

**REPORT DOCUMENTATION PAGE**

Form Approved  
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.  
**PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**

<b>1. REPORT DATE</b> (DD-MM-YYYY) 17/04/2020	<b>2. REPORT TYPE</b> Master of Military Studies (MMS) thesis	<b>3. DATES COVERED</b> (From - To) AY 2019-2020
--	--	---

<b>4. TITLE AND SUBTITLE</b> Artificial Intelligence Underestimated: Why Decision-Makers Must Prioritize AI-Focus	<b>5a. CONTRACT NUMBER</b> N/A
	<b>5b. GRANT NUMBER</b> N/A
	<b>5c. PROGRAM ELEMENT NUMBER</b> N/A

<b>6. AUTHOR(S)</b> Broersma, Andries T. (Major)	<b>5d. PROJECT NUMBER</b> N/A
	<b>5e. TASK NUMBER</b> N/A
	<b>5f. WORK UNIT NUMBER</b> N/A

<b>7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)</b> USMC Command and Staff College Marine Corps University 2076 South Street Quantico, VA 22134-5068	<b>8. PERFORMING ORGANIZATION REPORT NUMBER</b> N/A
--	--

<b>9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)</b> N/A	<b>10. SPONSOR/MONITOR'S ACRONYM(S)</b>
	<b>11. SPONSOR/MONITOR'S REPORT NUMBER(S)</b> N/A

**12. DISTRIBUTION/AVAILABILITY STATEMENT**  
Approved for public release, distribution unlimited.

**13. SUPPLEMENTARY NOTES**

**14. ABSTRACT**

In the next decade, the United States is at significant risk of losing AI leadership to China with a clear risk to international security and the balance of power. As current AI trends show revolutionary change in the military through autonomous systems, in the economy through risk of massive unemployment, and in the information environment through eroding social trust, policymakers must prioritize AI focus to prepare for a near, risky, and uncertain AI future. While modern AI has been around since the 1950s, future AI advancements is potentially so rapid that the next few decades will see an AI revolution with greater impact on life than nuclear and biological threats or the effects of climate change. This paper examines the driving forces behind evolving AI (i.e., computing power, AI experts, and data), some of the current and future trends of AI in the military, economic, and information domain, and the battle for AI leadership between the United States and China in the next decade. In the military domain, AI has the potential to revolutionize the way of war with intelligence analysis, improving situational awareness, autonomous systems, robotics, enhanced decision-making, and cyber capabilities. In the information domain, AI can become so capable in behavioral data analysis, pattern recognition, distributing disinformation, creating deep fakes, or manipulate content people see, that there is a risk for social trust to erode. In the economic domain, AI will radically alter how people work, and if the tempo of AI job replacement outpaces the creation of new jobs, there is a risk of massive unemployment and rise in inequality. All these effects impact international security and the balance of great powers in a future where China is likely to dominate the field of AI. This future prospect demands full attention of policymakers, decision-makers, and the wider community in general. While the United States currently enjoys AI leadership, China is catching up fast. Being in a more favorable position when it comes to abundant data, persistent entrepreneurs, and progressive government support, China is only outmatched by the United States in AI expertise itself. By analyzing multiple potential future worlds using AI leadership and the effect of AI on military power and unemployment as critical uncertainties, this paper shows that most scenarios describe a future where AI poses a risk to international security. Especially the scenarios where China leads the ever-more important field of AI calls for strategic decision-makers need to make AI a top priority and increase focus and resources. AI focus should include increased government funding for education and research, in part to mitigate most AI advancements originating in the commercial sector, increased government support for entrepreneurs to set the right conditions to advance and stay ahead of Chinese counterparts, improved policy on data collection to mitigate the disadvantage of ever-growing data abundance in China, and finally, increased international cooperation and agreements on AI policy.

**15. SUBJECT TERMS**  
Artificial Intelligence, Strategy, China, Superintelligence, Future,

<b>16. SECURITY CLASSIFICATION OF:</b>			<b>17. LIMITATION OF ABSTRACT</b>	<b>18. NUMBER OF PAGES</b>	<b>19a. NAME OF RESPONSIBLE PERSON</b>	
<b>a. REPORT</b>	<b>b. ABSTRACT</b>	<b>c. THIS PAGE</b>			USMC Command and Staff College	
Unclass	Unclass	Unclass	UU	44	<b>19b. TELEPHONE NUMBER</b> (Include area code) (703) 784-3330 (Admin Office)	

