

Stars, Stripes, and Circuits: The Current State of U.S. Robotics and Artificial Intelligence Policy

Bryan McMahon

Duke University, Department of Biology, Durham, NC 27708

Corresponding author: Bryan McMahon, bryan.mcmahon@duke.edu

Keywords: robotics, artificial intelligence, technology policy, U.S. policy, drones, self-driving cars, automation

Executive Summary: Innovations in robotics and artificial intelligence will soon reshape our economic, social, and political lives. However, the rapid advances of these technologies are outpacing policymakers' ability to fully address their effects on the general public. This paper examines current policy governing robotics and artificial intelligence, their historical precedent, and makes recommendations for our government given the current state of these technologies.

With effects already being felt in manufacturing, most jobs, including those performed by experts, will be affected as the technology progresses. Given their positive potential, it is vital to develop policy frameworks that respond to the particular challenges posed. The cost, complexity, and scale of regulations governing robotics and artificial intelligence has increased over time. Issues like cybersecurity cut across many agencies as the objects that surround us become networked. Multiple aspects of our economic, social, and political institutions will be affected at the hands of robotics and artificial intelligence technology as they reshape society as a whole. However, to date these technologies have been governed in a piecemeal fashion. Federal agencies have created a patchwork of regulations that narrowly govern robotics and artificial intelligence within a given agency's purview. This paper proposes two actions the federal government could undertake to address the coming disruptions posed by robotics and artificial intelligence. First, the Office of Technology Assessment, which provided expert, objective analysis of science and technology issues to Congress, needs to be revived. Second, Congress should create a robotics and artificial intelligence agency to unify the current patchwork of regulations and bring together the technological expertise required to govern and guide the coming waves of disruption.

Introduction

Robots, artificial intelligence, and automation dominate today's headlines. Since the beginning of the 20th century, people have dreamed of the role they will play in the future. The word robot, first used in the Czech play *R.U.R. (Rossum's Universal Robots)* in 1933, originally described human-like, plasma-filled beings that were designed for drudgery. These original robots were 'automatons'—efficient but emotionless and incapable of original thinking. In *R.U.R.*, robots were twice as efficient as a human worker but, "lacked a human soul."¹ The creation of these fictional robots reflected society's anxieties of a growing, mass-

produced culture, and fears of a future bereft of human inspiration and ingenuity.

Today, robot and artificial intelligence (AI) technologies look very different than what was imagined by fiction. Precisely defining them is difficult to do, in large part because many products we often see or interact with—like the Roomba—have some features that appear robotic, in this case cleaning without direction or intervention by a person, but they aren't considered robots. For the purpose of this paper, a robot is defined as "a machine with manipulators that can easily be programmed to do a variety of manual tasks automatically."² Essential to this definition of a robot are two elements: the robot's ability to affect and

interact with the physical world around it, and its programmable automation. Although it can be programmed to carry out automated tasks, the Roomba lacks instruments to manipulate and interact with its surroundings. While a “robot” has a degree of autonomy that reflects an ability to take in information about its surroundings and change its actions, “robotics” in turn refers to certain technologies that are part of a robot or allow a machine to act in some ways like a robot. Car factories, for example, use dozens of robots, and many cars are almost completely constructed by robotic workers. These machines can sense where they are and what they’re interacting with, have manipulators fashioned into tools specific for manufacturing, and can react to their environment, including recognizing when a human may be in danger.

Similarly, AI refers to “the branch of computer science devoted to programming computers to carry out tasks that if carried out by human beings would require intelligence,” including, “perception, reasoning, learning, understanding and similar cognitive abilities.”³ To date, no program or computer has reached this goal. AI is specifically not embodied, and is unable to affect its physical surroundings. It exists as an expansion of a computer’s capabilities, usually in the form of a computer program, that acts with intelligence. Throughout this paper, I will refer to any technologies not embodied that act with something similar to human intelligence as AI. Robotics and AI, as referred to throughout the paper, are specific technologies that represent facets or subtypes of a robot or AI.

While our present robots and artificial intelligence look different than science fiction predicted, many of the capabilities and concerns explored in these stories remain the same. The human-like robots imagined in fiction do not walk among us today, but modern artificial intelligences and robotics works on Wall Street, independently executing millions of trades a day without a banker’s say.⁴ They are being developed to drive cars, work in industrial settings, diagnose diseases from x-rays, conduct the discovery process for lawyers, and review the financial details of commercial-loan contract agreements for banks.^{5,6,7} The characteristic these technologies share is *automation*, either by machine, computer program, or data-crunching algorithm. Automation began as an outgrowth of Henry Ford’s

assembly line production method, in which each worker repeated their own specialized task as part of building an entire car. Today, aided by advances in computing power and access to vast quantities of data, automation is poised to affect almost every level of employment with the potential to completely eliminate some sectors.⁸ Upwards of 47% of jobs will be affected by automation, and the burgeoning Second Machine Age portends to reshape every aspect of our economic, social, and political lives.^{9 10}

Other studies, however, emphasize that while automation may take over specific parts of a job, many of the lost tasks are undesirable due to their rote nature (in robotics research, the responsibilities best suited for robots are often called “dirty, dull and dangerous”).¹¹ In this line of thinking, automation may free us from the repetitive parts of our jobs, and create positions to oversee automated systems.^{12 13 14} However, automation will also reach beyond physically intensive labor to affect expert fields. McKinsey’s Global Institute estimates that for 60% of all jobs, 30% of tasks can be automated with today’s robotic and AI.¹⁵ It appears inevitable that automation will reshape the job market, the skills required to stay employed, and the nature of the jobs themselves.

Given the possible scope of automation and its potential to reshape our future, it is vital to understand the current legislative, regulatory, and policy frameworks that exist involving robotics and AI technologies. Understanding the current landscape of policies and their origin is critical to support the creation of evidence-based policies to govern the future of these technologies. This paper summarizes the current state of robotics and AI policy in the United States, examines how such policies arose, and analyzes future directions for the regulation of robotics and AI. To do so, I first identify early government influences by tracing the historical development of robotics and AI. I then examine the federal regulations that currently make up the bulk of robotics and AI policy.

No broad, unified federal policy for robotics or AI exists. Instead, several federal agencies have responded to the adoption of various robotic or AI technologies with regulations restricted to specific domains. However, the scale and cost of implementing the regulations, technical complexity of objects being regulated, and number of technological objects is increasing rapidly, reflecting the growth in robotic or artificially intelligent

technologies and related industries. Parallel to this growth is an expansion in regulatory scope, as individual agencies are being tasked with regulating layers of issues entangled by increasingly complex technical details. In order to grapple with the current growth of these fields, I will evaluate the future implications of robotics and AI policy on our society, and recommend regulatory steps that should be enacted to help steer us through the coming waves of disruptive technologies.

History of Robotics and Artificial Intelligence

A useful guide to tracing the recent development of AI and robotics technologies is to follow International Business Machines Corporation's (IBM) research and computer product development. IBM has had a decades long involvement in both robotics and AI technologies, culminating in Watson, an AI program that leverages neural network architecture and machine learning to better approximate natural language. While the development of robotics and AI has risen and fallen along lines of progress and failure throughout the 20th century, a consistent theme emerges: the development of robotics and AI technologies were the result of strong public-private partnerships. Federal funding was responsible for much of the early research, and budgetary support is therefore among the earliest policy decisions that impacted the field.

The development of robotics and artificial intelligence share the same beginning: the invention of the computer¹⁶. Both a robot and an AI are in part defined by the execution of automated actions without human intervention, and this is achieved through the programmability of a computer. The first electronic, programmable computer was ENIAC (Electronic Numerical Integrator and Computer).¹⁷ ENIAC was a military project of the United States Army that was brought online in 1943 to calculate ballistic missile firing tables (projections of projectile flight paths) during World War II.¹⁸ ENIAC was the world's first operational, general purpose electronic digital computer. Most importantly, ENIAC was programmable.¹⁹ ENIAC had the unique ability to could carry out calculations or solve problems that required multiple steps. By breaking down problems into sequential bits that could be described in programming language, this paved the way for the general purpose computer we know today that can execute significantly more complex

tasks than solving calculations. The next step in computer development came in 1950 when Alan Turing, a British computer scientist and mathematician, described the conditions necessary for a 'thinking machine' in his paper, "Computing Machinery and Intelligence."²⁰ The paper also formulates a way to decide if a machine can be said to be intelligent, called the Turing Test, which has continued to be the standard to measure artificial intelligence's progress and ability to 'think.'²¹ Because thinking is hard to precisely define, Turing proposed a test in which a human interrogator questioned both a computer and another person, both of which were hidden from the questioner. A machine passes the test if the questioner cannot tell which respondent is the computer and which is the human. As of yet, no artificial intelligence has undisputedly passed the Turing Test. Turing's paper formally founded the field of AI as a branch of computer science, and inspired other mathematicians, computer scientists, and engineers to begin working on the design of an artificially intelligent machine. While ENIAC provided the necessary technology for AI and robotics to begin, Turing's formulation of the thinking machine gave the field its intellectual framing and direction.

