



July 22, 2016

Attn: Terah Lyons
Office of Science and Technology Policy,
Eisenhower Executive Office Building, 1650 Pennsylvania Ave. NW,
Washington, DC 20504

On behalf of the Center for Data Innovation (datainnovation.org), we are pleased to submit these comments in response to the Office of Science and Technology Policy's (OSTP) request for information on the overarching questions in artificial intelligence (AI).¹

The Center for Data Innovation is the leading think tank studying the intersection of data, technology, and public policy. With staff in Washington, DC and Brussels, the Center formulates and promotes pragmatic public policies designed to maximize the benefits of data-driven innovation in the public and private sectors. The Center is a non-profit, non-partisan research institute affiliated with the Information Technology and Innovation Foundation.

AI is a part of computer science devoted to creating computing machines and systems that perform operations analogous to human learning and decision-making.¹ Technological advancements over the past decade demonstrate that AI will become dramatically more powerful and effective at solving everything from mundane challenges, such as helping consumers figure out what to buy for the holidays, to the most pressing social and economic challenges, ranging from diagnosing and developing new treatments for devastating diseases to dramatically improving worker productivity. In this submission, we outline some of the most significant benefits and challenges of AI so that policymakers can take an active role in supporting the development of AI, as well as avoid succumbing to widespread yet unfounded alarmist narratives about how AI is a threat to economic and social well-being or even an existential threat to humanity.

¹ "Request for Information on Artificial Intelligence," Federal Register, June 27, 2016, <https://www.federalregister.gov/articles/2016/06/27/2016-15082/request-for-information-on-artificial-intelligence>.



Our responses to the relevant questions are in the attached document.

Sincerely,

Daniel Castro
Director
Center for Data Innovation
dcastro@datainnovation.org

Joshua New
Policy Analyst
Center for Data Innovation
jnew@datainnovation.org



THE LEGAL AND GOVERNANCE IMPLICATIONS OF AI

It was relatively clear how traditional software systems made decisions, as parameters were built in and largely understandable. In contrast, many AI systems make decisions based on complex models developed by an algorithm that continually adjust and improve based on experience. Since these adjustments may involve obscure changes in how variables are weighted in computer models, some critics have labeled these systems “black boxes” that are likely to create “algorithmic bias” that enables government and corporate abuse. These critics generally fall into two camps: those that believe companies or governments will deliberately “hide behind their algorithm” as a cover to exploit, discriminate, or otherwise act unethically; and those who argue that opaque, complicated systems will allow “runaway data” to produce unintended and damaging results.²

But resistance to AI because of these concerns fails to recognize a key point: AI, like any technology, can be used unethically or irresponsibly. AI systems are not independent from their developers and, more importantly, from the organizations using them. If a government or business wants to systematically discriminate against certain groups of persons, it does not need AI to do so. To put it simply, bad actors will do illegal things with or without computers.³

Nonetheless, many critics seem convinced that the complexity of these systems is responsible for any problems that emerge, and that mandating “algorithmic transparency” is necessary to ensure that the public can police against biased AI systems.⁴ Combatting bias and protecting against harmful outcomes is of course important, but mandating that companies open their proprietary AI software to the public would not solve these problems, but would create new ones. Consumers and policymakers are ill-equipped to actually understand the complicated decision-making processes of an AI algorithm, and AI systems can learn and change over time, making it difficult to measure unfairness by examining their underlying mechanics.⁵ Moreover, the economic impact of such a mandate would be significant, as it would prevent companies from capitalizing on their intellectual property and future investment and research into AI would slow.