*United States Marine Corps  
Command and Staff College  
Marine Corps University  
2076 South Street  
Marine Corps Combat Development Command  
Quantico, Virginia 22134-5068*

MASTER OF MILITARY STUDIES

---

**Artificial Intelligence Underestimated**

Why Decision-Makers Must Prioritize AI-Focus

SUBMITTED IN PARTIAL FULFILLMENT  
OF THE REQUIREMENTS FOR THE DEGREE OF  
MASTER OF MILITARY STUDIES

AUTHOR: Major Andries Broersma (RNLMC)

AY 2019-20

---

Mentor and Oral Defense Committee Member: Dr. Benjamin Jensen

Approved: //signed//

Date: April 17, 2020

Oral Defense Committee Members: Dr. Craig Hayden and Dr. Brandon Valeriano

Approved: //signed//

Date: April 17, 2020

DISCLAIMER

THE OPINIONS AND CONCLUSIONS EXPRESSED HEREIN ARE THOSE OF THE INDIVIDUAL STUDENT AUTHOR AND DO NOT NECESSARILY REPRESENT THE VIEWS OF EITHER THE MARINE CORPS COMMAND AND STAFF COLLEGE OR ANY OTHER GOVERNMENTAL AGENCY. REFERENCES TO THIS STUDY SHOULD INCLUDE THE FOREGOING STATEMENT.

QUOTATION FROM, ABSTRACTION FROM, OR REPRODUCTION OF ALL OR ANY PART OF THIS DOCUMENT IS PERMITTED PROVIDED PROPER ACKNOWLEDGEMENT IS MADE.

## Preface

To me, the concept of an Artificial Intelligence (AI) has always been fascinating. While many people misuse the term that does not have a universally accepted definition, I like to think of AI as an intelligence that is non-biological. Without thinking too philosophical, the non-biological part is relatively easy to grasp. Intelligence, however, is more difficult to describe. I like how Max Tegmark puts it as “the ability to accomplish complex goals,” where the term “complex” is deliberately somewhat vague. This vagueness adds to the fascination, where the pinnacle of AI will be a future superintelligence that can accomplish any task and outsmart any human on all levels. How soon this artificial general intelligence arrives, or whether it arrives at all, is debated amongst scientists and armchair philosophers. While I am in the camp that thinks we will reach that point or Singularity this century, I also think that the currently existing narrow AI evolution has the potential of far reaching effects on our way of life. Such effects will have great impact on the military and the way of war, the economy with AI replacing our jobs, and the international security and balance of power. As I feel that the consequences of rapidly evolving AI are potentially more impactful in the next two decades than even nuclear and biological threats, or the effects of climate change, I feel there is a disconnect between relevance and focus. That is, a safe AI future does not get enough attention from policymakers, decision-makers, the media, or the minds of my friends and colleagues. This thought drove me to analyze some of the current leading trends and thoughts on AI and justify why AI technology needs more of our attention.

## Table of Contents

Disclaimer .....	ii
Preface.....	iii
Table of Contents .....	iv
Executive Summary .....	v
<b>Introduction</b> .....	<b>1</b>
Thesis .....	2
Method .....	2
<b>History</b> .....	<b>4</b>
<b>Terminology</b> .....	<b>8</b>
<b>Trends and Driving Forces</b> .....	<b>11</b>
AI in the Military Domain.....	12
AI in the Information Domain.....	14
AI and the Economy.....	15
Rapid AI Advancement.....	18
AI Leadership: China and the United States .....	20
<b>Future Scenarios</b> .....	<b>23</b>
Mostly Positive Future .....	23
A Future with Risks.....	25
A Dangerous Future .....	26
<b>The Way Ahead</b> .....	<b>27</b>
<b>Conclusion</b> .....	<b>29</b>
Annex A: Three Future 2030 AI Worlds .....	30
Endnotes.....	31
Bibliography .....	314

## Executive Summary

**Title:** Artificial Intelligence Underestimated: Why Decision-Makers Must Prioritize AI-Focus.

**Author:** Major Andries T. Broersma, Royal Netherlands Marine Corps (RNLMC).

**Thesis:** In the next decade, the United States is at significant risk of losing AI leadership to China with a clear risk to international security and the balance of power. As current AI trends show revolutionary change in the military through autonomous systems, in the economy through risk of massive unemployment, and in the information environment through eroding social trust, policymakers must prioritize AI focus to prepare for a near, risky, and uncertain AI future.