Robotics

Soon after the invention of ENIAC, new developments added memory storage capabilities to computers, allowing them to execute more than one program without needing to be physically rearranged. This gave rise to the step-by-step programming of manufacturing robots. The leap from a programmable computer to a robotic manufacturing arm happened in 1969 at Stanford's Artificial Intelligence Lab. There, Victor Scheinman, with federal funding provided by The Defense Advanced Research Projects Agency (DARPA), designed the first Stanford Arm. This led to the "arm solution"—a robotics term that describes the complete programmability of the arm's movements in three dimensions with an attached mini-computer.

The use of robots in automated manufacturing processes rapidly spread to factories in the 1970s and early eighties following introduction of the Stanford Arm and related technologies. In general, robotics enjoyed continuous growth in its research and development because private companies were willing to take risks with the technology due to its potentially immediate

implementation. In contrast to AI's aforementioned boom and bust cycles, robotics received investments from strong public-private partnerships that bolstered its continuous growth and development. During the 1970s and 1980s, the movement, dexterity, and programmability of the Stanford Arm was improved, facilitating the spread of robotics out of the factory and into other industries. By 1984 automated robotic arms were capable of welding the Ford Fiesta's frame entirely without human intervention.²² As computer storage increased and computer programming became more intricate, the movements robotic arms and manipulators were capable of expanded in range and became increasingly refined. This has resulted in innovations such as the use of robotically assisted surgery (RAS) systems, in which a robotic arm originally designed for welding car chassis has been repurposed for use in delicate, non-laparoscopic brain biopsies.²³

While development of the Stanford Arm continued, increasing its dexterity and functional capabilities, military research opened parallel track of robotics technology development. A new approach in the 1980s, growing out of the resurgence of research on neural networks and improvements in sensor technology was the Autonomous Land Vehicle (ALV). The ALV grew out of DARPA's Strategic Computing Initiative (SCI). Initiated in 1983, the SCI funded research into advanced computer hardware and AI to modernize the military.²⁴ The research culminated in 1987 when the ALV successfully navigated over 2,000 miles on a rocky test course.²⁵ The ALV drove without the use of an internal map or GPS system, relying instead on sensors and a neural network architecture to process the data to complete the course. This system would go on to form the basis for contemporary self-driving cars.

Later, DARPA spurred robotics and autonomous systems research with a series of Grand Challenges. First offered in 2004, the DARPA Grand Challenge was a series of competitions designed to spur autonomous vehicle development of both hardware (the design of the car) and software (autonomous navigation).²⁶ A mixture of universities and private companies competed for the one-million-dollar prize, with the success of the first Grand Challenge leading to five competitions in total, each focused on developing different robotics and AI technologies. Many competitors in the DARPA Grand Challenge went on to work for the companies that are currently developing self-driving car technologies.

DARPA's Robotics Challenges was a concerted push to develop autonomous systems that could do more than drive: they could cross rubble, lift objects in the way, open a door, connect a firehose to a pipe and turn the valve, climb a ladder, and use a tool to break a concrete wall.²⁷ These autonomous systems and self-driving cars utilize a synthesis of robotics and artificial intelligence to sense their surroundings and respond based on the data accumulated from millions of miles on the road. They herald a new kind machine, a multipurpose, multifaceted robot that takes in data, processes it, and makes decisions contingent on both its past, current, and future environments and its programmed goals and objectives.

Artificial Intelligence

Following Turing's formulations of an intelligent machine, AI researchers Marvin Minsky and John McCarthy, along with senior scientists Claude Shannon and Daniel Rochester of IBM, organized the first gathering of AI researchers at Dartmouth College in 1956.²⁸ The conference formally named their field 'AI' and enumerated its mission, aiming to categorize and catalogue, "every aspect of learning or any other feature of intelligence [that] can be so precisely described that a machine can be made to simulate it."²⁹ Wild assertions and predictions, like that of two researchers claiming to have solved the philosophical mind/body problem,³⁰ abounded during this period following the enthusiasm borne out of the Dartmouth Conference. In particular, researchers predicted the invention of a general AI able to solve any problem a human could within 20 years.³¹ Eager promises such as these would set the stage for the boom and bust cycles of interest and funding experienced by AI and robotics over the years.

In 1963, seven years after the conference, DARPA funded research at three major AI labs, giving three million dollars to each per year and pushing AI research into two, discrete directions.³² The first, logic systems research, focused on rigid, step-by-step problem solving, and required the programs be fed simplified versions of the problems they aimed to solve. The other approach was called perceptrons, or neural networks, and featured a network of decision making agents within a larger system that was modeled after the human brain.³³ Early research favored the top-down approach of logic systems, and the idea of neural networks would not return until

the 1980s. However, the projects in the 1960s overpromised and raised unrealistically high expectations which could not ultimately be delivered. Falling short of the wild predictions posited by both the private and public sectors, government leaders lost interest in the field and reduced associated funding. The resulting shortfall in money and attention to AI, beginning with the cancelation of DARPA project funding in 1974, inaugurated the First AI Winter.³⁴

In addition to returning underwhelming results, AI researchers ran into technical limitations; at the time, computing power was not sufficient for sifting through the large volume of data necessary for AI. In contrast to robotics, AI's developments went through boom and bust cycles because it relied on the ebb and flow of government funding until businesses became interested and started investing in the applications of AI in the mid-2000s. AI research often promised unrealistic goals for the technology, which contributed to these boom and bust cycles. Federally funded U.S. AI research was not revived until the early 1980s, when the Strategic Computing Initiative (SCI) was started, largely in response to other countries' public announcements of AI projects. Initiated in 1983, the SCI funded research into advanced computer hardware and AI to modernize the military.³⁵ In total, the SCI funneled over \$1 billion dollars into university and government labs for AI research, but by 1987 the field again fell short of achieving general AI.³⁶ DARPA ended the SCI and, without a clear way forward for federally funded AI research, the Second AI Winter began.

The Second AI Winter signaled an end to government funded research projects in robotics or AI until DARPA's Grand Challenges of the 2000s. However, by the end of the 1980s, advances in computing power allowed private companies to lead their own AI projects and implement AI in a variety of applications for businesses. IBM took particular interest in marketable demonstrations of rapid advances in computer processing power. In 1997, their chess playing computer, Deep Blue, matched up against then world champion Garry Kasparov.³⁷ Over six games, Deep Blue won, 4-2. The highly-publicized match of "Man vs. Machine" brought AI to the general public, even though Deep Blue represented advances in computer processing more than advances in AI. The computer used a brute-force method of calculating the iterations of each move to

select the best option.³⁸ However, the advances in computer power that propelled Deep Blue to victory would be vitally important to the field of AI following the birth of the internet and would revive a method for designing AI.

Today's AI is a revival of early neural network architectures made possible by the marriage of increased computing power and a deluge of data collected from the internet to feed the programs.³⁹ IBM's Watson made history in 2011 by beating 74-time champion Ken Jennings on *Jeopardy!*, signaling that impressive new ground had been broken in the development of AI, as *Jeopardy's* puns, word games, and flipped answer-question format had previously stymied AI's natural language processing. AI's increased capacity to execute intricate decision-making, in particular those powered by machine learning, have recently learned to replicate other complex tasks as well. In 2013, Watson diagnosed lung cancer from interpreting x-rays, and, in 2016, was hired under the name ROSS by the law firm Baker&Hostetler to help read bankruptcy documents.^{40, 41}

This year Google's DeepMind AI, called AlphaGo, mastered a complex strategy game when it beat Go champion Ki Jie.⁴² This marked another high mark for artificial intelligence by showing AI's potential for strategic thinking; Go is estimated to have more possible moves than atoms in the universe, making it impossible for a program to run through every possible move set for a given board state. Instead, AlphaGo, developed novel strategies and unorthodox tactics, surprising the human players and opening new ground in the overall approach to Go.⁴³ AlphaGo was educated by being fed over 30 million human Go moves and then playing against a slightly different neural network structure of itself. The AI revolution has come full circle, now powered by technological advances in computing power and a flood of data connected through the internet and cloud computing.