Fortunately, many have recognized that embedding ethical principles into AI systems is both possible and effective. The White House’s framework for “equal opportunity by design” in algorithmic systems, as described by its report on the opportunities and challenges of big data and civil rights presents, is one promising method.⁶ This approach, described more generally by Federal Trade Commissioner Terrell McSweeney as “responsibility by design,” rightly recognizes that algorithmic systems can produce unintended outcomes, but does not demand a company waive rights to keep its software proprietary.⁷ Instead, the principle of responsibility by design



provides developers with a productive framework for solving the root problems of undesirable results in algorithmic systems: bad data as an input, such as incomplete data and selection bias, and poorly designed algorithms, such as conflating correlation with causation, and failing to account for historical bias.⁸ In particular, the federal government should help address the problem of data poverty, where a lack of high-quality data about certain groups of individuals puts them at a social or economic disadvantage.⁹

It also is important to note that some calls for algorithmic transparency are actually more in line with the principle of responsibility by design. For example, former chief technologist of the Federal Trade Commission (FTC) Ashkan Soltani said that although pursuing algorithmic transparency was one of the goals of the FTC, “accountability” rather than “transparency” would be a more appropriate way to describe the ideal approach, and that making companies surrender their source codes is “not necessarily what we need to do.”¹⁰ Rather than make companies relinquish their intellectual property rights, encouraging adherence to the principle of responsibility by design would allow companies to better police themselves to prevent unintended outcomes and still ensure that regulators could intervene and audit these systems should there be evidence of bias or other harms.¹¹ For example, policymakers could work with the private sector to develop a framework for predicting and accounting for disparate impact in AI systems. Companies would be more willing to deploy AI if they could clearly demonstrate they are acting in good faith by actively considering how their systems could produce discriminatory outcomes and taking steps to address these concerns.¹²

Figuring out just how to define responsibility by design and encourage adherence to it warrants continued research and discussion, but it is crucial that policymakers understand that AI systems are valuable because of their complexity, not in spite of it. Attempting to pull back the curtain on this complexity to protect against undesirable outcomes threatens the progress of AI.

THE USE OF AI FOR PUBLIC GOOD

AI systems can help organizations make better-informed, timelier decisions, as well as tackle complicated, data-intensive problems that humans are ill-equipped to solve. As such, the potential benefits of AI are will likely be quite large. There are already compelling examples of AI offering substantial benefits for civil rights, public health, conservation, energy efficiency, financial services, and healthcare.¹³ In addition, AI has the potential to make government substantially more efficient and citizen-friendly if government agencies adopt the technology.



THE SOCIAL AND ECONOMIC IMPLICATIONS OF AI

The criteria for measuring the success of AI should not be whether it is “perfect” but rather if it is an improvement over the status quo. AI offers an unprecedented opportunity to automate decision-making and reduce the influence of explicit and subconscious human bias that permeate every aspect of society and the economy.¹⁴ AI systems make decisions based on data, and quantifying and analyzing the decision-making process can both expose the underlying bias exhibited by human-made decisions, as well as prevent subjective and potentially discriminatory human decision-making from ever entering the equation.¹⁵

Regarding economic implications, one of the most widely-repeated warnings about AI is that it will lead to mass unemployment, as smart machines become increasingly adept at performing work normally carried out by humans in both blue and white collar jobs.¹⁶ However, the “AI will destroy jobs” argument is incorrect, for several reasons. First, most AI applications will not be able to fully replace human workers, but rather will automate particular tasks and allow a human worker to spend their time in more valuable ways.¹⁷ Very few jobs could conceivably be automated in the short or medium term, and automation will instead transform the function of many existing jobs rather than eliminate them.¹⁸

Second, in instances where AI does eliminate jobs, the jobs lost will be offset by the resulting productivity growth that leads to the creation of new jobs. When a business replaces a human worker with an AI system, it does so because the AI increases the business’s productivity by doing the job more effectively for lower cost. If jobs in one firm or industry are reduced or eliminated through higher productivity, then by definition production costs go down. These savings are passed on, often in the form of lower prices or higher wages. This money is then spent, which creates jobs in whatever industries supply the goods and services on which people spend their increased savings or earnings.¹⁹

