Statement of Greg Brockman
Co-Founder and Chief Technology Officer
OpenAI

before the

Subcommittee on Space, Science, and Competitiveness
Committee on Commerce, Science, and Transportation
United States Senate

Hearing on "The Dawn of Artificial Intelligence"

Thank you Chairman Cruz, Ranking Member Peters, distinguished members of the Subcommittee. Today's hearing presents an important first opportunity for the members of the Senate to understand and analyze the potential impacts of artificial intelligence on our Nation and the world, and to refine thinking on the best ways in which the U.S. Government might approach AI. I'm honored to have been invited to give this testimony today.

By way of introduction, I'm Greg Brockman, co-founder and Chief Technology Officer of OpenAI. OpenAI is a non-profit AI research company. Our mission is to build safe, advanced AI technology and ensure that its benefits are distributed to everyone. OpenAI is chaired by technology executives Sam Altman and Elon Musk.

The US has led the way in almost all technological breakthroughs of the last hundred years, and we've reaped enormous economic rewards as a result. Currently, we have a lead, but hardly a monopoly, in AI. For instance, this year Chinese teams won the top categories in a Stanford University-led image recognition competition. South Korea has declared a billion dollar AI fund. Canada produced some technologies enabling the current boom, and recently announced an investment into key areas of AI.

I'd like to share 3 key points for how we can best succeed in AI and what the U.S. government might do to advance this agenda. First, we need to compete on applications, but cooperate on open, basic research. Second, we need to create public measurement and contests. And third, we need to increase coordination between industry and government on safety, security, and ethics.

I. Competition and Cooperation

AI applications are rapidly broadening from what they were just a few years ago: from helping farmers decide which fields to seed, to warehouse robots, to medical diagnostics, certain AI-enabled applications are penetrating and enabling businesses and improving everyday life. These and other applications will create new companies and new jobs that don't exist today — in much the same way that the Internet did. But even discovering the full range of applications requires

significant scientific advances. So industry is not just working on applications: companies like Facebook, Google, and Microsoft are performing basic research as well, trying to create the essential AI building blocks which can later be assembled into products.

Perhaps surprisingly, the industry labs are publishing everything they discover. Publication allows them to pool their resources to create faster breakthroughs, and to attract top scientists, most of whom are motivated more by advancing society and improving the future, than personal financial gain.

Companies stay competitive by publishing their basic research, but not the details of their products. The inventor of a technique is usually the first to deploy it, as it has the right in-house infrastructure and expertise. For example, AI techniques developed by Google's subsidiary DeepMind to solve Atari video games were applied to increase the efficiency of Google's own data centers. DeepMind shared their basic techniques by publishing the Atari research papers, but did not share their applied work on data center efficiency.

Openness enables academia and industry to reinforce each other. Andrew Moore of Carnegie Mellon University says it's not unusual that between 10 and 20 percent of the staff he hires will take leaves of absence to work in industry or found a startup. Pieter Abbeel, a researcher at OpenAI, splits his time between OpenAI and the University of California at Berkeley; likewise, Stanford Professor Fei-Fei Li is spending time at both Stanford and Google; and many other companies and organizations work with academics. This ensures that the private sector is able to master the latest scientific techniques, and that universities are able to understand the problems relevant for industry.

Openness has concentrated the world's AI research activity around the US (including attracting many of the Canadian scientists who helped start the current AI boom), and allowed us to define its culture and values. Foreign firms like China's Baidu have opened US-based research labs and have also started publishing. As AI becomes increasingly useful, the pool of experts we're gathering will be invaluable to ensuring that its economic activity also remains centered on the US.

Recommendations -

We recommend the following, to ensure that our basic AI research community remains the strongest in the world:

A. Maintain or increase basic research funding for AI: In 2015 the government's unclassified investment in AI-related technology was approximately $1.1 billion, according to The National Artificial Intelligence Research and Development Strategic Plan report from the National Science and Technology Council [1]. As highlighted by Jason Furman, Chairman of the Council of Economic Advisers, there's evidence that the socially optimal level of funding for basic research is two to four times greater than actual spending[2]. Given that it only takes months for

a basic AI advance to result in new companies and products, usually by whoever made the advance, we support increasing funding for basic research in this domain. If we want these breakthroughs to be made in the US, we'll need to conduct basic research across a number of subfields of AI, and encourage the community to share their insights with each other. We'll need to allow our academics to freely explore ideas that go against consensus, or whose value has high uncertainty. This is supported by history: companies like Google and Microsoft rely on AI technologies that originated with a small group of maverick academics.

B. Increase the supply of AI academics: Industry has an insatiable demand for people with AI training, which will only increase for the foreseeable future. We need to grow the supply of people trained in AI techniques; this will let us make more research breakthroughs, give industry the people it needs to commercialize the basic science, and train the next generation of scientists. NSF could explore adjusting its policies to allow more competitive salaries for those working on federal academic grants.