Federal Agency Action: AI and Robotics Today

Until the past decade, robotics and AI technologies were not advanced enough to garner much political attention outside of funding. As these technologies advance, so do their impacts on society, economy, and our everyday lives. The economy will be reshaped as artificial intelligence and robotics automate more tasks and types of labor. AI is already being used on Wall Street to read market patterns, conduct high-frequency trading, and

execute long-term trading strategies. In law offices, AI can draw up contracts and carry out the discovery process while hospitals leverage machine learning tools like IBM Watson to assist in diagnoses or check for potential drug interactions. These algorithms can be too complex for their developers to be able to fully trace and explain how a program makes a decision or executes its commands. For example, an AI may rely on too many variables to carry out its goal for its developers to be able to track and follow. In addition, machine learning programs often develop their own ways of weighing variables or arriving at a decision as they are fed data, so it's not always possible to 'open up' the program to understand how it is working. This problem, called the 'black box' problem, poses an asymmetric risk as automation becomes more centralized and automates more labor.⁴⁴ As the technologies spread across industries, the federal frameworks that develop must address the black box problem and legal remedies need to adapt to accommodate harms carried out by autonomous agents. Finally, manufacturing jobs are already being displaced by robots, and as drones and self-driving cars develop, the labor market in many industries, from delivery and taxi services to knowledge and expert services, is poised to become increasingly tumultuous. Job and skills retraining and workplace oversight will need to respond to the new economy.

Our political lives are being changed as well. Two recent, unpredicted electoral events—the election of Donald Trump as President and Britain's "Brexit" vote, in which the country voted to leave the European Union—boast deep ties to artificial intelligence.⁴⁵ Both relied in part on Cambridge Analytica for its advertising, a company that boasts an algorithm that can predict the personality of a given voter and, through machine learning, test and tailor a political ad to that audience (called microtargeting). Campaign advertising and online news algorithms that shape the public's images of and information about politicians and policy are driven by, and are by and large unregulated, AI. Bots—an automated computer program—also play a role in today's politics as they automatically push a message on social media while posing as real people.⁴⁶ As these technologies continue to penetrate our political life, policies must be developed to ensure the health of our political processes and institutions, including the media ecosystem and our elections.

Finally, robotics and artificial intelligence technologies bring unique concerns for policy-makers. First, cybersecurity is an increasing risk as robotic and AI technologies are networked together. In particular, smart homes—a house in which many of the essential functions, like heating and air conditioning or even the locks on the house, are networked together and connected to the internet—and self-driving cars present a challenge to secure. Given the potential for harm to both property and persons, the cybersecurity of self-driving cars and smart homes as well as the public infrastructure like the roads and the networks they rely on, is paramount. A related set of issues arise as robots and AI are increasingly connected to the internet. Privacy and data ownership and sharing will continue to be fundamentally changed as these technologies become networked. In the above examples, the privacy of ourselves in our homes is potentially threatened, while new gushes of data—ranging from GPS locations in our cars to monitoring how we use our devices in our homes—will spring from self-driving cars and smart homes. The confluence of these issues when combined with the increasing spread of networked AI and robotics technologies present challenges for policy-makers.

The following section highlights the current federal regulations and policy frameworks, how they developed, and the gaps in these frameworks that are poised to widen as robotics and AI technologies mature and spread. While to date Congress has not passed comprehensive laws regulating the broad field of robotics and AI, we have reached a point wherein federal regulation is necessary. Without legislation passed by Congress uniformly addressing robotics and AI, the regulatory responsibility has been distributed piecemeal across various federal agencies. As discussed in the introduction, robotics and AI, as they exist today, are multifaceted groupings of specific technologies and areas of research, like automation, machine learning, or computer vision. Current regulations of robotics and AI target specific facets of robotics or AI as they are used in real-world applications and products. Thus, an attempt to regulate automation leads to disparate policies enacted by agencies as diverse as the Occupational Safety and Health Administration, when workplace safety is affected, or the Food and Drug Administration (FDA), when the safety and efficacy of medical devices must be determined. Given the variation among technologies that could

be categorized as robotics or AI technologies—from drones to medical diagnostic tools to automated industrial machines— current federal policies regulate robotics not as a broad class of objects but in specific domains of use.⁴⁷ This framework is intended to balance the need to support innovative development of new technologies with safety, cybersecurity, privacy and data sharing protections for the public. As I will argue, this siloed approach to regulation is leading to redundant and ineffective policies compared to the issues these technologies raise. This section summarizes the state of robotics and AI and details the resulting policy frameworks of the following federal agencies as they encountered robotics or AI: the Occupational Safety and Health Administration (OSHA), Federal Rail Administration (FRA), Federal Aviation Administration (FAA), National Highway Transportation Safety Administration (NHTSA), and the Food and Drug Administration (FDA).

Automation and Regulation in Industrial Manufacturing

Cost is a major motivation driving the adoption of factory robots. The initial, nonrecurring capital investment to buy the robot can be markedly cheaper than paying an employee wages and benefits over his or her lifetime.⁴⁸ This trend continues into the present, with more facets of human labor in manufacturing continuing to be automated. For example, in the Fremont California Tesla car manufacturing plant, opened in 2013, robots are employed to do multiple tasks. Working in coordination, these robots can, “autonomously swap the tools wielded by their robotic arms in order to complete a variety of tasks. The same robot installs the seats, retools itself, and then applies adhesive and puts the windshield into place.”⁴⁹ ⁵⁰ Improvements in industrial robots’ functional capabilities, like increasing the total range of motion, improving fine motor control of the machine’s “hand,” increasing the computing power and overall programmability, synchronizing multiple robots, or augmenting their autonomous capacities by installing visual systems or wireless communication systems, have greatly expanded the tasks and types of labor these machines can reliably automate which contributed to their broader adoption.

The rapid and widespread adoption of robotic and automated technology has resulted partly from a ‘hands-off’ regulatory approach, particularly at the

federal level. Though robots were first installed to take over the most dangerous parts of the manufacturing process, their presence posed its own risk. By 1983, 66,000 automated industrial robots had been installed, and OSHA began investigating workplace-robot safety accidents.⁵¹ An agency within the U.S. Department of Labor, OSHA, sometimes in partnership with state government, oversees workplace safety standards and regulations. OSHA’s authority to oversee the health and safety of workplaces is mainly accomplished through compliance officers who visit and inspect workplace environments to verify that they are in compliance with OSHA’s safety standards and workplace regulations, levying fines on facilities that violate OSHA’s requirements. The first accident occurring between a human worker and an automated machine was reported on November 12th, 1984, noted as a worker’s “right leg caught in a welding apparatus.”⁵² The machine, a Hitachi robot welding the frames of General Motors vans, could either be automated or manually controlled and was in its automated mode when it caught the leg of an employee who had been working with the robot for nearly two months. The machine was in its automated mode, and other workers failed to stop the machine as no override switch had been installed.⁵³ More accidents followed, with increasing consequences. In 1987, two industrial robots crushed two workers to death. The first death was caused by a robotic arm that had automated grabbing and welding metal plates and crushed the worker between two plates.⁵⁴ The second worker was killed by an automated lathe.⁵⁵

In response to these accidents, the Office of Technology Assessment (OTA), a government body that existed from 1972 until 1995 to provide expert analysis on scientific and technological issues, developed a set of recommended safety standards and workplace regulation requirements for manufacturers using automated industrial robots. The guidelines, adopted by OSHA in 1987, described necessary safety precautions to help prevent further accidents between workers and automated machines as robots continued to be adopted and grew increasingly complex in shape and function.⁵⁶ The guidelines emphasized the dangers different robots may pose and describe what physical and procedural precautions are needed to ensure their safe operation.

At present, robots struggle to recognize the intentions and body language of humans, and fail to recognize if a worker is in danger unless its interactions are clearly defined.⁵⁷ The risk to the worker presented by a robot varies greatly according to the robot's degree of automation, the level of, frequency and kind of interaction with a human, and the type the work performed.⁵⁸ Given the rapid adoption and evolution of industrial robots, OSHA regularly updates the OSHA Technical Manual. Section IV, Chapter 4 of the manual describes robot types, robot system types, and the variety of operations carried out by industrial robots.⁵⁹ ⁶⁰ In addition, OSHA's training manuals and educational materials have been expanded to reflect the increasingly technical nature of the industrial workplace. The pace of industrial robotics adoption has only increased as robots have become increasingly adept, networked, and capable of learning from past mistakes, and regulatory agencies must be able to follow.⁶¹

Robotically-Assisted Surgery Systems

Medicine, like many other professional fields, is poised to be reshaped by robotics and AI. Created in 1906, the FDA is responsible for the protection and promotion of public health by ensuring the safety and efficacy of drugs, medical devices, the food supply, cosmetics, and biological products.⁶² The Food, Drug, and Cosmetics Act of 1937 expanded FDA's authority to include a pre-market evaluation of the safety and efficacy of any drugs, medical devices, or biological products. Under this authority, robotically-assisted surgical (RAS) devices are regulated as a medical devices. A RAS system is an interface between the surgeon and the patient that, "enables surgeons to use computer technology to control and move surgical instruments through tiny incisions in patient's body."⁶³ Despite incorporating the term "robotic," RAS are teleoperated; a surgeon is always in control of the surgical tools. The surgeon performs the procedure with machine-steadied tools while looking at a screen that displays the surgery site. RAS systems are used almost exclusively for laparoscopic procedures as the robotically-supported surgical instruments may provide greater precision during laparoscopic procedures than the surgeon's hands alone.⁶⁴

