Commentary on Artificial Intelligence – the Revolution Hasn't Happened Yet by Michael J. Jordan, University of California, Berkeley

Maria Fasli

Published on: Jun 22, 2019

Updated on: Dec 10, 2019

DOI: [10.1162/99608f92.0c23c330](https://doi.org/10.1162/99608f92.0c23c330)

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

Professor Jordan’s article brings to the forefront important issues in what is commonly referred to as “Artificial Intelligence” (AI) nowadays, and is both instructive and thought-provoking. AI is indeed the mantra of our times. Until a few years ago, it was almost considered an anathema to say that you work in AI, but recent proliferation of data as well as the exponential growth of computational power have fuelled a resurgence in the field. Everyone is doing AI or desperately wants to jump on the AI bandwagon. Capitalizing on the newfound fame and popularity of AI, research foundations, including governments, around the world put millions of dollars in funding and initiatives for further progress to be made in the field.

True, there have been some advances in techniques, but not so much so as one would have thought. Neural networks are not new—they have been around for a long time—but it is only now that we have the “machinery” to make them work, both in terms of computational power and the enormous amounts of data to train them. Yet we have not broken down any significant barriers to creating AI as the original proponents of the term imagined.

We should not underestimate the confusion among both scientists/academics and non-academics alike, which could be severely detrimental to the whole endeavour of the field. The distinction that Jordan draws between Intelligence Augmentation (IA) and Intelligent Infrastructure (II) is a useful one and illustrates the point of how much of what is rendered or presented as AI today can be labelled differently.

What Jordan calls “human-imitative AI” or “human-like AI”—recreating our own intelligence in a machine, but perhaps an improved version of us, devoid of the computational boundaries and the limitations of how much data we can process to reach decisions—has served as the driving force behind AI. This drive to create artifacts and even mechanistic versions of ourselves or animals is not new: in Greek mythology for instance, Hephaestus would create automata in his workshop.

Is the elusive goal of generating human-like AI distracting us? Yes and no. On one hand, this insatiable quest to create human-like AI led to the development of the field, the realization that ‘intelligence’ itself is far more complicated than originally thought—and not just the matter of being able to do fast calculations. AI thrives across diverse fields in which researchers drive progress. Had it not been for this, AI may not have developed in the same way that it has (for better or worse). It is a distraction, on the other hand, to imagine an AI that will be more like us and presume that AI systems in realising their superiority in comparison to humans and in order to survive they will wipe out everything in their way—much in the same way that we are currently doing to other species around the world and even to our planet itself.

Despite our flaws, we are good decision-makers, we deal with uncertainty well, and we are, astonishingly, able to draw generalizations and apply learnings from one problem to a completely different domain. This is where we have yet to make any progress in AI systems—they can learn, and improve, and you may be able to apply the same system in a similar domain, but it is nearly impossible to apply this to a completely different domain. This high-level reasoning extends beyond existing techniques, irrespective of how much data we have available. This has nothing to do with data; humans are able to perform this kind of reasoning with very few examples indeed. Human-level intelligence also entails social interaction which we have yet to “engineer” or adequately build in current AI systems.

Interestingly, like Jordan, John McCarthy defined AI as “the science and engineering of making intelligent machines.” Jordan advocates for a new engineering discipline to develop AI systems. To this extent, he posits that, similarly to civil engineering and chemical engineering, which did not focus on the creation of an artificial bricklayer or chemist, we should move the emphasis away from the creation of human-like AI as it distracts from the challenges that IA and II pose. One could suggest, though, that an artificial bricklayer or chemist still falls under the realm of AI versus civil or chemical engineering.

Another interesting parallel exists between the quest for the development of human-like AI, and the much-needed engineering discipline behind AI systems and chemical engineering. Chemistry and chemical engineering have, to a certain extent, roots in alchemy: the ultimate goal was to create gold out of base metals or to find a universal elixir. This parallels our quest for the elusive human-like AI. Alchemy ultimately failed, but in the process, alchemists developed basic laboratory techniques, experimental methods, theories, and techniques that led the foundations for modern science. Alchemy also had quasi-mystical elements, similarly to what some ascribe to AI, albeit laced with more jargon. Could the quest for human-like AI be more like alchemy? Perhaps not delivering the “gold” or transmutation that we are after, but helping us lay the foundations for the new engineering discipline required along the way.

The new engineering discipline should be human-centric, and principles from ethics, human rights, fairness, transparency, and accountability should be among the foundational pillars. The current level of hype is damaging and may have long-term, detrimental consequences. Unless we help dispel confusion and manage expectations on the uses and abuses of AI, we risk plunging into another AI winter.

Discussion

Further commentary by:

- [Rodney Brooks](#) (MIT)
- [Emmanuel Candes, John Duchi, and Chiara Sabatti](#) (Stanford University)
- [Greg Crane](#) (Tufts University)
- [David Donoho](#) (Stanford University)
- [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)

This article is © 2019 by Maria Fasli. 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.