Game Changers: Artificial Intelligence Part III

House Oversight Committee, Subcommittee on IT

Prepared Testimony and Statement for the Record of Ben Buchanan

Postdoctoral Fellow, Belfer Center Cybersecurity Project
Harvard University

Thank you, Chairman Hurd and Ranking Member Kelly, for holding this important hearing and for inviting me to testify.

My name is Ben Buchanan. I am a fellow at Harvard University's Belfer Center for Science and International Affairs, and a Global Fellow at the Woodrow Wilson International Center for Scholars. My research specialty is examining how nations deploy technology, and I especially focus on cybersecurity and artificial intelligence. Recently, with Taylor Miller of the Icahn School of Medicine at Mount Sinai, I co-authored a paper entitled "Machine Learning for Policymakers."[1]

To help open today's hearing on artificial intelligence, I'd like to make three points: one on privacy, one on cybersecurity, and one on economic impact.

First, to simplify a bit, we can think of most modern artificial intelligence systems as relying on a triad of pillars: data, computing power, and learning algorithms. While we have seen remarkable advances in computer hardware and machine learning software, for many policy purposes it is the role of data that is most vital to understand. Data is the fuel of machine learning systems; without it, these systems produce embarrassingly poor results. Gathering relevant and representative data for training, development, and testing purposes is a key part of building modern artificial intelligence technology. On balance, the more data that is fed into a machine learning system, the more effective it will be.[2] It is no exaggeration to say that there are probably many economic, scientific, and technological breakthroughs that have not yet occurred because the right data sets have not yet been assembled.

However, there is a catch, and a substantial one: much of the data that might—and I emphasize *might*—be useful for future machine learning systems is intensely personal, revealing, and appropriately private. Too frequently, the allure of gathering more data to feed a machine learning system distracts from the harms that collecting that data brings. There is the risk of breaches by hackers, of misuse by those who collect it or access it, and of secondary use—in which data collected for one purpose is later re-appropriated for another. Frequently, attempts at anonymization do not work nearly as well as promised. It suffices to say that, in my view, any company or government agency collecting large amounts of data is assuming an enormous

---

[1] Buchanan, Ben and Taylor Miller. "Machine Learning for Policymakers." *Belfer Center for Science and International Affairs* (2017), https://www.belfercenter.org/sites/default/files/files/publication/MachineLearningforPolicymakers.pdf.

[2] Michele Banko and Eric Brill, 'Scaling to Very Very Large Corpora for Natural Language Disambiguation' (paper presented at 'Proceedings of the 39th Annual Meeting on Association for Computational Linguistics', 2001).

responsibility. Too often, these collectors fall far short of meeting that responsibility. And yet, in an era of increased artificial intelligence, the incentive to collect ever more data is only going to grow.

Technology cannot replace policy, but some important innovations offer some mitigation to this problem. Technical approaches such as differential privacy ensure that the particular data of any one individual is obscured but that large data sets retain almost all of their value. On-device processing reduces the amount of data transmitted back to central servers and makes interception and aggregation of private information harder. But these technological advances are too infrequently deployed, especially when they conflict with short-term financial interests. This is an area in which much remains to be done.

Second, AI is poised to make a significant impact in cybersecurity, potentially redefining key parts of the industry. Automation on offense and defense is an area of enormous significance. The most high-profile example of this is the DARPA Grand Cyber Challenge, performed live at the DEF CON hacking conference in 2016, in which automated computer systems played both offense and defense against one another in a hacking competition. In the long run, it is uncertain whether increased automation will give a decisive cybersecurity advantage to hackers or to defenders, but there is no doubt of its immediate relevance.[3]

AI systems also pose new kinds of cybersecurity challenges. Most significant among these is the field known as adversarial learning, in which the learning mechanisms of algorithms can be misled. This can cause AI systems to make bizarre and unpredictable decisions. The more powerful the system, the potent such a mistake can be. Cybersecurity is challenging enough to begin with; adding in modern AI technology such as deep learning only makes it more so. Research in the world of AI system security is fairly early-stage, especially compared to the much more developed body of mainline cybersecurity research and best practices.

A more general security concern is AI safety. AI safety is the field of research and development that ensures that AI systems, once deployed, remained aligned with the original interests of their designers and do not pose unanticipated threats. This is not a question of Terminator scenarios. Rather it is usually far subtler, but vitally important and too frequently neglected.[4] I think it is fair to say we have barely scratched the surface of the important safety and basic security research that can be done in AI, and that the United States should be a leader in these areas.

Third, AI will have significant economic effects. Some of these, to be sure, will be positive. On the other hand, some will be quite negative. The drumbeat of beneficial technological progress has always brought some upheaval with it. Nonetheless, my concern in this instance is that the further development and integration of AI will occur at such a rapid rate that its economic impacts might be hard to anticipate and counteract where appropriate. Some research already clearly supports this view: one major Oxford study cited in a landmark White House report suggests that 47% of jobs

---

[3] Schneier, Bruce. "Artificial Intelligence and the Attack/Defense Balance." *IEEE Security and Privacy* (2018) https://www.schneier.com/essays/archives/2018/03/artificial_intellige.html.

[4] For one of the leading works in this area, see Amodei, Dario, et. al. "Concrete problems in AI safety." *arXiv preprint:1606.06565* (2016). For further development of these ideas and others, see "Example Topics: AI Alignment." *Open Philanthropy* (2017). https://www.openphilanthropy.org/focus/global-catastrophic-risks/potential-risks-advanced-artificial-intelligence/open-philanthropy-project-ai-fellows-program#examples.

could be threatened by machines, while another McKinsey analysis places the number at 30%.[5] Other studies bear out a worrying trend: even when economic theory predict jobs lost to automation would be replaced, later empirical analysis suggests that this replacement did not occur in practice.[6]

We must make sure that our current and future workforce is competitive in an age in which AI will be ubiquitous, when it is as broadly integrated throughout our society as electricity, according to Stanford professor and leading AI researcher Andrew Ng. We should recognize that many American students and workers will need to have an understanding of computer science, statistics, and machine learning in order to be most competitive in the global economy; too often these subjects are not taught in our schools, or not taught well and with appropriate resources. Innovations like Massive Open Online Courses, or MOOCs, have been immensely important in helping Americans acquire these valuable skills, but government and formal educational settings have important roles to play as well. Other nations, such as China, have begun to dramatically invest in these areas of education and research, and we must do the same.

Simply put, AI is exciting, economically vital, and geopolitically important. Maximizing its potential will require a whole of society effort. I appreciate your efforts in holding this series of hearings, and I look forward to your questions.

---

[5] Frey, Carl and Michael Osborne, "The Future of Employment: How Susceptible Are Jobs to Computerization," Oxford University (2013) (http://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf. See also, Manyika, James, et. al. "Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation." *McKinsey Global Institute* (2017), https://www.mckinsey.com/~/media/McKinsey/Global%20Themes/Future%20of%20Organizations/What%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/MGI-Jobs-Lost-Jobs-Gained-Report-December-6-2017.ashx.

[6] Acemoglu, Daron, and Pascual Restrepo. "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment", *NBER* (2016), https://www.nber.org/papers/w22252.pdf. Acemoglu, Daron, and Pascual Restrepo. "'Robots and Jobs: Evidence from US Labor Markets." *NBER* (2017) https://www.nber.org/papers/w23285.