Preparing for the Future: An Assessment of Emerging Cyber Threats

House Homeland Security Committee, Subcommittee on Cybersecurity, Infrastructure Protection, and Innovation

Prepared Testimony of
Ben Buchanan

Assistant Teaching Professor, School of Foreign Service
Senior Faculty Fellow, Center for Security and Emerging Technology
Georgetown University

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

My name is Ben Buchanan. I am an Assistant Teaching Professor at the School of Foreign Service and a Senior Faculty Fellow at the Center for Security and Emerging Technology, both at Georgetown University. I am also a Global Fellow at the Woodrow Wilson International Center for Scholars, where I teach introductory classes on artificial intelligence and cybersecurity for congressional staff. My research specialty is examining how cybersecurity and AI shape international security. --I co-authored a paper entitled "Machine Learning for Policymakers."[1]

I will confine my opening remarks to the impact of artificial intelligence on cybersecurity, since I think it is the emerging technology poised to have the most significant effect in this area. While there is an enormous amount of hype and debate around AI in general, the intersection of AI and cybersecurity is understudied and underappreciated.

At least three dimensions of this problem deserve analysis:

First and most significant is the cybersecurity of AI systems themselves. AI systems are just as likely to be susceptible to the kinds of software vulnerabilities that are present in other kinds of computer code. As we have seen for decades, hackers can exploit these vulnerabilities for their own ends. There is no reason to think that hackers will not try to do the same to AI systems, and there is no reason to think that they will not at times succeed. This possibility is particularly worrying given the high stakes of some AI applications; it is not a reason to avoid using AI, but vigilance is imperative to preserve cybersecurity.

---

[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.

But to stop our analysis at just the traditional kinds of software vulnerabilities is to miss a great deal of the cybersecurity risk that AI systems pose. The neural network architecture that underpins a lot of modern AI is immensely powerful but presents a new class of cybersecurity risks that we are only beginning to uncover. We call this field adversarial learning.

Using adversarial learning, hackers can cause neural networks to make bizarre errors, causing systems that rely on those networks to fail or to reveal confidential information. This is a field that requires a great deal more attention.

Second, AI can also change traditional offensive cyber attacks against regular computer systems. Modern hackers in many cases do not need artificial intelligence to achieve their ends. That said, I think it is noteworthy that some of the most potent cyber attacks we have seen--including Stuxnet, the 2016 blackout in Ukraine, and the 2017 attack known as NotPetya that caused at least ten billion dollars in damage--feature some forms of automated propagation and attack capability. I can imagine a world in which future cyber operations will use more sophisticated automated capabilities to achieve particular tasks, such as vulnerability discovery, target selection, command and control, and attack execution.

I suspect that such automation could offer significant upsides to sophisticated hackers faced with complex targets. In some respects, the possible upside to automation is higher in this area than in physical warfare; whether a plane is operated by a person or a human, the laws of physics still apply, but it is likely that automated cyber capabilities--if sophisticated enough--could operate much faster than their human-directed counterparts. I stress, however, that we have not seen this come to fruition yet.

This leads to the third area of analysis: the possibility that AI might help on cyber defense. This idea is also the subject of a lot of hype and a lot of venture capital investment. There seem to be discrete ways in which AI can indeed help secure computer systems, both in discovering vulnerabilities

before hackers do and also in detecting the presence of malicious code. However, we must be careful not to let the hype outrun the reality on this front. In evaluating cybersecurity advances in this area, we should be careful to compare them to the baseline of technologies we already use-- many of which already involve automation--and understand how, if at all, artificial intelligence improves our defenses.

I do believe that AI-enabled tools are likely to be a fundamental part of modern and future cyber defense; the scale, size, and speed of cyber operations will make this inevitable, and it is imperative that we develop these tools. That said, we must not forget that cyber operations, no matter how sophisticated, are still fundamentally human operations. For as much as we will talk about technology today, we must remember that the people in our organizations are key to addressing these threats.

I look forward to your questions.