Who Gets to Decide? — Designing Human-AI Teaming for Security (HATS)

Call for Papers
SOUPS 2026 Workshop on August 23, 2026

Overview

Traditionally, organizational cybersecurity has relied on systems that place the human explicitly in control. When a security alert fired or an access request was flagged, a human analyst or employee ultimately made the call. However, the rapid integration of agentic AI into cyber-defense and daily enterprise workflows is challenging this principle. As AI systems shift from advisory tools to autonomous agents capable of taking action, guaranteed human oversight is at threat of erosion.

Whether employees are reviewing AI chatbot outputs for sensitive data leaks, software engineers are evaluating AI-generated code, or SOC analysts are triaging automated alerts, this shifting balance of autonomy creates critical, underexplored vulnerabilities. Inappropriate reliance on AI, ranging from under-reliance due to broken trust to over-reliance resulting in automation bias, can severely compromise organizational security.

Topics of Interest

To address this, we urgently need robust design principles that ensure appropriate AI autonomy and meaningful human agency in security decision-making. This workshop invites researchers and practitioners from diverse backgrounds to help us identify the practical barriers to human-AI teaming, define the metrics needed to evaluate it, and co-create design principles to optimize it.

Topics of interest include, but are not limited to:

- Designing for appropriate automation levels and agentic AI in cybersecurity.
- Automation bias, trust calibration, and employee perceptions of AI decisions.
- Human-centric frameworks and workflows for shared decision-making in security teams.
- AI-augmented threat hunting, signal processing, and SOC triage.
- Technical, psychological, interpersonal, or societal risks associated with human-AI security workflows.
- Error handling in AI-driven security tools and its impact on organizational dynamics.
- Practical barriers and facilitators for deploying human-AI collaboration in enterprise security.
- Opportunities for human-AI teaming to address cybersecurity workforce shortages.

Submission Guidelines

To participate, please submit a **two- to four-page position paper** (including references and any appendices) presenting your perspective on decision-making and human-AI teaming in the context of cybersecurity. All papers must follow the SOUPS formatting template. Submissions will be reviewed with acceptance criteria focused on: (1) relevance towards the workshop themes, (2) diversity of perspectives, and (3) space limitations. Further, we review scholarly rigour, originality, and clarity.

We plan to accept up to 30 participants, and at least one author of each accepted paper must attend the workshop in person. Accepted papers will be published as workshop proceedings. During the workshop, you will present your ideas and take part in collaborative group activities aimed at co-creating actionable design principles for human-AI interaction.

Important Details:

- **Submission Deadline:** Monday, June 01, 2026 AoE (Anywhere on Earth)
- **Author Notification:** Sunday, June 14, 2026 AoE (Anywhere on Earth)
- **Camera-ready Deadline:** Thursday, June 25, 2026 AoE (Anywhere on Earth)
- **Workshop Website & Further Information:** `https://hats-soups26.github.io/`
- **Submission:** `https://hats2026.hotcrp.com`
- **Formatting Instructions:** Available on the SOUPS website

Let's come together and co-create actionable principles to inform future research on human-AI teaming for cybersecurity decision-making, as we tackle the question: "Who gets to decide?"