CENTER FOR DATA INNOVATION

September 1, 2017

Re: Response to the Call for Evidence by the House of Lords Select Committee on Artificial Intelligence

Dear Members of the Committee,

On behalf of the Center for Data Innovation, we are pleased to submit the following response to the call for evidence by the select committee on artificial intelligence.

The nonprofit, nonpartisan 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. It educates policymakers and the public about the opportunities and challenges associated with data, as well as technology trends such as artificial intelligence, data analytics, and the Internet of Things. In our answers to the Committee's questions, there are two particularly salient points we wish to emphasize.

First, there is little to no evidence to support the hyperbolic fears about AI, such as the suggestion that the technology will cause cataclysmic job destruction, loss of privacy, bias and abuse, and even human extinction or enslavement. The notion that AI raises such grave concerns that policymakers should take a precautionary regulatory approach to limit the damages it could allegedly cause is both wrong and harmful to societal progress. However, there is substantial evidence of AI's economic benefits. Thus, rather than attempt to limit AI, the role of policy should be to accelerate its development and adoption.

Second, over the long-term the potential benefits of AI are largely dependent on an adequate supply of data. Policymakers should therefore ensure they do not constrain the supply of data—such as by enacting overzealous data protection regulations—which would limit the positive impact of AI in jurisdictions where they apply, not to mention limit the growth of AI firms. Furthermore, policymakers should also work to close the "data divide"—the social and economic inequalities that may result from an insufficient collection or use of data about individuals or communities.

Yours faithfully,

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

Nick Wallace
Senior Policy Analyst
Center for Data Innovation
nwallace@datainnovation.org

What is the current state of artificial intelligence and what factors have contributed to this? How is it likely to develop over the next 5, 10, and 20 years? What factors, technical or societal, will accelerate this development?

AI is a field of computer science devoted to creating computing machines and systems that perform operations analogous to human learning and decision-making.[1] The Association for the Advancement of Artificial Intelligence describes AI as "the scientific understanding of the mechanisms underlying thought and intelligent behavior and their embodiment in machines."[2] AI does not necessarily imply machines with human-level intelligence, or machines that think in a human-like way. In fact, the very term "artificial intelligence" is a misnomer. Rather, AI describes a broad range of systems designed to behave in ways that humans think of as intelligent, and the level of intelligence in any given implementation of AI can vary greatly.

Contemporary AI systems generally exhibit one or more of the following functions: monitoring data to identify anomalies and patterns; extracting insights from large datasets in order to discover new connections and stimulate new solutions; predicting how trends are likely to develop; interpreting unstructured data that was hitherto difficult to classify; interacting with connected sensors and actuators in the physical environment; and interacting, communicating, and collaborating with humans and other machines. Practical applications of AI may involve just one or two of these functions, or may involve a complex array of algorithms performing near enough all of them, such as in autonomous vehicles.[3]

AI as a field of computer science began during the aftermath of the Second World War. Despite considerable excitement that major breakthroughs were just a few years away, the field showed only modest progress. However, the scientific and technological breakthroughs that have spurred its recent advancement, and made it more commercially viable during the last few years, are much more recent. Of particular importance is machine learning, a method whereby developers write algorithms that autonomously and iteratively build new analytical models in response to new data, without programming the solutions. Prior to this breakthrough, computer scientists had to laboriously pre-program outwardly intelligent behavior. The underlying factors that enabled the development of machine learning include better, cheaper computer hardware, particularly faster processing power and higher-capacity storage, as well as a greater supply of machine-readable data and better algorithms.[4]

Machine learning will produce ever more advanced algorithms that can interpret and respond to more complex data in more sophisticated and more reliable ways. This will expand the variety and complexity of tasks to which computer scientists can dedicate AI tools in a reliable and

commercially viable way. But contrary to speculation by some vocal critics of AI, the current progress of algorithmic development does not point towards the development of artificial consciousness, human-level or human-like artificial intelligence, sometimes called artificial general intelligence (AGI), anytime in the foreseeable future. Many of the dystopian fears about AI stem from the notion that AGI is imminent, feasible, or uncontrollable. In the 1960s, technologists began predicting that AGI was just a few years away. Since then, AI has progressed dramatically, and the underlying technology that supports it has developed even faster than predicted, yet AGI is likely just as far away today as it was 50 years ago. There is a very significant difference between the rapidly-advancing ability of machines to solve very specific problems in response to a narrow array of data supplied by humans, and a machine that can find solutions on its own to an infinite number of unpredictable and hitherto unknown problems with zero indication of what information might be pertinent to it.[5] This difference is akin to that between a jet that can fly at the speed of sound and a spacecraft that travel at warp speed.