**Discussion:** While modern AI has been around since the 1950s, future AI advancements is potentially so rapid that the next few decades will see an AI revolution with greater impact on life than nuclear and biological threats or the effects of climate change.

This paper examines the driving forces behind evolving AI (i.e., computing power, AI experts, and data), some of the current and future trends of AI in the military, economic, and information domain, and the battle for AI leadership between the United States and China in the next decade.

In the military domain, AI has the potential to revolutionize the way of war with intelligence analysis, improving situational awareness, autonomous systems, robotics, enhanced decision-making, and cyber capabilities. In the information domain, AI can become so capable in behavioral data analysis, pattern recognition, distributing disinformation, creating deep fakes, or manipulate content people see, that there is a risk for social trust to erode. In the economic domain, AI will radically alter how people work, and if the tempo of AI job replacement outpaces the creation of new jobs, there is a risk of massive unemployment and rise in inequality. All these effects impact international security and the balance of great powers in a future where China is likely to dominate the field of AI. This future prospect demands full attention of policymakers, decision-makers, and the wider community in general.

**Conclusion:** While the United States currently enjoys AI leadership, China is catching up fast. Being in a more favorable position when it comes to abundant data, persistent entrepreneurs, and progressive government support, China is only outmatched by the United States in AI expertise itself. By analyzing multiple potential future worlds using AI leadership and the effect of AI on military power and unemployment as critical uncertainties, this paper shows that most scenarios describe a future where AI poses a risk to international security. Especially the scenarios where China leads the ever-more important field of AI calls for strategic decision-makers need to make AI a top priority and increase focus and resources.

AI focus should include increased government funding for education and research, in part to mitigate most AI advancements originating in the commercial sector, increased government support for entrepreneurs to set the right conditions to advance and stay ahead of Chinese counterparts, improved policy on data collection to mitigate the disadvantage of ever-growing data abundance in China, and finally, increased international cooperation and agreements on AI policy.



## Introduction

Most people are to some degree familiar with the concept of Artificial Intelligence (AI), but have different understandings of what it is. This is partly due to the invisibility of AI; it is usually not something you can directly see or touch. To the general public, AI is often a smart robot or computer that can perform human-like tasks. Within this realm—strengthened by Hollywood and popular science fiction—this robot or computer is often imagined as Stanley Kubrick’s HAL 9000 in *2001: A Space Odyssey* (1968), James Cameron’s *The Terminator* (1984), or the human-like robot in Alex Proyas’ *I, Robot* (2004). A popular example of a real-life robot, as seen in various scientific and entertainment media, is Hanson Robotics’ Sophia.<sup>1</sup> Given further thought, the general public recognizes that other popular examples of AI are Tesla’s self-driving cars, autonomous military drones, Facebook’s facial image recognition, or Apple’s assistant Siri—all with varying degrees of controversy.

It is relatively easy to recognize and understand the controversy of AI in self-driving cars or autonomous drones. Far more difficult to comprehend or grasp is AI’s potential future effect on the stability of the economy, the capabilities of the military, and security in the information domain. Best case, the effects are mostly positive, with AI making life safer, the economy more productive, and the military more efficient. However, as the future and impact of evolving AI is difficult to predict, the effects may also be far more negative and therefore severely impact international security. If AI alters the conduct of war and significantly increases military capabilities, if it heavily impacts the economy and leads to a large loss of jobs, and if AI technology is dominated by an adversary of the United States, e.g., China, a safe future international environment is in jeopardy. Furthermore, as AI advancement is potentially



exponentially fast, its effects are felt within the next few decades—sooner and with greater impact on life than climate change for example.

How AI will evolve and what its impact will be is debated amongst scientists and thinkers, where an AI future ranges from utopia to dystopia. This paper examines the driving forces behind evolving AI, some of the current and future trends of AI in the military, economic, and information domain, and the battle for AI leadership between the United States and China in the next decade. Then, by creating multiple potential future worlds using these trends as critical uncertainties, this paper shows that most scenarios describe a future where AI poses significant risk to international security. Where climate change is a hot topic in the media today, AI is far less at the forefront. The aim of this paper is to increase AI awareness of the reader and argue that strategic decision-makers need to make AI a top priority and increase focus and resources in the AI field.

### *Thesis*

In the next decade, the United States is at significant risk of losing AI leadership to China with a clear risk to international security and the balance of power. As current AI trends show revolutionary change in the military through autonomous systems, in the economy through risk of massive unemployment, and in the information environment through eroding social trust, policymakers must prioritize AI focus to prepare for a near, risky, and uncertain AI future.

