

Building Machine Learning Systems that Understand

Jeff Dean
Google Research
jeff@google.com

ABSTRACT

Over the past five years, deep learning and large-scale neural networks have made significant advances in speech recognition, computer vision, language understanding and translation, robotics, and many other fields. Deep learning allows the use of very raw forms of data in order to build higher-level understanding of data automatically, and can also be used to learn to accomplish complex tasks. In the next decade, it is likely that a fruitful direction for research in data management will be in how to seamlessly integrate these kinds of machine learning models into systems that store and manage data. In this talk, I will highlight some of the advances that have been made in deep learning and suggest some interesting directions for future research.

CCS Concepts

• Theory of computation~Online algorithms • Theory of computation~Online learning theory

Short Bio

Jeff joined Google in 1999 and is currently a Google Senior Fellow in Google Research, where he leads the Google Brain team, Google's deep learning research team, working on machine learning systems and their use for speech recognition, computer vision, language understanding, and various predictive tasks. He has co-designed/implemented five generations of Google's crawling, indexing, and query serving systems, and co-designed/implemented major pieces of Google's initial advertising and AdSense for Content systems. He is also a co-designer and co-implementor of Google's distributed computing and storage infrastructure, including the MapReduce, BigTable, Spanner and TensorFlow systems, protocol buffers, LevelDB, and a variety of internal and external libraries and developer tools.



He received a Ph.D. in Computer Science from the University of Washington in 1996, working with Craig Chambers on compiler techniques for object-oriented languages. He is a Fellow of the ACM, a member of the U.S. National Academy of Engineering and the American Academy of Arts and Sciences, and a recipient of the Mark Weiser Award and the ACM-Infosys Foundation Award in the Computing Sciences.

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). Copyright is held by the owner/author(s).

SIGMOD'16, June 26–July 1, 2016, San Francisco, California, USA.

ACM 978-1-4503-3531-7/16/06.

<http://dx.doi.org/10.1145/2882903.2932259>