

2017-2018 NRT-DESE Seminar Series

Presented by the National Science Foundation, NSF Research Traineeship Data-Enabled Science and Engineering (NRT-DESE) award for Graduate Training in Data-Enabled Research into Human Behavior and its Cognitive and Neural Mechanisms

Friday, October 6, 2017
11:00 am - 12:00 pm Talk
Wegmans Hall 1400 (Auditorium)

Deep Reinforcement Learning for Sequential Decision Making Tasks with Natural Language Interaction

Abstract: The success of Deep Learning (DL) on visual perception has led to rapid progress on Reinforcement Learning (RL) tasks with visual inputs. More recently, Deep Learning is showing promise at certain kinds of supervised natural language problems and this too is making its way into helping on RL tasks with natural language inputs. In this talk, I will describe two projects in this direction from my group. The first (url 1 below) involves learning to query, reason, and answer questions on simple forms of ambiguous texts designed to focus on a specific problem that occurs in dialog systems. The second (url 2 below) involves zero shot generalization to unseen instructions in a 3d maze navigation task for which we develop a hierarchical DeepRL architecture.

1. <http://web.eecs.umich.edu/~baveja/Papers/GuoICLR2017.pdf>

2. <http://web.eecs.umich.edu/~baveja/Papers/task-generalization.pdf>

Satinder Singh, PhD - University of Michigan, Professor of Computer Science & Engineering; Director of the Artificial Intelligence Laboratory



Biography: Dr. Satinder Singh is a Professor of Computer Science & Engineering at the University of Michigan where he also served as the Director of the Artificial Intelligence Laboratory from 2006 to 2017. He is also a co-founder and Chief Scientist of Cogitai, Inc. Dr. Singh's research interests focus on the field of Reinforcement Learning, i.e., on building algorithms, theory, and architectures for software agents that can learn how to act in uncertain, complex, and dynamic environments. Specific interests include building models of dynamical systems from time-series data, learning good interventions in human-machine interaction, dealing with partial observability and hidden state in sequential decision-making, dealing with the challenge of exploration-exploitation and

delayed feedback, explaining animal and human decision making using computational models, and optimal querying in semi-autonomous agents based on value of information. He is interested in applications from healthcare, robotics, and game-playing. He is a Fellow of the Association for the Advancement of Artificial Intelligence, was Program Co-Chair of AAI 2016, has received an outstanding faculty award from his department, and has published over 150 papers in his field.

