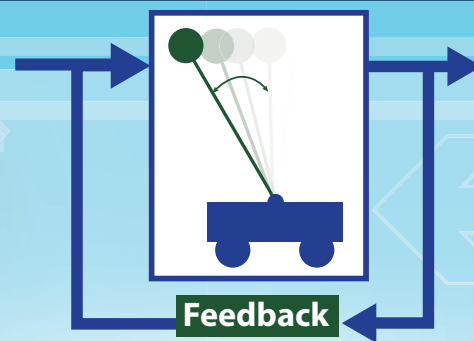


COLLEGE OF ENGINEERING

Control Seminar



Sponsored by: Bosch, Ford, and Toyota

Learning in Non-Stationary Environments: Near-Optimal Guarantees



LILLIAN RATLIFF

Assistant Professor

University of Washington, Department of Electrical Engineering

Friday, March 29, 2019

3:30 pm – 4:30 pm 1500 EECS

ABSTRACT: Motivated by scenarios in which heterogeneous autonomous agents interact, in this talk we present recent work on the development of learning algorithms with performance guarantees for both simultaneous and hierarchical decision-making. Adoption of new technologies is transforming application domains from intelligent infrastructure to e-commerce, allowing operators and intelligently augmented humans to make decisions rapidly as they engage with these systems. Algorithms and market mechanisms supporting interactions occur on multiple time-scales, face resource constraints and system dynamics, and are exposed to exogenous uncertainties, information asymmetries, and behavioral aspects of the human decision-making process. Techniques for synthesis and analysis of decision-making algorithms, for either inference or influence, that fundamentally depend on an assumption of environment stationarity often breakdown in this context. For instance, humans engaging with platform-based transportation services make decisions that are dependent on their immediate observations of the environment and past experience, both of which are functions of the decisions of other users, multi-timescale policies (e.g., dynamic pricing and fixed laws), and even environmental context that may be non-stationary (e.g., weather patterns or congestion). Implementation of algorithms designed assuming stationarity may lead to unintended or unforeseen consequences. Stochastic models with high-probability guarantees that capture the dynamics and the decision-making behavior of autonomous agents are needed to support effective interventions such as physical control, economic incentives, or information shaping mechanisms. Two fundamental components are missing in the state-of-the-art: (i) a toolkit for analysis of interdependent learning processes and for adaptive inference in these settings, and (ii) certifiable algorithms for co-designing adaptive influence mechanisms that achieve measurable improvement in system-level performance while ensuring individual-level quality of service through design-in high-probability guarantees. In this talk, we discuss our work towards addressing these gaps. In particular, we provide (asymptotic and non-asymptotic) convergence guarantees for simultaneous play, multi-agent gradient-based learning (a class of algorithms that encompasses a broad set of multi-agent reinforcement learning algorithms) and performance guarantees (regret bounds) for hierarchical decision-making (incentive design). In addition, we draw connections between several existing algorithmic frameworks for multi-agent learning showing a common derivation for such schemes and we discuss interpretations that suggest interesting approaches for learning approaches to multi-agent settings with agents of bounded rationality. Building on insights from these results, the talk concludes with a discussion of interesting future directions in the design of certifiable, robust algorithms for adaptive inference and influence.

BIO: Lillian J. Ratliff is an Assistant Professor in the Department of Electrical Engineering at the University of Washington. Prior to joining UW she was a postdoctoral researcher in EECS at UC Berkeley (2015-2016) where she also obtained her PhD (2015) under the advisement of Shankar Sastry. She holds a MS (UNLV 2010) and BS (UNLV 2008) in Electrical Engineering as well as a BS (UNLV 2008) in Mathematics. Her research interests lie at the intersection of game theory, optimization, and learning.

M | **ECE** ELECTRICAL &
COMPUTER ENGINEERING
UNIVERSITY OF MICHIGAN

BOSCH



Questions? Contact: Judi Jones asap@umich.edu