Harvard Data Science Review • 1.1

Challenges for Human-Level Intelligence

Rodney Brooks

Published on: Jun 23, 2019

Updated on: Dec 10, 2019

DOI: [10.1162/99608f92.7ef56abe](https://doi.org/10.1162/99608f92.7ef56abe)

This piece is a commentary on the article: [“Artificial Intelligence—The Revolution Hasn’t Happened Yet.”](#)

Data science and machine learning are important to Artificial Intelligence (AI) and will likely be at the nexus of advancement of AI for the next several years. As Michael Jordan says, there is a general confusion where common parlance treats data science and machine learning as all of AI.

When he coined the term ‘artificial intelligence’ in a 1955 funding proposal, John McCarthy based the concept on “the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it.” Just a few sentences later he adds, “If a machine can do a job, then an automatic calculator can be programmed to simulate the machine” (1955).

From its very beginnings, AI appeared as both a computational endeavor and as an attempt to mimic all of human intelligence.

The capabilities that Jordan describes, which data science and machine learning have recently delivered when trained on very large curated data sets, were largely unexpected just a few years ago. They have been good enough to be at the core of applications that deliver new technological capabilities that we now use on a daily basis.

However, they leave out a lot of what will be necessary in order to have human-level intelligence that can operate in a general way. Human-level intelligence (and, in some cases, animal-level intelligence) has drives and goals that are maintained over long periods of time. It has episodic memory and can change perceptions based on the introduction of a single new, symbolically-expressed fact. Human-level intelligence has models of the short-term future based on physics and of the longer-term future based on both regularities in the world and psychological models of other actors in the world. Human-level intelligence can reason about categories and hypotheticals involving those categories and can reinterpret long ago events based on new facts. Machine learning has not tackled any of these examples in a meaningful way.

Let us return to what data science and machine learning can currently do.

When data science or machine learning duplicate something that humans can readily do (e.g., reliably extract a stream of phonemes from a human speaking, recognize human faces), the energy requirements for the learning is usually far greater than that for an individual human to learn the same task over a period of years. Of course, the machine learning system produces a result, which has a very low marginal cost in energy to have it installed in thousands or millions of end products. This tells us that the way our data science and machine learning systems currently work are not the only possible ways.

More than this, we also know how humans are able to infer complex categories from just a handful of exposures to instances of the category. We use other background information and knowledge to infer the boundaries of categories from very little information—generalization capabilities that data science and machine learning have not yet explored. Even for reinforcement learning and game playing, humans necessarily work very differently, as they see only a tiny fraction of the board instances that the recent heroes of machine game playing have seen.

Just as there are lots of aspects of AI that current data science and machine learning cannot tackle, there are ways in which the current machine learning techniques cannot match human performance. There is room in data science and machine learning for much more science that will extend the collection of algorithms at our disposal.

Discussion

Further commentary by:

- [Emmanuel Candes, John Duchi, and Chiara Sabatti](#) (Stanford University)
- [Greg Crane](#) (Tufts University)
- [David Donoho](#) (Stanford University)
- [Maria Fasli](#) (UNESCO)
- [Barbara Grosz](#) (Harvard University)
- [Andrew Lo](#) (MIT)
- [Maja Mataric](#) (USC)
- [Brendan McCord](#) (Tulco Labs)
- [Max Welling](#) (University of Amsterdam)
- [Rebecca Willett](#) (University of Chicago)

Rejoinder by: [Michael I. Jordan](#) (UC Berkeley)

References

McCarthy, J., Minsky, M., Rochester, N., and Shannon, C. (2006). A proposal for the Dartmouth Summer Research Project in Artificial Intelligence. *AI Magazine*, Vol 27, Number 4, pp12-14.
<https://doi.org/10.1609/aimag.v27i4.1904>

This article is © 2019 by Rodney Brooks. The article is licensed under a Creative Commons Attribution (CC BY 4.0) International license (<https://creativecommons.org/licenses/by/4.0/legalcode>), except where otherwise indicated with respect to particular material included in the article. The article should be attributed to the author identified above.