How can the general public best be prepared for more widespread use of artificial intelligence?

The coming changes are not so dramatic as to require government to prepare the general public. The incoming wave of AI applications, though socially and economically important in its benefits, does not threaten to cause revolutionary social upheaval, especially not quickly. In fact, even the most socially consequential applications—such as the gradual emergence of autonomous vehicles—seems mundane in comparison to technological revolutions we have already seen, such as the rise of the Internet, not to mention the automobile itself, around which much of our urban infrastructure has been built. If policymakers act on the baseless assumption that AI has implications so dramatic as to require the public to be prepared, they risk creating undue panic, in turn generating political pressure for hasty policy decisions based on fear rather than fact and likely intended to slow down adoption. We have already seen such fears turn into ill-advised proposals to regulate and tax smart robots.

The general public does not need special preparation, though the continued evolution of the workforce requires government to maintain and strengthen programs that offer job retraining and other supports for dislocated workers. AI will cause economic disruption in some sectors, but this disruption will come slower and affect fewer sectors than many popular commentators allege.[6] For example, most doctors will not be replaced by AI, nor will nurses, journalists, civil servants, paramedics, or police officers. Taxi and bus drivers, airline pilots, and even lorry drivers will likely remain employable for the medium term due to remaining technological hurdles, consumer demand, public opinion, and public policy. However, light rail train drivers may face changes far

sooner, as autonomous trains are already commercially viable and in use in urban subway systems around Europe.[7] Policymakers should be prepared to help those who face such disruption retrain and find new career paths.

In most other fields, workers are more likely to find themselves working with AI than replaced by it. This will stimulate some demand for new skills, but the necessary experience will often be contingent upon industry-specific expertise that workers in that sector already have. Doctors, for example, will have to learn how to use some AI applications responsibly—but just as they will not be replaced by machines, nor will they be replaced by AI experts who are not doctors, and the AI tools they will use will have been designed with doctors in mind.

Who in society is gaining the most from the development and use of artificial intelligence and data? Who is gaining the least? How can potential disparities be limited?

As with most technology-driven efficiency gains, AI will benefit consumers and workers through increased productivity that will lead to greater choice and cheaper products and services and higher wages. This is particularly critical for the UK, which is suffering from an unprecedented productivity crisis, with productivity stagnant over the last decade. Unless Britain can find a way to boost productivity, social and political crises will increase as incomes stagnate, especially in the face of the increased proportion of retirees. AI will also lead to benefits for UK residents in a range of other areas, including healthcare, transportation, and environment. Those who stand to gain the least are people subject to types of social exclusion that restrict the supply of data pertinent to them, which in turn diminishes the relevance of AI tools to their circumstances.

Early uses of AI in healthcare are beginning to benefit patients, such as by helping doctors to identify problems in medical imaging and test results far earlier and more consistently than they might have otherwise.[8] But the benefits of these tools remain as limited as health services' readiness or ability to deploy them—the technology, however, is out there already.

Businesses that invest in AI sooner will enjoy earlier rewards than those who turn to it later: for while AI will yield large returns for successful developers of it, like many other kinds of information technology, AI will help businesses in virtually all sectors to become more efficient and productive, and competitive markets will mean that workers and consumers will benefit. Citizens also benefit from early implementations of AI, such as autonomous vacuum cleaners, personal virtual assistants, and personalized language learning.[9] AI is beginning to help lenders and insurers calculate risks more accurately using an unprecedented supply of data.[10] This helps to accept or reject applications more wisely, and fine-tune premiums, interest rates, repayment periods, excess, and the quantity lent or value insured. For applicants and wider society, this promises to improve access to these financial services.

