Editor's Column

AI and the Humanities

AT FIRST GLANCE IT SEEMED NO DIFFERENT FROM ANY OTHER MLA session: in a midsize room at the Washington State Convention Center, well attended but not quite filled to capacity, with people leafing through their programs, checking their phones, drifting in and out. It was session 388, Being Human, Seeming Human. Arranged by the Office of the Executive Director, it was the first of its kind. Four of the six speakers were from Microsoft,[1] expressly invited to start a conversation about what it means for those who self-identify as human to share the planet with those who seem to be.

The Microsoft representatives talked about OpenAI's GPT-2, a widely used text generator. Can we always tell that the writer is an algorithm? And should we object if it happens to write like us?[2] These playful conundrums were exactly right for the occasion, but just the tip of the iceberg. Indeed, "seeming human" might turn out to be one of the less scary things AI can do. Replacing, supplanting, and eliminating human beings are also in the cards.

From self-driving cars to facial recognition biometrics to drones carrying out remote assassinations, AI is poised to transform the fabric of life and the future of work. The Brookings Institution, drawing on a study by the Stanford graduate student Michael Webb,[3] reported on 20 November 2019 that, unlike the automation enabled by robots such as those in Amazon warehouses (Edwards), which mostly affects low-paying jobs, the predictive and decision-making powers of AI—in the form of machine-learning algorithms—will affect every employment sector, hitting educated workers the hardest (Muro et al.). Those "with graduate or professional degrees will be almost four times as exposed to AI as workers with just a high school degree" (Muro et al.). This point was reiterated in a 29 January 2020 update from the Brookings Institution, further underscoring the

https://doi.org/10.1632/pmla.2020.135.3.449 Published online by Cambridge University Press

difference between AI and robots and the extent to which the former might render educated human beings superfluous (Gaynor).

Warnings have come from those who know the technology from the inside. Stephen Hawking, suffering from amyotrophic lateral sclerosis and dependent on an Intel-designed AI system to talk and write, famously told the BBC in 2014 that "the development of full artificial intelligence could spell the end of the human race" (qtd. in Cellan-Jones). Also in 2014, Elon Musk, the CEO of Tesla and SpaceX, warned his audience at the Massachusetts Institute of Technology (MIT) that AI is humanity's "biggest existential threat," akin to "summoning the demon" (qtd. in Gibbs). Earlier this year, Sundar Pichai, the CEO of Google, wrote in the *Financial Times* that "we need to be clear-eyed about what could go wrong."

Two complementary (and controversial) gifts, one to MIT and the other to the University of Oxford, made by Stephen Schwarzman, the CEO of the Blackstone Group, the world's largest private equity firm, were meant to stave off any apocalyptic scenario. At MIT, a donation of $350 million launched the $1 billion Stephen A. Schwarzman College of Computing, created to promote a model of AI held accountable by other disciplines. "As computing reshapes our world, M.I.T. intends to help make sure it does so for the good of all," said MIT's president, L. Rafael Reif (qtd. in Adams). Artificial intelligence is ordinarily designed only by computer scientists and engineers; at MIT, it will henceforth be informed by "humanistic perspectives" ("Computing"). At Oxford, meanwhile, a parallel donation of $188.75 million (£150 million) created the Stephen A. Schwarzman Centre for the Humanities, which will house all the humanities departments as well as the new Institute for Ethics in AI (Dawkins).

Schwarzman, a billionaire dealmaker and Trump insider, is an incendiary figure in his own right.[4] Outcry from faculty members and students against the donor eclipsed the announced objective of the donations (Karr; Tandon). But the potential threat of artificial intelligence remains a live issue. AI design is key to mitigating this threat. How can we create algorithms that would complement rather than replace human beings, help rather than destroy us? Many top-ranked computer science schools have brought in the humanities to help tackle the problem. At Carnegie Mellon, the Human-Computer Interaction Institute, founded in 1993, includes faculty members from the School of Fine Arts and the Dietrich School of Humanities and Social Sciences. At the University of California, Berkeley, the Human-Compatible Artificial Intelligence Institute involves faculty members from philosophy, psychology, and gender and women's studies. At the University of Southern California, the cofounder and director of the Center for Artificial Intelligence in Society is a professor from the Suzanne Dworak-Peck School of Social Work.

