House Committee on Oversight and Government Reform
Subcommittee on Information Technology
February 14, 2018

"Game Changers: Artificial Intelligence Part I"

Witness: Charles Isbell, Georgia Institute of Technology

Chairman Hurd, Ranking Member Kelly, and distinguished members of the subcommittee, my name is Dr. Charles Isbell and I am a Professor and Executive Associate Dean for the College of Computing at Georgia Tech. Thank you for the opportunity to appear before this Subcommittee to discuss the development, uses, barriers to adoption, and potential challenges and advantages of government use of artificial intelligence.

By way of explaining my background, let me note that while I tend to focus on statistical machine learning, my research passion is actually artificial intelligence. I like to build large integrated systems, so I also tend to spend a great deal of my time doing research on autonomous agents, interactive entertainment, some aspects of human-computer interaction, software engineering, and even programming languages

I think of my field as interactive artificial intelligence. My fundamental research goal is to understand how to build autonomous agents that must live and interact with large numbers of other intelligent agents, some of whom may be human. Progress towards this goal means that we can build artificial systems that work with humans to accomplish tasks more effectively; can respond more robustly to changes in environment, relationships, and goals; and can better co-exist with humans as long-lived partners.

As requested by the Subcommittee, my testimony today will focus on the potential for artificial intelligence and machine learning to transform the world around us. I will:

1. Situate recent developments in AI in the larger context of developments in computing more generally;
2. Explore the potential uses and misuses of this technology by focusing on the human-machine loop; and
3. Discuss the gaps in education and training that threaten to minimize participation in the field.

As the members of this Subcommittee well know, there has been an explosion in the development and deployment of what we might call AI technology. With that explosion has come a corresponding explosion in interest in AI.

In any discussion—particularly technical ones—it helps to define our terms. There are many potential definitions of AI. My favorite one is that it is "the art and science of making computers act like they do in the movies." In the movies, computers are often semi-magical and anthropomorphic; they do things that, if humans did them, we would say they required intelligence. This definition is borne out in our use of AI in the everyday world. We use the infrastructure of AI to search billions upon billions of documents to find the answers to a staggering variety of questions—often expressed literally as questions. We use automatically tagged images to organize our photos, and we use that same infrastructure to plan optimal routes for trips—even altering our routes on-the-fly in the face of changes in traffic. We are able to detect automatically the presence of tumors from x-rays, even those trained doctors find difficult to see. We let computers finish our sentences as we type texts and use search engines, sometimes facilitating a subtle shift from prediction of our behavior to influence over our behavior. Often we take advantage of these services by using our phones (our phones!) to interpret a wide variety of spoken commands.

This basic definition, of course, ignores what is going on underneath the hood. Perhaps a somewhat better way of grappling with AI is to understand that it is a set of computing techniques and approaches that attempt to solve exponentially hard problems in reasonable time by cheating in clever ways. In other words, at bottom, AI is about *computing methods* for automated understanding and reasoning, especially ones that *leverage data* to adapt their behavior over time.

That AI is really about computing is important. What has enabled many of the advances in AI over the last decade is the stunning increase of computational power combined with the ubiquity of that computing. That AI also leverages data is equally important. The advances in AI over the last decade are also due in large part to the even more stunning increase in the availability of data, again made possible by the ubiquity of the internet, social media, and relatively inexpensive sensors (including cameras, GPS, and the computer itself) that track our every move.

Consider the problems above: Google leverages arrays of server farms to index and search documents now available digitally; neural networks use millions of examples of pictures of human faces to perform the hundreds of millions of calculations necessary to do face-tagging in the same way that we can do phoneme and word detection from audio; our navigation apps like Waze use both the digital expression of maps to sort through millions or even billions of possible paths from one place to another, as well as the ubiquity of GPS in other vehicles to detect changes in traffic; medical prediction software can build tumor detectors by leveraging decades of data on x-rays and ground-truth labels of cancer; and the same techniques are used to crowd-source likely completions to search queries.

