"It's going to kill us!" and Other Myths About the Future of Artificial Intelligence

by Robert D. Atkinson, Ph.D.

Editor's note: The following article is reprinted with permission from the author and appeared originally in the ITIF Update.

Given the promise that AI holds for economic growth and societal advancement, it is critical that policymakers not only avoid retarding the progress of AI innovation, but also actively support its further development and use

The past decade has seen important advancements in computer science that enable software systems to compile and process new information to continually improve the way they function. This improved artificial intelligence is enabling computers to become an ever more powerful and valuable complement to human capabilities: improving medical diagnoses, weather prediction, supply-chain management, transportation, and even personal choices such as where to go on vacation or what styles of clothes to buy.

> Many now argue that AI will power a productivity explosion so great that it will destroy jobs faster than the economy can keep up, creating an unemployed underclass that will be dominated by an elite class of "machine owners." These are not new predictions, and they are as wrong today as they have been in years past.

Although artificial intelligence has become commonplace—most smartphones contain some version of AI, such as speech recognition—the public still has a poor understanding of the technology. As a result, a diverse cast of critics, driven by fear of technology, opportunism, or ignorance, has jumped into the intellectual vacuum to warn policymakers that, sooner than we think, AI will produce a parade of horribles: mass unemployment, abuse from "algorithmic bias," the end of privacy, an atrophying of human agency, and even the destruction of humanity as "Skynet"-like machines decide the world is better off without us. Indeed, these voices have grown so loud, espousing a message that a click- hungry media eagerly amplifies, that we are very near

the point where these narratives may be accepted as truth. Needless to say, when AI is so vociferously demonized (indeed, the engineering magnate Elon Musk has explicitly warned that AI could be "the demon" that threatens our existence, especially if actions are not taken to design systems that can remain under human control), there is a real risk that policymakers will seek to retard its progress.

This would be a terribly unfortunate outcome, because the truth is that AI systems are no different than shovels or tractors: They are tools in the service of humans, and we can use them to make our lives vastly better. Given the promise that innovation in AI holds for economic growth and societal advancement, it is critical that policymakers actively support its further development and use. The cost of not developing artificial intelligence, or developing it more slowly, would be enormous: lower growth in per-capita incomes, slower progress in areas such as health care and the environment, and reduced quality improvement in a wide array of public and private goods and services. This report provides a primer on artificial intelligence and debunks five prevailing myths that, if left unchecked, could undermine its progress. Rather than give in to fear, policymakers should be doing everything possible to accelerate the progress of AI innovation.

MYTH 1: AI WILL DESTROY MOST JOBS

Reality: AI will be like past technologies, modestly boosting productivity growth and having no effect on the overall number of jobs or unemployment rates.

Many now argue that AI will power a productivity explosion so great that it will destroy jobs faster than the economy can keep up, creating an unemployed underclass that will be dominated by an elite class of "machine owners." These are not new predictions, and they are as wrong today as they have been in years past.

The apocalyptic views that AI will kill jobs suffer from two major errors. The first is that they vastly overestimate the capabilities of AI to replace humans. It is actually quite hard for technology, AI or otherwise, to eliminate jobs, as evidenced by the fact that U.S. productivity has been growing at a historically slow pace. And it is particularly hard to automate large numbers of jobs with AI, because virtually all AI is "narrow AI," designed to focus on doing one thing really well. So, in many occupations, the introduction of AI may not lead to job loss at all; it may instead increase output, quality, and innovation.

The second reason is that even if AI were more capable, there still would be ample job opportunities, because if jobs in one firm are reduced through higher productivity, then costs go down. These savings are recycled though lower prices or higher wages. This puts more money into the economy, and the money is then spent creating jobs in whatever industries supply the goods and services that people demand as their incomes go up. This is why, historically, there has been a negative relationship between productivity and unemployment rates.

MYTH 2: AI WILL MAKE US STUPID

Reality: AI will help us make smarter decisions.

Even beyond the unfounded fear that smart machines will take our jobs, some dystopians assert that AI will turn us into helpless automatons who are bound to become overly dependent on the machines and in so doing lose our own native skills—so when the machines occasionally fail, we'll be ill-equipped to take back control. To be sure, some skills may become less necessary as AI is able to handle routine tasks that humans used to do—just as machines like the automobile made it unnecessary for most people to know how to ride a horse—but it will open up new areas of skill. And the issue is not whether these systems won't make errors; it is whether on net they will make fewer errors than human-controlled activities. The answer is yes; they will make fewer errors—otherwise they will not be used—and that will be a boon to mankind.

MYTH 3: AI WILL DESTROY OUR PRIVACY

Reality: AI will have no effect on privacy, since most information practices are bound by laws and regulations.

