NISKANEN CENTER

Public Interest Comment

*Submitted to the Office of Science and Technology Policy in the Matter of:*

# A Request for Information on Artificial Intelligence


## Ryan Hagemann

Technology and Civil Liberties Policy Analyst
The Niskanen Center



Submitted: July 22, 2016
Docket No. 2016-15082
Word count: 1862


## Executive Summary


The long-awaited promise of artificial intelligence (AI) is beginning to materialize. Powerful AIs, such as IBM's Watson and Google's Deepmind, which has bested the world's Go champion, herald the "springtime" of AI research and development. However, some find the flowering of the technology alarming, and wonder aloud whether AI may lead to a Terminator-style future in which incomprehensibly intelligent computers destroy human civilization. Even moderate critics of AI warn that we now stand on the verge of a mass labor dislocation in which up to half of all jobs may be taken by machines. For now, however, these worries are extremely speculative, and the alarm they cause can be counterproductive.

In order to maximize the benefits associated with ongoing developments in AI, we recommend that policymakers and regulators:

(1) avoid speaking of hyperbolic hypothetical doomsday scenarios, and

(2) embrace a policy of regulatory restraint, intervening in the development and use of AI technology only when and if the prospect of harm becomes realistic enough to merit government intervention.

# Introduction

As the renowned science fiction author Isaac Asimov once wrote, "Any sufficiently advanced technology is indistinguishable from magic." For many people, the seemingly magical nature of unfamiliar technology invites wild speculation about the human implications of its development and adoption. Nowhere is this more true than in artificial intelligence (AI). In the interest of brevity, these comments will address just one of the topics raised in the Office of Science and Technology Policy's request for information: (4) the social and economic implications of AI.

# Social Implications

AI is unlikely to herald the end times. It is not clear at this point whether a runaway malevolent AI, for example, is a real-world possibility. In the absence of *any* quantifiable risk along these lines government officials should refrain from framing discussions of AI in alarming terms that suggest that there *is* a known, rather than entirely speculative, risk.[1] Fanciful doomsday scenarios belong in science fiction novels and high-school debate clubs, not in serious policy discussions about an existing, mundane, and beneficial technology. Ours is already "a world filled with narrowly-tailored artificial intelligence that no one recognizes. As the computer scientist John McCarthy once said: 'As soon as it works, no one calls it AI anymore.'"[2]

The beneficial consequences of advanced AI are on the horizon and potentially profound. A sampling of these possible benefits include: improved diagnostics and screening for autism;[3] disease prevention through genomic pattern recognition;[4] bridging the genotype-phenotype divide in genetics, allowing scientists to glean a clearer picture of the relationship between genetics and disease, which could introduce a wave of more effective personalized medical care;[5] the development of new ways for the sight- and hearing-impaired to experience sight and sound.[6] To be sure, many of these developments raise certain practical, safety, and ethical concerns. But there are already serious efforts underway by the private ventures developing these AI applications to anticipate and responsibly address these, as well as more speculative, concerns.

---

[1] On the distinction between quantifiable risk and uncertainty, or unquantifiable risk, see Frank H. Knight, *Risk, Uncertainty, and Profit, Houghton-Mifflin,* 1921, http://www.econlib.org/library/Knight/knRUP.html.

[2] Ryan Hagemann, "The Decentralization, Distribution, and Democratization of Artificial Intelligence," Niskanen Center, June 6, 2016, https://niskanencenter.org/blog/the-decentralization-distribution-and-democratization-of-artificial-intelligence/.

[3] Daniel Bone, Somer Bishop, Matthew P. Black, Matthew S. Goodwin, Catherine Lord, Shrikanth S. Narayanan, "Use of machine learning to improve autism screening and diagnostic instruments: effectiveness, efficiency, and multi-instrument fusion," *Journal of Child Psychology and Psychiatry*, 2016; DOI: 10.1111/jcpp.12559.

[4] David Cyranoski, "Exclusive: Genomics pioneer Jung Wang on his new AI venture," *Nature*, July 28, 2015, http://www.nature.com/news/exclusive-genomics-pioneer-jun-wang-on-his-new-ai-venture-1.18091.

[5] David Beyer, "The future of machine intelligence," *O'Reilly*, March 28, 2016, https://www.oreilly.com/ideas/the-future-of-machine-intelligence/page/4/brendan-frey-deep-learning-meets-genome-biology.

[6] John Markoff, "Israeli Start-Up Gives Visually Impaired a Way to Read," *New York Times*, June 3, 2013, http://www.nytimes.com/2013/06/04/science/israeli-start-up-gives-visually-impaired-a-way-to-read.html.