Consider the technology behind them: Deep learning is an update on an algorithm whose modern expression was known about the time of my birth. It uses layers of artificial "neurons" to map from a set of features (*e.g.,* pixels, sounds, financial information, and so on) to more abstract concepts (*e.g.,* names of objects, words, credit-worthiness, and so on). As recently as twenty years ago, computational power was such that one could only build one or two layers. Performance required highly trained humans hand-tuning both network structure and the form of features themselves. Now, with both cheap, fast computing power and an abundance of data, the structure and features can also be learned, freeing computing professionals to develop better techniques that take advantage of this newfound power. Accordingly, the new systems work far better than we had available even a few years ago.

So, in some very important sense, AI already exists. It is not the AI of science fiction, neither benevolent intelligences working with humans as we traverse the galaxy, nor malevolent AI that seeks humanity's destruction. Nonetheless, we are living every day with machines that make decisions that, if humans made them, we would attribute to intelligence. And the machines often make those decisions faster and better than humans would.

Importantly, each of the examples we consider above is a distinctly human-centered problem. It is human-centered both in the sense that these systems are trying to solve problems that humans deal with every day—question answering, symptom evaluation, navigation—but also human-centered in the sense that humans have or currently perform some of those tasks. Presumably, these developments are all to the good. We are living up to the promise of technology that allows us to automate away work that is dirty, dangerous, or dull, freeing up human capital to be more productive, and, hopefully, for humans to be more fulfilled. The social and economic benefits are potentially immense.

There are also some reasons for concern. Those concerns also have social and economic components. I will focus briefly on two potential issues: the opaqueness of our intelligent machines and the potential impact on jobs.

We are increasingly using our AI systems to make real decisions, and we do not necessarily understand those decisions. In particular, these decisions can have severe impacts. For example, according to the Marshall Project, almost every state uses some form of "risk assessment" at some stage in the criminal justice system.

> Risk assessments have existed in various forms for a century, but over the past two decades, they have spread through the American justice system, driven by advances in social science. The tools try to predict recidivism — repeat offending or breaking the rules of probation or parole — using statistical probabilities based on factors such as age, employment history, and prior criminal record. They are now used at some stage of the criminal justice process in nearly every state. Many court systems use the tools to guide decisions about which prisoners to release on parole, for example, and risk as-

> sessments are becoming increasingly popular as a way to help set bail for inmates awaiting trial.

Consider the automation of this process, relying on an algorithm in lieu of a judge's discretion. As noted by Cathy O'Neil, author of *Weapons of Math Destruction*, the data used by these algorithms to build models are sometimes suspect. Worse, we treat the output as "objective" without understanding that the data are themselves not objective. In this particular case, we set out to predict recidivism as if that means *the chance of committing a crime again* when in fact we are predicting *the chance of being arrested and convicted again.* It does not take much imagination to see how being from a heavily policed area raises the chances of being arrested again, being convicted again, and in aggregate leads to even more policing of the same areas, creating a feedback loop. One can imagine similar issues with determining fit for a job, or credit-worthiness, or even face recognition and automated driving. In computing, we call this garbage-in-garbage-out: an algorithm is only as good as its data. This saying is certainly true, and especially relevant for AI algorithms that learn based on the data they are given.

Luckily, one way to address these issues is straightforward: to increase transparency. An AI algorithm should inspectable. The kind of data the algorithm uses to build its model should be available. The decisions that such algorithms make should be inspectable. In other words, as we deploy these algorithms, each algorithm should be able to explain its output. "This applicant was assigned high risk because…" is more useful than, "This applicant was assigned high risk."

