahmed Mohammed, Shames Al Mandalawi, Mert Can Cakmak, and John Talburt

Is Weight Tying Still Beneficial for Decoder-Only LLMs in Private Settings Under DP-SGD?

Razan El Mais, Ali Chehab, Ibrahim Issa, and Razane Tajeddine

On the Robustness of Small Language Models for Smishing Detection Against Adversarial Text Perturbations

Seyed Mohammad Sanjari Pirmahalleh, Mir Mehedi Ahsan Pritom, and Maanak Gupta

PaTAS: A Framework for Trust Propagation in Neural Networks Using Subjective Logic

Koffi Ismael Ouattara, Ioannis Krontiris, Theo Dimitrakos, Dennis Eisermann, Houda Labiod, and Frank Kargl

CogMI Research Session 3: Memory, Retrieval, and Knowledge-Guided Reasoning

Time: 10:00AM – 12:00PM

Room: [Oak](#)

A Two-Layer Event Memory with Reconstruction-Aware Forgetting

Artur Kalimullin

Selective Reconsolidation in Persistent Agent Memory

Pranava Kumar, Sushaan Kandukoori, and Manolis Kellis

Platoon-RAG: Consensus-Driven Multi-Agent Retrieval-Augmented Generation via Semantic Headway Optimization

Naveen Kumar Rajarajan

EpiWarrant: Defeasible Learner-Knowledge Attribution

Hasan Jamil

Cerno: A Declarative Language for Knowledge-Guided Image Analysis with Guaranteed Dispatch or Refusal

Hasan Jamil

Invited Session 3: Cybersecurity, Privacy Preservation and Secure AI Systems

Time: 10:00AM – 12:00PM

Room: [Willow](#)

Session Chair: Danda B Rawat (Howard University)

Evaluation, Certification, and Accreditation of CyberAI Systems

Bhavani Thuraisingham

Aegis: Towards Governance, Integrity, and Security of AI Voice Agents

Wenqi Wei

Secure Timed Data Release Using Blockchains

Balaji Palanisamy

Integration of AI, Software Engineering, and Security

Bhavani Thuraisingham

Information-Theoretic Bounds on Privacy Leakage in Federated Learning at the Tactical Edge: Analysis With and Without Packed Homomorphic Encryption

Danda Rawat

RISC Session 3: AI, Anomaly Detection & Resilient Autonomy for Space Systems

Time: 10:00AM – 12:00PM

Room: [Juniper](#)

Session Chair: Hoda Mehrpouyan (Boise State University)

A Categorical Framework for Modeling Agentic AI in Cyber-Physical Systems: Semantics, Attack Modeling, and Fault Detection (*research*)

Mohamed Gebril (George Mason University)

Autonomous Edge Inference in Extreme Environments: A 4H-SiC Spiking Neural Network Coprocessor
(research)

Thomas Coverstone (Johns Hopkins University Applied Physics Laboratory)

Securing the Intelligent Edge: LLM Attack Taxonomy and AI Governance Gaps in Space and Critical Systems
(research)

Zeba Wahab (Independent Researcher)

The Deception Arms Race: A Survey of Honeypots and Anti-Honeypot Techniques for Resilient Networks and Critical Systems (application)

Teresita Noelia Nunez Migliorisi (University of Delaware), Michael De Lucia (University of Delaware), and Andrew Novocin (University of Delaware)

Telemetry Integrity Under Test: A Reproducible Stress Test for Spacecraft Anomaly Detectors (application)

Swastik Sengupta (Las Positas Community College)

CIC Application Session 1: Intelligent Data Systems

Time: 10:00AM – 12:00PM

Room: Cedar/Pine

Session Chair: Firoj H. Shaikh (Meharry Medical College)

Auditing Search Before Content Verification: A Manifest-Gated Study of YouTube API Availability and Stability

Liang Hu (Columbia University)

Unified Cross-Layer Log Analytics for Anomaly Detection in Enterprise IT Systems

Haitham Khalifa (D3m Tech Company), Ahmed Kamal (Tawuniya Insurance Company) and Yasser Youssef (California State University, Bakersfield)

Ask, Answer, or Abstain: Auditing Ambiguity in Chinese Automotive-Parts Queries

Liuye Hua (Xinjiang University), Leixin Zhang (Xinjiang University), Jiong Zheng (Xinjiang University) and Haoze Zheng (Xinjiang University)

