

主催: 東京工業大学 工学院 システム制御系
共催: 統計数理研究所 統計的機械学習研究センター

A Machine Learning Approach to Optimal Control

日時: 13:30–15:00 2019年11月26日(火)

場所: 東京工業大学 大岡山キャンパス
西8号館E棟10階大会議室

Dr. Marc Deisenroth

DeepMind Chair in Artificial Intelligence,
University College London



Abstract: Optimal control has seen many success stories over the past decades. However, when it comes to autonomous systems in open-ended settings, we require methods that allow for automatic learning from data. Reinforcement learning is a principled mathematical framework for autonomous learning of good control strategies from trial and error. Unfortunately, reinforcement learning suffers from data inefficiency, i.e., the learning system often requires collecting much data before learning anything useful. This extensive data collection is usually not practical when working with mechanical systems, such as robots. In this talk, I will outline two approaches toward data-efficient reinforcement learning, and I will draw connections to the optimal control setting. First, I will detail a model-based reinforcement learning method, which exploits probabilistic models for fast learning. Second, I will discuss a model-predictive control approach with learned models, which allows us to provide some theoretical guarantees. Finally, I will discuss some ideas that allow us to learn good predictive machine learning models that obey the laws of physics. This geometric approach finds physically meaningful representations of high-dimensional time-series data. With this, we can learn long-term predictive models from a few tens of image observations.

Key references:

- Marc P. Deisenroth, Dieter Fox, Carl E. Rasmussen, Gaussian Processes for Data-Efficient Learning in Robotics and Control, IEEE Transactions on Pattern Analysis and Machine Intelligence, volume 37, pp. 408–423, 2015
- Sanket Kamthe, Marc P. Deisenroth, Data-Efficient Reinforcement Learning with Probabilistic Model Predictive Control, Proceedings of the International the Conference on Artificial Intelligence and Statistics (AISTATS), 2018
- Steindór Sæmundsson, Alexander Terenin, Katja Hofmann, Marc P. Deisenroth, Variational Integrator Networks for Physically Meaningful Embeddings, arXiv:1910.09349

主催: 東京工業大学 工学院 システム制御系

共催: 統計数理研究所 統計的機械学習研究センター

Dr. Marc Deisenroth

DeepMind Chair in Artificial Intelligence,
University College London



Bio: Marc Deisenroth is the DeepMind Chair in Artificial Intelligence at University College London. He also holds a visiting faculty position at the University of Johannesburg. From 2014 to 2019, Marc was a faculty member in the Department of Computing, Imperial College London. Since September 2016, Marc has also been an advisor to PROWLER.io, a Cambridge-based startup. Marc's research interests center around data-efficient machine learning, probabilistic modeling and autonomous decision making. Marc was Program Chair of EWRL 2012, Workshops Chair of RSS 2013 and received Best Paper Awards at ICRA 2014 and ICCAS 2016. In 2019, Marc co-organized the Machine Learning Summer School in London with Arthur Gretton. In 2018, Marc has been awarded The President's Award for Outstanding Early Career Researcher at Imperial College. He is a recipient of a Google Faculty Research Award and a Microsoft PhD Grant. In 2018, Marc spent four months at the African Institute for Mathematical Sciences (Rwanda), where he taught a course on Foundations of Machine Learning as part of the African Masters in Machine Intelligence. He is co-author of the book Mathematics for Machine Learning (Cambridge University Press, 2020).