The FDA regulates the safety and efficacy of medical devices like RAS systems through pre-market evaluations and post-market monitoring. Medical

devices are classified into three categories based on their expected levels of risk. RAS systems are regulated as Class II, 510(k) devices due to their "moderate risks and moderate regulatory controls."⁶⁵ 510(k) pre-market notifications require manufacturers to submit a letter showing the, "device to be marketed is at least as safe and effective, that is, substantially equivalent, to a legally marketed device,"⁶⁶ in lieu of the FDA itself testing the device. In order to demonstrate substantial equivalence and qualify for a 510(k) notification, the RAS must have the same intended use, technological characteristics and safety and efficacy as the predicate device already on the market.⁶⁷ To date, the FDA has made no independent standards specific to RAS devices, instead relying on predicate devices to guide RAS's safety and effectiveness. IBM's Watson, which is being developed to diagnose illness and recommend the appropriate medications, is not currently regulated by the FDA. This medical AI package is considered a clinical decision support tool, rather than a medical device. Clinical decision support tools retrieve existing clinical information about the patient from electronic medical records or about their medical tests, conditions, and drugs from accessing the medical literature online. In contrast, medical devices conduct diagnostic tests, and thus fall under the regulatory purview of the FDA.

Self-Driving Cars

Few robotics or AI developments have captured the public's interest as much as self-driving cars. While ideas for self-driving cars go back as far as the 1930s, the autonomous vehicles of today are a direct result of DARPA's Grand Challenges and companies investing in autonomous driving technologies in the mid 2000s. These vehicles are a synthesis of robotics and artificial intelligence technologies. Equipped with sensors that read their environment and machine learning algorithms that react depending on the millions of miles of driving they've been fed, the vehicles respond to their environment, manipulating themselves to navigate. The quick maturation of these technologies and the regulatory issues they bring with them have catalyzed the need for government action. Stakeholders recognized the need to create safety guidelines, an upgraded infrastructure to integrate self-driving technology, and ongoing public-private collaboration. In addition, privacy concerns surrounding data ownership and the cybersecurity of the car itself present a need for

an efficient and complete framework to guide self-driving cars to market.

Complicating the regulatory picture further, not all companies are interested in fully automated (i.e., completely without human intervention) vehicles. Tesla, for example, combined automatic braking, rearview cameras, and distance sensors that other car companies introduced to create an 'Autopilot' mode.⁶⁸ Through Autopilot, Tesla's cars maintain appropriate speed, distance from other cars and stay between lane markings; however, the driver's attention and intervention are still required.⁶⁹ Other car companies have also been incrementally introducing features that assist or fully automate some parts of driving. Many vehicles now use radar sensors to warn the driver if he or she is too close to another car, and some engage the brakes automatically if the driver fails to respond. Mirroring the various approaches to a self-driving car, the current laws governing them are a complicated patchwork that often differ according to the type of self-driving car.

Much of the policy concerning self-driving cars written to-date has been a binary of whether or not autonomous vehicles are even allowed on the road. There is an ongoing competition to be the leader in setting policy between car companies and tech companies at the state level. Tennessee, Georgia, Illinois, and Maryland have crafted bills to allow only traditional automakers access to test and develop their self-driving cars, while Nevada's Department of Transportation granted Google's fully autonomous vehicles the first license for a self-driving car in 2012.⁷⁰ Uber's fleet of self-driving cars left San Francisco after the California Department of Motor Vehicles revoked their registrations because no fully automated vehicles had been approved by the municipality, but in Pittsburgh, self-driving Ubers have been offering rides since 2016.⁷¹ The current piecemeal approach will lead to more confusion because no fewer than twenty companies, ranging from Silicon Valley startups to traditional car manufacturers like Ford and BMW, have entered the self-driving car market, and the number continues to grow.⁷²

Facing a groundswell of self-driving technologies that redefine what 'driving' means and who does it, a wider framework of policies has been called for by many of the companies themselves, safety and consumer advocates, and regulators. Prototype cars and trucks, automated or not, now regularly cross in

and out of municipalities and state lines and drive on roads maintained by state and local infrastructure funds. Given the number of cars that could become autonomous and the need for a seamless integration of self-driving technology into the existing infrastructure for drivers and pedestrians alike, a larger and coordinated framework of policies are needed to satisfy the public and private interests touched by automated driving. As in the case of automated industrial robots, existing safety regulations were amended to accommodate the arrival of robotics and AI to the roads. The National Highway Transportation Safety Agency (NHTSA) oversees the national safety regulations for motor vehicles, setting the safety standards to which all motor vehicles must comply. In September 2016, NHTSA released a set of voluntary guidelines for automated driving technology called, "Federal Automated Vehicles Policy".⁷³ This is a set of policy recommendations to expand the agency's mission of enhancing safety on the roads. By only releasing recommendations, the NHTSA is encouraging an ongoing conversation between stakeholders and an in-depth review of proposed regulations before the final policies are put into place. This document breaks regulation of self-driving cars, called Highly Automated Vehicles (HAVs), into four broad parts: vehicle performance guidelines for automated vehicles, model state policy, NHTSA's current regulatory tools, and recommended new tools and authorities.⁷⁴ These regulations apply to systems that are considered fully automated, i.e., that can complete the entire driving task without human intervention, or semi-automated, like Tesla's Autopilot. All HAVs that will drive on public roads, including HAVs that are in the testing and development phase, need to be covered under NHTSA's guidelines as well. Importantly, regulation of automated vehicles, despite the expectation for them to significantly disrupt industries and transform infrastructure, cities, and personal mobility, is primarily motivated by according to the potential harm they may pose to the public.

NHTSA's guidelines outline the performance standards of HAVs to meet its safety regulations for vehicle manufacturers⁷⁵. These include how well a car must protect the passengers during a crash, as well as how it should function when driven. For HAVs, the safety standards must cover automated driving and all the technologies that contribute, making it a highly technical review process. A new

fifteen-point standard would amend NHTSA's current vehicle manufacturer safety standards, the "Federal Motor Vehicle Safety Standards and Regulations", as authorized by the National Traffic and Motor Vehicle Safety Act in 1967, to include safety regulations explicitly for HAVS.

Although the amended safety standards are currently just guidelines, NHTSA has requested for vehicle safety reports to be voluntarily submitted, so that the effectiveness of the proposed regulations may be tracked and be used in later rule-making by the agency. These recommended safety amendments reflect the unique challenge automated cars may bring regulators, as 'safety' has expanded into the domains of data sharing between customers, companies, and third-parties, the resulting tension between data sharing and personal privacy, cybersecurity, and complex ethical situations revolving around questions like who a vehicle will choose to hit during a crash that would inevitable involve multiple people. These are complex questions to address, let alone regulate, without impeding the development of autonomous vehicles, and this complexity is reflected in the diversity of the stakeholders participating the regulatory process and the multifaceted framework that has so far developed.

The proposed guidelines would follow traditional regulatory jurisdictional responsibilities: states oversee vehicle licensing and registration, traffic laws, and motor vehicle insurance and liability structures while federal agencies oversee the regulation, like safety, of motor vehicles. The Department of Transportation be in charge of licensing HAVs as the 'driving' in fully automated vehicles is not done by individuals but by the vehicle itself; therefore, separate registration and licensing processes to accommodate HAVs need to be added.⁷⁶ The guidelines also propose restructuring insurance and liability rules to spread risk among owners, operators, passengers, and manufacturers. Emphasized throughout the recommendations for states is the need to avoid a patchwork of laws from state to state or laws different enough that encourage competing interests from car makers. Public roads exist within and across state lines and therefore require state and federal policy to mesh and leave no regulatory gaps.

The final two sections of NHTSA's guidelines review its current authorities and regulatory tools and recommend where they will need to be expanded.