Because of the important impact on productivity growth from AI, virtually all UK residents will benefit. However, those who stand to gain the least are those living in a state of what one can call "data poverty."[11] These are social groups about whom little data is ever collected, which limits the extent to which data-driven services can be of use to them. These tend to be groups that are already marginalized in myriad other ways too, such as refugees. To give one example of the potential dangers of data poverty: we already know that some societal groups experience higher rates of certain diseases than others, for reasons that are in some cases fully understood by the medical profession, and in others less so.[12] This means that a paucity of data on a given

community limits the usefulness to those communities of AI tools intended to help tackle such problems.

The most important thing policy makers can do is to communicate a message around AI that highlights the progressive forces AI represent. Just as UK policy makers have supported technological change from the steam engine to the Internet, and not given in to the demands of Luddites, they need to do the same today. Exaggerations about the impact of AI have led to many harmful policy recommendations, particularly the claim that automation bolsters the case for a universal basic income (UBI). This claim begs the question, because it assumes, contrary to evidence, that high productivity from automation will cause joblessness. UBI would increase social exclusion and unemployment, and reduce living standards, because it is not time-limited, which distorts incentives.[13] The UK government also should not succumb to techno-panic by following the path of some who propose harmful polices like taxing AI, regulating smart robots, or significantly limiting data access on which so much AI depends.

The UK government should support public R&D into AI, to help the UK become a global leader in this emerging field. At the same time, it should ensure that the education system produces more data scientists and computer scientists with a understanding of AI.

Finally, the government should take steps to address data poverty. Data poverty is not usually an isolated problem, but a symptom of broader social exclusion. As data becomes more important in the economy, there is a real danger that the economic consequences of social exclusion could become more severe. Attempts to tackle social exclusion, therefore, must be combined with more ambitious approaches to the collection and use of data in public policy and public administration.[14]

Should efforts be made to improve the public's understanding of, and engagement with, artificial intelligence? If so, how?

Public understanding, and even demand for, artificial intelligence, can help accelerate its adoption. Policymakers can facilitate this understanding by doing three things:

1. Inform themselves about what AI is and what it is not, and use this information to speak and argue more intelligently and more honestly in policy debates pertaining to AI.
2. Promote data skills throughout the education system, particularly as part of vocational and professional training in fields where data and AI are likely to play an important role, such as medicine.

3. Encourage the use of AI in public services and ensure out-of-date regulations do not become an unnecessary barrier. For example, UK medical regulations currently pose challenges for testing AI with patient data.[15]

How can the data-based monopolies of some large corporations, and the 'winner- takes-all' economies associated with them, be addressed? How can data be managed and safeguarded to ensure it contributes to the public good and a well-functioning economy?
There are no "data-based monopolies," and the winner does not take all. Data is non-rivalrous: customers who give their personal data to one company can provide it again to another. There are thousands of companies developing AI tools using large datasets. Accumulating personal data confers economic benefits on a company, but it does not automatically create a monopoly.[16] However, policymakers can boost competition by encouraging the free flow of data. For example, the law should extend the data portability rights of personal data subjects to users of systems (such as cars) that generate non-personal data, allowing those users to share that raw machine data with third parties, such as insurance companies.

What are the ethical implications of the development and use of artificial intelligence? How can any negative implications be resolved?
The ethical implications of AI are specific to the circumstances in which AI is deployed. For example, lethal autonomous weapons will demand a more robust ethical framework than autonomous vacuum cleaners.[17] Moreover, many ethical dilemmas commonly associated AI are independent of the technology. For example, a popular question is what a self-driving car should do when forced to choose between equally lethal alternatives. This is not a dilemma caused by AI, but by cars. Cars are dangerous machines that kill a staggering number of people with great frequency. AI will mitigate this dilemma by significantly reducing the number of accidents, but the fundamental implications of sitting inside a metal object and hurtling it forward at considerable speed remain the same.

