Response to “Artificial Intelligence—The Revolution Hasn’t Happened Yet”

Gregory Crane

Published on: Jun 23, 2019

Updated on: Dec 10, 2019

DOI: [10.1162/99608f92.e32f6dec](https://doi.org/10.1162/99608f92.e32f6dec)

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

I face two challenges in responding to Michael Jordan’s views on Artificial Intelligence and the still-unfulfilled revolution. First, I agree almost entirely with his analysis, and so my comments will build upon, rather than provoke controversy about, his piece. AI as a process to imitate human consciousness appeals to the popular imagination, but long-evolving methods under the broad rubric of machine learning provide the substance behind the hype. In particular, I find appealing Jordan’s emphasis on two different, two-letter acronyms: IA (intelligence augmentation) and II (intelligent infrastructure). Both IA and II are fundamental to the humanities, not simply insofar as specialists teach students in formal academic settings and produce publications for each other (although IA and II are relevant to both), but to the way human beings are able to understand the cultural record and each other.

Second, my task is to respond as a humanist, but the humanities are so heterogeneous that I cannot imagine providing even a broad—much less comprehensive—summary of how AI (or, more properly, a range of technologies under the rubric of machine learning) may affect the humanities as a whole. When two of my students in the Bachelor of Science program in Digital Humanities at Leipzig developed a report on AI, they focused first on the relationship between AI and human consciousness, and then upon the expanding range of ethical questions raised by autonomous cars, humanlike personal assistants and coaches, and the potential of independent weapon systems choosing whom they would kill. Such perspectives are essential, but pose challenges if we wish to speak of AI and humanities; they are both too broad (they do not connect to much of the work done in the humanist) and too narrow (as they connect to specialized work in ethics and philosophy). For now, I doubt that I could even formulate a description of the humanist that would not provoke controversy. Ultimately, everyone—whether they think of themselves as a humanist or not—is living a kind of Platonic dialogue, where we struggle to develop some form of the good life.

If I were to summarize the humanities in its broadest sense and as it most fully bears upon our understanding of technology, I might go back to a phrase attributed to the fifth century Greek thinker Protagoras and that can be fairly paraphrased: human beings are the measure by which we evaluate everything.¹ We assess the past, present, and potential of anything—including technological innovations—by the impact that something has upon human beings. In this broadest and fullest of senses, the humanities begin whenever we move from the question of how to that of why. We are all, in this sense, humanists in at least some measure. Otherwise, we fall prey to the charge Oscar Wilde levelled on cynics and which many level on economists: we may know the price (or some other numerical metric) of everything, but we know the value of nothing.

Hype and AI, then and now

Jordan’s insights encouraged me to reflect on my own perceptions of AI and its potential implications, as well as what computation might mean for the study of the past. I began reflecting on the potential of Artificial Intelligence almost as soon as I began to consider what computation might mean for the study of the past. If the AI revolution has not taken place in 2019, we should remember that an AI revolution was expected a generation ago. While various applications from machine learning are already having serious and disruptive impacts, my own experience makes me more skeptical about the implications of a more general Artificial Intelligence. Reading the fifth-century Greek author Herodotus, the first person to produce a book-length prose history (his account of Persian/Greek relations), would be useful for a variety of reasons. In particular, Herodotus emphasizes that, over long periods of time, entities rise and fall in importance. Even the most powerful at any given time vanish in the end.

In 1982, when I first began exploring the implications of computation for the study of the ancient world in particular and of the humanities in general, AI was already in the news and destined to transform society. I was a graduate student at Harvard and, across town, MIT hosted a revolution of advancements in AI research, including a cluster of startup firms seeking to commercialize AI. Researchers in Cambridge nicknamed the area around Kendall Square “AI Alley.”² Over the next several years, I would make a number of pilgrimages to talk with potential collaborators—some friendly and curious, but none able to dedicate time to research in the humanities with an unfunded Ph.D. student. On one memorable visit in 1984, two graduate students in political science, who had gotten access to high end equipment, told me that I could easily collaborate—I just had to install some code on my Lisp machine. The affable—and far better funded—MIT graduate students were surprised that I did not have one and informed me that \$100,000 would get me a fine system. In 1985, Xerox established an equipment grant program, and my first successful application as an assistant professor elect was for a room full of Lisp machines. My plan was to use NoteCards, a pioneering hypertext system implemented by Randall Trigg and Frank Halasz in Lisp, on Xerox D workstations. NoteCards wasn’t AI, of course, but Lisp was developed in large measure to support research in AI.

