Harvard Data Science Review • 1.1

Artificial Intelligence versus Intelligence Engineering

Max Welling

Published on: Jun 22, 2019

Updated on: Dec 10, 2019

DOI: [10.1162/99608f92.364ce476](https://doi.org/10.1162/99608f92.364ce476)

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

I congratulate Prof. Jordan with his astute commentary on contemporary trends in the research field known as AI. Under the field’s current hype, AI has become a container term confusing experts and laymen alike. Prof. Jordan argues that this state of affairs is not helping the field focus on the truly important tasks that lay ahead of us, namely to turn the new AI technology into safe, privacy preserving, practical, human-centric tools that benefit society. He proposes to disentangle the field into (at least) three subfields: Human-Imitative AI, Intelligence Augmentation, and Intelligent Infrastructure, and argues that the latter two are at least as important as the first one despite its receiving all the attention in the media.

I tend to agree with Prof. Jordan that the term AI is overloaded, the field is overhyped, and the media could be more objective about the way they report. It’s perhaps interesting to note that these attention biases are not unique to AI (please excuse me for using the term myself as a container). In physics, a disproportionate fraction of the brightest minds works in string theory because it speaks to the imagination to do research on the ‘theory of everything.’ For similar reasons I believe many young researchers choose to work on Human-Imitative AI rather than Intelligence Augmentation or Intelligent Infrastructure: building human-like superintelligence has a certain appeal. This is also the reason why the media overemphasizes this type of AI despite the fact that it may not actually be the largest contributor to useful and reliable AI.

Prof. Jordan points out that the field is nowhere as far as we are led to believe from the media. I agree here as well. Talking about a *Terminator* scenario or a singularity is, in my opinion, like worrying about dementia at the age of a young child. Deep learning, the most important technology that has fueled recent progress in the field, is at this point not much more than a (very powerful) glorified signal processing tool for pattern recognition. An actual deep understanding of the world, including its causal and physical structures, is still lacking. Current AI techniques are not yet capable of high-level reasoning. At the same time, progress is sometimes surprisingly fast, and I continue to be amazed at systems like Deepmind’s AlphaZero or Google Duplex that seem to be making rather serious progress in that direction.

There is one key factor that sets the current developments apart from what we have witnessed thus far, namely that it is fueled by a huge interest from industry. I believe that the emergence of AI labs in industry, the rapid transfer of academic ideas to industry, and the massive level of investments in AI by industry are the true revolution that we are witnessing. Technologies such as speech recognition, machine translation, and face recognition make it into products and services on a very short timescale. Academic papers end up on a company desk the day after publication and may influence products and services the same year, which will in turn lead to new investments. It’s that positive feedback loop that

will keep propelling the field forward. Even though investment decisions are also subject to hype and “fear of missing out,” I still believe this level of investment is a clear sign that the current AI technology is making a real impact in society, and will continue to do so in the foreseeable future. In this sense, the hype is more rooted in reality than one might imagine.

The fact that industry is accelerating the AI flywheel is also the reason that I am less concerned about the perceived lack of acknowledgement for the importance and urgency of topics such as Intelligence Augmentation and Intelligent Infrastructure. The amount of engineering spent behind the scenes to make AI systems work in practice is likely larger than the time spent on algorithmic development. Companies understand the significance of ‘Intelligence Engineering’ very well, as they are faced with it every day. In fact, these topics remain heavily researched in academic circles also, though they are not as much in the media’s spotlight as Human-Imitative AI.

There are also great reasons to be interested in Human-Imitative AI. Current deep learning algorithms consume too much energy to be profitable for many applications, such as ranking recommendations in a web shop or ranking posts on a social network. Moreover, in small (edge) devices, such as smartphones, mobile health devices and so on, we face the challenge that these devices cannot dissipate the heat created by the power-hungry deep learning models without getting too hot. Researchers will be looking again to natural intelligence for inspiration. We will ask how the brain manages to squeeze a factor 100,000 more intelligence out of the same power budget. These engineering questions are mainly driven by companies pursuing profits rather than researchers pursuing the coolest research topics. Market forces will outweigh sentiments amplified by the media about the importance or dangers of imitating human intelligence.

Despite biased and fearmongering reporting by the media, I continue to marvel at the awe-inspiring results that a decade of deep learning has brought us, and will likely continue to bring us in the future. I marvel at deep learning-based algorithms that beat human doctors at certain, admittedly narrowly-defined, medical diagnosis tasks in pathology, dermatology, retinopathy, etc. I continue to look in awe how systems such as AlphaZero can train themselves in a matter of hours to beat the best human beings in complex games such as Chess or Go, or how virtual assistants like Google Duplex can reserve a table without the human noticing it was talking to a chatbot. All of these applications are enabled by deep learning, combined with a whole lot of complex engineering in the background.

In conclusion, there are many exciting and important endeavors in the broad spectrum of activities that we call AI. I subscribe to the importance of human-centric AI, intelligent infrastructure and engineering. I agree that we need a broad palette of disciplines to work together to develop complex AI systems that are safe, useful, private, fair, transparent and so on. And as anyone who has developed a real system will tell you, the engineering efforts behind the scenes often represent the bulk of the

hard work. Thus, a new field of “Intelligence Engineering” may indeed be necessary and may in fact already be under construction behind the scenes.

Discussion

Further commentary by:

- [Rodney Brooks](#) (MIT)
- [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)
- [Rebecca Willett](#) (University of Chicago)

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

This article is © 2019 by Max Welling. 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.