### *Method*

In order to show the importance of preparing for a future in which AI plays an important—if not the most important—role, this paper creates three potential future worlds for

the next decade using the scenario-developing method developed by Peter Schwartz, one of the world's leading futurists. These potential worlds help visualize the impacts and implications of the focal issue or decision of this paper: do decision-makers need to increase focus and resources on AI development and policy? In his book *The Art of the Long View*, Schwartz identifies choosing the focal decision as the first step in an eight-step process that helps inform that decision.<sup>2</sup>

Next, this paper provides a brief overview of AI history and examines key AI terminology before proceeding to the second step: listing the key factors that influence the success or failure of the focal decision—key trends that decision-makers need to understand in order to make the right decision. The scope of this paper is to examine four trends: China is likely to dominate the AI field, AI is increasingly impacting military power, AI is increasingly affecting the information domain, and AI is increasingly capable of replacing human jobs. Together with step three—listing the driving forces that influence these four trends—these are the two most research-intensive steps of the process. The focus of those driving forces will be on what drives the speed of AI progress: increasing computing power, the importance of engineering talent, new scientific breakthroughs, academics and large companies, large data sets, the importance of experts and entrepreneurs, and the role of government support.

In step four, all the factors, trends, and driving forces are ranked in order of importance. The top few are then used in step five—selecting scenario logics—as axes along which the eventual scenarios will differ. Here, we will use three axes: will China or the United States exercise AI leadership, to what extent will AI impact military power and capabilities, and to what extent will AI cause unemployment on a global scale? By looking at the extremities of each axis (e.g., limited versus large negative effects on the human job market), this leads to eight

possible scenarios (two to the power of three) that will be grouped into three possible worlds: a mostly optimistic one, a future with risk and challenges to international security, and a dangerous future. In step six, these worlds are fleshed out using analysis of the other trends and driving forces.

Lastly, looking at the way ahead, step seven will look at some implications of a poor decision on the focal issue of step one, and help stress the importance and urgency of increased strategic focus on AI. Step eight will identify leading indicators and signposts to keep informing the decision-making process of step one in predicting the world that is most likely to unfold.

## History

Prior to analyzing current trends and building future scenarios, it is useful to explore the history of AI. While previously mentioned technologies and applications are relatively new, AI itself is not. Although the ancient Greeks already wrote about artificially created life-forms, the field of modern AI started in the 1950s. In his paper *Computing Machinery and Intelligence* (1950), Alan Turing proposed the question “Can Machines Think?”<sup>3</sup> By formulating a game, which he called “the imitation game,” Turing constructed a method to test a computer’s “intelligence.” The goal of the test was to convince human interrogators they were communicating with another human instead of a computer. Many new improved variants of this test are now still referred to as “The Turing Test.”<sup>4</sup> Soon after—intrigued by the question whether machines can model intelligent life—scientists made subsequent developments in this field, including a neural network machine modeling the behavior of a rat in a maze.<sup>5</sup> In 1952, Arthur Samuel built the first self-learning Checkers-playing program.<sup>6</sup>

All these successes inspired scientists such as John McCarthy (1927-2011), computer scientist and professor at Stanford University. He first coined the term “Artificial Intelligence” in 1956 at a conference at Dartmouth College in Hanover, New Hampshire.<sup>7</sup> McCarthy wrote the proposal for the conference, stating that he aimed "to proceed on the basis of the conjecture that every aspect of learning or any other feature of intelligence can in principle be so precisely described that a machine can be made to simulate it."<sup>8</sup> Together with a small team of scientists, including Marvin Minsky, MIT cognitive scientist, Nathaniel Rochester, designer of the IBM 701, and Claude Shannon, founder of the digital circuit design theory, he was very optimistic of the rapid progress of AI. In the same proposal, they forecasted that during that summer of 1956, the small selected group of scientists could find out how to create machines that improve themselves and solve problems that were then only reserved for humans. This optimistic view of rapid AI-progress continued for almost two decades, with initial developments in image processing, recognition of speech and written language, and the introduction of Rosenblatt’s Perceptron.<sup>9</sup>

However, creating an AI proved difficult. With the lack of understanding on how to imitate the complex human brain with machine learning programs—so called neural networks—and the limited computing power available, AI-progress stalled in the mid-1960s. As a result of disappointing advancements in machine translation, i.e., obstacles hindering the progress of the perceptron and realistic evaluations such as the critical *Lighthill Report* (1973), government funding decreased significantly.<sup>10</sup> The United States Defense Advanced Research Agency (DARPA) now funded less fundamental work and focused more on applied AI projects. The period that followed became known as the first AI winter and lasted until 1980.