The tools would support the expanded authority of NHTSA to regulate autonomous vehicles. In particular, allowing NHTSA to require pre-market testing and affirmative approval for all HAVs, a change from its current purview.⁷⁷ New proposed statutory amendments would grant NHTSA the further authority to regulate HAVs in three main areas. First, given the potential scope of an autonomous system, NHTSA would be able to require manufacturers to "cease and desist" their operations in emergency sections. A single city or manufacturer may be responsible for a live network of HAVSs, and so immediate intervention in the network in case of a systemic error or in times of an emergency is necessary. Second, the number of vehicles exempted from some federal regulations for research purposes would be expanded. Research into total integration of the autonomous vehicles into the larger infrastructure needs to be expanded, and the exemptions allow testing to progress. Finally, 'over the air,' or wireless, software updates would be first checked for safety and errors by NHTSA.⁷⁸ As outlined above, part of the recommended expansion of federal authority is in response to the highly technical nature of carrying out basic regulations about cars, and cities or states may not have the money or ability to support the necessary highly technical staff. Another reason an expansion of federal regulatory authority may be needed is the sheer scope of implementing of a new, potentially disruptive technology across the United States.

Autonomous vehicle technology is not yet ready for the leap to fully automated driving. More time and better synchronized testing and development between public and private stakeholders is needed. Policies on the books now may miss what HAVs and their infrastructure will look like and cost; worse, regulations could limit or all-together prevent HAVs from hitting the road by regulating before the technology is mature. NHTSA must ensure the public interest is served by increasing safety associated with driving and expanding access to transportation. Given the scope of implementation and potentially transformative nature of HAVs, NHTSA's guidelines allow for an open and private-public approach to policy-making that sets the stage for well-targeted regulations in near future.

Drones

Unmanned aircraft systems (UAS), more widely known as drones, have resulted in significant

pressure on the current piecemeal approach to AI and robotics in several areas. Drones challenge our current regulatory frameworks by combining robotic and artificial intelligence technologies and overwhelming the current regulatory network by sheer numbers and speed of development. Combined with the steadily declining costs to purchase, equip and use drones, an estimated 2.5 million commercial or personal drones were in flight by 2016, which will increase to an estimated 7 million by 2020 according to an FAA report.⁷⁹ This growth is the result of next-generation drone technology sensors, high-resolution cameras, and the ease of remote piloting by smartphone. The five largest sectors predicted to incorporate drones are real-estate/aerial photography, insurance, agriculture (crop inspection and pesticide spraying), industrial inspection (bridge, powerline, and antennae inspections), and government (primarily law enforcement).

Regulation of Commercial and Personal Drones

In an attempt to keep up with the rising tide of drones, several key policies have been enacted in the last five years. Importantly, in contrast to industrial robots and PTC, which both represented regulatory responses to accidents, UAS policies grew out of conversation between multiple stakeholders and the FAA. Private companies, university researchers, hobbyists, and safety and privacy advocates all contributed to a proactive public-private approach. Two reasons made the public-private approach possible. First, drone hobbyists had a prior history of self-regulation under their umbrella organization, the Academy of Model Aeronautics (AMA), created in 1940.⁸⁰ Second, the widespread adoption of drones brought together the numerous and diverse stakeholder interests whose vision for drone uses intersected and overlapped, triggering the effort to develop a framework of policies that could best accommodate all parties.

In 2012, facing a groundswell of drone technology and stakeholder interest, Congress passed the FAA Modernization and Reform Act to allow the FAA to develop a framework and infrastructure to integrate new technologies, like UASs.⁸¹ The basic requirements for a commercial aircraft—a certified and registered craft, a licensed pilot, and operation approval—continued to apply to UAS. But, as the framework to integrate drones was developed, the act also created the Section 333 exemption, which granted the FAA the authority to determine the need

for a commercial drone to obtain an airworthiness certificate on a case by case basis. However, the Section 333 exemption was slow to get approved and proved to be too narrow and restrictive a pathway for companies. In fact, the difficulties caused by the Section 333 exception laid the groundwork for a revised drone rule to be issued later. The FAA Modernization and Reform Act formally regulated personal or hobbyist drones by the FAA, in place of self-regulation according to previously-coordinated agreements with the AMA.⁸² Drone use had expanded to the point that the FAA decided a unified and codified set of federal regulations needed to replace the system of self-regulation. Section 336 of the FAA Modernization and Reform Act codified existing operating procedures and included language that UASs for personal use must still stay within visual-line-of-sight (VLOS) of the operator at all times. This language in particular emphasized the reason for federal regulation of hobbyist drones because UAS pose a serious risks to privacy, safety, and an integrated national airspace if left unrestricted by VLOS rules. Section 336 also exempted personal drones from future FAA rules, setting the piece of the larger framework in place. As the larger framework and infrastructure was developed, the FAA worked closely to best accommodate the diverse and growing uses of drones. Drone delivery services, represented by Amazon and Google, researchers, farmers, and many other types of industries, coordinated priorities and concerns to the developing regulations, working closely with the FAA regulators. Issues beyond those of the private stakeholders included privacy concerns for the public due to surveillance drones, raised by the Electronic Privacy Information Center, were a part of the conversation.

The resulting policy framework was the FAA's final rule published on June 28th, 2016, known as the "Department of Transportation (DOT) Operation and Certification of Small Unmanned Aircraft Systems."⁸³ The Rule put into place firm regulations for the type of drone allowed in commercial operations. In Section 333, the Small UAS Rule clarified and enabled drone use to ensure commercial drone use could expand without endangering the public. First, the drone must be in constant VLOS of either its operator or another person in communication with the operator. This restriction, while burdensome for many of the drones' uses, like flying over large areas

of crops, simplifies and lowers the risks of UAS operation beyond VLOS in the National Airspace. Second, much to the chagrin of companies developing drone delivery systems, each UAS must be controlled by one person; this one-to-one system prevents large scale automation of a drone's flying, surveilling, or delivering. Third, the drone's pilot must have a remote pilot certification. Finally, the drone and its payload cannot exceed 55 pounds nor be operated more 400 feet above the ground or travel more than 100 miles per hour during its flight path. In this final rule, the FAA balanced the need for a well-regulated National Airspace and risks associated with drones operating outside of VLOS of its operators against the opportunity for the rapid introduction of commercial drones across a range of industries. The rule, while prohibiting drone delivery services like Amazon and Google for the time being, will allow commercial drones to create an estimated 100,000 jobs and \$82 billion dollars in new services over the next decade while preserving the integrity of the National Airspace.⁸⁴ The FAA Modernization and Reform Act and FAA "Operation and Certification of Small Unmanned Aircraft Systems" Final Rule respond to and unify diverse and consequential uses for UAS, particularly in commerce and research with a framework of policies moving beyond self-regulation to ensure a well-integrated National Airspace.

Regulation of Military Drones

Drones are challenging our regulatory frameworks in other spaces, too. While the FAA continues to develop a comprehensive framework for public and private drone flights, the expanding capabilities and application of drones in combat has developed largely out of public sight. The military first used drones for surveillance and reconnaissance missions starting soon after World War II. The Predator drone, primarily used by the United States Air Force and Central Intelligence Agency (CIA), was first flown in 1995 for missions in Serbia and remained active as a reconnaissance and surveillance aircraft.⁸⁵ By 2000, a satellite telecommunications connection allowed truly remote piloting, where the pilot could control the drone from anywhere, instead of radio-controlled flights that required the pilot to be at least in range of the radio. In 2001, the U.S. Airforce equipped Predator drones with missiles and developed remote aircraft strategies for targeting and striking enemy combatants, providing more mobile strike forces while keeping soldiers out of

danger.⁸⁶ These two technological advances made the modern military drone, enabling U.S. forces to remotely fly, target, and kill around the globe from a base in the United States.

In February of 2002, the CIA committed the first lethal drone strike undertaken separately from any ongoing military operation in Afghanistan to kill a "tall man" believed to be Osama Bin Laden.⁸⁷ Lethal drone strikes were authorized under the Authorization of Military Force (AUMF) Congress signed after the 9/11 attacks to fight "a global war on terror."⁸⁸ Specifically, the AUMF authorized lethal force against Al-Qaeda or other terrorist networks deemed to help plan or aid the 9/11 attackers or their organization. Importantly, President Bush, empowered by the language of the AUMF, claimed the power to classify anyone, even American citizens, as "enemy combatants" as defined by the Geneva Convention.⁸⁹ Enemy combatants carry a precise definition under the Geneva Convention, which governs the internationally accepted rules of war. Since the AUMF, the President has claimed executive power to determine anyone, if known or suspected to be a terrorist or part of a terrorist organization, an enemy combatant subject to lethal force carried out by the United States government. Under the Obama administration, drone strikes have spread outside of the original conflict zones in Afghanistan, and the executive branch expanded its claims under the AUMF to include persons not in active combat zones. Rather than adhering to the stricter classification of a combatant as outlined in the Geneva Convention, the Obama administration maintained "kill lists" of known or suspected terrorists and their associates.⁹⁰ Such strikes, including two that killed three American citizens, cleric Anwar al-Awlaki, his son, and magazine editor Samir Khan, expanded outside of combat zones bringing into question the domestic and international legality of both the targeting those who are not "enemy combatants" and killing outside of combat zones.

