

Presenter Name: Hao-Ping (Hank) Lee

Talk Title: Privacy in the age of AI: What has changed and what should we do about it?

Presenter Short Bio: Hao-Ping (Hank) Lee is a PhD student at the Human-Computer Interaction Institute at Carnegie Mellon University, advised by Professors Sauvik Das and Jodi Forlizzi. His research lies at the intersection of usable privacy security, human-computer interaction (HCI), and human-centered AI. He studies and builds tools that enable practitioners to identify, reason about, and mitigate AI-entailed privacy risks during the development of consumer AI products. His research has been published at top privacy & security and HCI conferences such as IEEE S&P, USENIX, and CHI.

Short Talk Abstract: How does AI change privacy? Are the designers, engineers, and technologists who create AI technologies equipped to recognize and mitigate the unique privacy risks entailed by the AI products and services they create? Addressing these questions is crucial to steer the development of AI products and services toward their promise and away from privacy invasions. In this presentation, I will detail my PhD research on privacy in AI. I will begin by introducing a taxonomy of AI privacy risks we created, highlighting how AI changes the landscape of privacy by introducing risks not previously accounted for and amplifying the existing ones. Following this, I will discuss insights from our interviews with 35 AI practitioners, which reveal their practices of AI privacy work and the barriers therein — awareness, motivation, and ability. Lastly, I will introduce our ongoing efforts to develop a tool designed to assist AI practitioners in identifying privacy risks associated with their early-phase AI product ideas. This tool aims to enhance practitioners' understanding of the trade-offs between AI utility and privacy, promoting more informed and responsible AI development.