Stanford's Institute for Human-Centered Artificial Intelligence (HAI) is perhaps the most thoroughgoing. "As artificial intelligence becomes better and better, the technologies that take advantage of it can also seem to become more and more 'human,'" writes Persis Drell, Stanford's provost. The "interdisciplinary nature" of HAI "really forces us to address deep questions on 'what it means to be human?' and what distinguishes technology from humanity." John Etchemendy, a codirector of HAI, agrees. "AI is going to transform every discipline, including and especially the humanities," because it will generate "fundamental questions that the humanities are best equipped to answer":

> These range from the obvious questions in ethics that arise when we create systems controlled by AI to fundamental questions about our conception of the human. Ever since Aristotle, rationality has been considered the defining feature of humans that distinguishes us from other animals. What happens when

https://doi.org/10.1632/pmla.2020.135.3.449 Published online by Cambridge University Press

> this defining feature is shared, or even exceeded, by artifacts that we humans create?

To pursue these questions, HAI is directed by a philosopher, Etchemendy, and a computer scientist, Fei-Fei Li. The institute's faculty includes English professors, anthropologists, an ethicist, historians, linguists, musicologists, philosophers, psychologists, and law professors. "If AI is to serve the collective needs of humanity," the directors say in their welcome letter, "the creators and designers of AI must be broadly representative of humanity. This requires a true diversity of thought—across gender, ethnicity, nationality, culture and age, as well as across disciplines" (Li and Etchemendy).

There is no guarantee, of course, that these diverse disciplines would be able to work meaningfully together or that the AI design they come up with would translate into meaningful action. A recent seminar, AI for Earth and Environment, part of the AI for Good series sponsored by HAI and Stanford's Institute for Computational and Mathematical Engineering, highlights this point. Featuring Lucas Joppa, Microsoft's chief environmental officer, and Stefano Ermon, a computer scientist at HAI, the livestreamed event was conducted in a language decidedly nontechnical. Joppa, much in the news lately because of Microsoft's initiative to set aside $1 billion for carbon removal in order to become carbon negative by 2030,[5] chose on this occasion to focus on something equally visionary: a collective effort to fill in data gaps, surprisingly common in an age where data are supposed to be overabundant. "We live in the most narcissistic definition of an information age," he said. "We have so much information about ourselves and so little information about the planet we call home" (qtd. in Deb).

Joppa mentioned in passing an AI project yet to be implemented—a "planetary computer"—which he had proposed last year in a *Scientific American* article. This computer, a cloud-based AI platform collecting data about the health of ecosystems around the world, would also serve as a "decision engine," formulating large-scale policies based on its data analytics:

> A planetary computer will borrow from the approach of today's internet search engines, and extend beyond them in the form of a geospatial decision engine that supports queries about the environmental status of the planet, programmed with algorithms to optimize its health. Think of this less as a giant computer in a stark white room and more as an approach to computing that is planetary in scale and allows us to query every aspect of environmental and nature-based solutions available in real time. (Joppa, "Planetary Computer")[6]

The foremost task for a planetary computer is to monitor earth's environment and outline actions that would preserve its integrity. Why dream of such a complex AI platform, so clearly beyond our reach at the moment? And why hand over collective decision-making to this machine, when human beings have for so long enjoyed a laissez-faire existence?

Joppa refers us to the humanities for an answer, in the shape of an essay by Jonathan Franzen, "What If We Stopped Pretending?," published in the *New Yorker* in September 2019. "Climate apocalypse is coming. To prepare for it, we need to admit that we can't prevent it," Franzen begins. He goes on to lay out what exactly is in store for us: "If you're younger than sixty, you have a good chance of witnessing the radical destabilization of life on earth—massive crop failures, apocalyptic fires, imploding economies, epic flooding, hundreds of millions of refugees fleeing regions made uninhabitable by extreme heat or permanent drought. If you're under thirty, you're all but guaranteed to witness it." Is there any chance that such scenarios could be averted, that human beings would clean up their act in time to prevent global temperatures from rising another two degrees? Not by any stretch of the imagination, Franzen says:

https://doi.org/10.1632/pmla.2020.135.3.449 Published online by Cambridge University Press

> As a non-scientist, I do my own kind of modelling. I run various future scenarios through my brain, apply the constraints of human psychology and political reality, take note of the relentless rise in global energy consumption. . . . Call me a pessimist or call me a humanist, but I don't see human nature fundamentally changing anytime soon. I can run ten thousand scenarios through my model, and in not one of them do I see the two-degree target being met.

