The Next Generation of Neural Networks

Geoffrey Hinton

Google Research & Vector Institute

Toronto

geoffrey.hinton@gmail.com

ABSTRACT

The most important unsolved problem with artificial neural networks is how to do unsupervised learning as effectively as the brain. There are currently two main approaches to unsupervised learning. In the first approach, exemplified by BERT and Variational Autoencoders, a deep neural network is used to reconstruct its input. This is problematic for images because the deepest layers of the network need to encode the fine details of the image. An alternative approach, introduced by Becker and Hinton in 1992, is to train two copies of a deep neural network to produce output vectors that have high mutual information when given two different crops of the same image as their inputs. This approach was designed to allow the representations to be untethered from irrelevant details of the input.

The method of optimizing mutual information used by Becker and Hinton was flawed (for a subtle reason that I will explain) so Pacannaro and Hinton replaced it by a discriminative objective in which one vector representation must select a corresponding vector representation from among many alternatives. With faster hardware, contrastive learning of representations has recently become very popular and is proving to be very effective, but it suffers from a major flaw: To learn pairs of representation vectors that have N bits of mutual information we need to contrast the correct corresponding vector with about 2^N incorrect alternatives. I will describe a novel and effective way of dealing with this limitation. I will also show that this leads to a simple way of implementing perceptual learning in cortex.

CCS CONCEPTS

• **Computing methodologies** → **Unsupervised learning; Neural networks;**

KEYWORDS

deep learning, unsupervised learning, neural networks

ACM Reference Format:

Geoffrey Hinton. 2020. The Next Generation of Neural Networks. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '20)*, June 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3397271.3402425>

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(s).
SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8016-4/20/07.

<https://doi.org/10.1145/3397271.3402425>

BIOGRAPHY

Geoffrey Hinton received his PhD in Artificial Intelligence from Edinburgh in 1978. After five years as a faculty member at Carnegie-Mellon he became a fellow of the Canadian Institute for Advanced Research and moved to the Department of Computer Science at the University of Toronto where he is now an Emeritus Distinguished Professor. He is also a Vice President & Engineering Fellow at Google and Chief Scientific Adviser of the Vector Institute.

He was one of the researchers who introduced the backpropagation algorithm and the first to use backpropagation for learning word embeddings. His other contributions to neural network research include Boltzmann machines, distributed representations, time-delay neural nets, mixtures of experts, variational learning and deep learning. His research group in Toronto made major breakthroughs in deep learning that revolutionized speech recognition and object classification.

REFERENCES

- [1] Suzanna Becker and Geoffrey E. Hinton. 1992. Self-organizing neural network that discovers surfaces in random-dot stereograms. *Nature* 355, 6356 (1992), 161–163. <https://doi.org/10.1038/355161a0>
- [2] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A Simple Framework for Contrastive Learning of Visual Representations. (2020). [arXiv:cs.LG/2002.05709](https://arxiv.org/abs/2002.05709)
- [3] Alberto Paccanaro and Geoffrey E. Hinton. 2001. Learning Distributed Representations of Concepts Using Linear Relational Embedding. *IEEE Trans. on Knowl. and Data Eng.* 13, 2 (March 2001), 232–244. <https://doi.org/10.1109/69.917563>