The 1980s saw a revival of interest in the field of AI, this time with a focus on commercial products. Additionally, both the British and United States government increased funding to compete with the Japanese efforts on their Fifth Generation Computer System (FGCS) and its potential threat to Western technological dominance.<sup>11</sup> In 1983, the Strategic Computing Initiative of the United States government funded various projects with the goal of creating AI software in a ten-year time frame; the Department of Defense (DoD) spent over one billion dollars on the project.<sup>12</sup> However, new spending was cancelled in 1988 as it became apparent again that expectations on machine intelligence were too high. As director of DARPA's Information Science and Technology Office (ISTO) from 1986-1989, Jacob Schwartz stated that AI research had "very limited success in particular areas, followed immediately by failure to reach the broader goal at which these initial successes seem at first to hint." Again, although less severe than in 1980, general interest and funding decreased, and many companies even rebranded their company names to something seemingly less associated with AI to maintain their funding.<sup>13</sup>

In the mid-1990s, research and interest picked up, with achievements such as Cyberknife in 1994, a robot that could surgically remove tumors, and Sojourner in 1997, a robot built by NASA that could semi-autonomously operate on Mars.<sup>14</sup> Perhaps the most well-known AI achievement in the 90s is that of IBM's supercomputer Deep Blue beating chess grandmaster Garry Kasparov in 1997. This popularized application of man against machine continued in the next century with IBM's computer system Watson beating its human competitors in the television quiz *Jeopardy!* in 2011, and recently Google's AlphaGo beating Korean Go player Lee Sedol in 2016. AlphaGo's achievement was mainly possible due to great progress in the field of "deep learning," a subclass of machine learning. Discovered by neural-nets engineer

Geoffrey Hinton, deep learning significantly increased the performance of neural networks that had disappointed in the past. In October 2012, Hinton's neural network won the prominent ImageNet computer-vision contest by identifying objects with almost twice the accuracy as its opponents. The result that neared human performance was a turning point for AI research, now with a focus on deep learning, and "convinced lots and lots of people who had been very skeptical before."<sup>15</sup> Just a few years later in 2015, the results even surpassed human performance.

With the new developments and achievements, government funding also again increased. In 2007, DARPA started funding a program to create a Cognitive Technology Threat Warning System (CT2WS), a system that would integrate a human brain (e.g., an infantryman) with computer vision to spot danger more efficiently. Funding for an Autonomous Real-Time Ground Ubiquitous Surveillance-Imaging System (ARGUS-IS), an advanced AI-enhanced camera system, started that same year. Similarly, the European Union (EU) started funding AI research programs from 2007-2013 through its Framework Programs for Research and Technological Development (EU-FP7) with a budget of over 50 billion USD.<sup>16</sup> Today, AI continues to make progress but the rate of development in the future is debated, including by leading AI experts. Nonetheless, government involvement and funding in AI development have only increased. The United States mentions a new category for AI in the President's Budget Request for 2020, with a budget of one billion USD for non-defense R&D purposes.<sup>17</sup> The increased funding is in part a response to China, the other "superpower" now heavily investing in AI.

Since that conference in Dartmouth in 1956, the field of AI has seen significant progress. It helped solve various "problems", varying in complexity and usefulness, from spam blocking to voice and image recognition, and self-driving cars to tumor recognizing systems. And although

many of the possibilities envisioned by Turing and McCarthy seem to be in the distant future, progress continues rapidly. With continuous exponential growth of computing power, and ever-growing data sets, the large impact of AI in the future is difficult to deny.

### **Terminology**

The term Artificial Intelligence, also referred to as machine intelligence or computational intelligence, does not have one universal definition. The *English Oxford Living Dictionary* defines AI as “The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision-making, and translation between languages.”<sup>18</sup> The online dictionary *Merriam-Webster* defines it as “A branch of computer science dealing with the simulation of intelligent behavior in computers,” or “The capability of a machine to imitate intelligent human behavior.”<sup>19</sup> As elements of these definitions are subject to change due to the evolution of AI, Max Tegmark, in his book *Life 3.0*, defines AI more broadly as “Non-biological intelligence,” where intelligence is the “Ability to accomplish complex goals.”<sup>20</sup> The word “complex” is deliberately somewhat broad, as it is difficult to quantify and compare complexity and intelligence. Comparing the intelligence of a system that excels in speech recognition to one that beats any human in a game of chess makes no direct sense. It is, however, useful to make a distinction between “narrow” and “broad” intelligence.