Veracity-Aware Machine Learning Data Exchange: Grounding JSON-LD in Subjective Logic

Muntaser Syed (Florida Institute of Technology), Sheikh Abujar (The University of Alabama at Birmingham) and Marius Silaghi (Florida Institute of Technology)

From Ambiguous Requirements to Hallucinated Tests: A Multi-Domain Study of the Mirror Effect in LLM-Assisted Software Testing

Kiran Paul Kanikaram (Independent Researcher)

An Intelligent Agentic AI Framework for Autonomous Network Performance Diagnosis and Self-Healing Using AIOps and Machine Learning

Mohit Bajpai (Independent Researcher)

CogMI Research Session 4: NeuroAI, Human Cognition, and Social Robotics

Time: 10:00 AM – 12:00 PM

Room: Redwood

Session Chair: **Reza Zafarani (Syracuse University)**

Composability and Profitability Boundaries of Temporal–Distributed Parallelism for Scalable Brain-Inspired Spiking Neural Networks

Yunfei Chen, Hong Shi, and Zhenyu Zhang

The Value of Spike Timing: A Leakage-Resistant Benchmark of SNN Design Choices for Network Intrusion Detection

Raj Patel, Shaswata Mitra, David Amebley, Taye Akinrele, Sayanton Dibbo, and Shahram Rahimi

Self-Supervised Temporal Convolutional Network Modeling of Eye Gaze Reveals Implementation Strategies in Proactive Cognitive Control

Jason Reimer, Dylan Diaz, Vishwa Bhatt, Jonathan Martinez, Jennifer Jin and Kevin Rosales.

Multimodal Prediction of Student Belonging Using EEG Responses to Mindset-Based Stimuli in Computing Education

Andrew Martinez, Meora Giusiano, Joel Saju Puthankalam, Pallavi Das, Zoe Wood, Julie Bettergarcia, and Sumona Mukhopadhyay

Evaluating Design Guidelines for Social Robots: Connecting People and Building Relationships

Daniel Saba, Kathryn Degnan, Michael Baggot, and Jeffrey W. Herrmann

TPHAC Session 1

Time: 10:00 AM – 12:00 PM

Room: **Maple**

Session info: <https://sites.google.com/pitt.edu/tphac/home?authuser=0>

Lunch Break (provided by conference)

12:00 PM – 1:00 PM (Room: **Atrium**)

Keynote 3 (Room:)

01:00 PM – 2:00 PM

Keynote 3: Irina Dolinskaya, Directorate Head, NSF Directorate for Computer and Information Science and Engineering (CISE)
(Room: **Willow/Oak/Cypress**)

Title: Advancing the frontiers of computing: CISE research opportunities and emerging efforts.

(Chair: James Joshi, (University of Pittsburgh, USA)

Panel 2 (Room: Willow/Oak/Cypress)

2:00 PM – 3:30 PM

IEEE CogMI Panel: Pacing the Frontier: Prudence, Paralysis, or Pipe Dream?

Panelists: Elisa Bertino (Purdue University), Padmanabhan (Babu) Pillai (Carnegie Mellon University), Calton Pu (Georgia Institute of Technology), KC (Casey) Santosh (Oregon Institute of Technology)

Moderator: *Paolo Boldi, Università degli Studi di Milano, Italy*

Coffee Break (15 min) (Room: Atrium)

TPS Research Session 4: AI-Assisted Cyber Defense, Malware Deception, and Reverse Engineering

Time: 03:45 PM – 05:45 PM

Room: Cypress

Session Chair: Seyed Mohammad Sanjari (Tennessee Technological University)

Fine Tuning LLM-Based Decompiler Backend for Output Code Quality and Trustworthiness

Muqi Zou, Hongyu Cai, Hongwei Wu, Zion Leonahenahe Basque, Arslan Khan, Z. Berkay Celik, Dave Tian, Antonio Bianchi, Ruoyu Wang, and Dongyan Xu

File Localization in Firmware Image for Agentic Reverse Engineering

Nathan Harrison and Samuel Mulder

ORCAGen: Orchestrating Context-Aware Malware Deception with RAG-Guided Generative AI

Shihab Ahmed, Md Sajidul Islam Sajid, Teryl Taylor, Frederico Araujo, and Tariqul Islam

IRENA: Intrusion Response by Expected-Utility Negotiation over Attack–Defense Trees

