

Learning to Learn and Compositionality with Deep Recurrent Neural Networks

Nando de Freitas

Google DeepMind and University of Oxford

ABSTRACT

Deep neural network representations play an important role in computer vision, speech, computational linguistics, robotics, reinforcement learning and many other data-rich domains. In this talk I will show that learning-to-learn and compositionality are key ingredients for dealing with knowledge transfer so as to solve a wide range of tasks, for dealing with small-data regimes, and for continual learning. I will demonstrate this with several examples from my research team: learning to learn by gradient descent [1], neural programmers and interpreters [3], and learning communication [2].

Keywords

Deep Learning; Recurrent Neural Networks; Compositionality; Transfer Learning; Learning to Learn, Abstraction, Algorithm Induction

Bio

Nando de Freitas is a machine learning professor at Oxford University and a senior staff research scientist at Google DeepMind. He is a fellow of the Canadian Institute For Advanced Research (CIFAR) in the successful Neural Computation and Adaptive Perception program, and an action editor for the Journal of Machine Learning Research.

He received his PhD from Trinity College, Cambridge University in 2000 on Bayesian methods for neural networks. From 1999 to 2001, he was a postdoctoral fellow at UC Berkeley in the artificial intelligence group of Stuart Russell. He became a professor at the University of British Columbia from 2001 to 2014. Nando has spun-off a few technology companies, and received several awards, including a few best paper awards, the Charles A. McDowell Award for Excellence in Research, and the Mathematics of Information Technology and Complex Systems Young Researcher Award.

Bibliography

- [1] M. Andrychowicz, M. Denil, S. Gomez, M. W. Hoffman, D. Pfau, T. Schaul, and N. de Freitas. Learning to learn by gradient descent by gradient descent. ArXiv preprint 1606.04474, 2016.
- [2] J. N. Foerster, Y. M. Assael, N. de Freitas, and S. Whiteson. Learning to communicate with deep multi-agent reinforcement learning. ArXiv preprint 1605.06676, 2016.
- [3] S. E. Reed and N. de Freitas. Neural programmer - interpreters. In International Conference on Learning Representations, 2016.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).

KDD'16, August 13-17, 2016, San Francisco, CA, USA

ACM 978-1-4503-4232-2/16/08.

<http://dx.doi.org/10.1145/2939672.2945358>