Narrow AI excels in one or a small set of specific tasks. The term narrow refers to the limited set of predefined tasks the AI is programmed to do, not the ability to accomplish them, or how well it does so in comparison to human intelligence. Some systems are designed for one specific task, such as earlier mentioned Google’s AlphaGo. Others accomplish a slightly broader

range, e.g., Google’s DeepMind DQN system that can play a wide range of Atari games—often better than humans. Yet others go slightly further and combine multiple narrow systems to solve more broader tasks, such as self-driving cars. However, all current AI is technically narrow, or sometimes called “weak AI”; it is merely responding to how it is programmed and does not understand or derive meaning from its task.

Broad AI, better known as Artificial General Intelligence (AGI) refers to a system that, when maximally broad, can accomplish any goal, including learning. The term AGI was popularized in 2002 by Ben Goertzel, Shane Legg, and Mark Gubrud, as Goertzel was unsatisfied with the controversial title “Real AI” for his new book.<sup>21</sup> Alternative terms are “human-level AI” and “strong AI.” However, as Tegmark points out, these terms may cause confusion. A simple calculator performs human-level tasks, but in a very narrow sense. And, as the antonym of strong AI is weak AI, that results in Google’s supercomputer AlphaGo unnecessarily being called weak.<sup>22</sup> However, while currently many narrow AI systems surpass human capabilities, current technology has not yet achieved human-level AGI, a system with the ability to accomplish any cognitive task at least as well as humans, or even a Superintelligence, an AGI far beyond human level.

All current powerful AI systems are based on artificial neural networks. Older AI systems were often rule-based, where programmers encoded logical rules (IF-THEN statements) into a system. Ideally this led to every situation (IF x) resulting in the correct answer or action (THEN y). For simple problems, such as winning at a game of tic-tac-toe, this works well, even with very little computing power. For more complex problems, e.g., winning a game of chess, the correct answer, or move, is not always obvious, and calculating all possible moves by brute force proves impossible.<sup>23</sup> The addition of human knowledge from experts (e.g., chess masters) into



the system greatly enhances the capabilities. Combined with great computing power, these “expert systems” or “symbolic systems” can achieve impressive results, with earlier mentioned IBM’s Deep Blue as an example.

Machine learning (ML), an AI subfield, uses a different approach. Instead of relying on pre-programmed rules, ML uses artificial neural networks that mimic the architecture of the organic (human) brain, with multiple layers of connected neurons representing biological neurons and synapses. The system does not use hand-coded software programs, but makes predictions after learning from, and identifying patterns in, large data sets. As an example, to distinguish a handwritten “X” from an “O,” a rule-based system requires pre-programmed rules that describe the characteristics of those characters. ML identifies those characteristics itself by processing millions of images labeled “X” or “O” (by humans) respectively. However, despite the explosion of digital data as a result of the internet, and increasing computing power, ML networks in the early 2000s were still limited in their capabilities. As complex problems required many additional layers of artificial neurons, researchers lacked an efficient method of training these layers. A breakthrough occurred when Geoffrey Hinton found a way to efficiently train those new layers and hence significantly improved the power of neural networks—now known as deep learning.

Deep learning is an advanced variant and subclass of ML, optimized for complex calculations and less suited for simple ones. Here, the algorithm trains itself by using large amounts of data and recognizing deep patterns and correlations, including ones that seem irrelevant to human interpreters. Ideally the data input is labeled and structured, as described above, but this is not required as different layers within the network also classify the multiple levels of output. Today, deep learning is a promising technique for AI in solving real-world

problems. However, it currently requires a narrow domain, a concrete goal, and large amounts of relevant data to be successful. The first two requirements indicate that deep learning is still mostly valid in the domain of narrow AI. The latter requirement is less of a problem as relevant data quickly becomes more abundant in many fields. In fact, the large amount of structured and unstructured data is so large, that it is difficult—or impossible—to process by traditional methods. Both this large volume of data, as well as the field that deals with analyzing and extracting useful information from it, is referred to as “big data.”

