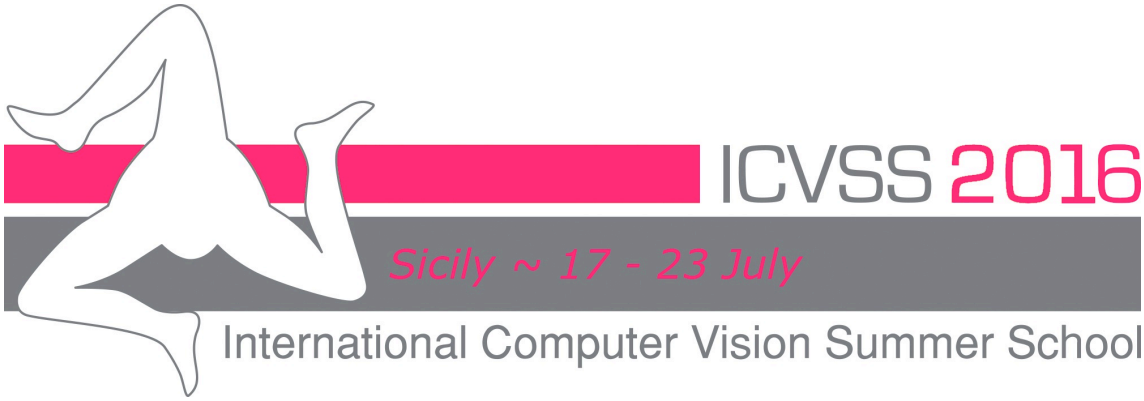

Deep Learning for Robotics: Learning Actionable Representations

Sergey Levine
UC Berkeley, University of Washington, Google, USA

Abstract

Deep learning methods have had a transformative effect on supervised machine perception fields, such as vision, speech recognition, and natural language processing. However, extending the successes of deep learning to robotics has proven more challenging. Robotics, as well as other domains that combine perception with active decision making and control, introduce substantial additional challenges, since each decision that is taken effects the next input to the system. In this lecture, I will provide an overview of current methods that combine deep learning with decision making and control, discuss applications of deep learning to robotic control and perception, and present an in-depth discussion of several topics that are core to "active" deep learning: reinforcement learning, optimal control, and dynamic programming. I will conclude by discussing how techniques from robotic deep learning can be extended to other fields, including computer vision, where they can be used to enable active perception, active data collection, and end-to-end training that combines both perception and decision making into a single architecture that is optimally adapted to the task at hand.

Keywords
Robotics, deep learning, reinforcement learning