Stefano Iannucci, Shahram Rahimi, and Ioana Banicescu

Shortcut Learning in Early Intrusion Detection: Header Fingerprints Across Four Datasets

Michele Guida, Stefano Iannucci, Raj Patel, Shahram Rahimi, Sudip Mittal, and Paolo Merialdo

CogMI Application Session 1: Explainable and Human-Centered AI

Time: 03:45 PM – 05:45 PM

Room: Oak

Session Chair: Sheikh Rabiul Islam (University at Albany)

How Explanation Intent Shapes Human Use of XAI for Neural-Network Predictions

Muhammad Umair Danish, Memoona Aziz, Umair Rehman, and Katarina Grolinger

Generalization-Aware and Explainable AI for EEG-Based ADHD Classification

Nazia Natasha, Stephen Gravereaux, Akash Pedapaga, Sammed Admuthe, and Sheikh Rabiul Islam

When Can a System Trust Its Own Disagreement? Characterizing a Multimodal Affective Signal

Benjamin Hurt

Does AI Benefit Measure the User? Case-Level Persistence and Construct Validity of Person-Specific AI Assistance Effects

João Borges

From Macro to Micro: Intent-Aware Micro-Assessment Generation

Aashish Pandey and Ram Dantu

Leveraging an LLM-Based Evaluator with Code Analysis for Student Code Evaluation

Hadel Alhawasi and Abdou Youssef

Invited Session 4: Trustworthy AI, Explainable Learning and Responsible AI Governance

Time: 03:45 PM – 05:45 PM

Room: Willow

Session Chair: James Joshi (University of Pittsburgh)

Safety Tax: Safety Alignment Makes Your Large Reasoning Models Less Reasonable

Ling Liu

Explainable Continual Learning: Memory Organization and Evolution

KC Santosh

Agentic AI Governance Agent-in-the-Loop Governance Framework for Agentic AI Ecosystem

Meruyert Karzhaubayeva, James Joshi

M-Fibration Theory with Applications to Neural Network Compression

Paolo Boldi

RISC Poster Session: Space Systems

Time: 03:45 PM – 05:45 PM

Room: Juniper

Session Chair: Rob Cunningham (University of Pittsburgh), Zeba Wahab (Independent)

SPACE MEDICINE

A Cardiovascular Digital Twin for Personalized Prediction of Microgravity Adaptation and Acute Gravity Transitions

Lakshya Dharwal (Somnium Biolabs), Bhavna Malladi (University of California Berkeley), Anna Wu (Carnegie Mellon University), and Chintan Dave (Somnium Biolabs)

Dynamic Hippo-YAP Modeling for Osteocyte Mechanotransduction in Spaceflight

Adithi Kona (UC Riverside)

SPACE RESILIENCE

Design of a Digital Forensics Framework for Spacecraft Flight Software

Logan Flowers (New Mexico State University) and Efrén López-Morales (New Mexico State University)

Does Drift Severity Predict Forecaster Risk in Spacecraft Telemetry?

Korede Bishi (University of Georgia), Adetayo Okunoye (University of Georgia), and Dominic Okonkwo (University of Georgia)

CAPSTONE as a Software-Defined Satellite: Demonstrating Multi-Level On-Orbit Software Updates at Lunar Distance

Elizabeth Geist (NASA), William Stanton (NASA), Paul Swenson (NASA/Vantage Systems Inc.), and Theresa Beech (NASA)

CogMI Application Session 2: Reasoning and Reliability in Large Language Models

Time: 03:45 PM – 05:45 PM

Room: Cedar/Pine

Session Chair: Reza Zafarani (Syracuse University)

The Long Road to the Same Answer: Cognitive Bias Under Escalating Reasoning Budgets in Large Language Models

Obada Kraishan

Cognitive Stability of Chain-of-Thought Reasoning in Large Language Models Under Semantic Perturbation

John Kingsley Arthur, Eric Appiah Mantey, Adu Baffour, Michael Kelly Osei, and Dominic Dalyngton Damoah

Measuring Compositional Reasoning in Deployable LLMs on Financial Payoffs with Risk-Decomposed Mispricing

Varun Viswapriyan

When Gold Labels Move: Evidence-State Heterogeneity and Oracle Sensitivity in Multi-Source LLM Evaluation

Byungsik Seo

Loud Failures, Quiet Failures: Fault Detection and Recovery in Tool-Using Language Model Agents

Obada Kraishan

TPS Research Session 5: Access Control, Authentication, Blockchain, and Data Provenance

Time: 03:45 PM – 05:45 PM

Room: Redwood

Session Chair: Asma Jodeiri Akbarfam (Washington State University)

Dynamic Trust Based Attribute-Centric Access Control for Smart Homes

Hritik Thapa and Smriti Bhatt

Enhancing Security of Template-Less Biometric Authentication Against Geometric Inversion Attack

Mahafujul Alam, Lan Zhang, and Bertrand Cambou

Pre-Authorization Screening of Ethereum EIP-7702 Account Delegation under an Execution-Audited Adaptive Adversary

M M Imran, Anik Tahabilder, Changxin Bai, Qiang Zhu, and Shiyong Lu

Cross-Chain Detection of Pump-and-Dump Attacks in Permissionless Blockchains

Abdulrahman Alhaidari, Bhavani Kalal, Balaji Palanisamy, and Shamik Sural

WMDS: Contract-Linked Layered Watermarking for Attributable Provenance in Data Spaces

Anhelina Kovach, Shuwen Liu, Yucong Li, Jorge Lanza, Leticia Montalvillo, Aitor Urbietta, and George C. Polyzos

TPHAC Session 2

Time: 03:45 PM – 05:45 PM

Room: TBA

Session Chair: TBA

Banquet (provided by conference)

Room: Redwood/Maple

06:45 PM – 09:45 PM

Conference Day 3: November 6, 2025

Registration (Room: Foyer) & Continental Breakfast (provided by conference) (Atrium – hotel will have private signs)

7:15 AM - 8:15 AM

Welcome and Opening Remarks

8:15 am – 8:20 am

All Participants and Chairs

Room -Willow/Oak/Cypress

Keynote 4

08:20 AM – 9:20 AM

Keynote 4: Ramesh Raskar, MIT Media Lab, USA

(Room: Willow/Oak/Cypress)

Title: Billions of Personal Agents vs. AGI: Which Future Wins?

(Chair: Danda B. Rawat, Howard University, USA)

Keynote 5

09:20 AM – 10:20 AM

Keynote 5: Craig Knoblock, University of Southern California, USA

(Room: Willow/Oak/Cypress)

Title: From Papers to Knowledge: Building Knowledge Graphs from Technical Literature

(Chair: Reza Zafarani, Syracuse University, USA)

Coffee Break (10 min)

Room: Atrium

TPS Application Session 1: Secure Agentic AI and AI-Assisted Vulnerability Analysis and Forensics

Time: 10:30 AM – 12:30 PM

Room: Cypress)

Session Chair: Shaswata Mitra (University of Alabama)

Runtime Scope Composition for Secure Tool Invocation in Stateless Model Context Protocol Agents

Vishnu Gatla and Pushpjeet Shrivastava

Livelock in the Agent Harness: Delegation-Induced Liveness Failure in LLM Systems

Mohammad Hasan, Aciful Khan, Hadi Rashid, Md Motasim Billa, Shafkat Islam, and Mariofanna Milanova

VEXA IoT: Autonomous IoT Vulnerability EXPloitation using AI Agents

Katherine Swinea, Kshitiz Aryal, Lopamudra Praharaj, and Maanak Gupta

Identifying Multiple Security Patches for Complete Vulnerability Fixes

Aarya Niraula, Weiliang Qi, W. Eric Wong, and Xinda Wang

Evidence-Grounded LLM-Assisted Android Mobile Forensic Reporting

Madhab Chandra Das, Sajid Ahmed, Tasnia Tabassum, and Qingzhong Liu

CogMI Application Session 3: Efficient, Resilient, and Secure AI Systems

Time: 10:30 AM – 12:30 PM

Room: Oak

Session Chair: Hoda Mehrpouyan (Boise State University)

HIPQ-R: Interface-Sensitive Risk-Constrained Compression of Hybrid ANN–SNNs for Resource-Constrained Neuromorphic Perception

Xueyi Wang and Zhenyu Zhang

Adaptive Transformer Neural Architecture Search via Federated Reinforcement Learning and Distillation-Based Architecture Evaluation

Sudha Pelluri, Navya Chilaka, and Preethi Taralpalli

Toward Voice-Driven Edge Intelligence: Lightweight LLM Deployment for Interactive Application Generation