### **Trends and Driving Forces**

The field of AI is not new. It continues to evolve, has more and more useful applications, receives increasing funding, has government and commercial interest, and is likely to have a growing impact on everyday life—with positive and negative effects. Already founded in 1979, the Association for the Advancement of Artificial Intelligence (AAAI) promotes responsible use of AI, and aims to “increase public understanding of AI, improve the teaching and training of AI practitioners, and provide guidance for research planners and funders concerning the importance and potential of current AI developments and future directions.”<sup>24</sup> To prepare for an AI future, various countries—including the United States and China—recently developed specific strategic AI policies, dealing with both current and future AI implications. Although the timeframe and severity of AI impact is debated, an increasing number of scientists agree that it requires effort and time to well prepare for an AI future. As the impact is potentially large, so is the relevance of its discussion—labeled by Tegmark as “[T]he Most Important Conversation of Our Time.”<sup>25</sup> In 2015, Tegmark and other scientists founded the Future of Life Institute (FLI) with a mission to “catalyze and support research and initiatives for safeguarding life and developing optimistic visions of the future, including positive ways for humanity to steer its own course considering

new technologies and challenges.”<sup>26</sup> Next to biotechnology, nuclear technology, and climate change, AI is the focus of technologies with the largest impact on the future of mankind. A key argument for relevance is the notion that AI does not merely impact a few aspects of life, but numerous, including the military, information, and economic domain—affecting national and international security and stability. And, although scientific progress is often hyped and overestimated, significant effects may arrive sooner than many anticipate.

### *AI in the Military Domain*

In the United States DoD, DARPA, and the Intelligence Advanced Research Projects Agency (IARPA), both research organizations within the different services, initiate and manage most AI applications. A recent investment is Project Maven that enhances existing DoD systems with AI technology. With a focus on intelligence processing (e.g., drone imagery), this project vastly improves situational awareness (SA) and enhances the decision-making process. Another impactful application is gaining and exploiting dominance in the electromagnetic spectrum. A self-improving AI can send signals through contested environments while concurrently jamming that same signal. By using deep learning, the result is a means to dominate the spectrum. A third example is an application in command and control (C2). As intelligence collection systems (through AI) continue to faster gather increased amounts of data, AI can help process that same data—transforming it into information. This will aid commanders or leaders in better understanding their environment, and therefore make better informed decisions—some of which can even be delegated to an automated system without human intervention or approval.

As AI technology advances, the implications for military superiority are potentially large. Two key fields of military capabilities that AI applications will affect are those of cyber and

robotics. In 2016, Michael Rogers, director of the National Security Agency (NSA), stated that AI is “foundational to the future of cyber security.”<sup>27</sup> That same year, DARPA began Project Voltron, using autonomous cyber security systems to find and fix vulnerabilities in military systems.<sup>28</sup> As for robotics, with the cost decreasing, the demand and size of the market grows at an almost exponential rate. In combination with 3D printing technology, robotics enables states to operate millions of autonomous micro-drones. Unfortunately, such capabilities are not exclusive to the United States and its allies, and subsequently create opportunities for its adversaries—including hostile non-state actors. To illustrate this, Stuart Russel, professor at the University of California-Berkeley and member of the scientific advisory board of the FLI, presented a fictional warning video in November 2017, at an event hosted by the Campaign to Stop Killer Robots during the UN Convention on Conventional Weapons. The short video depicts a disturbing future in which lethal autonomous drones (i.e., weapons) are cheap and ubiquitous.<sup>29</sup> Although fictional, Russel warned that the technology is already available, and timely action is needed to prevent the dire prospect.

One possible consequence of cheap lethal robotics is a diffusion of power, comparable to the impact of cyber operations in the past. With cyber, smaller actors less powerful in a conventional way, including terrorists and criminals, acquired a relatively cheap means of affecting the world to a much larger extent than before; a similar trend is possible with evolving robotics. Whereas now a nation-state’s military potential is mainly determined by the number of young males in its population, this correlation might be far less relevant in the future. Through AI, robotics, and cyber, a state’s military power is likely to grow disconnected from its population size and economic power.

