Deep Learning and Cultural Evolution

Yoshua Bengio
Computer Science and Operations Research
University of Montreal, Canada
http://www.iro.umontreal.ca/~bengioy/yoshua_en/index.html

Abstract

We propose a theory and its first experimental tests, relating difficulty of learning in deep architectures to culture and language. The theory is articulated around the following hypotheses: learning in an individual human brain is hampered by the presence of effective local minima, particularly when it comes to learning higher-level abstractions, which are represented by the composition of many levels of representation, i.e., by deep architectures; a human brain can learn such high-level abstractions if guided by the signals produced by other humans, which act as hints for intermediate and higher-level abstractions; language and the recombination and optimization of mental concepts provide an efficient evolutionary recombination operator for this purpose. The theory is grounded in experimental observations of the difficulties of training deep artificial neural networks and an empirical test of the hypothesis regarding the need for guidance of intermediate concepts is demonstrated. This is done through a learning task on which all the tested machine learning algorithms failed, unless provided with hints about intermediate-level abstractions.

Categories and Subject Descriptors
I.2.6 Computing Methodologies, ARTIFICIAL INTELLIGENCE, Learning: Connectionism and neural nets

Keywords
deep learning, cultural evolution

Short Bio

Yoshua Bengio received a PhD in Computer Science from McGill University, Canada in 1991. After two post-doctoral years, one at M.I.T. with Michael Jordan and one at AT&T Bell Laboratories with Yann LeCun and Vladimir Vapnik, he became professor at the Dept of Computer Science and Operations Research at Université de Montréal. He is the author of two books and around 200 publications, the most cited being in the areas of deep learning, recurrent neural networks, probabilistic learning algorithms, natural language processing and manifold learning. He is among the most cited Canadian computer scientists and is or has been associate editor of the top journals in machine learning and neural networks. Since '2000 he holds a Canada Research Chair in Statistical Learning Algorithms, since '2006 an NSERC Industrial Chair, since '2005 his is a Fellow of the Canadian Institute for Advanced Research. He is on the board of the NIPS foundation and has been program chair and general chair for NIPS. He has co-organized the Learning Workshop for 14 years and co-created the new International Conference on Learning Representations. His current interests are centered around a quest for AI through machine learning, and include fundamental questions on deep learning and representation learning, the geometry of generalization in high-dimensional spaces, manifold learning, biologically inspired learning algorithms, and challenging applications of statistical machine learning. In Octobr 2013, Google Scholar finds more than 14500 citations to his work, yielding an h-index of 52.

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.
ACM 978-1-4503-2662-9/14/07.
http://dx.doi.org/10.1145/2598394.25983965