Of course, as we make our AI better and easier to understand, it is difficult not to imagine that AI will do more and more for us. In today's climate, we are imagining not only robots that assemble our cars, but that those cars will drive themselves. We can see a world where we will not only have algorithms that allow us to watch the stock market but will do a faster, better job buying and selling stocks than stockbrokers do. We may soon trust the x-ray machine itself to tell us if we have a tumor more than we trust a doctor. I am skeptical that we will create such AI machines in the near future, but it does seem that we are making inexorable progress toward that end. We may not replace all truck drivers and taxi cab drivers, but we may replace many of them. We may not replace all cashiers, but we may replace many of them. In a country where there are nearly 3 million truck drivers and more than 3 million cashiers, one can imagine what a significant impact such automation will have on the economy and on the job force.

Luckily again technology and automation does not simply destroy jobs, it creates them. In this particular case, it creates jobs that require technological sophistication and understanding. Here, it is important to return to our definitions. AI is about computing methods for automated understanding and reasoning, especially ones that leverage data to adapt their behavior over time. Thus, the future belongs to those who are not simply highly literate but compurate; that is, those who understand computing and how it fits into problem solving will be most productive and impactful in the future.

We can see in the current data that our fellow citizens understand this reality. At Georgia Tech, we launched an affordable online master's degree in Computer Science four years ago. We are currently enrolling 6,365 students, 70% of whom are US citizens or permanent residents. Across the country, undergraduate computer science enrollments are at an all-time high at Research I universities, growing 113% between 2009 and 2015. From 2006 to 2015, the average number of CS majors increased for large departments (10+ faculty) from 320 to 970 and for small departments from 160 to 500 majors. The overall numbers are significantly higher than at the height of the dot-com boom.

At the same time, non-majors are increasingly taking upper-division computing courses for use in their own fields. According to Generation CS, the number of non-majors in courses intended for majors is increasing at a rate equal to or higher than that of majors. We are also seeing increasing interest in AI. For example, at Georgia Tech, 43% of our CS minors are focused on Artificial Intelligence. This year, our peers are reporting record numbers of graduate student applicants in machine learning and artificial intelligence.

Even more telling, institutions have been forced to cap the number of students who major in a program. This throttling of support suggests that demand may be even higher than it seems, but it also suggests that we are not capable of responding to this demand even as we need to educate more and more students in the area. The number of Ph.D. graduates in computer science going into higher education is dropping significantly. Further, this issue is not limited to those seeking undergraduate and advanced degrees. We are seeing an increasing need to educate students at the high school level as well and a corresponding lack of teachers available who are qualified to teach foundational computer science in K-12. Given the slow pace of production and the lack of an incentive structure for graduates in computer science to become teachers, the country will not be able to produce enough CS teachers quickly enough to meet demand.

In Georgia, for example, there are approximately 519,000 high school students. Only 29,000 of them are enrolled in computing of any kind—less than 6 percent. According to the Professional Standards Commission, the governing body over teacher certification in the state, there were only 93 credentialed teachers in 2017. The majority of the computing courses in the state are being taught by approximately 400-500 engaged and committed teachers who are not certified to do so. The state is in its nascent stages of offering a framework to guide what "high school-level CS" actually means. For now, the curricula and quality of the CS courses vary tremendously. The College Board's Advanced Placement Computer Science A exam is more formalized and demonstrates the magnitude of the problem for rigorous computing. Data from the College Board suggest that, in 2017, only 125 of the 500 high schools in the state offered AP Computer Science. In Atlanta Public Schools, which is in the heart of Georgia's technology hub, there are only two high schools that offer Advanced Placement Computer Science.

Under these circumstances, possibly the only way to deploy this subject broadly is to offer blended learning courses. The core content of computational courses will ultimately have to be delivered through online platforms in close conjunction with classroom teachers who can be present and facilitate the actual process of learning.

In conclusion, I am excited by these hearings. Advances in AI are central to our economic and social future. The issues are being raised here can be addressed with thoughtful support for robust funding in basic research in artificial intelligence—including research in how to engage in education; support for that education throughout the pipeline; and in developing standards for the proper use of intelligent systems.

I thank you very much for your time and attention today, and I look forward to working with you in your efforts to understand how we can best develop these technologies to create a future where we are partners with intelligent machines.

Thank you. This concludes my testimony.