The Xerox equipment grant landed me space in a building next to the Harvard Divinity School, reportedly the site of a World War II temporary structure with hardened points for anti-aircraft guns on the roof, and the surviving Harvard twin to MIT’s now-demolished Building 20, famous as a skunkworks. The Lisp machines also served as excellent heaters—a crucial advantage when winter closed in on our poorly insulated and unevenly heated lab—but we ended up shifting our work to mainline Unix systems. When I developed linguistic analysis software for ancient Greek, I used Lisp via either a DEC Vax minicomputer or a Sun workstation, two genres of computer and two once-mighty companies that are now extinct.

Organizations may vanish forever and with almost no trace, but ideas may return from weakness or even oblivion. AI is back, but AI is both more and less what it was understood to be a generation ago. A generation ago, AI often focused on modelling human decision processes and even consciousness. Marvin Minsky once suggested that the time would come when no one would believe that the books in the library did not talk to each other (Stefik, 1997). He was imagining, if I understand correctly, that the books would embody some form of intelligence and interact in an ongoing fashion, growing ever more intelligent. That may or may not happen at some point. My biggest doubt about whether the Minsky quote fully comes true lies elsewhere: I wonder how long our students will understand the references to books and to libraries.

The years after 1985 were not good for AI. While the results of Google’s Ngram Viewer should always be used cautiously, the trend for AI is clear and negative.

Figure 1. Google Ngram results for “AI” and “Artificial Intelligence” from 1800 through 2000. <https://books.google.com/ngrams>

A decade later when, as a professor at Tufts, I sat in on our AI class, there were as many students as I had in my first-year class of Ancient Greek (six, if I recall correctly, although whatever the precise number, it was identical). Jim Schmolze, whom we lost far too young to cancer, commented morosely that every time AI produced something useful, mainline computer science would appropriate the method and leave AI to intractable problems and oversold promises. Google Ngram needs to be used with caution, but it certainly provides quick and striking evidence that the AI bubble a generation ago was swift and disastrous.

When companies go to zero or vanish after being acquired, they are generally gone forever. Sun Microsystems, founded in 1982, was arguably the most innovative and successful of the Unix workstation companies that sprang up in the 1980s, but was also one of the companies that suffered most from overvaluations of the tech bubble. Not only did its stock price crash, but its business evaporated. Oracle acquired Sun Microsystems in 2010 and Sun exists as a web page and corporate trophy on the Oracle website.³

The fall of Digital Equipment Corporation (or DEC as it was commonly known) was arguably more dramatic and, to me at least, more chilling. My first real exposure to computing came in 1982 at the Computer Based Lab (CBL) in the Harvard Psychology Department. CBL was a DEC shop. I learned to program in C with the Kernighan and Ritchie “white book” using a PDB 11/44; I even got some exposure to machine language via an old PDP 4. Washing machine-sized hard disks with 80 mbytes each spun over the raised floor in the machine room. A few hundred meters away, computer science had Vax 750 and 780 mini-computers, but with no networking aside from phone lines, these were fabulous instruments that I could occasionally see, but from the standpoint of computation, might as well have been on the other side of the planet. When work began in 1987 on Perseus⁴, our digital library of Greek, Latin, and other literatures, DEC employed over 140,000 people worldwide. It was the second-biggest computer company in the world at the time after IBM, but where the tech wreck sent Sun Microsystems into a death spiral, DEC did not even survive as a separate entity long enough to experience the Internet bubble. Compaq acquired DEC in 1998. While in Maynard a few years ago, after hunting at some length, I found a single historical plaque that recalled the existence of DEC, remains less imposing than the shattered visage of Shelley’s *Ozymandias*.

Voice recognition in flagship products such as Amazon’s Alexa and Apple’s Siri bring the popular imagination back to the talking computers of the Enterprise and 2001 and, in this respect, they embody the image of disembodied consciousness of at least some serious AI research in the 1980s. However, Alexa and Siri build on automatic speech recognition and on question-answering, a branch of information retrieval focused on converting natural language questions into structured queries. In the early 1970s, William Woods⁵ developed a question-answering system to manage natural language queries about moon rocks brought back by the Apollo missions, and the Text Retrieval Conference sponsored by the National Institute of Standards (NIST) first offered a question-answering contest track in 1999 (Woods, 1973).⁶ Both speech recognition and question-answering systems address challenging technical problems, but neither seeks to model human consciousness. Insofar as each reflect artificial intelligence, they exploit techniques from machine learning and thus seek to exploit statistical analysis of data to generate decisions in well-specified domains.

Technological Evolution and Humanist Thinking

This long digression on the vicissitudes of AI as a label (if not as a well-defined technical term) emphasizes one intellectual feature that can distinguish humanistic thinking: the emphasis on framing things in a longer-term perspective. This perspective entails the lived recognition—a tangible sense that shapes our daily perceptions—that we comprise tenuous links in a chain that extends backwards into the past and outwards to the end of intelligent life on this planet.

