

Introduction to the Special Theme

Modern Machine Learning: More with Less, Cheaper and Better

by Sander Bohte and Hung Son Nguyen

While the discipline of machine learning is often conflated with the general field of AI, machine learning specifically is concerned with the question of how to program computers to automatically recognise complex patterns and make intelligent decisions based on data. This includes such diverse approaches as probability theory, logic, combinatorial optimisation, search, statistics, reinforcement learning and control theory. In this day and age with an abundance of sensors and computers, applications are ubiquitous, ranging from vision to language processing, forecasting, pattern recognition, games, data mining, expert systems and robotics.

Historically, rule-based programs like the Arthur Samuel checkers-playing program were developed alongside efforts to understand the computational principles underlying human learning, in the developing field of neural networks. In the '90s, statistical AI emerged as a third approach to machine learning, formulating machine learning problems in terms of probability measures. Since then, the emphasis has vacillated between statistical and probabilistic learning and progressively more competitive neural network approaches.

The breakthrough work by Krizhevsky, Sutskever & Hinton [1] on deep neural networks in 2012 has been a catalyst for AI research by demonstrating a step function in performance on the Imagenet computer vision competition. For this, they used a deep neural network trained exhaustively on 'GPUs': a garden-variety parallel computing hardware used for video-games. Similar advances were then quickly reported for speech recognition and later for machine translation and natural language processing. In short order, big companies like Google, Microsoft and Baidu established large machine learning groups, quickly followed by essentially all other big tech companies. Since then, with the combination of big data and big computers, rapid advances have been reported, including the use of machine learning for self-driving cars, and consumer-grade real-time speech-to-speech translation. Human performance has even been exceeded in some specialised domains. It is probably safe to say that at present, machine learning allows for many more applications than there are engineers capable of implementing them.

These rapid advances have also reached the general public, with often alarming implications: think tanks are declaring that up to 70% of all presently existing jobs will disappear in the near future, and serious attention is being given to potentially apocalyptic futures where AI capabilities exceed human intelligence. We believe, however, that it is safe to say that this will not happen in the next five years, as machine learning still faces some serious obstacles before reaching human levels of flexible intelligence.

Some of the current challenges in machine learning are reflected in the articles presented in this special issue: the much glorified deep learning approaches all rely on the availability of massive amounts of data, often needing millions of correctly

labelled examples. Many domains, however, including some important areas such as health care, will never have such massive labelled datasets. Similarly, robots cannot be trained for millions of trials, simply because they wear out long before. The question is thus how to learn more with less. Here, statistics and prior knowledge will likely play a big role, and some promising work is presented in this issue – see for example the articles by Mouret and by Welling. Some work is also examining whether quantum computing can help reduce the computational complexity of machine learning, as explained in the article by Wittek. At the same time, massive data streams generate problems of their own: Cieliebak and Benczur each present work on how to deal with huge torrents of streaming data.

Apart from these technical challenges, we, as a community, need to train the future experts in machine learning, such that the wider industry and society can benefit from what is currently already possible. Some of the exciting applications are laid out in this issue: Kappen for example showcases the Bonaparte Disaster Victim Identification system, which uses Bayesian statistical modelling to identify victims based on their DNA and that of next of kin. Potamias presents a wonderful application of machine learning in ‘Pharmagenomics’, where the challenge is to determine which genes interact with which drugs. This is a key determining factor in the efficacy of these drugs and central to the future of personalised medicine.

A separate line of ongoing research is the link between the one working example of intelligence, the brain, and learning principles. This relates to such diverse questions as ‘How can goals be selected in an autonomous fashion?’ and ‘How can we optimise over many different learning problems with one system?’, but also in reverse: ‘What can the success of deep neural networks tell us about the brain?’. Articles by Alexandre and Oudeyer cover current efforts on these topics.

The study of machine learning has thus grown from the efforts of a handful of computer engineers exploring whether computers could learn to play games and mimic the human brain, and a field of statistics that largely ignored computational considerations, to a booming discipline that is actively transforming the world in which we live.

Reference:

[1] Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." *Advances in neural information processing systems (NIPS-2012)*. Pages 1097-1105, 2012.

Please contact:

Sander Bohte, CWI, The Netherlands
sander.bohte@cwi.nl

Hung Son Nguyen, University of Warsaw, Poland
son@mimuw.edu.pl