To be a humanist in the twenty-first century is to be clear-eyed about the limitations of our species, clear-eyed about the slim chance that we would be able to change our ways in time. Throughout history human beings have shown little ability to care for the planet; our track record has been especially abysmal in the past two hundred years. If *humanism* equals *pessimism* at this moment, it is because it seems unlikely that we would be able to rewire our brains fast enough to guarantee the survival of our species, let alone the survival of other species. At no other point in time is this fatal flaw more evident than now.

Joppa's planetary computer is the logical answer. Setting aside illusions about who or what we are, this blueprint for the future requires a division of labor between human beings and machines. "We need complementary decision making power that can optimize over the dimensions of space and time that our brains aren't naturally evolved to do, but which our computers can strongly complement us in," Joppa says (E-mail). "It's the job of *Homo sapiens* to describe our overall objective for the Earth. Then it's the job of computers to produce optimization results that are aligned with the human-defined objective" (qtd. in Strickland). Replacing human beings here is life-saving rather than life-destroying. "We need to do some radical things, and we need to do them now," Joppa maintains (qtd. in Trotman).

The idea is not as far-fetched as it might seem. Already defense departments around the world are lining up AI to make life-and-death decisions in war (Chandler). During a pandemic such as COVID-19, an AI platform coordinating global health data can offer early warnings, monitor the spread of the disease, help develop vaccines and predict the efficacy of existing drugs, and model long-term recovery as well as urgent decisions in the emergency room (Wu; Broad). Yet, even if we were to invest vast resources to build this planetary computer, there is no guarantee that the policies it formulates could be implemented to avert fast-approaching catastrophes. Joppa recognizes this problem. "Technology is only part of the toolbox, not a silver bullet," he says. "Human behavior got us to the present day, and it will take a massive effort from people to get us out of the crisis" ("Planetary Computer").

How to promote human activism? If a division of labor between human beings and machines is indeed the best way to go forward, a planetary computer is only half the answer. The other half would have to come from us. To maximize grassroots participation, Joppa stresses the importance of democratizing technology, making broadly available AI tools that would allow everyone to input data with ease. He points to iNaturalist, a global biodiversity monitoring system that uses computer-vision algorithms to allow its 575,000 members to upload more than 7 million photos to date ("Case"), crucial as a safeguard against future pandemics (Scott). These photos used to require expert verification; deep-learning neural networks make that unnecessary. Thanks to these user-friendly AI systems, crowdsourcing is now the backbone of a crisis-responsive "citizen science" (Joppa, "Case").

But how to make human beings change their ways beyond this broad-based data collection? A planetary computer can only formulate policy; its data analytics have teeth only if there are people to enforce them. For that to happen, AI literacy would seem important for a much larger segment of the

https://doi.org/10.1632/pmla.2020.135.3.449 Published online by Cambridge University Press

population than we think. "We need a public sector knowledge infrastructure in place," Etchemendy says. He points to Regulations Laboratory, or RegLab, a project undertaken by Dan Ho, a law professor and one of HAI's associate directors. "RegLab has a formal partnership with EPA to design, pilot, and evaluate AI-based interventions, which EPA has never done." The lab is about to host an AI boot camp for thirty-five state EPAs. In addition, it is working with "a ProPublica journalist to write a piece timed with the release of a dataset on facilities evading permitting requirements under the Clean Water Act."

EPA, ProPublica, Stanford Law School—these entities have standing only if the general public respects the rule of law, a free and independent press, and educational institutions charged with the dissemination of actionable knowledge. These civic infrastructures are never more important than when society is under the stress of catastrophes. And the actionable knowledge they disseminate needs to include not only AI literacy but also literacy about the human species. We need to know what kind of a species we have been, our culpability as well as vulnerability throughout history. We need to know the large-scale consequences of our lifestyle choices and what exactly the future would look like if we stick to our dietary and energy habits.