C. Enhance the professional diversity of the AI field: Today, AI consists mostly of individuals with degrees in computer science, mathematics, and neuroscience, with a significant gender bias towards men. As AI increases its societal impact, we need to increase the diversity of professional views within the AI community. Government can explore making more interdisciplinary research grants available to incentivize experts in other fields, such as law or agriculture or philosophy, to work with AI researchers. We also support the White House's Computer Science for All initiative, and the OSTP's recommendation that government should create a federal workforce with diverse perspectives on AI.

II. The Need For Public Measurement and Contests

Objective measures of progress help government and the public distinguish real progress from hype. It's very easy to sensationalize AI research, but we should remember that advanced AI has seemed just around the corner for decades. Good policy responses and a healthy public debate hinge on people having access to clear data about which parts of the technology are progressing, and how quickly. Given that some AI technologies, such as self-driving cars, have the potential to impact society in a number of significant ways, we support OSTP's recommendation that the government keep a close watch on the advancement of specific AI technologies, and work with industry to measure the progression of the technology.

Also, having a measurable goal for AI technologies helps researchers select which problems to solve. In 2004, DARPA hosted a self-driving car competition along a 150-mile course in the Mojave Desert — the top competitor made it only seven miles. By 2007, DARPA hosted an Urban Challenge to test self-driving cars on a complex, urban environment, and six of the eleven teams completed the course. Today, Uber, Google, Tesla, and others are working on commercializing self-driving car technology.

Similarly, when Fei-Fei Li and her collaborators at Stanford launched the image recognition ImageNet competition in 2010, it was designed to be beyond the capabilities of existing systems. That impossibility gave the world's research community an incentive to develop techniques at the very edge of possibility. In 2012, academics won first place using a neural network-based approach, which proved the value of the technique and kickstarted the current AI boom. The winning ImageNet team formed a startup and were subsequently hired by industry to create new products. One member, Ilya Sutskever, is one of my co-founders at OpenAI, and the other two members work at Google. This shows how competitions can provoke research breakthroughs, and translate into an economic advantage for industry.

We're moving from an era of narrow AI systems to general ones. Narrow AI systems typically do one thing extremely well, like categorize an image, transcribe a speech, or master a computer game. General AI systems will contain suites of different capabilities; they will be able to solve many tasks and improvise new solutions when they run into trouble. They will require new ways to test and benchmark their performance. Measuring the capabilities of these new multi-purpose systems will help government track the technology's progress and respond accordingly.

Recommendations -

Government can create objective data about AI progress in the following ways:

A. Modern competitions: AI systems have often been measured by performance on a static dataset. Modern systems will act in the real world, and their actions will influence their surroundings, so static datasets are a poor way to measure performance. We need competitions which capture more of the complexity of the real world, particularly in developing areas such as robotics, personal assistants, and language understanding. The government can continue designing competitions itself, as DARPA did recently with the Cyber Grand Challenge, or support others who are doing so.

B. Government information gathering: Government should gather information about the AI field as a whole. Researchers tend to focus on advancing the state of the art in one area, but the bigger picture is likely to be crucial for policymakers, and valuable to researchers as well. The government can invest in careful monitoring of the state of the field, forecasting its progress, and predicting the onset of significant AI applications.

III. Increase coordination between industry and government on safety, security, and ethics

The Internet was built with security as an afterthought, rather than a core principle. We're still paying the cost for that today, with companies such as Target being hacked due to using insecure communication protocols. With AI, we should consider safety, security, and ethics as early as possible, and bake these into the technologies we develop.

Academic and industrial participants are starting to coordinate on responsible development of AI. For example, we recently worked with researchers from Stanford, Berkeley, and Google to lay out a roadmap for safety research in our paper "Concrete Problems in AI Safety"[3]. Non-profit groups like the Partnership on AI and OpenAI are forming to ensure that research is done responsibly and beneficially.

Recommendations -

Industry dialog: Government can help the AI community by giving feedback about the what aspects of progress it needs to understand in preparing policy. As the OSTP recommended in its report, Preparing for the future of Artificial Intelligence[4], the NSTC Subcommittee on Machine Learning and Artificial Intelligence should meet with industry participants to track the progression of AI. OpenAI and our peers can use these meetings to understand what we should monitor in our own work to give government the telemetry needed to calibrate policy responses.

Accenture recently reported that AI has the potential to double economic growth rates by 2035, which would make it the engine for our future economy. Having the most powerful economy in the world will eventually require having the most AI-driven one, and the US accordingly must lead the development and application of AI technologies along the way. The best way to ensure a good future is to invent it.

Thank you for your time and focus on this critical topic. I am pleased to address any questions.

[1] National Science and Technology Council, Networking and Information Technology Research and Development Subcommittee. 2016. "The National Artificial Intelligence Research and Development Strategic Plan" report: https://www.whitehouse.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf
[2] Furman, Jason. 2016. "Is This Time Different? The Opportunities and Challenges of Artificial Intelligence" report: https://www.whitehouse.gov/sites/default/files/page/files/20160707_cea_ai_furman.pdf
[3] Amodei, Dario et al. 2016. "Concrete Problems in AI Safety" research paper: https://arxiv.org/abs/1606.06565
[4] Executive Office of the President, National Science and Technology Council Committee on Technology. 2016. "Preparing for the future of artificial intelligence" report: https://www.whitehouse.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf