

“Reinforcement Learning”

Pieter Abbeel & Jan Peters

U.C. Berkeley / Technische Universität Darmstadt

Friday, 14 May 2012

0900-1530

Room 11 - St. Paul River Centre

Abstract: This all-day tutorial introduces the audience to reinforcement learning. Prior experience in this area is not assumed. In the first half of this tutorial we will cover the foundations of reinforcement learning: Markov decision processes, value iteration, policy iteration, linear programming for solving an MDP, function approximation, model-free versus model-based learning, Q-learning, TD-learning, policy search, the likelihood ratio policy gradient, the policy gradient theorem, actor-critic, natural gradient and importance sampling. In the second half of this tutorial we will discuss example success stories and open problems.

Bios: Pieter Abbeel received a BS/MS in Electrical Engineering from KU Leuven (Belgium) and received his Ph.D. degree in Computer Science from Stanford University in 2008. He joined the faculty at UC Berkeley in Fall 2008, with an appointment in the Department of Electrical Engineering and Computer Sciences. He has won various awards, including best paper awards at ICML and ICRA, the Sloan Fellowship, the Air Force Office of Scientific Research Young Investigator Program (AFOSR-YIP) award, the Okawa Foundation award, and the 2011's TR35. He has developed apprenticeship learning algorithms which have enabled advanced helicopter aerobatics, including maneuvers such as tic-tocs, chaos and auto-rotation, which only exceptional human pilots can perform. His group has also enabled the first end-to-end completion of reliably picking up a crumpled laundry article and folding it. His work has been featured in many popular press outlets, including BBC, New York Times, MIT Technology Review, Discovery Channel, SmartPlanet and Wired. His current research focuses on robotics and machine learning with a particular focus on challenges in personal robotics, surgical robotics and connectomics.

Jan Peters is a computer scientist (Dipl.-Inform. Fernuni Hagen; M.Sc., Ph.D., Univ.S.California), an electrical engineer (Dipl.-Ing. TU Muenchen), and a mechanical engineer (M.Sc. Univ.S.California). Jan has held visiting research positions at ATR, Japan and at National University of Singapore. Currently, Jan Peters is a full professor (W3) at Technische Universität Darmstadt heading the FG Intelligente Autonome Systeme while at the same time leading the Robot Learning Lab at the Max Planck Institute for Intelligent Systems. Jan Peters' research interests span a large variety of topics in robotics, machine learning and biomimetic systems with a strong focus on learning of motor skills. He has (co-)authored 165 papers and received the 2007 Dick Volz Best US Robotics Ph.D. Runner-Up Award.