THE WEIZMANN INSTITUTE OF SCIENCE
FACULTY OF MATHEMATICS AND COMPUTER SCIENCE
Faculty Seminar

Room 155 ,Ziskind Building
on Sunday, Nov 17, 2019at 11:15

Sarah Keren Harvard University

Automated Design of Better Environments for Intelligent Agents

Abstract:

Automated agents and humans increasingly interact and collaborate: at home, at the workplace, on the road, and in many other everyday settings. In all these settings, effectively recognizing what users try to achieve, providing relevant assistance (or, depending on the application, taking relevant preventive measures), and supporting an effective collaboration in the system are essential tasks. All these tasks can be enhanced via efficient system redesign, and often even subtle changes can yield great benefits. However, since these systems are typically complex, hand crafting good design solutions is hard.

Utility Maximizing Design (UMD) addresses this challenge by automating the design process. It does so by offering informed search strategies to automatically and efficiently find optimal design solutions for maximizing a variety of targeted objectives. One example is Goal Recognition Design (GRD), which seeks a redesign to an environment that minimizes the maximal progress an agent can make before its goal is revealed. A second is Equi-Reward Utility Maximizing Design (ER-UMD), which seeks to maximize the performance of a planning agent in a stochastic environment. The third, Reinforcement Learning Design (RLD), parallels ER-UMD, but considers an environment with a reinforcement learning agent.

Among the different ways that may be available to change a setting, the talk will focus on information shaping, which consists of selectively revealing information to a partially informed agent in order to maximize its performance. As an example application, I will demonstrate how information shaping can be applied to enhance the performance of a robot that is navigating in an unfamiliar environment. The talk will conclude with a discussion of future extensions and various applications related to the agenda of Utility Maximizing Design, including human-robot collaboration, intrusion detection, and assisted cognition.