Ethical concerns about explaining algorithmic decisions in human terms have led the EU to legislate for a "right to explanation" in the General Data Protection Regulation. Whether an individual has a right to have a decision explained depends on the decision, not the technology used to make it. The auditing of algorithms should be appropriate to the decisions they make, and not to a separate standard that applies solely to algorithms, such as that set out in the GDPR. Moreover, such approaches assume that human decision making is objective, transparent, and unbiased, something research has consistently shown is often wrong.

What lessons can be learnt from other countries or international organizations (e.g. the European Union, the World Economic Forum) in their policy approach to artificial intelligence?

The Chinese government is striving for leadership in AI by pushing state-controlled businesses to invest in developing and implementing the technology. Countries that would compete with China, or that would prefer not to see it dominate AI development, should put in place strategies that both support AI research and identify opportunities to deploy AI in industry, without mirroring China's mercantilism.[18]

Japan is a useful example in demonstrating how this might be achieved. The Japanese government has developed a roadmap for the commercialization of AI tools, which complements the AI development funding the government provides.[19] Admittedly, Japan's rapidly ageing population means the country has less to fear from claims of job destruction, and might be expected to take a more proactive approach to deploying AI in its industries. But as mentioned above, these claims are exaggerated, so the UK would do well to formulate a similar strategy that ties investment in AI research to social and economic gains—not least because the British government already sponsors AI research anyway.

The EU is more of a cautionary example: it has tried to regulate AI too early, imposing rules that address theoretical concerns without respect for evidence. The aforementioned right to explanation will not guarantee accountability in algorithmic decisions, because it isolates individual decisions, making it harder to identify algorithmic bias, at the same time as imposing pointless costs on business. Statistical auditing is a more practical way to root-out bias in automated decisions.[20] Furthermore, the European Parliament has endorsed a report that calls for the regulation of robots and speculates wildly about their capabilities and risks.[21] Just as over-regulation of biotechnology during the 1980s allowed the United States to take the lead, the new regulations threaten to have similar effects on AI, ceding leadership to other regions.

The World Economic Forum is also a poor example, as it too has largely succumbed to the "AI is out of control" narrative. Klaus Schwab, head of WEF, writes that "We stand on the brink of a technological revolution that will fundamentally alter the way we live, work, and relate to one another. In its scale, scope, and complexity, the transformation will be unlike anything humankind has experienced before."[22] As we highlighted above, there is simply no evidence for such hyperbolic claims. Viewing the development of AI in these overblown terms is virtually guaranteed to lead to bad policy.

---

[1] Daniel Castro and Joshua New, "The Promise of Artificial Intelligence" (Center for Data Innovation, October 10, 2016), https://www.datainnovation.org/2016/10/the-promise-of-artificial-intelligence/.
[2] "AI Overview: Broad Discussions of Artificial Intelligence," *AITopics*, accessed September 29, 2016, http://aitopics.org/topic/ai-overview.