The expansion of these strikes brings into question whether the AUMF can still be said to be the guiding policy framework for lethal drone strikes. In response, the executive branch has crafted rules and policy guidances to regulate the use of drone strikes and have allowed the military to attack and kill using drones remotely and autonomously with little public, congressional, or judicial oversight. In 2013, President Obama released a Presidential Policy

Guidance (PPG), which was partially redacted, outlining the rules and guidances that underwrite the drone program's targeting and killing of people in Afghanistan, Iraq, Syria, Libya, Yemen, Pakistan, and Somalia.⁹¹ One of the guiding principles outlined in the PPG is the lawful use of force against an imminent threat, as described in the Geneva Convention. The PPG therefore considers any strike against known or suspected terrorists to be lawful, having previously characterized terrorists anywhere an imminent threat after 9/11.⁹²

The legality of such strikes under international and domestic law continues to be contested. Internationally, a 2013 report by Amnesty International and Human Rights Watch argues the targeting and killing contravene the laws of armed conflict, international law, and Obama's own policies.⁹³ The report, detailing several CIA drone strikes, including one in Pakistan that killed a grandmother gathering vegetables, charges that the U.S. officials who ordered the strike should stand trial for war crimes. Domestically, federal courts considering the cases of Anwar Al-Awlaki and his son Abdulrahman, ruled in favor of the U.S. government on procedural and jurisdictional grounds that it could kill its own citizens without a trial if they are believed to be terrorists.⁹⁴ To date, no court has adjudicated the legality of these drone strikes, Congress has made no law, no regulations have been passed into the U.S. Title Code, and no federal agency has reviewed or developed policy framework governing drones using lethal force.

An unclear set of executive rules and guidances authorize the current lethal drone strikes, and the regulatory picture stands to become more complicated soon. Clearer policy underpinnings and legal structures will be needed as drones are poised to become one of many autonomous weapons. For example, currently the U.S. military is experimenting with autonomous drones using artificial intelligence to fly, identify, target, and, potentially, kill.⁹⁵ The potential of lethal force being carried out by autonomous systems only places greater pressure on the murky policies currently guiding drone strikes. In addition, missiles that can decide what to attack and ships that autonomously track and hunt submarines for thousands of miles are being developed.⁹⁶ Internationally, in what is perhaps the clearest example of the potential of autonomous weapons, an autonomous robot patrols the demilitarized zone between North and South Korea

that can spot and track a person for 2 miles and then kill him or her with a high-powered automatic rifle.⁹⁷ The robot has two settings: one requires human sign off to kill and one that is automated.⁹⁸ While current guidances still require a human to pull the trigger with a drone, autonomous weapons are only years away, not decades.

In response to the potential of autonomous weapons, over 3,000 AI researchers signed and published an open letter warning against a global AI or autonomous weapons arms race.⁹⁹ The letter emphasized the unique threat autonomous weapons would pose to our international ethical and legal frameworks and the treaties and conventions governing the rules of war. Specifically, the researchers warned of autonomous weapons being "the Kalashnikovs [the original AK-47] of tomorrow" because, "Unlike nuclear weapons, they require no costly or hard-to-obtain raw materials, so they will become ubiquitous and cheap for all significant military powers to mass-produce... Autonomous weapons are ideal for tasks such as assassinations, destabilizing nations, subduing populations and selectively killing a particular ethnic group."¹⁰⁰ Autonomous weapons are likely to cause a major shift in not just the weapons used in war, but the way war is conducted as well. While relying on autonomous weapons can protect soldiers from being in direct combat, it may also thereby encourage conflicts by reducing the risks sustained by fighting.¹⁰¹ In light of the threat of a global AI and autonomous weapons arms race, the United Nations is working on a treaty to ban lethal autonomous and artificial intelligence weapons.¹⁰² Experts gathered by the United Nations Convention on Conventional Weapons will meet twice in 2017 to build a framework for a broad ban on autonomous weapons, although, even if signed, such a treaty may not be effective in banning the full range of autonomous weapons possible. Given the contentious international and domestic challenges to the current drone program, the lack of a comprehensive policy framework governing use, limits, and authorization for autonomous weapons poses an imminent threat to our national security and international law. Furthermore, the authority for a machine to autonomously kill a person, combatant or otherwise, raises profound ethical and legal questions about robots. The introduction of autonomous without proper policy guidance, oversight, and public debate poses grave risk, and at present, no existing policy

structure is poised to address the ethical, legal, or regulatory issues autonomous weapons raise. In summary, autonomous weapons challenge our domestic and international legal and ethical frameworks. Due to the serious risks they pose and the current lack of oversight and regulation, autonomous weapons require the highest level of policy-making processes in complex, international treaties and frameworks to appropriately govern these tools.

Conclusion and Future Implications

The rise of robotics and AI and government policy-making have been intertwined since the birth of the computer during World War II. At first, regulators encouraged the development of these technologies through decades of funding research projects, universities, and holding prize competitions to spur innovation. As the technologies spread and matured, federal agencies established regulatory frameworks in response to safety needs. Recently, however, robotics and AI have progressed to a point that is straining the patchwork of federal agency regulations due to the technology's ubiquity, technical complexity, and potential capabilities. Automation and the continuing evolution of robotic and AI technologies will undoubtedly reshape our economic, social, and political lives.

One of the primary factors underlying the current state of robotics and AI policy is the increase in scale, cost, and complexity of enacting policy frameworks for rapidly evolving robotic and AI technologies. OSHA was the first agency to intersect with and regulate robotics by amending and updating its existing safety standards. The FDA was able to accommodate RAS systems into its current regulatory structure for medical devices, but has no authority over AI diagnosis support tools as it does not regulate the practice of medicine. Finally, both the FAA and NHTSA require expanded authorities and regulatory tools to put into place the necessary policy frameworks for drones and HAVs. The policy frameworks to regulate robotic or AI technologies within an agency's domain have necessarily become more complex over time.

The evolving nature of robotic and AI technologies is multiplying the policy concerns and compounding the regulatory tools needed by any single agency.¹⁰³ For example, to carry out its mission with HAV's, NHTSA must have the technical expertise to determine the crashworthiness of both the physical

car and the safety of the self-driving technologies. Implicit in the range of technologies an HAV uses to drive, however, are concerns beyond NHTSA's traditional ones of safety and crashworthiness. NHTSA will also need to regulate the data collection, sharing, storage inherent to each car, the privacy of the driver, and the cybersecurity of the HAV. The issue extends well beyond self-driving cars, however. Smart homes connected through household objects, networked factories, networked trains, and even future medical devices will all be connected to the internet and susceptible to being hacked. Autonomous weapons and surveillance robots are powerful tools that could bring greater security and lower the human cost in wars; they could also make conflict more violent, damaging, and pervasive or empower an authoritarian government to watch, record, and weaponized every moment of our lives. Legal and ethical frameworks addressing data ownership, privacy, cybersecurity, and the use of lethal force must be established to govern robots or AI domestically and internationally and demand a unified approach to regulation.

The current patchwork of federal agencies is reacting independently to developments in robotics and AI in each specific domain, and will not be able to adequately address the changes and challenges brought by the Second Machine Age.¹⁰⁴ Instead, a unified approach that puts robotic and AI technology first is needed. There are two steps Congress can quickly take to accomplish this. The first step is to restore the Office of Technology Assessment (OTA). Before it was defunded in 1995, the OTA produced nonpartisan analysis by experts in science and technology fields, lending their expertise on legislation, regulations, or through reports on science and technology topics. Guided by leading experts in their field, the OTA kept the government informed of new and cutting-edge science and technology developments. OTA produced a total of 750 reports on topics ranging from acid rain and global climate change to healthcare and polygraphs. Its closing prompted Republican House Representative Amo Houghton to say, "We are cutting off one of the most important arms of Congress when we cut off unbiased knowledge about science and technology."¹⁰⁵ Reinstating the OTA would provide members of Congress with objective and authoritative scientific and technological analysis on complex issues for a bill, report, or set of regulations being considered.

Reviving the OTA is relatively simple to do because the Office was never abolished. It was defunded and stripped of its experts, but no new legislation recreating the office is needed to bring it back. Its annual budget at its closing in 1995 was \$20 million, a small drop in the \$3.2 billion budget Congress has for internal matters.¹⁰⁶ Robotics and AI will soon reshape our economy, our social lives, and our political institutions, and Congress needs scientific and technological expertise to craft the right legislation for our future.

In addition to bringing rigorous scientific and technological analysis back to Congress, a new federal agency dedicated to robotics and AI is needed. When the radio was invented, the technology proved to be transformative for our society and required infrastructure that spanned the country, the Federal Radio Commission was created. Now the FCC, the agency smoothed the integration of radio, cell phones, television, and the internet into society in ways that aimed to serve the public interest. Robotics and AI stand as transformative technologies, but the incorporation has been far from smooth. Economic dislocation has already begun with little policy oversight, and creating an agency to effectively govern a complex set of regulatory needs is historically the rule not the exception.