Already in 1982 I had the opportunity to embrace the digital future. The hardware and software tools were certainly primitive and most of our effort focused on low-level tasks such as automated print typesetting and replacing print indices with computer searches. But I was working with scholarship that had been composed more than a century before and the study of Greek literature as a whole extended thousands of years into the past. It made no difference to me whether the shift to a digital world took a decade, a generation, or a century. The shift was inevitable. It would be irreversible and it would fundamentally recast how we lived our lives and viewed the world. However, the greater the potential change, the greater the potential opportunity for us to realize and revitalize our larger goals as students of the past. When the revolution would take place was less important than that the revolution was inevitable. The more I felt anchored in a tradition that extended centuries and millennia into the past, the more essential I felt that I should embrace long-term change.

Whenever—or even whether—the AI revolution takes place, the two separate topics that Michael Jordan cites will remain fundamental to humanists: intelligence augmentation (IA) and intelligent infrastructure (II). With II, Jordan points out that “a web of computation, data and physical entities exists that makes human environments more supportive, interesting and safe.” Within this category I would include the ever larger and more heterogeneous digital libraries with which we can develop questions about human history and culture, present as well as past. Such a digital library certainly extends beyond enhanced metadata and, in its ideal form, incorporates the full record of humanity: not just traditional texts, but every video on YouTube, and not just the linguistic production of humanity, including musical and mathematical idioms, but the physical traces of humanity’s impact on earth in the present and the past.

For the humanist, intelligence augmentation must now and forever be our goal. Machine learning only matters insofar as it makes us fundamentally more intelligent and deepens our understanding. Here again, the contrast between prices (and quantifiable metrics of all kinds) and value is central. Intelligence balances and embraces both the how and why. Such a statement may seem obvious, but for me as a professional humanist, the consequences are profound. Insofar as Intelligent Infrastructure and Intelligence Augmentation have any non-trivial meaning, they combine to disrupt

my audience as a humanist: I can go beyond mute translations for texts or subtitles for film and build an intelligent reading environment that invites readers and audiences. Such new modes of reading may initially appear for smaller, well-studied and stable areas such as Classical Greek and Classical Chinese, but the same principles will support audiences watching video series produced in Hindi, Turkish, or Colombian Spanish on Netflix. Whether we are reflecting on sources thousands of years old or delving into the background of a new video on Netflix or YouTube, my goal as a humanist is to challenge myself and others to push as deeply as we wish into the language and culture that they are viewing. My own goal is to develop increasingly intelligent reading environments that provide as much immediate and personalized understanding as possible and that enable audiences to develop, for any language and culture, an empathetic understanding that goes as deeply as their determination and their abilities can take them.

Discussion

Further commentary by:

- [Rodney Brooks](#) (MIT)
- [Emmanuel Candes, John Duchi, and Chiara Sabatti](#) (Stanford 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

- Stefik, M. (1997). *Internet Dreams: Archetypes, Myths, and Metaphors*. Cambridge, MA: MIT Press.
- Woods, W. A. (1973). Progress in natural language understanding. *Proceedings of the June 4-8, 1973, National Computer Conference and Exposition on AFIPS '73*, 441. doi:10.1145/1499586.1499695
-

This article is © 2019 by Gregory Crane. 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.

Footnotes

1. The phrase is familiar in English as “man is the measure of all things,” but *anthrôpos*, the Greek word translated as “man”, can be applied to men and women and can designate a human being of either (or any) gender. Plato alludes to this saying in the Theaetetus (152a), while Sextus Empiricus provides a direct quotation (Adv. Math. 7.60). [↵](#)
2. Echoes of this now largely-forgotten nickname can (as of December 2018) be found at <http://news.mit.edu/2011/kendall-square-part1-1122> and <https://www.linkedin.com/pulse/reviving-bostons-ai-alley-peter-polachi>. [↵](#)
3. Oracle does maintain a web page on “Oracle and Sun Microsystems”: <https://www.oracle.com/sun/>; this web page does contain link to “Sun Servers” and “Sun Storage and Tape” but these links resolve to a page touting Oracle Cloud servers, <https://www.oracle.com/servers/>, and Engineered Storage. Neither page mentions Sun Microsystems. [↵](#)
4. <http://www.perseus.tufts.edu/hopper/> [↵](#)
5. I was able to gain my first exposure to natural language processing (NLP) in a classroom setting when, in spring 1985, Harvard hired Bill Woods to teach NLP as a professor of the practice. He was the one who told me about the Xerox equipment grant program and made it possible to begin my own career as an assistant professor of Classics in fall 1985 with a full computer lab of my own. [↵](#)
6. <https://trec.nist.gov/data/qa.html>. [↵](#)