[3] Daniel Castro and Joshua New, "The Promise of Artificial Intelligence" (Center for Data Innovation, October 10, 2016), https://www.datainnovation.org/2016/10/the-promise-of-artificial-intelligence/.
[4] Ibid.
[5] Ibid.
[6] Ibid, Robert Atkinson and John Wu, "False Alarmism: Technological Disruption and the U.S. Labor Market, 1850-2015" (Information Technology and Innovation Foundation, May 8, 2017), https://itif.org/publications/2017/05/08/false-alarmism-technological-disruption-and-us-labor-market-1850-2015.
[7] "One billion travelers on Europe's automated metro systems" (Allianz pro Schiene November 30, 2016), https://www.allianz-pro-schiene.de/en/pressemitteilung/overview-automated-metro-systems-europe/.
[8] Nick Wallace, "5 Q's for Matteo Carli, Chief Technology Officer and Founder of xbird" (Center for Data Innovation, June 29, 2017), https://www.datainnovation.org/2017/06/5qs-for-matteo-carli-chief-technology-officer-and-founder-of-xbird/.
Nick Wallace, "5 Q's for Eyal Toledano, Chief Technology Officer at Zebra Medical Imaging" (Center for Data Innovation, June 8, 2017), https://www.datainnovation.org/2017/06/5qs-for-eyal-toledano-co-founder-of-zebra-medical-vision/.
[9] Nick Wallace, "5 Q's for Mait Müntel, Co-Founder of Lingvist" (Center for Data Innovation, July 17, 2017), https://www.datainnovation.org/2017/07/5qs-for-mait-muntel-co-founder-of-lingvist/.
[10] Nick Wallace, "5 Q's for David Hand, Emeritus Professor at Imperial College, London" (Center for Data Innovation, January 30, 2017), https://www.datainnovation.org/2017/01/5-qs-for-david-hand-emeritus-professor-at-imperial-college-london/.
[11] Daniel Castro, "The Rise of Data Poverty in America" (Center for Data Innovation, September 10, 2014), https://www.datainnovation.org/2014/09/the-rise-of-data-poverty-in-america/.
[12] Erick Forno and Juan C. Celedón, "Asthma and Ethnic Minorities: Socioeconomic Status and Beyond" *Current Opinion in Allergy and Clinical Immunology*, 9(2):154-60, April 2009, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3920741/.
Pellowski et al, "A pandemic of the poor: social disadvantage and the U.S. HIV epidemic" *American Psychologist*, 68(4):197-209 May-June 2013, https://www.ncbi.nlm.nih.gov/pmc/articles/PMC3700367/.
"Why are some genetic conditions more common in particular ethnic groups?" *Genetics Home Reference*, (U.S. National Library of Medicine) last modified August 22, 2017, https://ghr.nlm.nih.gov/primer/inheritance/ethnicgroup.
[13] Robert Atkinson, "Robots, Automation, and Jobs: A Primer for Policymakers" (Information Technology and Innovation Foundation, May 8, 2017), https://itif.org/publications/2017/05/08/robots-automation-and-jobs-primer-policymakers.
[14] Daniel Castro, "Europe Should Promote Data For Social Good" (Center for Data Innovation, October 3, 2016), https://www.datainnovation.org/2016/10/europe-should-promote-data-for-social-good/.
[15] Nick Wallace, "UK Medical Regulations Need an Update to Make Way for Medical AI" (Center for Data Innovation, August 12, 2017), https://www.datainnovation.org/2017/08/uk-regulations-need-an-update-to-make-way-for-medical-ai/.
[16] Joe Kennedy, "The Myth of the Data Monopoly: Why Antitrust Concerns About Data Are Overblown" (Information Technology and Innovation Foundation, June 6, 2016), https://itif.org/publications/2017/03/06/myth-data-monopoly-why-antitrust-concerns-about-data-are-overblown.
[17] Daniel Castro, "'Ban the Killer Robots' Movement Could Backfire" *Computerworld*, November 13, 2015, https://www.computerworld.com/article/3005204/emerging-technology/ban-the-killer-robots-movement-could-backfire.html.

[18] Joshua New, "How Governments Are Preparing for Artificial Intelligence" (Center for Data Innovation, August 18, 2017), https://www.datainnovation.org/2017/08/how-governments-are-preparing-for-artificial-intelligence/.
"Why China's AI push is worrying" *The Economist*, July 27, 2017, https://www.economist.com/news/leaders/21725561-state-controlled-corporations-are-developing-powerful-artificial-intelligence-why-chinas-ai-push.
[19] Joshua New, "How Governments Are Preparing for Artificial Intelligence" (Center for Data Innovation, August 18, 2017), https://www.datainnovation.org/2017/08/how-governments-are-preparing-for-artificial-intelligence/.
[20] Nick Wallace, "EU's Right to Explanation: A Harmful Restriction on Artificial Intelligence" *TechZone360*, January 25, 2017, http://www.techzone360.com/topics/techzone/articles/2017/01/25/429101-eus-right-explanation-harmful-restriction-artificial-intelligence.htm.
[21] "Civil Law Rules on Robotics - TEXTS ADOPTED - European Parliament resolution of 16 February 2017 with recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL))" http://www.europarl.europa.eu/sides/getDoc.do?pubRef=-//EP//NONSGML+TA+P8-TA-2017-0051+0+DOC+PDF+V0//EN.
[22] Klaus Schwab, "The Fourth Industrial Revolution: what it means, how to respond" (World Economic Forum, January 14, 2016) https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/.