HUMANITY CENTERED ROBOTICS INITIATIVE
RHODE ISLAND

BROWN

Humanity Centered Robotics Initiative Talk

Scott Niekum, University of Texas

Wednesday, September 18 at 2pm
CIT 368

"Scaling Probabilistically Safe Learning to Robotics"

Abstract: Before learning robots can be deployed in the real world, it is critical that probabilistic guarantees can be made about the safety and performance of such systems. In recent years, so-called "high-confidence" reinforcement learning algorithms have enjoyed success in application areas with high-quality models and plentiful data, but robotics remains a challenging domain for scaling up such approaches. Furthermore, very little work has been done on the even more difficult problem of safe imitation learning, in which the demonstrator's reward function is not known. This talk focuses on new developments in three key areas for scaling safe learning to robotics: (1) a theory of safe imitation learning; (2) scalable inverse reinforcement learning in the absence of models; (3) efficient policy evaluation. The proposed algorithms offer a blend of safety and practicality, making a significant step towards high-confidence robot learning with modest amounts of real-world data.

Scott Niekum is an Assistant Professor and the director of the Personal Autonomous Robotics Lab (PeARL) in the Department of Computer Science at UT Austin. He is also a core faculty member in the interdepartmental robotics group at UT. Prior to joining UT Austin, Scott was a postdoctoral research fellow at the Carnegie Mellon Robotics Institute and received his Ph.D. from the Department of Computer Science at the University of Massachusetts Amherst. His research interests include imitation learning, reinforcement learning, and robotic manipulation. Scott is a recipient of the 2018 NSF CAREER Award.

Host: George Konidaris/HCRI