Hyun Suk Nam and Andrew Lee

CascadeAuth: Separating the Resource Knob from the Decision Boundary in Risk-Adaptive Authorization Cascades

Lovina Dmello

AWARE-RG: Abstraction Aware Explanation of Resiliency Graphs for Cyber Physical Systems

Shadaab Kawnain Bashir, Rakesh Podder, Philip D. Hopkins, Sarath Sreedharan, Indrakshi Ray, and Indrajit Ray

TPS Application Session 2: AI Assurance, Auditing, Red-Teaming, and Supply-Chain Security

Time: 10:30 AM – 12:30 PM

Room: Oak

Session Chair: Hoda Mehrpouyan (Boise State University)

VerifyBench: Auditing Public Verifiability in Responsible-AI Benchmarks for Foundation Models

Anshima Srivastava

RLMS-Shield: The Abstention Frontier and Peer-Prediction Auditing for Automated LLM Red-Teaming

Matias Zabaljauregui

What's Actually In the Models You Pull? A Registry-Wide Audit of Behavioral Supply-Chain Integrity in the Consumer GGUF Ecosystem

Mike Wagner

SAFESHIELD: A Decision-Organization Framework for Deployment-Time Safety of Small Language Models

Xingru Zhou, Luis Sentis, and Aarti Choudhary

Trust-Boundary Architectures for Auditable LLM-Assisted Contract Risk Analysis: Threat Model, Enforcement, and Empirical Evaluation

Santhosh Guntupalli

TPS Application Session 3: Privacy-Preserving Data Usage, Synthetic Data, and Policy Governance

Time: 10:30 AM – 12:30 PM

Room: [Juniper](#)

Session Chair: Ram Dantu (University of North Texas)

Auditable Yet Confidential: Machine-Verifiable Trust in Cross-Organizational Digital Ecosystems

Christoph Stach, Ioulia Litou, and Clémentine Gritti

Auditable Privacy-Preserving Data Usage for Trustworthy Intelligent Data Services

Wei-Yu Liu, Ming-Yu Ho, and Gwan-Hwan Hwang

Truth, Permissiveness, and Silence: What a Database May Say When It Cannot Say Everything

Hasan Jamil

Identity-by-Descent Membership Inference Attacks on Synthetic Genomic Data

Xinyue Wang and Jaideep Vaidya

ONSET: Ontology-Supervised Extraction and Translation for Secure Natural-Language Policy Automation

Indravadan Patel and Ram Dantu

Government/Industry Track 1: Governance, Trust, and Security for AI and Cloud Systems 1

Time: 10:30 AM – 12:30 PM

Room: [Cedar/Pine](#)

Session Chair: Zeba Wahab (Independent Researcher)

Code Over Prompts, Gates Over Judges: Deploying and Blindly Re-Evaluating an Autonomous Security-Bug Remediation Agent

Shubhankar Unhale and Purshotam Shah

Tamper-Verifiable Audit Logs for Hardware Security Modules Using Merkle-Tree Integrity Anchored to Vendor Device Identity

Neelakantam Gaddam

SCOPE: Semantic & Context Oriented Prompted Extraction

Ashwin Parameswaran and Soamar Homsai

Agent-Warden: eBPF-Based Kernel-Native Process-File Provenance Tracking for LLM Agents

Dongxu Cui, Zhichao Gu, Ping Zheng, Simeng Han and Yong Liao

A Threat Model for Agentic DevOps Assistants: Attack Surface, Trust Boundaries, and Least-Privilege Guardrails for LLM Agents in Cloud-Native Operations

Parth Haresbhai Sabhadiya, Fenil Kanubhai Godhani and Vaikunth Desai

EEE CIC NextGen Program: Juniors at Conference.

website: <https://cic.ieee-cs.org/2026/nextgen-cs-juniors/>

Time: 8:30 AM – 4:30 PM

Room:

Session Chair:

NAND Workshop Session 1:

Time: 10:30 AM – 12:30 PM

Room:

Session Chair: TBA

Lunch Break

12:30 PM – 1:30 PM (Room: Atrium)

TPS Application Session 4: Trusted and Secure Systems, Networks, and Critical Infrastructures

Time: 01:30 PM – 03:30 PM

Room: Juniper)

Session Chair: Sayantan Ghosh (LinkedIn)