If smart machines can crunch massive amounts of data, then surely they will destroy our privacy. Or so AI dystopians warn us. But there are several reasons why these opponents are wrong. First, while AI systems have the ability and even the need to collect and analyze more information, the threat to privacy is little greater than in non-AI systems, which already collect and analyze large amounts of information. Moreover, the rules that already govern data use and protect privacy today will cover data analyzed by AI, too.

In short, this is basically a policy question, not a technology question. If we don't want government agencies to collect certain data, then Congress can require that and courts enforce it. Whether agencies have or do not have machine-learning systems is irrelevant. In addition, many, if not most of the benefits of AI-enabled data analysis can be obtained without the need to risk disclosing personally identifiable information.

MYTH 4: AI WILL ENABLE BIAS AND ABUSE

Reality: In most cases, AI will be less biased than humans.

Machine-learning systems are more complex than traditional software systems. It was relatively clear how the older rules-based expert systems made decisions. In contrast, machine-learning systems continuously adjust and improve based on experience. Some critics claim this level of complexity will result in "algorithmic bias" that promotes government and corporate abuse, whether unintentional or deliberate, because organizations will hide behind their algorithms and use the algorithms' complexity as a cover to justify exploitation, discrimination, or other types of unethical or damaging behavior.

It is certainly true that AI systems, like any technology, can be used unethically or irresponsibly. But those who resist AI based on this concern fail to recognize a key point: Machine-learning systems are not independent from their developers or the organizations using them. If an organization wants to systematically discriminate against certain groups, it doesn't need AI to do so. Furthermore, if an algorithmic system produces unintended and potentially discriminatory outcomes, it's not because the technology itself is malicious; it's because it simply follows instructions set by human decision-making or, more often, relies on real-world data sets that may reflect bias. Finally, in most cases these systems are less biased than human decision making, where subconscious or overt biases permeate every aspect of society.

MYTH 5: SMART MACHINES WILL TAKE OVER AND POTENTIALLY EXTERMINATE THE HUMAN RACE

Reality: We will be lucky if smart machines become smart enough to make us a sandwich.

Finally, some argue that machines will become super-intelligent and decide they are better off without humans. It's a sad commentary that the public has become so technophobic that we are even taking these sci-fi claims seriously. The view that smart machines will kill us overstates the pace of technological progress, particularly because the processing power of silicon computer chips is slowing down and progress in AI outside of deep learning is relatively modest. Moreover, machines and the human mind are completely different systems, and even major advances in computing are highly unlikely to produce a machine with humanity's intellectual capacity, imagination, or adaptability. As MIT computer scientist Rodney Brooks puts it, "We generalize from performance to competence and grossly overestimate the capabilities of machines—those of today and of the next few decades." Just as importantly, even if human-level intelligent machines could be built, which is unlikely, they will remain under the control of humans, because we would never build them unless they are largely safe, with the benefits outweighing the costs (just as we do with all technologies in the marketplace).

What Policymakers Should and Should Not Do

Making sure that societies capture the full benefits of AI requires accelerating, rather than restricting the technology's development and adoption. So policymakers should not give in to an AI techno-panic; they should instead embrace future possibilities with optimism.

We should operate on the innovation principle, not the precautionary principle. In other words, we should proceed on the assumption that AI will be fundamentally good, and while it will present some risks—as every new technology does—we should focus on addressing risks if and when they arise, rather than slowing the technology. Living by the innovation principle also means understanding that AI will involve both Type I and Type II errors, which is to say it will produce some errors, but it will also reduce or even eliminate many others.

We should of course be clear-eyed in identifying and resolving challenges as we pursue the vast opportunities AI offers. Indeed, we need open and frank discussion about the potential challenges and benefits, if only because many, if not most, of the voices dominating the AI debate are spreading techno-panic. But governments should look at AI rationally and calmly.

Finally, policymakers need to significantly increase support for research on AI development, including on making AI more powerful and effective, but also safer, more secure, and more transparent. Policymakers also should support companies and government agencies in using AI to better accomplish their tasks. In short, technological progress has been and will remain key for future progress, and AI is poised to play a key role in that progress, provided we do not give in to fear.

Robert D. Atkinson, Ph.D., is founder and president of the Information Technology and Innovation Foundation (ITIF). He is the author of Innovation Economics: The Race for Global Advantage (Yale, 2012), Supply-Side Follies: Why Conservative Economics Fails, Liberal Economics Falters, and Innovation Economics is the Answer (Rowman & Littlefield, 2006), and The Past And Future Of America's Economy: Long Waves Of Innovation That Power Cycles Of Growth (Edward Elgar, 2005).