Conference on AI and Agentic Systems

Association for Computing Machinery

MEDIA ADVISORY

Contact:
Jim Ormond
ACM Media Relations
212-626-0505
ormond@hq.acm.org

AI Agents, "Vibe Coding," and the Future of Real-World AI Systems Take Center Stage at Inaugural ACM CAIS Conference

New peer-reviewed research examines whether AI agents work in practice—and what happens when they don't

SAN JOSE, Calif. — May 20, 2026 — AI agents are quickly becoming one of the technology industry's biggest bets, but researchers gathering at the inaugural ACM Conference on AI and Agentic Systems (CAIS 2026) believe critical questions remain about whether these systems are truly reliable, secure, and ready for real-world deployment.

From May 26–29 in San Jose, California, CAIS 2026 will bring together researchers, engineers, and practitioners presenting new peer-reviewed research examining how these systems behave in practice—from AI coding agents and multi-agent collaboration to cybersecurity misuse, safety failures, enterprise deployment and AI social behavior. The conference will feature 61 peer-reviewed research papers alongside 45 live demonstrations of emerging agentic AI systems and workflows.

Presented by ACM, the Association for Computing Machinery, the conference focuses on the architectures and engineering practices behind modern AI systems, including retrieval-augmented generation (RAG), multi-agent systems, tool-using AI agents, and large-scale deployment infrastructure.

A limited number of press passes are available for journalists interested in attending CAIS 2026 in person. Reporters interested in attending should email josh@cmpnd.ai with a brief statement of interest.

Selected Research Highlights

- "Why Johnny Can't Use Agents: Industry Aspirations vs. User Realities with AI Agents": Researchers at Carnegie Mellon University evaluated 102 commercial AI agents against real-world user tasks and identified major gaps between how AI agents are marketed and what users can successfully accomplish in practice.

- "ViBench: A Benchmark on Vibe Coding": One of the first rigorous evaluations of vibe coding—building applications through natural-language prompting alone—found that even frontier AI systems struggle with realistic end-to-end software development

workflows.

- "The Cost of Consensus: Isolated Self-Correction Prevails Over Unguided Homogeneous Multi-Agent Debate" - Researchers from the Jozef Stefan Institute found that groups of AI agents can amplify errors rather than correct them, often collapsing into false consensus instead of effective peer review.

- "Does Socialization Emerge in AI Agent Society? A Case Study of Moltbook" - Researchers studying an AI-agent social network observed human-like social behaviors emerging among autonomous agents, including influence persistence and collective consensus formation.

- "Malice in Agentland: Down the Rabbit Hole of Backdoors in the AI Supply Chain" - Researchers from ServiceNow Research, Mila, and Polytechnique Montréal demonstrate how AI agent systems can be compromised through hidden backdoors embedded in training data, base models, or deployment environments.

- "When Harmful Intent Dissolves into Technical Detail: How Safe Are Coding Agents Against Cyber Misuse?" - Researchers at Purdue University found that coding agents can execute harmful cyber tasks when malicious intent is distributed across seemingly harmless technical steps.

- "The Verifier Tax: Horizon Dependent Safety–Success Tradeoffs in Tool Using LLM Agents" - Researchers quantified how adding runtime safety enforcement to AI agents can significantly reduce task success rates as interactions become more complex.

Featured Demonstrations

- "From Bug Report to Pull Request" - An autonomous AI agent pipeline capable of tracing production software errors, identifying root causes, and automatically opening code fixes

- "Agent 4: Teamwork and Collaboration for Vibe-Coding" - A multi-agent system designed to collaboratively build applications through natural-language prompting workflows

- "TRACE: A Multi-Agent System for Natural Language-Driven Social Graph Investigation" - A multi-agent system for social graph forensics that uses natural-language behavior detection and LLM-driven graph exploration, achieving 10 times network expansion and 91.9% discovery of unknown suspicious entities.

Keynote Speakers

- **Andy Konwinski**, co-founder of Databricks and Perplexity AI and founder of Laude Institute
- **Thariq Shihipar,** Member of Technical Staff working on Claude Code at Anthropic

- **Percy Liang**, Stanford University professor and director of the Center for Research on Foundation Models (CRFM)

Conference Details

- **Dates:** May 26–29, 2026
- **Location:** San Jose, California
- **Website:** https://caisconf.org

About ACM

ACM, the Association for Computing Machinery, is the world's largest educational and scientific computing society, uniting educators, researchers, and professionals to inspire dialogue, share resources, and address the field's challenges. ACM strengthens the computing profession's collective voice through strong leadership, promotion of the highest standards, and recognition of technical excellence. ACM supports the professional growth of its members by providing opportunities for life-long learning, career development, and professional networking.

###