AppScopes: Fine-grained Access Control for App Collaboration on Android

Sanket Goutam, Kasra Kakaee, and Amir Rahmati

Shortcut or Signature? A Confound-Controlled Reassessment of Network-Traffic Honeypot Detection

Teresita Noelia Nunez Miglitoris, Michael De Lucia, and Andrew Novocin

Beyond Timing Evasion: A Machine Learning Analysis of Evolving Internet-Wide Reconnaissance

Citlalli Martinez Cervantes, Ramy Othman, Alex Carbajal, Shereen Ismail, and Asma Jodeiri Akbarfam

Cyber-Attack Containment in Heat-Recovering Data Centers: Shifting the Resilience Frontier

Hoda Mehrpouyan

CROWN: Carbon- and Price-Aware Multi-Objective Reinforcement Learning for Battery Energy Management under Integrity Attacks

Shengwen Ding and Mai Abdelhakim

CogMI Application Session 4: Visual Perception and Vision-Language Models

Time: 1:30 PM – 3:30 PM

Room: Willow/Oak/Cypress

Session Chair: KC Santosh (Oregon Institute of Technology)

An Accurate and Efficient Two-Stage Framework for Gun Detection in Videos

Badhan Chandra Das, M. Hadi Amini, and Yanzhao Wu

Beyond Pairwise Feedback: Listwise Vision-Language Supervision for Preference-Based Reward Learning

Srivalli Katkuri, Maxwell Kawada, and Juan Wachs

Cheap to Verify, Expensive to Fix: Deterministic Structural Verification as a Routing Signal for Matched-Cost VLM Table Extraction

Michael Hamaty and Amith Kamath Belman

Neuro-Symbolic Graph Reasoning for Explainable Medical Image Retrieval

Divyesh Patel, Rishab Darshan Shylendra, Piyush Gulhane, Pavithran Gnanasekaran, and Alina Vereshchaka

Declarative Smart Image Analysis: From Analytical Intent to Verified Execution with CERNŌ

Moumita Chanda, Alexander Maas, and Hasan M. Jamil

Invited Session 5: AI for Societal Impact, Knowledge Integration and Information Integrity

Time: 1:30pm – 3:30pm

Room: Willow

Session Chair: Danda B Rawat (Howard University)

A Global Switchboard for the Agentic Web: Connecting Discovery Islands Across Enterprises, SMBs, and Individuals

Ramesh Raskar, Pradyumna Chari, Luca Muscariello, Samuel Sharaf, Junjie Bu, Karan Bharadwaj, Vijay Pandey

Rethinking Threat Modeling for Agentic AI Systems

Andrew Cahan, Elisa Bertino

Challenges and Opportunities in Research on Nonprofits: Measuring the Countable and Assessing the Noncountable

Calton Pu

Fake News Theories: Harnessing Disciplinary Insights for Computational Modeling, Detection, and Explanation

Reza Zafarani

Government/Industry Track 2: Governance, Trust, and Security for AI and Cloud Systems 2

Time: 01:30 PM – 03:30 PM

Room: Cedar/Pine

Session Chair: Julian Jarrett (Lutron Electronics)

Topology-Preserving Obfuscation of Network Configurations for LLM Analysis

Giannis Tziakouris and Nadhem Al-Fardan

Unified Risk Budgets for Reliability, Security, and Compliance in Cloud-Native Operations

Deepak Bhaskaran

Scalable Detection of Adversarial Synthetic Slop and Coordinated Media Abuse: A LoRA-Enabled Multimodal Defense System

Abhinav Mathur, Claire Liu, Kelvin Tan and Yifei Liu

DynaToken: An AI-Driven Context-Aware Framework for Dynamic Least-Privilege Access Tokens

Aniket Tendulkar and Kuldeep Marganache

Permission Anomaly Detection in Salesforce Access Graphs: A Synthetic Benchmark and a Per-Archetype Diagnostic

Sharan Sukesh

IEEE CIC NextGen Program: Juniors at Conference.

website: <https://cic.ieee-cs.org/2026/nextgen-cs-juniors/>

Time: 8:30 AM – 4:30 PM

Room:

Session Chair:

NAND Workshop Session 1:

Time: 01:30 PM – 03:30 PM

Room:

Session Chair: TBA

Closing Remarks

Room: Willow/Oak/Cypress

3:30 PM – 4:00 PM