For this kind of species literacy, humanists, schooled by long records from the past and the outpouring of apocalyptic fiction right now, are uniquely suited. Literature from *Gilgamesh* on has taught us about the human assault on the nonhuman world. It has also taught us the art of assisted survival by making kin with nonhuman beings. The emergence of AI at this moment of crisis makes that art all the more urgent.

Wai Chee Dimock

Notes

1. The four speakers from Microsoft were Donald Brinkman, Bill Dolan, Jonathan Grudin, and Allison Hegel.

2. On the issues at stake, see Lea.

3. Webb's study was funded by Stanford's Institute for Human-Centered Artificial Intelligence.

4. Schwarzman chaired Trump's short-lived Strategic and Policy Forum, a sixteen-member business forum created to advise the president on productivity and economic growth.

5. See "Microsoft's Billion-Dollar Investment"; Vetter; Gold.

6. See also "A Planetary Computer for a Sustainable Future."

Works Cited

Adams, Susan. "Stephen Schwarzman Makes Anchor Gift for New $1 Billion School of Artificial Intelligence at MIT." *Forbes*, 15 Oct. 2018, www.forbes.com/sites/susanadams/2018/10/15/stephen-schwarzman-makes-anchor-gift-for-new-1-billion-school-of-artificial-intelligence-at-m-i-t/#478c3e48ed99.

Broad, William J. "A.I. versus the Coronavirus." *The New York Times*, 26 Mar. 2020, www.nytimes.com/2020/03/26/science/ai-versus-the-coronavirus.html.

Cellan-Jones, Rory. "Stephen Hawking Warns Artificial Intelligence Could End Mankind." *BBC*, 2 Dec. 2014, www.bbc.com/news/technology-30290540.

Chandler, Simon. "How Artificial Intelligence Will Make Decisions in Tomorrow's War." *Forbes*, 20 Jan. 2020, www.forbes.com/sites/simonchandler/2020/01/20/how-artificial-intelligence-will-make-decisions-in-tomorrows-wars/#9c60d8f4e011.

"Computing and Artificial Intelligence: Humanistic Perspectives from MIT." *MIT News*, 24 Sept. 2019, news.mit.edu/2019/computing-and-ai-humanistic-perspectives-0924.

Dawkins, David. "Billionaire Stephen Schwarzman Gifts $188 Million to Oxford University." *Forbes*, 19 June 2019, www.forbes.com/sites/daviddawkins/2019/06/19/billionaire-stephen-schwarzman-gifts-188-million-to-oxford-university/#41fd92bb6ab7.

Deb, Sarina. "'AI for Good' Talk Pushes Tech Usage to Mitigate Humans' Environmental Impact." *Stanford Daily*, 28 Jan. 2020, www.stanforddaily.com/2020/01/27/ai-for-good-talk-pushes-tech-usage-to-mitigate-humans-environmental-impact/.

Drell, Persis. E-mail to Wai Chee Dimock. 31 Jan. 2020.

Edwards, David. "Amazon Now Has 200,000 Robots Working in Its Warehouses." *Robotics and Automation*, 21 Jan. 2020, roboticsandautomationnews

https://doi.org/10.1632/pmla.2020.135.3.449 Published online by Cambridge University Press

.com/2020/01/21/amazon-now-has-200000-robots-working-in-its-warehouses/28840/.

Etchemendy, John. E-mail to Wai Chee Dimock. 1 Feb. 2020.

Franzen, Jonathan. "What If We Stopped Pretending?" *New Yorker*, 8 Sept. 2019, www.newyorker.com/culture/cultural-comment/what-if-we-stopped-pretending.

Gaynor, Michael. "Automation and AI Sound Similar, but May Have Vastly Different Impacts on the Future of Work." *Brookings*, 29 Jan. 2020, www.brookings.edu/blog/the-avenue/2020/01/29/automation-and-artificial-intelligence-sound-similar-but-may-have-vastly-different-impacts-on-the-future-of-work/.

Gibbs, Samuel. "Elon Musk: Artificial Intelligence Is Our Biggest Existential Threat." *The Guardian*, 25 Oct. 2014, www.theguardian.com/technology/2014/oct/27/elon-musk-artificial-intelligence-ai-biggest-existential-threat.

Gold, Russell. "Microsoft Strives for a Carbon Free Future." *Wall Street Journal*, 30 Jan. 2020, www.wsj.com/articles/microsoft-strives-for-a-carbon-free-future-a-setback-in-fargo-shows-the-hard-reality-11580380200.

