MEDIA ADVISORY

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

Tom Romanoff
ACM Policy Director
212-626-0543
romanoff@hq.acm.org

ACM Technology Policy Council: Agentic AI Is Outpacing the Laws & Safeguards Designed to Govern It

New Techbrief Examines Legal Liability, Security Risks, and Workforce ImpactsaAs Autonomous AI Systems Move Into Mainstream Deployment

New York, NY — June 11, 2026 — AI systems are increasingly browsing the web, executing code, managing files, and sending messages without step-by-step human approval, raising new risks in the process, according to a new TechBrief from the Association for Computing Machinery's Technology Policy Council (TPC) on the rise of agentic AI.

The TechBrief, "Agentic AI: Autonomy, Opportunities, and Challenges of Action-Taking AI Systems," examines AI systems that plan and execute multi-step tasks toward a user-defined goal. Such agentic systems are being rapidly adopted by enterprises and consumers. A 2025 survey of more than 500 technology leaders found that 48% are already deploying or adopting agentic AI. The TechBrief finds that this acceleration is outpacing the legal, regulatory, and technical frameworks designed to govern it.

The brief arrives as governments are beginning to respond. On May 1, 2026, CISA and five allied national cybersecurity agencies published the first coordinated multinational security guidance specifically targeting agentic AI, just 42 days after the White House released its National Policy Framework for Artificial Intelligence on March 20. Yet neither effort fully resolves the accountability questions agentic systems raise.

"Many people are rapidly adopting agentic AI systems for their businesses and personal lives. They know that these systems can cause great harm when they misbehave, but the short-term advantages of deploying them and hoping for the best are nearly impossible to resist," said Simson Garfinkel, Chief Scientist at BasisTech and Chair of the ACM TechBriefs Committee. "These systems can offer tremendous advantages to their users, but anyone deploying them today is taking on real risk with very little legal protection. When something goes wrong with a system that takes actions on a user's behalf, it can be genuinely difficult to determine who is responsible. Existing law simply doesn't answer this question."

The TechBrief identifies four key policy dimensions where existing frameworks fall short:

- **Legal liability without a clear accountable party:** When an AI agent causes harm, as in a documented case where one deleted a company's entire production database, responsibility may fall to the model provider, the framework developer, the person or company that deployed the agent, or the end user. No person made the decision, yet harm was done. No case law currently exists that can help resolve liability.
- **Serious and underappreciated security risks:** Because Large Language Models (LLMs) process text as both data and commands, agentic systems cannot reliably distinguish legitimate content from embedded malicious instructions. Documented incidents include an AI agent that exposed private Slack data after processing a message containing hidden instructions, and consumer agent marketplaces found to contain malicious extensions reaching hundreds of thousands of users.
- **Lack of consumer transparency and recourse:** Users often cannot determine what systems an agent can access, what actions it can take without confirmation, or how to revoke its permissions. No standardized data format currently exists to disclose or control agent authority, a concern the brief flags as particularly acute in high-stakes settings such as healthcare.
- **Workforce disruption outpacing evidence:** A 2025 Gartner survey found that 55% of supply chain leaders expect agentic AI to reduce entry-level hiring. Yet the productivity claims driving those decisions have not been independently verified at scale, and the long-term effects on skill formation and labor markets remain unmeasured.

"Who is legally responsible when autonomous systems cause harm?" Garfinkel continued. "Today's license agreements just point fingers elsewhere. But if there is no underlying technology that can make these systems reliably follow the policies they are given, that finger-pointing may not be legally binding. This is not a problem specific to the US, Europe, or any single Asian country. We need to work on both law and technology to provide strong assurances to businesses and consumers throughout the entire industrialized world."

The TechBrief concludes that addressing these challenges will require defined authentication and delegation standards, robust audit trails, standardized consumer disclosures, and sector-specific guidance in areas such as healthcare, financial services, and critical infrastructure, where existing law assumes a human decision-maker.

Read the full TechBrief:

ACM's TechBriefs are designed to complement ACM's activities in the policy arena and to inform policymakers, the public, and others about the nature and implications of information technologies. Earlier ACM TechBriefs have covered topics such as vibe coding, buying vs building LLMs, automated speech recognition, governmental digital transformation, accessibility, and generative artificial intelligence among others.

About the ACM Technology Policy Council

ACM's global Technology Policy Council sets the agenda for global initiatives to address evolving technology policy issues and coordinates the activities of ACM's regional technology policy committees in the US and Europe. It serves as the central convening point for ACM's interactions with government organizations, the computing community, and the public in all matters of public policy related to computing and information technology. The Council's members are drawn from ACM's global membership.

About ACM

ACM, the Association for Computing Machinery, is the world's largest educational and scientific computing society, uniting computing 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.