Consider OpenAI, "a non-profit artificial intelligence research company." OpenAI's goal "is to advance digital intelligence in the way that is most likely to benefit humanity as a whole, unconstrained by a need to generate financial return."[7] AI researchers are already thinking deeply and carefully about AI decision-making mechanisms in technologies like driverless cars, despite the fact that many of the most serious concerns about how autonomous AI agents make value-based choices are likely many decades out.[8] Efforts like these showcase how the private sector and leading technology entrepreneurs are ahead of the curve when it comes to thinking about some of the more serious implications of developing true artificial general intelligence (AGI) and artificial superintelligence (ASI). It is important to note, however, that true AGI or ASI are unlikely to materialize in the near-term, and the mere possibility of their development should not blind policymakers to the many ways in which artificial narrow intelligence (ANI) has already improved the lives of countless individuals the world over. Virtual personal assistants, such as Siri and Cortana, or advanced search algorithms, such as Google's search engine, are good examples of already useful applications of narrow AI.

## Economic Implications

The extent to which AI's may "disrupt" labor markets is difficult to measure. It is clear that as AI becomes more advanced, it will result in the increased automation of work. This trend may or may not result in mass job dislocation. However, some low-skilled jobs are clearly vulnerable to automation and improvements in AI technologies will certainly result in the loss of some of these jobs. It's important to recognize that AI is like many, many other technological developments that have led to the replacement of labor by machines. What's new is the kinds of jobs AI will allow to be automated. The negative impact for certain workers in certain fields should not blind us to the likely benefits of increased productivity in terms of economic performance and job-creation elsewhere in the economy. Government policies that both promote economic growth and help dislocated workers with unemployment insurance, retraining, and other forms of public assistance can facilitate disruptive innovation while protecting the welfare of those most likely to lose jobs to AI technology. If policymakers get these policies right, advanced AI and increasing automation will help bring about rising, broad-based prosperity.[9]

Policies to ameliorate negative consequences of increased automation in the economy must be informed by empirical research. Some researchers have suggested that traditional measurements, such as gross domestic product per capita, may not accurately capture the true scope of the costs and

[7] OpenAI, About, https://openai.com/about/.

[8] Nate Soares and Benya Fallenstein, "Agent Foundations for Aligning Machine Intelligence with Human Interests: A Technical Research Agenda," *Machine Intelligence Research Institute*, technical report 2014-08, 2014, https://intelligence.org/files/TechnicalAgenda.pdf.

[9] Particular policy recommendations can run the gamut, such as: (1) investing in technologies that reduce the friction and transaction costs associated with properly "matching" workers to ideal employment opportunities, (2) supporting expanded funding for research and development initiatives, and (3) taking advantage of more targeted pigovian taxes made possible by better data analytics (i.e. taxing traffic congestion or taking advantage of smart metering for dynamic utility-usage pricing) that can help fund a more robust social safety net for workers who experience the most acute effects of labor dislocation resulting from automation. For more detailed recommendations, we suggest: Erik Brynjolfsson and Andrew McAfee, The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies, (W.W. Norton & Company: New York, NY), 2014, pgs. 188-247.

benefits of AI.[10] As such, further research assessing more appropriate metrics for quantifying the effects of AI and related automation will be needed in order to clarify policymakers' options for dealing with the negative implications of continued advances in the technology.

The Future of Life Institute has observed that "our civilization will flourish as long as we win the race between the growing power of technology and the wisdom with which we manage it. In the case of AI technology … the best way to win that race is not to impede the former, but to accelerate the latter, by supporting AI safety research."[11] Government can play a positive and productive role in ensuring the best economic outcomes from developments in AI by promoting consumer education initiatives. By working with private sector developers, academics, and nonprofit policy specialists government agencies can remain constructively engaged in the AI dialogue, while not endangering ongoing developments in this technology.

## General Policy Recommendations

*Recommendation #1 (social): Because doomsday scenarios overstate the known risks of AI, official discussion of AI policy should be conducted in measured and moderate terms, and focus on actual or predictable risks of existing or emerging technology rather than on unfettered speculation about unknowable future developments.*

Many of the worst-case scenarios associated with AI are fueled by hyperbolic references to the potential for a Terminator-style apocalypse. But existing AI is in an early, almost childlike stage of development. Current AI technology doesn't remotely approach the level of sophistication that would merit the level of concern some critics have encouraged. It is encouraging that the White House recognized this in its original announcement of an interagency working group, which formed the foundation for the Office of Science and Technology Policy's call for comments. Ed Felten, the Deputy U.S. Chief Technology Officer, pointed out that current "AI is confined to narrow, specific tasks, and isn't anything like the general, adaptable intelligence that humans exhibit."[12] If left unchallenged, exaggerated worries that implicitly misrepresent the nature of current and near-term AI capabilities could impede development in this nascent field of science and technology.

