

SEMLA 2019

The 2nd International Symposium on Software Engineering for Machine Learning Applications

May 23th – 24th 2019, Montréal, QC, Canada



Stephen J. Eglash is Director of Research at the Stanford Institute for Human-Centered AI (Stanford HAI) where his group is responsible for research programs and external partnerships. Stanford HAI works to advance AI technology and applications, as well as to understand and influence the interaction of AI with society. Stanford HAI conducts research in AI technology, policy, and societal impact; convenes stakeholders from academia, industry, government, and civil society; and educates students, professionals, and others. Previously at Stanford, Steve developed and managed research programs in computer science, energy, and sustainability. Prior to joining Stanford, Steve was president and CEO of the solar energy company Cyrium Technologies, consultant for the National Renewable Energy Laboratory and the US Department of Energy, venture capitalist at Worldview Technology Partners, vice president at SDL (JDSU), and member of the technical staff at MIT Lincoln Laboratory. Steve received a PhD and MS from Stanford University, and a BS from the University of California at Berkeley, all in Electrical Engineering.

Keynote Title: Human-Centered AI Systems

Abstract: Practitioners and researchers in software engineering and machine learning are increasingly considering human factors and societal impact when developing sophisticated applications. These human-centered systems need to work in ways that are safe, reliable, and reasonable, even in expected situations and with culturally diverse users. Applications like insurance decision making, autonomous vehicles, commercial aircraft collision avoidance, the electric grid, financial transactions, healthcare, and controlling massive chemical plants are particularly unforgiving. These are applications where errors can lead to serious and catastrophic consequences. Industry, research, government, and society need AI systems that are verifiable, reliable, robust against adversarial attacks, auditable, explainable, and unbiased. This talk will describe human-centered AI, present current research on safe and reliable systems, and discuss how developers can make use of emerging solutions.

What is SEMLA?

SEMLA aims at bringing together leading researchers and practitioners in software engineering and machine learning to reflect on and discuss the challenges and implications of engineering complex data-intensive software systems.

Full Program: <https://semla.polymtl.ca/program/>

Registration: <https://semla.polymtl.ca/registrationpage/>