As robotic and AI technologies continue to evolve, the policy frameworks needed to balance public and private interests will grow increasingly complex and depend on highly technical knowledge. Relying on individual agencies to regulate these technologies in their specific domains blocks the necessary accrual of policy expertise a consolidated federal agency could solve. Furthermore, the U.S. government's historical approach of recruiting the "best and brightest" has failed to bring in talented computer scientists, engineers, and technology policymakers to work for the government.¹⁰⁷ Instead, the best have gone to work for tech companies like Google, Amazon, and Facebook where they work on cutting-edge technologies and are rewarded with high salaries. In addition to being better positioned to recruit the best talent, a unified federal agency would be better suited to deal with the issues, like cybersecurity, that cut across multiple technology domains. By consolidating technology issues that impact multiple regulatory domains—cybersecurity, ownership of data, and automated decision-making more generally—under the purview of a new

robotics and AI agency, the resulting regulations would be leaner and less restrictive. The current model forces agencies without the requisite technical experience, knowledge, or policy-makers to address these issues under a given agency's narrower regulatory authority, which has historically been safety of these technologies. While NHTSA grapples with the cybersecurity of self-driving cars, the FAA writes rules to protect drones from being hacked. Thus, the regulations strain the regulatory authority of the agencies and creates redundant policies for problems that cut across multiple domains. Much of the current public policy debate surrounding these technologies splits along the lines of whether they require more, or less regulation. The former argues that unregulated technologies can pose a danger to the public and harm the public good unless brought under appropriate regulatory guidance, while the latter argues that policy-makers are too quick to regulate these technologies and threaten to strangle further growth and innovation. The creation of an AI or robotics agency would enact both fewer regulations and a more comprehensive and efficient policy framework.

Finally, a federal agency is needed to ensure the public is protected in a world with expanded AI and automation. Currently, policies are being negotiated between a distributed collection of stakeholders, but because the agencies are currently constrained in their regulatory authority, the public interest and the voice of citizens are similarly constrained. A federal robotics and AI agency is a critical need to help steer the United States through the disruptive waves to ensure we are all best served as we enter the Second Machine Age.

References

- ¹ Roberts, Adam. *The History of Science Fiction*. (New York: Palgrave MacMillan, 2006), 168.
- ² Gevarter, William. *Intelligent Machines*. (Upper Saddle River: Prentice-Hall Publications, 1985), 8.
- ³ Gevarter, William. *Intelligent Machines*. 150.
- ⁴ Popper, Nathaniel. "The Robots are Coming for Wall Street." *New York Times*. <https://www.nytimes.com/2016/02/28/magazine/the-robots-are-coming-for-wall-street.html>. Accessed 4/16/17
- ⁵ Ford, Martin. *The Rise of the Robots*. 1-12.

- ⁶ Sills, Anthony. "IBM Ross and Watson Tackle the Law." IBM (blog). <https://www.ibm.com/blogs/watson/2016/01/ross-and-watson-tackle-the-law/>. Accessed 3/21/17.
- ⁷ Son, Hugh. "JP Morgan Does in Seconds What Took Lawyers 360,000 Hours." Bloomberg News. <https://www.bloomberg.com/news/articles/2017-02-28/jpmorgan-marshals-an-army-of-developers-to-automate-high-finance>. Accessed 4/15/17.
- ⁸ McKinsey Global Institute. "A Future that Works: Automation, Employment and Productivity". January 2017.
- ⁹ Brynjolfsson, Erik and McAfee, Andrew. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. (New York: W. W. Norton Company, 2014), 2-12.
- ¹⁰ McKinsey Global Institute. "A Future that Works: Automation, Employment and Productivity".
- ¹¹ Leila Takayama, Wendy Ju, and Clifford Nass. "Beyond dirty, dangerous and dull: what everyday people think robots should do." *Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction* (2008), 25-32.
- ¹² Association for Advancing Automation. "Robots Fuel the Next Wave of US Productivity". White Paper. October 2015.
- ¹³ A Roadmap for Robotics: From Internet to Robotics. 2016 Edition.
- ¹⁴ Ibid.
- ¹⁵ McKinsey Global Institute. "A Future that Works: Automation, Employment and Productivity".
- ¹⁶ McCorduck, Pamela. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. (Nantick: AK Peters Ltd, 2004), 111-136.
- ¹⁷ William T. Moyer. "The Army Sponsored Revolution." U.S. Army Archives. January 1996. <http://ftp.arl.mil/~mike/comphist/96summary/index.html>. Accessed 4/01/17.
- ¹⁸ McCartney, Scott. *ENIAC: The Triumphs and Tragedies of the World's First Computer*. (New York: Walker & Co., 1999), 7-18.
- ¹⁹ William T. Moyer. "The Army Sponsored Revolution."
- ²⁰ Turing, Alan. "Computer Machinery and Intelligence." *Mind*. 59 (1960): 433-460.
- ²¹ Turing, Alan. "Computer Machinery and Intelligence." *Mind*. 59 (1960): 433-460.
- ²² Ibid
- ²³ Samadi, David. "History of Robotic Surgery." Blog. <http://www.roboticoncology.com/history/>. Accessed 4/14/17
- ²⁴ Dr. Robert Leighty. "DARPA ALV Summary". U.S. Army Engineers Topographic Laboratories. Report #ETL-R-085. 1986
- ²⁵ Pomerleau, Dean A. "An Autonomous Land Vehicle in a Neural Network". Carnegie Mellon University, 1989.
- ²⁶ <http://archive.darpa.mil/grandchallenge/>
- ²⁷ Tactical Technology Office. "Broad Agency Announcement: DARPA Robotics Challenge." DARPA. DARPA BAA-12-39. April 10, 2012.
- ²⁸ McCorduck, Pamela. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. 111-136.
- ²⁹ McCarthy, John, et. al. *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*. 31 August 1955. retrieved April 17th, 2017.
- ³⁰ McCorduck, Pamela. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. 111-136.
- ³¹ Crevier, Daniel. *AI: The Tumultuous History of the Search for Artificial Intelligence*. (New York: Basic Book, 1993), 52-107.
- ³² Adjusted for inflation, that means each lab received approximately 24 million dollars in today's money per year.
- ³³ McCorduck, Pamela. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. 111-136.
- ³⁴ Crevier, Daniel. *AI: The Tumultuous History of the Search for Artificial Intelligence*. 52-107.
- ³⁵ Dr. Robert Leighty. "DARPA ALV Summary". U.S. Army Engineers Topographic Laboratories. Report #ETL-R-085. 1986
- ³⁶ Ibid.
- ³⁷ Hsu, Feng-Hsiung. "Behind Deep Blue: Building the Computer that Defeated the World Chess Champion". (Princeton: Princeton University Press, 2002), 92-97.
- ³⁸ Hsu, Feng-Hsiung. "Behind Deep Blue: Building the Computer that Defeated the World Chess Champion". 107.
- ³⁹ Kaplan, Jerry. *Humans Need Not Apply: A Guide to Wealth and Work in the Age of Artificial Intelligence*. (New Haven: Yale University Press, 2015), 23-28.
- ⁴⁰ Cha, Arianna. "IBM Supercomputer Watson's Next Feat? Taking on Cancer." Washington Post. <http://www.washingtonpost.com/sf/national/2015/06/27/watsons-next-feat-taking-on-cancer/>. Accessed 4/18/17.
- ⁴¹ Jesus, Cecille. "IBM Lawyer Ross Has Been Hired By Its First Official Law Firm." Futurism. <https://futurism.com/artificially-intelligent-lawyer-ross-hired-first-official-law-firm/>. Accessed 4/7/17
- ⁴² Mozur, Paul. "Google's AlphaGo Defeats Chinese Go Master in Win for AI." New York Times. <https://www.nytimes.com/2017/05/23/business/google-deepmind-alphago-go-champion-defeat.html>. Accessed 8/7/17.
- ⁴³ Mozur, Paul. "Google's AlphaGo Defeats Chinese Go Master in Win for AI."

44 Rankin, Keith. "The Dark Secret at the Heart of AI." MIT Technology Review.
<https://www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/>. Accessed 9/1/17.

47 Calo, Ryan. "Why We Need a Federal Agency on Robotics." *Scientific American*.
<https://www.scientificamerican.com/article/why-we-need-a-federal-agency-on-robotics>. Accessed 3/2/17.