Joppa, Lucas. "The Case for Technology Investments in the Environment." *Nature*, 19 Dec. 2017, www.nature.com/articles/d41586-017-08675-7.

———. E-mail to Wai Chee Dimock. 3 Feb. 2020.

———. "A Planetary Computer to Avert Environmental Disaster." *Scientific American*, 19 Sept. 2019, www.scientificamerican.com/article/a-planetary-computer-to-avert-environmental-disaster/.

Karr, Tiffany. "Schwarzman College at MIT Spurs Outcry by Students, Faculty." *Bloomberg News*, 18 Feb. 2019, www.bloomberg.com/news/articles/2019-02-18/schwarzman-college-at-mit-spurs-backlash-as-group-scolds-school.

Lea, Richard. "If a Novel Was Good, Would You Care If It Was Created by Artificial Intelligence?" *The Guardian*, 27 Jan. 2020, www.theguardian.com/commentisfree/2020/jan/27/artificial-intelligence-computer-novels-fiction-write-books.

Li, Fei-Fei, and John Etchemendy. "Welcome to the Stanford Institute for Human-Centered Artificial Intelligence." *Stanford University*, hai.stanford.edu/about/welcome.

"Microsoft's Billion-Dollar Investment in Carbon Removal." *NPR Marketplace*, 17 Jan. 2020, www.marketplace.org/shows/marketplace-tech/microsofts-billion-dollar-investment-in-carbon-removal/.

Muro, Mark, et al. "Which Jobs Are Most Affected by AI? Better-Paid, Better-Educated Workers Face the Most Exposure." *Brookings*, 20 Nov. 2019, www.brookings.edu/research/what-jobs-are-affected-by-ai-better-paid-better-educated-workers-face-the-most-exposure/.

Pichai, Sundar. "Why Google Thinks We Need to Regulate AI." *Financial Times*, 19 Jan. 2020, www.ft.com/content/3467659a-386d-11ea-ac3c-f68c10993b04.

"A Planetary Computer for a Sustainable Future through the Power of AI." *Microsoft*, innovation.microsoft.com/a-planetary-computer-for-earth/.

Scott, John. "How Biodiversity Loss Is Hurting Our Ability to Combat Pandemics." *World Economic Forum*, 9 Mar. 2020, www.weforum.org/agenda/2020/03/biodiversity-loss-is-hurting-our-ability-to-prepare-for-pandemics/.

Strickland, Eliza. "The Ultimate Optimization Problem." *IEEE Spectrum*, 20 Sept. 2020, spectrum.ieee.org/tech-talk/energy/environment/the-ultimate-optimization-problem-how-to-best-use-every-square-meter-of-the-earths-surface.

Tandon, Arya. "Societies, Academics and Students Condemn Schwarzman Donation." *The Oxford Student*, 15 Oct. 2019, www.oxfordstudent.com/2019/10/15/societies-academics-and-students-condemn-schwarzman-donation/.

Trotman, Andy. "'We Need to Do Some Radical Things, and We Need to Do Them Now': Microsoft's Chief Environmental Officer Reveals Why He Thinks the World of AI for Earth." *Microsoft*, news.microsoft.com/en-gb/features/we-need-to-do-some-radical-things-and-we-need-to-do-them-now/.

Vetter, Jonathan. "Microsoft Is Going Carbon Negative: What Does That Mean?" *Forbes*, 20 Jan. 2020, www.forbes.com/sites/davidrvetter/2020/01/20/microsoft-is-going-carbon-negative-what-does-that-mean/#38f2d1dab130.

Webb, Michael. "The Impact of Artificial Intelligence on the Labor Market." *SSRN*, 15 Nov. 2019, papers.ssrn.com/sol3/papers.cfm?abstract_id=3482150.

Wu, Jun. "How Artificial Intelligence Can Help Fight Coronavirus." *Forbes*, 19 Mar. 2020, www.forbes.com/sites/cognitiveworld/2020/03/19/how-artificial-intelligence-can-help-fight-coronavirus/#180a6d4e4d3a.

https://doi.org/10.1632/pmla.2020.135.3.449 Published online by Cambridge University Press