

Scaling Up Deep Learning

Yoshua Bengio
University of Montreal
Montreal, Canada
yoshua.bengio@umontreal.ca

Deep learning has rapidly moved from a marginal approach in the machine learning community less than ten years ago to one that has strong industrial impact, in particular for high-dimensional perceptual data such as speech and images, but also natural language. The demand for experts in deep learning is growing very fast (faster than we can graduate PhDs), thereby considerably increasing their market value. Deep learning is based on the idea of learning multiple levels of representation, with higher levels computed as a function of lower levels, and corresponding to more abstract concepts automatically discovered by the learner. Deep learning arose out of research on artificial neural networks and graphical models and the literature on that subject has considerably grown in recent years, culminating in the creation of a dedicated conference (ICLR). The tutorial will introduce some of the basic algorithms, both on the supervised and unsupervised sides, as well as discuss some of the guidelines for successfully using them in practice. Finally, it will introduce current research questions regarding the challenge of scaling up deep learning to much larger models that can successfully extract information from huge datasets.

I.2.6 []: Learning—*Connectionism and neural nets*

deep learning; neural networks; conditional computing; big data

Yoshua Bengio is Full Professor of the Department of Computer Science and Operations Research, head of the Machine Learning Laboratory (LISA), CIFAR Fellow in the Neural Computation and Adaptive Perception program, Canada Research Chair in Statistical Learning Algorithms, and he also holds the NSERC-Ubisoft industrial chair. His main research ambition is to understand principles of learning that yield intelligence. He teaches a graduate course in Machine Learning (IFT6266) and supervises a large group of graduate students and post-docs. His research is widely cited (over 16000 citations found by Google Scholar in early 2014, with an H-index of 55).

Yoshua Bengio is currently action editor for the Journal of Machine Learning Research, editor for Foundations and Trends in Machine Learning, and has been associate editor for the Machine Learning Journal and the IEEE Transactions on Neural Networks.

Yoshua Bengio was Program Chair for NIPS'2008 and General Chair for NIPS'2009 (NIPS is the flagship conference in the areas of learning algorithms and neural computation). Since 1999, he has been co-organizing the Learning Workshop with Yann Le Cun, with whom he has also created the International Conference on Representation Learning (ICLR). He has also organized or co-organized numerous other events, such as the ICML'2012 Representation Learning Workshop, the NIPS'2011 Deep Learning and Unsupervised Feature Learning Workshop, the NIPS'2010 Deep Learning and Unsupervised Feature Learning Workshop, the ICML'2009 Workshop on Learning Feature Hierarchies and the NIPS'2007 Deep Learning Workshop.

Permission to make digital or hard copies of all or part 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. Copyright 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'14, August 24–27, 2014, New York, NY, USA.
ACM 978-1-4503-2956-9/14/08.
<http://dx.doi.org/10.1145/2623330.2630802>