48 Association for Advancing Automation. "Robots Fuel the Next Wave of US Productivity".

49 Ford, Martin. *The Rise of the Robots*. Pg. 32-45.

50 Lavrinc, Davon. "A Peek Inside Tesla's Robotic Factory". *Wired*. July 16, 2013.
<https://www.wired.com/autopia/2013/07/tesla-plant-video/>. Accessed 2/13/17.

51 Occupational Safety and Health Administration Accidents Database. Keyword "Robot."

52 Ibid.

53 Ibid.

54 Lathe is a machine used to shape wood, metal, or other materials by a rotating drive

55 Occupational Health and Safety Administration Accidents Database. Keyword "Robot".

56 Occupational Safety and Health Administration, Safety Standards. 01-12-002 [PUB 8-1.3].

57 Lipson, Hod and Kurman, Melba. "Driverless: Intelligent Cars and the Road Ahead." (Cambridge: MIT Press, 2016), 6-10.

58 Occupational Safety and Health Administration. Technical Manual.
https://www.osha.gov/dts/osta/otm/otm_iv/otm_iv_4.html

59 Ibid.

60 To date, 38 incidents have been reported to OSHA relating to a robot with 25 resulting in deaths.

61 Ford, Martin. *The Rise of the Robots*. 23-24.

62 Food & Drug Administration. <https://www.fda.gov/>

63 FDA. "Products and Medical Procedures."
<https://www.fda.gov/MedicalDevices/ProductsandMedicalProcedures/SurgeryandLifeSupport/ComputerAssistedSurgicalSystems/default.htm>

64 Whether the advertised greater precision of surgical robots translates into better health outcomes is still being investigated. See: Payne TN, Dauterive FR. "A Comparison of Total Laparoscopic Hysterectomy to Robotically Assisted Hysterectomy: Surgical Outcomes in a Community Practice." (2008) and Barbash, Gabriel, and Glied, Sherry. "New Technology Healthcare Costs—The Case of Robotically Assisted Surgery." (2013) for contrasting views.

65 Food and Drug Administration. "Device Regulation and Guidance."
<https://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/Overview/>

66 21CFR807.92(a)(3))

67 Food and Drug Administration. "Device Regulation and Guidance."
<https://www.fda.gov/MedicalDevices/DeviceRegulationandGuidance/HowtoMarketYourDevice/PremarketSubmissions/PremarketNotification510k/default.htm>

68 <https://www.tesla.com/autopilot>. Accessed 3/22/17.

69 <https://www.tesla.com/blog/all-tesla-cars-being-produced-now-have-full-self-driving-hardware>. Accessed 3/22/17.

70 Bhuiyan, Johana. "A Series of US State Laws Could Prevent Google and Uber from Operating Self-Driving Cars". *Recode*.
<https://www.recode.net/2017/2/25/14738966/self-driving-laws-states-gm-car-makers>. Accessed 4/27/17.

71 Mele, Christopher. "In a Retreat, Uber Ends Its Self-Driving Car Experiment in San Francisco." *New York Times*.
<https://www.nytimes.com/2016/12/21/technology/san-francisco-california-uber-driverless-car.html>. Accessed 3/21/17.

72 Muoio, Danielle. "These 19 Companies are Racing to Build Self-Driving Cars in the Next 5 Years." *Business Insider*. <http://www.businessinsider.com/companies-making-driverless-cars-by-2020-2017-1>. Accessed 5/3/17.

73 NHTSA. "Federal Automated Vehicles Policy." *Policy Guidance*. September 2016.

74 Ibid.

75 Ibid.

76 Ibid.

77 Ibid.

78 Ibid.

79 Federal Aviation Administration, *FAA Aerospace Forecast*, FAA-1636.
https://www.faa.gov/data_research/aviation/aerospace_forecasts/media/FY2016-36_FAA_Aerospace_Forecast.pdf. Accessed 3/06/17.

80 Model Aircraft Museum. "Aeromodeling History."

81 *FAA Modernization and Reform Act*, Public Law 112-95, US Statutes at Large 126 (2012).

82 Ibid.

83 Federal Aviation Agency. "Operation of Small Unmanned Aircraft Systems, Final Rule." *Federal Register* 14 CFR 21. 2016.

84 Ibid.

85 Sifton, John. "A Brief History of Drone Strikes." *The Nation*. <https://www.thenation.com/article/brief-history-drones/>. Accessed 8/8/16.

86 Ibid.

87 Ibid.

88 Currier, Cora. "Drone Strikes Test Legal Grounds for War on Terror."
<https://www.propublica.org/article/drone-strikes-test-legal-grounds-for-war-on-terror>. Accessed 8/9/17.

⁸⁹ Ibid.

⁹⁰ Jaffer, Jameel. "How the U.S. Government Justifies Drone Strikes: Targeted Killing, Secrecy, and the Law." *The Guardian*. <https://www.theguardian.com/us-news/2016/nov/15/targeted-killing-secrecy-drone-memos-excerpt> Accessed 8/8/17.

⁹¹ White House Memorandum. "Presidential Policy Guidance." ACLU. <https://www.aclu.org/foia-document/presidential-policy-guidance>. Accessed 8/9/17.

⁹² Hussain, Murtaza. "American Drone Assassinations May Violate International Law, Experts Say." *The Intercept*. <https://theintercept.com/2015/10/29/american-drone-assassinations-may-violate-international-law-experts-say/>. Accessed 8/9/17.

⁹³ Boone, Jon. "US Drone Strikes Could be Classed as War Crimes, says Amnesty International." *The Guardian*. <https://www.theguardian.com/world/2013/oct/22/amnesty-us-officials-war-crimes-drones>. Accessed 8/9/17.

⁹⁴ Jaffer, Jameel. "How the U.S. Government Justifies Drone Strikes: Targeted Killing, Secrecy, and the Law." *The Guardian*.

⁹⁵ Markoff, John and Rosenberg, Mathew. "The Pentagon's 'Terminator Conundrum': Robots that Could Kill on Their Own."

<https://www.nytimes.com/2016/10/26/us/pentagon-artificial-intelligence-terminator.html>. Accessed 8/7/17.

⁹⁶ Ibid.

⁹⁷ Chappell, Bill. "Researchers Warn Against Autonomous Weapons Arms Race." NPR.

<http://www.npr.org/sections/thetwo-way/2015/07/28/427189235/researchers-warn-against-autonomous-weapons-arms-race>. Accessed 8/8/17.

⁹⁸ Chappell, Bill. "Researchers Warn Against Autonomous Weapons Arms Race." NPR.

⁹⁹ Russell, Staurt, et. al. "Autonomous Weapons: An Open Letter From AI & Robotics Researchers." *The Future of Life Institute*. <https://futureoflife.org/open-letter-autonomous-weapons/>. Accessed 8/8/17.

¹⁰⁰ Ibid.

¹⁰¹ Ibid.

¹⁰² Morris, David. "UN Moves toward Possible Ban on Autonomous Weapons." *Forbes*. <http://fortune.com/2016/12/24/un-ban-autonomous-weapons/>. Accessed 8/9/17.

¹⁰³ Kaplan, Jerry. 24-26.

¹⁰⁴ Brynjolfsson, Erik and McAfee, Andrew. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. 2-12.

¹⁰⁵ Nader, Ralph. "Nader Proposes Reviving Office of Technology Assessment". *Independent Political Report*. <http://independentpoliticalreport.com/2010/06/nader-proposes-reviving-congressional-office-of-technology-assessment/>. Accessed 4/26/17.

¹⁰⁶ Ibid.

¹⁰⁷ Calo, Ryan. "Why We Need a Federal Agency on Robotics."

Author Biography

Bryan McMahon is a recent graduate from Duke University, where he majored in Biology and went to every Duke basketball game he could. In his last year and a half at Duke, he worked closely with Duke's new Science and Society department studying the growing intersections of science, technology, and society. This work included conducting original research as a member of Duke's Science, Law, and Policy, Lab looking at the evolving legal and policy issues of emerging science and technologies like 3D printing, direct-to-consumer genetic and neuroscientific technology, and the growing impact of neuroscience in the courtroom. More recently, Bryan had the opportunity to be a Science Policy Fellow at Duke and track and analyze upcoming science and technology policy through weekly published briefs on scipol.duke.edu. He specifically concentrated on artificial intelligence and robotics policy and hopes to continue working in this area in the future. He now resides in Arlington, Virginia with his wonderful girlfriend, Hayley, their cat, and four geckos.

Acknowledgements:

I would like to thank Dr. Michael Clamann for his deep knowledge of the subject matter and his guidance and advisement throughout the course of this paper. I would also like to thank Dr. Nita Farahany for first opening the door allowing me to study the intersections of science and society. In addition, I would like to acknowledge Thomas Williams for his help and guidance with all things science policy. And finally, I would like to thank my girlfriend, Hayley, for being an editing extraordinaire.