

Deep Learning and Machine Learning in Robotics

By Fabio Bonsignorio, David Hsu, Matthew Johnson-Roberson, and Jens Kober

Deep learning has gone through massive growth in recent years. In many fields—computer vision, speech recognition, machine translation, game playing, and others—deep learning has brought unprecedented progress and become the method of choice. Will the same happen in robotics and automation?

In a sense, it is already happening. Today, deep learning is often the most common keyword for work presented at major robotics conferences. At the same time, robots, as physical systems, pose unique challenges for deep learning in terms of sample efficiency and safety in real-world robot applications. With robots, data are abundant, but labels are sparse and expensive to acquire. Reinforcement learning in principle does not require data labeling but does require a significant number of iterations on real robots. Transferring the capabilities learned in simulation to real robots and collecting sufficient data for practical robot applications both present major challenges. Further, mistakes by robot learning systems are often much more costly than those by their counterparts in the virtual world. These mistakes may cause irreversible damage to robot hardware or, even worse, loss of human lives. Safety is thus paramount for robot learning systems. A related issue is interpretability, i.e., the ability for humans to understand the learning process. We need shared supervision and control approaches

where humans have the option to integrate with machine learning systems and overcome the systems' cognitive limitations. This requires that humans understand the decision processes of the machine in sufficient detail.

This special issue of *IEEE Robotics and Automation Magazine* focuses on approaches that have been validated on real-world robots, scenarios, and automation problems. It features an exceptionally large number of scientific articles, 11 in total, in keeping with the recent massive growth in the field of robot learning. The articles were carefully selected by the guest editors in two rounds of reviews and provide a good account of the depth and breadth of the research in this area. The methods range from algorithms based on probability theory to deep learning. The whole spectrum of robotics—from manipulators to legged robots to drones to autonomous vehicles—is covered, with applications ranging from rehabilitation to cargo transportation. Additionally, this issue contains two other loosely connected regular articles.

“Movement Primitive Learning and Generalization” by You Zhou et al. considers the representation of movements capable of generalizing according to variations in the environment and task, rather than simply reproducing a demonstration. The focus of this article is the ability to represent behaviors that have multiple modes (i.e., different strategies) in a joint representation. The final evaluation is a throwing task where the robot decides whether to bounce the ball off a wall or not

depending on the target's location. “Gaussians on Riemannian Manifolds” by Sylvain Calinon provides an overview of the use of Riemannian geometry in robotics. The article then illustrates the use of Gaussians on Riemannian manifolds for movement generation in more detail. This is important as a possible way to reduce the computational burden of learning strategies by considering the manifold structure of the data coming from robot motions. Similarly, the following article proposes a structured approach to deal with temporal features, another characteristic aspect of robotic systems. “Interactive Learning of Temporal Features for Control” by Rodrigo Pérez-Dattari et al. explores how robots can be taught new skills interactively by human teachers. The authors focus on tasks where the learning agent needs to be equipped with memory to succeed.

“Multifingered Grasp Planning via Inference in Deep Neural Networks” by Qingkai Lu et al. employs deep learning to predict grasp success, using a voxel representation of the object and the grasp pose as inputs. This model, jointly with a prior over-grasp pose, is then used to infer the grasp with the highest chance of success on unseen, novel objects. “Optimal Deep Learning for Robot Touch” by Nathan Lepora and John Lloyd shows how deep learning methods can be employed to estimate object poses based on input from optical tactile sensors. The authors also demonstrate how the method can then be employed for 3D object exploration.

“Machine Learning for Active Gravity Compensation in Robotics” by Axier Ugartemendia et al. improves the performance of mechatronic systems used in rehabilitation, bringing their capabilities closer to those of human physiotherapists. The article compares various methods for learning to compensate for gravity. “Decoding Motor Skills of Artificial Intelligence and Human Policies” by Kai Yuan et al. offers the opposite perspective. It shows how using machine learning to discover how a humanoid robot’s push recovery strategies can be used to better understand and study human balance control. Those discovered strategies can then be employed for control design.

“Assured Runtime Monitoring and Planning” by Esen Yel et al. addresses one of the biggest open questions in robot learning: safety and verification. This article proposes to speed up online reachability analysis with a pretrained neural network where the neural network itself has been verified. The approach is demonstrated experimentally with quadrotors.

“Multifidelity Reinforcement Learning With Gaussian Processes” by Varun Suryan et al. shows how data collected on less realistic simulations can be used to minimize the amount of data that need to be collected on a real robot. The

approach is validated with a navigation task on a Pioneer robot. “Unsupervised Pedestrian Pose Prediction” by Xiaoxiao Du et al. addresses a problem that is highly relevant for autonomous driving, i.e., not only detecting pedestrians but also predicting their future behavior. In particular, the authors present a method that does not require labeled pedestrian data for training, which has been a major bottleneck in the practical application of such methods. “Deep Learning-Based Localization and Perception Systems” by Zhe Liu et al. presents a system, based on deep learning, that can very accurately localize itself. This is an essential building block toward enabling autonomous navigation, in this case in the Hong Kong International Airport, without requiring modifications to the environment.

As a whole, the 11 articles in this special issue cover a broad range of challenges, ranging from crucial theoretical problems (manifold structure of data, temporal features, unsupervised learning, and multifidelity reinforcement learning) to key transversal needs (assurance and verification of autonomous operations, gravity compensation, and human skill decoding) while covering a wide and diverse set of applications and robotic morphological structures. In addition, the article on the CUDA

parallel computing platform and programming model by Enric Cervera suggests some interesting technical insights that can be useful for the practical implementation of machine/deep learning systems. Further, the article by David St-Onge et al. on planetary exploration provides an example of an application where those approaches could have huge benefits in the future.

We hope that, while providing a rich picture of the state of the art in the application of machine and deep learning in intelligent robotics, this special issue will inspire further research from both the theoretical and application standpoints.

Fabio Bonsignorio is with Heron Robots, Genoa, Italy. Email: fabio.bonsignorio@heronrobots.com.

David Hsu is with the National University of Singapore. Email: dyhsu@comp.nus.edu.sg.

Matthew Johnson-Roberson is with the University of Michigan, Ann Arbor, United States. Email: mattjr@umich.edu.

Jens Kober is with the Delft University of Technology, The Netherlands. Email: j.kober@tudelft.nl.



IEEE connects you to a universe of information!

As the world’s largest professional association dedicated to advancing technological innovation and excellence for the benefit of humanity, the IEEE and its Members inspire a global community through its highly cited publications, conferences, technology standards, and professional and educational activities.

Visit www.ieee.org.



Publications / IEEE Xplore® / Standards / Membership / Conferences / Education

IMAGE LICENSED BY INGRAM PUBLISHING