In a recent report on AI, Robert Atkinson, the president of the Information Technology and Innovation Foundation, put it best:

> *Making sure that societies receive the full economic and social benefits that AI has to offer first and foremost requires accelerating, rather than restricting the technology's development and adoption. And that in turn requires that policymakers resist an AI*

[10] Stuart Russell, Daniel Dewey, Max Tegmark, "Research Priorities for Robust and Beneficial Artificial Intelligence," Association for the Advancement of Artificial Intelligence, *AI Magazine*, Winter 2015, http://futureoflife.org/data/documents/research_priorities.pdf.

[11] Benefits and Risks of Artificial Intelligence, Future of Life Institute, http://futureoflife.org/background/benefits-risks-of-artificial-intelligence/.

[12] Ed Felten, "Preparing for the Future of Artificial Intelligence," White House blog, May 3, 2016, https://www.whitehouse.gov/blog/2016/05/03/preparing-future-artificial-intelligence.

*techno-panic; they must instead embrace future possibilities with optimism and hope.*[13]

Avoiding apocalyptic rhetoric will help ensure a reasonable and practical policy discussion about AI. Policymakers and regulators who give in to the temptation to dabble publicly in speculation about cinematic worst-case scenarios invite reckless and counterproductive regulation unmoored from realistic cost-benefit analyses. Government agencies and lawmakers would do well to avoid discussing AI in hyperbolic terms, lest we delay or altogether lose out on the many great benefits AI can offer.

*Recommendation #2 (economic): AI's full potential can only be actualized if government embraces a policy of regulatory restraint.*

Private stakeholders are well-poised to explore and manage the costs and risks associated with ongoing developments in AI. The government can be an effective partner and collaborator, but regulators should stand down for the time being. Regulating too early, or on the basis of knee-jerk reactions to merely hypothetical doomsday scenarios, will hinder technological progress and innovation. Restraint and realism are especially important to encourage ongoing private capital investments in AI research and development.[14]

The general regulatory framework the Niskanen Center recommended in response to the National Telecommunications Information Administration's request for comments on the Internet of Things (IoT) should also be applied to the field of AI.[15] Indeed, the IoT is a nexus of developments in AI, big data collection and analysis, and robotics and automation. Because all these technologies are interrelated, a lack of regulatory forbearance in one area will have negative consequences that reverberate through the entire emerging technology ecosystem.

## Conclusion

As AI research and development continues, regulators and policymakers must remain realistic about the nature and size of potential costs and weigh them responsibly against actual and probable benefits. Speaking of developments in AI in apocalyptic and eschatological terms distracts from the real and important issues facing this nascent field. The benefits of narrowly-tailored AI can already be seen all around us, and much greater benefits are on the horizon. Meanwhile, it is not presently clear whether the AI technology that might lead to a doomsday scenario is even possible. For now, these

---

[13] Robert D. 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?mc_cid=54d778623c&mc_eid=77888594ca.

[14] Tom Simonite, "Inside Facebook's Artificial Intelligence Engine Room," *MIT Technology Review*, July 13, 2016, https://www.technologyreview.com/s/601882/inside-facebooks-artificial-intelligence-engine-room/?utm_campaign=add_this&utm_source=twitter&utm_medium=post.

[15] Ryan Hagemann, *comments submitted to the National Telecommunications Information Administration in the matter of: The Benefits, Challenges, and Roles for the Government in Fostering the Advancement of the Internet of Things*, Docket No. 160331306-6306-01, submitted May 23, 2016, https://www.ntia.doc.gov/files/ntia/publications/niskanencenter_ntia_iot_comments.pdf. (In particular, we would refer OSTP to the "Framework for Global Electronic Commerce" in the annex: pgs. 13-36. Although this framework was promulgated as the Clinton Administration's official policy towards the commercial Internet, the theme of regulatory forbearance and humility in the face of decentralized innovation can aptly apply to policies related to the development of AI and machine learning.)

scenarios should be approached with an air of dismissive skepticism. It is often possible to *imagine* catastrophic consequences of new technologies. But it's neither rational nor responsible to take nightmares about costs into account alongside real benefits that have already begun to accrue.

As AI develops, government can be a valuable ally in promoting engagement between researchers and academics, the private sector, government agencies, and civil society organizations. However, this engagement should avoid conjuring the specter of legislation or regulatory action that may hinder the important work being done in this field. Unless a clear need for intervention can be established with a cost-benefit analysis that balances *real* harms against *real* benefits, regulators and policymakers should remain on the sidelines and keep watch over ongoing developments.

We thank the Office of Science and Technology Policy for the opportunity to provide these comments.