

Artificial intelligence: How can we regulate without stifling innovation?

Some people are afraid that [heavily armed artificially intelligent robots](#) might take over the world, enslaving humanity – or [perhaps exterminating us](#). These people, including tech-industry billionaire Elon Musk and eminent physicist Stephen Hawking, say artificial intelligence technology [needs to be regulated](#) to [manage the risks](#). But Microsoft founder [Bill Gates](#) and Facebook's [Mark Zuckerberg](#) disagree, saying the technology is not nearly advanced enough for those worries to be realistic.

As someone who researches how AI works in robotic decision-making, [drones](#) and [self-driving vehicles](#), I've seen how beneficial it can be. I've developed AI software that lets robots working in teams [make individual decisions](#), as [part of collective efforts](#) to explore and solve problems. Researchers are already subject to existing rules, regulations and laws designed to protect public safety. Imposing further limitations risks reducing the potential for innovation with AI systems.

How is AI regulated now?

ai-1518-2 While the term “artificial intelligence” may conjure fantastical images of human-like robots, most people have encountered AI before. It helps us [find similar products](#) while shopping, offers [movie and TV recommendations](#) and helps us [search for websites](#). It [grades student writing](#), provides [personalized tutoring](#) and even [recognizes objects](#) carried through [airport scanners](#).

In each case, the AI makes things easier for humans. For example, the AI software I developed could be used to plan and execute a search of a field for a plant or animal as part of a science experiment. But even as the AI frees people from doing this work, it is still basing its actions on human decisions and goals about where to search and what to look for.

In areas like these and many others, AI has the potential to do far more good than harm – if used properly. But I don't believe additional regulations are currently needed. There are already laws on the books of nations, states and towns governing civil and criminal liabilities for harmful actions. Our drones, for example, must obey FAA regulations, while the self-driving car AI must obey regular traffic laws to operate on public roadways.

Existing laws also cover what happens if a robot injures or kills a person, even if the injury is accidental and the robot's programmer or operator [isn't criminally responsible](#). While lawmakers and regulators may need to refine responsibility for AI systems' actions as technology advances, creating regulations beyond those that already exist could prohibit or slow the development of capabilities that would be overwhelmingly beneficial.

Potential risks from artificial intelligence

It may seem reasonable to worry about researchers developing very advanced artificial intelligence systems that can operate [entirely outside human control](#). A common thought experiment deals with a [self-driving car forced to make a decision](#)

ai 1 5 18 3

about whether to run over a child who just stepped into the road or veer off into a guardrail, injuring the car's occupants and perhaps even those in another vehicle.

Image not found or type unknown

Musk and Hawking, among others,

worry that hypercapable AI systems, no longer limited to a single set of tasks like controlling a self-driving car, might decide it doesn't need humans anymore. It might even look at human stewardship of the planet, the interpersonal conflicts, theft, fraud and frequent wars, and [decide that the world would be better without people](#).

Science fiction author Isaac Asimov tried to address this potential by proposing [three laws](#) limiting robot decision-making: Robots cannot injure humans or allow them "to come to harm." They must also obey humans – unless this would harm humans – and protect themselves, as long as this doesn't harm humans or ignore an order.

But [Asimov himself knew](#) the three laws were not enough. And they don't reflect the complexity of human values. What constitutes "harm" is an example: Should a robot protect humanity from suffering related to overpopulation, or should it protect individuals' freedoms to make personal reproductive decisions?

We humans have already wrestled with these questions in our own, nonartificial intelligences. Researchers have proposed [restrictions on human freedoms](#), including reducing reproduction, to control people's behavior, population growth and environmental damage. In general, society has decided against using those methods, even if their goals seem reasonable. Similarly, rather than regulating what AI systems can and can't do, in my view it would be better to [teach them human ethics and values](#) – like parents do with human children.

Artificial intelligence benefits

People already benefit from AI every day – but this is just the beginning. AI-controlled robots could assist law enforcement in responding to human gunmen. Current police efforts must focus on preventing officers from being injured, but [robots could step into harm's way](#), potentially changing the outcomes of cases like the [recent shooting of an armed college student at Georgia Tech](#) and an [unarmed high school student in Austin](#).^{1 5 18 4}

Image not found or type unknown

Intelligent robots can help humans

in other ways, too. They can perform repetitive tasks, like [processing sensor data](#), where human boredom may cause mistakes. They can limit human exposure to dangerous materials and dangerous situations, such as when [decontaminating a nuclear reactor](#), working in areas humans can't go. In general, AI robots can provide humans with more time to pursue whatever they define as happiness by freeing them from having to do other work.

Achieving most of these benefits will require a lot more research and development. Regulations that make it more expensive to develop AIs or prevent certain uses may delay or forestall those efforts. This is particularly true for small businesses and individuals – key drivers of new technologies – who are not as well equipped to deal with regulation compliance as larger companies. In fact, the biggest beneficiary of AI regulation may be large companies that are used to dealing with it, because startups will have a harder time competing in a regulated environment.

The need for innovation

Humanity faced a similar set of issues in the early days of the internet. But the United States [actively avoided](#) regulating the internet to avoid stunting its early growth. Musk's PayPal and numerous other businesses helped build the modern online world while subject only to regular human-scale rules, like those preventing theft and fraud.

Artificial intelligence systems have the potential to change how humans do just about everything. Scientists, engineers, programmers and entrepreneurs need time to develop the technologies – and deliver their benefits. Their work should be free from concern that some AIs might be banned, and from the delays and costs associated with new AI-specific regulations.

Jeremy Straub is an Assistant Professor in the Department of Computer Science at the North Dakota State University. Straub's research spans the gauntlet between technology, commercialization and technology policy. In particular, his research has recently focused on

robotic command and control, aerospace command and 3D printing quality assurance.

A version of this article was originally published on the Conversation's website as ["Does regulating artificial intelligence save humanity or just stifle innovation?"](#) and has been republished here with permission.