To be sure, there are winners and losers in the process of productivity improvement: Some workers will lose their jobs, and it is appropriate for policymakers to help those workers quickly transition to new employment. But there is simply no merit in the belief that productivity growth will reduce the overall number of jobs.²⁰



THE MOST IMPORTANT RESEARCH GAPS IN AI THAT MUST BE ADDRESSED TO ADVANCE THIS FIELD AND BENEFIT THE PUBLIC

Ensuring that AI systems produce fair, unbiased, and safe results without mandating algorithmic transparency can pose complicated technical challenges that warrants further research. For example, Carnegie Mellon researchers have developed a method for determining why an AI system makes particular decisions without having to divulge the underlying workings of the system or code.²¹ This research will be useful to addressing regulators' concerns about discrimination, as well as helping companies that want to ensure they are acting ethically.

Federal research programs should avoid working in social engineering into AI research. For example, the National Science Foundation's National Robotics Initiative focuses on accelerating the development of robotic systems, but only systems that work beside or cooperatively with humans.²² While AI-powered worker assistance applications will be beneficial, limiting the focus of this research in such a manner precludes opportunities to develop AI systems that could replace workers, which history has shown to produce greater economic benefits in the moderate and long run.

CONCLUSION

AI has the potential to generate substantial benefits to the economy, society, and overall quality of life, and it is encouraging to see OSTP proactively working to better understand the technology, promote its research and development, and set the record straight about the potential opportunities associated with the technology. OSTP should also play an active role in dispelling the prevalent alarmist myths about AI, particularly concerns that AI will lead to higher rates or unemployment and even eradicate the human race, which, besides being wrong, threaten the acceptance and advancement of this technology.



REFERENCES

1. Robert Atkinson, "'It's Going to Kill Us!' and other Myths About the Future of Artificial Intelligence," Information Technology and Innovation Foundation, June 2016, <http://www2.itif.org/2016-myths-machine-learning.pdf>.
2. Tim Hwang and Madeleine Clare Elish, "The Mirage of the Marketplace," *Slate*, August 9, 2015, http://www.slate.com/articles/technology/future_tense/2015/07/uber_s_algorithm_and_the_mirage_of_the_marketplace.html and David Auerbach, "The Code We Can't Control," *Slate*, January 14, 2015, http://www.slate.com/articles/technology/bitwise/2015/01/black_box_society_by_frank_pasquale_a_c_hilling_vision_of_how_big_data_has.html.
3. Katherine Noyes, "The FTC Is Worried About Algorithmic Transparency, and You Should Be Too," *PC World*, April 9, 2015, <http://www.pcworld.com/article/2908372/the-ftc-is-worried-about-algorithmictransparency-and-you-should-be-too.html>
4. For example, the Electronic Privacy Information Center argues that "Entities that collect personal information should be transparent about what information they collect, how they collect it, who will have access to it, and how it is intended to be used. Furthermore, the algorithms employed in big data should be made available to the public." See Mark Rotenburg, "Comments of The Electronic Privacy Information Center to The Office of Science and Technology Policy, Request for Information: Big Data and the Future of Privacy" (Electronic Privacy Information Center, April 4, 2014), <https://epic.org/privacy/big-data/EPIC-OSTP-Big-Data.pdf>; See also David Auerbach, "The Code We Can't Control," *Slate*, January 14, 2015, http://www.slate.com/articles/technology/bitwise/2015/01/black_box_society_by_frank_pasquale_a_c_hilling_vision_of_how_big_data_has.html.
5. Lauren Smith, "Algorithmic Transparency: Examining from Within and Without," IApp, January 28, 2016, <https://iapp.org/news/a/algorithmic-transparency-examining-from-within-and-without/>.
6. The White House Executive Office of the President, "Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights," May 2016, https://www.whitehouse.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf.
7. Terrell McSweeney, "Tech for Good: Data for Social Empowerment" (keynote remarks of Commissioner Terrell McSweeney at Google DC Tech Talk Series, Washington District of Columbia, September 10, 2015), https://www.ftc.gov/system/files/documents/public_statements/800981/150909googletechroundtable.pdf.
8. The White House Executive Office of the President, "Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights," May 2016, https://www.whitehouse.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf.
9. Daniel Castro, "The Rise of Data Poverty in America," Center for Data Innovation, September 10, 2014, <http://www2.datainnovation.org/2014-data-poverty.pdf>.
10. Christopher Zara, "FTC Chief Technologist Ashkan Soltani On Algorithmic Transparency and the Fight Against Biased Bots," *International Business Times*, April 9, 2015, <http://www.ibtimes.com/ftc-chieftechologist-ashkan-soltani-algorithmic-transparency-fight-against-biased-1876177>.

