

Kush R. Varshney, Distinguished Research Staff Member and Manager with IBM Research, Thomas J. Watson Research Center

Title: Developing Trust in Artificial Intelligence and Machine Learning for High-Stakes Applications

Abstract: In recent times, we often hear a call of the governance of AI systems, but what does that really mean? In this talk, I will first adopt a control theory perspective to explain governance that includes diverse stakeholders determining the reference input via value alignment, data scientists acting as the controller to meet the values in a machine learning system, and facts captured in transparent documentation as the feedback signal. I will later go into further depth on value alignment via CP-nets and performance metric elicitation, and transparency via factsheets. The talk will assume some familiarity with algorithmic fairness, robustness, and explainability.

Bio: Dr. Varshney is a distinguished research staff member and manager with IBM Research at the Thomas J. Watson Research Center, Yorktown Heights, NY, where he leads the machine learning group in the Foundations of Trustworthy AI department. He was a visiting scientist at IBM Research - Africa, Nairobi, Kenya in 2019. He is the founding co-director of the IBM Science for Social Good initiative. He applies data science and predictive analytics to human capital management, healthcare, olfaction, computational creativity, public affairs, international development, and algorithmic fairness, which has led to recognitions such as the 2013 Gerstner Award for Client Excellence for contributions to the WellPoint team and the Extraordinary IBM Research Technical Accomplishment for contributions to workforce innovation and enterprise transformation, and Harvard Belfer Center Tech Spotlight runner-up for AI Fairness 360. He conducts academic research on the theory and methods of trustworthy machine learning. His work has been recognized through best paper awards at the Fusion 2009, SOLI 2013, KDD 2014, and SDM 2015 conferences and the 2019 Computing Community Consortium / Schmidt Futures Computer Science for Social Good White Paper Competition. He is currently writing a book entitled 'Trust in Machine Learning' with Manning Publications. He is a senior member of the IEEE.

Photo:

