

Guest editorial: Special issue on robot learning, Part A

Jan Peters · Andrew Y. Ng

Received: 7 May 2009 / Accepted: 19 May 2009 / Published online: 28 May 2009
© Springer Science+Business Media, LLC 2009

1 Introduction

Creating autonomous robots that can assist humans in unpredictable situations of daily life has been a long standing vision of robotics, artificial intelligence, and the cognitive sciences. With the current rise of physical humanoid and other highly mechanically capable robots in robotics research labs around the globe, we have come a step closer to this aim. Thus, it has become essential to create robot systems that learn to accomplish a multitude of different tasks, triggered by environmental context or higher level instruction. Only if machine learning succeeds at making robots fully adaptive, is it likely that we will be able to take real robots out of research labs into human inhabited environments. To do so, future robots will need to make proper use of perceptual stimuli such as vision, proprioceptive and tactile feedback, and translate these into motor commands. In order to close this complex loop from perception to action, machine learning will be needed in various stages such as sensor-based action determination, high-level plan generation and torque-level motor control. Among the important problems hidden in these steps are perceptuo-action coupling, imitation learning, movement decomposition, probabilistic planning, motor primitive learning, reinforcement learning, model learning and motor control.

J. Peters (✉)

Department of Empirical Inference, Max Planck Institute for Biological Cybernetics, Spemannstrasse 38, 72076 Tübingen, Germany
e-mail: mail@jan-peters.net

A.Y. Ng

Department of Computer Science, Stanford University, Room 156, Gates Building 1A, Stanford, CA 94305-9010, USA
e-mail: ang@cs.stanford.edu

Driven by high-profile competitions such as RoboCup and the DARPA LAGR & Learning Locomotion challenges, as well as the growing number of robot learning research programs funded by governments around the world, robot learning has become a central research problem in many labs. While the interest in machine learning and statistics within robotics has increased substantially, robot applications have become important for motivating new algorithms and formalisms in the machine learning community. As a result the interest in robot learning has reached an unprecedented high point as underlined both by the large number of robot learning workshops and the number of published robot learning papers at important robotics conferences such as ICRA, IROS and RSS as well as at major machine learning conferences such as NIPS and ICML.

In this Autonomous Robots Special Issue on Robot Learning, we intend to highlight recent successes in the application of domain-driven machine learning methods to robotics. For obtaining a representative sample of interesting recent results and important current work, papers were solicited with an open call that was circulated in the four months preceding the deadline. As a result we have received 46 submissions to the special issue which were rigorously reviewed by up to eight reviewers as well as by at least one of the guest editors—on average, all papers had 4.5 reviewers. As a result, a small sample of eight papers has been selected for the special issue. The first four papers form Part A of the special issue. They are described next.

2 Guide to Part A of the special issue

In the first paper on *on-line learning and modulation of periodic movements with nonlinear dynamical systems*, Gams et al. report on recent results in imitation learning of rhythmic

movements using motion primitives. They present a two-layered system that is capable of learning nonlinear, oscillatory motion without explicit knowledge of frequency or wave form of the modeled movement. The resulting system can be learned rapidly, is robust towards perturbations and its' implementation has only low computational requirements. The authors are able to show that periodic motion can be modulated with an external signal as the motor primitives are represented by dynamical systems. The authors show results in bimanual drumming with a HOAP-2 humanoid robot as well as several examples of rhythmic movements learned by the same system.

Ratliff et al. present a new approach to imitation learning based on functional gradient techniques that is called *learning to search*. Instead of taking a direct imitation approach as Gams et al. do in the article before, the authors champion the idea of rather taking an inverse optimal control approach that recovers the cost function underlying the demonstration of the teacher and subsequently makes use of planning techniques in order to find good policies that imitate the teacher's intent rather than the teacher's actions. Their work centers around a maximum margin planning approach which has resulted in the efficient and easy to implement algorithm called LEARCH (LEARNING to seaRCH). The presented method has proven to be highly successful in the DARPA Legged Locomotion and DARPA UPI programs. Its applications include autonomous navigation on the NREC Crusher mobile platform, legged locomotion on a BDI Little Dog as well as in robot grasping.

In the paper *reinforcement learning for robot soccer* by Riedmiller et al. the authors reveal the secret of the success of their team BrainStormers at the soccer world championship RoboCup. Using a powerful reinforcement learning algorithm called Neural Fitted Q-Iteration, the authors were able to learn a variety of important behaviors that enabled them to win the world cup both in 2006 and 2007 in the Middle Size League and in 2005, 2007 and 2008 in the simulation league. Among the learned motor programs was an aggressive defensive behavior that provided them with decisive performance gain when facing the team WrightEagle in the final matches of both the world-cup 2007 and 2008. As current three wheel omnidrive RoboCup-style mobile robots are hard to control in the presence of noise due to the uncertainty of the essential ground reaction forces, the authors also discuss how they can learn control policies that perform particularly accurate motor speed control. As a third application, they show how dribbling skills can be learned which significantly outperform any handcrafted dribbling policy. That learned behavior helped them win in 2007 in the Middle Size League.

In the last paper of Part A of the special issue, Ko et al. discuss the use of Bayesian filtering in the context of state estimation for robotics. To achieve this goal, the authors

show in the paper on *GP-BayesFilters* how Bayesian filtering can be achieved using Gaussian process prediction and observation models. In this paper, the authors first discuss how such models are learned and, subsequently, they integrate them with various Bayesian filters, particularly particle filters and extended and unscented Kalman filters. A major advantage of the presented approach is that the uncertainty about the model can be treated in the same framework as the state estimation. The authors underline the functionality of the presented algorithms by results on synthetic data as well as on tracking an autonomous blimp.

All four papers present essential developments in robot learning and we hope that you will enjoy reading them as much as we did.

Acknowledgements The Guest Editors wish to express their gratitude to Gaurav Sukhatme, the AUTONOMOUS ROBOTS Editor-in-Chief, as well as Editorial Assistant Valerie Schofield and JEO assistant Rolan Profugo, for their constant support. We are also enormously grateful to the more than a hundred colleagues who agreed to provide reviews for the 46 papers submitted to this Special Issue.



Jan Peters is a senior research scientist and heads the Robot Learning Lab (RoLL) at the Max Planck Institute for Biological Cybernetics (MPI) in Tuebingen, Germany. He graduated from University of Southern California (USC) with a Ph.D. in Computer Science. He holds two German M.S. degrees in Informatics and in Electrical Engineering (from Hagen University and Munich University of Technology) and two M.S. degrees in Computer Science and Mechanical Engineering from USC. Jan Peters has been a visiting researcher at the Department of Robotics at the German Aerospace Research Center (DLR) in Oberpfaffenhofen, Germany, at Siemens Advanced Engineering (SAE) in Singapore, at the National University of Singapore (NUS), and at the Department of Humanoid Robotics and Computational Neuroscience at the Advanced Telecommunication Research (ATR) Center in Kyoto, Japan. His research interests include robotics, nonlinear control, machine learning, reinforcement learning, and motor skill learning.



Andrew Y. Ng received his B.Sc. from Carnegie Mellon University, his M.Sc. from the Massachusetts Institute of Technology, and his Ph.D. from the University of California, Berkeley. He is an Assistant Professor of Computer Science at Stanford University, and his research interests include machine learning, robotic perception and control, and broad competence AI. His group has won best paper/best student paper awards at ACL, CEAS and 3DRR. He is also a recipient of the Alfred P. Sloan Fellowship.