11. Katherine Noyes, "The FTC Is Worried About Algorithmic Transparency, and You Should Be Too," *PC World*, April 9, 2015, <http://www.pcworld.com/article/2908372/the-ftc-is-worried-about-algorithmictransparency-and-you-should-be-too.html>.
12. Travis Korte and Daniel Castro, "Disparate Impact Analysis is Key to Ensuring Fairness in the Age of the Algorithm," Center for Data Innovation, January 20, 2015, <https://www.datainnovation.org/2015/01/disparate-impact-analysis-is-key-to-ensuring-fairness-in-the-age-of-the-algorithm/>.
13. Hope Reese, "AI App Uses Social media to Spot Public Health Outbreaks," *TechRepublic*, March 9, 2016, <http://www.techrepublic.com/article/ai-app-uses-social-media-to-spot-public-health-outbreaks/>, "Energy Savings from the Nest Learning Thermostat: Energy Bill Analysis Results," Nest Labs, February 2015, <https://nest.com/downloads/press/documents/energy-savings-white-paper.pdf>, Derrick Harris, "How PayPal Uses Deep Learning and Detective Work to Fight Fraud," *Gigaom*, March 6, 2015, <https://gigaom.com/2015/03/06/how-paypal-uses-deep-learning-and-detective-work-to-fight-fraud/>, and Alexander Kostura, "Artificial Intelligence is Key to Using Data for Social Good," Center for Data Innovation, July 11, 2016, <https://www.datainnovation.org/2016/07/artificial-intelligence-is-key-to-using-data-for-social-good/>.
14. Joshua New, "It's Humans, Not Algorithms, That Have a Bias Problem," Center for Data Innovation, November 16, 2015, <https://www.datainnovation.org/2015/11/its-humans-not-algorithms-that-have-a-bias-problem/>.
15. Joshua New, "The White House is Starting to Get it Right on Big Data," Center for Data innovation, June 13, 2016, <https://www.datainnovation.org/2016/06/the-white-house-is-starting-to-get-it-right-on-big-data/>.
16. Robert Atkinson, "'It's Going to Kill Us!' and other Myths About the Future of Artificial Intelligence," Information Technology and Innovation Foundation, June 2016, <http://www2.itif.org/2016-myths-machine-learning.pdf>.
17. Ibid.
18. Michael Chu, James Manyika, and Mehdi Miremadi, "Four Fundamentals of Workplace Automation," *McKinsey Quarterly*, November 2015, <http://www.mckinsey.com/business-functions/business-technology/our-insights/four-fundamentals-of-workplace-automation>.
19. Robert Atkinson, "'It's Going to Kill Us!' and other Myths About the Future of Artificial Intelligence," Information Technology and Innovation Foundation, June 2016, <http://www2.itif.org/2016-myths-machine-learning.pdf>.
20. Ibid.
21. "Carnegie Mellon Transparency Reports Make AI Decision-Making Accountable," Carnegie Mellon University, May 26, 2016, <https://www.ecnmag.com/news/2016/05/carnegie-mellon-transparency-reports-make-ai-decision-making-accountable>.
22. "National Robotics Initiative (NRI)," National Science Foundation, accessed July 20, 2016, <http://www.nsf.gov/pubs/2016/nsf16517/nsf16517.htm>.