Next to the United States and China, another prominent nation investing in the application of AI, cyber, and robotics in the military is Russia. President Putin stated in 2017 that the nation that leads in AI will be the ruler of the world.<sup>30</sup> The Russian Military Industrial Committee recently approved a strategic plan stating that by 2025, a third of Russia's combat power consists of autonomous systems.<sup>31</sup> Through the Foundation for Advanced Studies, a defense research organization similar to DARPA, their focus is on robotics, including weaponized unmanned ground vehicles (UGV). Despite continuous cuts in defense budgets, the Uran-6, a demining-robot, and the Uran-9, a Lethal Autonomous Weapon System (LAWS) designed for combat operations, are in their final stage of development.<sup>32</sup> Achieving a leading role in their active pursuit of LAWS gives them a unique advantage. Even without the dominant role, disruptive capabilities must not be underestimated, as shown by Russia's influence in cyberspace without ever leading in internet technology.<sup>33</sup>

### *AI in the Information Domain*

In the information age, where both humans and machines access and process more information every day, AI has significant implications on information security. By exploiting behavioral data and recognizing patterns, autonomous programs, or bots, already target specific audiences, distribute disinformation (or propaganda), create fake news, and therefore influence and manipulate the content people see. This is arguably less impactful when used in tailoring advertisements for consumer products, but harmful when used in political elections, for example. In 2018, bots made up nearly 40 percent of internet traffic—of which half were bad bots (i.e., bots with a malign purpose).<sup>34</sup> On Twitter, despite machine learning models to counter them, between 9 and 15 percent of accounts are estimated to be bots—with increasing complexity.<sup>35</sup>

Deep fakes are another impactful phenomenon, where AI can mimic any person's voice, image, or video with a large enough data set. For voice, a recorded sample of a few minutes is enough to convincingly create new fake audio. When hundreds of hours of data are available—as is often the case with public figures—the result is often indistinguishable from authentic material for humans. Countermeasures driven by AI exist, but are not flawless. As it remains difficult to create universal filters, human intervention is often still necessary. Recently, the International Fact-Checking Network (IFCN) partnered with Google and created an algorithm with an 80 percent success rate in recognizing fake news.<sup>36</sup> To counter deep fakes, a partnership between Microsoft, Facebook, MIT, and others, initiated an industry-wide effort called the Deep Fake Detection Challenge (DFDC) to detect manipulated media.<sup>37</sup> Without such countermeasures, the reliability of information will decrease and inevitably erode social trust.

Another foreseeable improvement through AI in the information domain is that of data collection and analysis capabilities in the security and intelligence realm. AI technology can both increase the amount of data collected, as well as help analyze it. With big data, there already is substantially more data available than humans can analyze within a meaningful timeframe, and the growth rate is exponential. The International Data Corporation (IDC) estimates an annual overall data growth of 61 percent until 2025.<sup>38</sup> For national intelligence services, this is both an opportunity and a challenge, as nation-states with far outperforming data-analyzing capabilities, are at a major advantage.

### *AI and the Economy*

In the economic or financial domain, AI is used in a variety of ways. One application is in law enforcement by countering illicit financing operations that often support criminal

activities and terrorism. The United Nations (UN) estimates that current efforts only seize about one percent of illicit funds.<sup>39</sup> AI will improve these results by analyzing big data and identifying correlations between seemingly innocent behavior and fraud—far better than humans can. This very trend of AI outperforming human capabilities is one of controversy. As increasingly more jobs are better performed by machines than humans, this could disrupt the economic system and lead to an AI-induced crisis of unemployment and increased inequality.

PricewaterhouseCoopers predicts that by 2030, AI may contribute almost 16 trillion USD to the global economy.<sup>40</sup> This is a productivity growth of over one percent per year. To put this into perspective, the introduction of steam engines in the nineteenth century increased production only by about 0.3 percent annually.<sup>41</sup> In contrast to less impactful innovations that perform single tasks (e.g., the typewriter), innovations such as the steam engine, electricity, and the internet, called General Purpose Technologies (GPT), radically altered how people work. Such innovations are scarce, and AI can be added to this list. The question arises whether this extra wealth and productivity through AI progress leads to either additional jobs and increased human welfare, or massive unemployment and a rise in inequality. Erik Brynjolfsson, director of the MIT initiative on the Digital Economy and author of *The Second Machine Age*, calls it “the biggest challenge of our society for the next decade.”<sup>42</sup> Throughout history, technological progress never led to a rise of unemployment in the long term. During the industrial revolution, the introduction of the tractor decreased the requirement of farmers, but work displacement was temporary as factories created new jobs in the now industrial economy. Optimists say that AI will be no different as automation decreases prices, increases disposable income and the demand for more and newer products, and eventually increases and creates new jobs to meet this demand.

However, AI might be different as it can replace both manual and cognitive jobs, at a much higher pace than previous GPT.

Where the power loom in the early industrial revolution replaced high skilled weavers but subsequently created jobs for many more unskilled workers, AI will not facilitate the deskilling of economic production. Instead, AI may replace jobs without either creating new ones, or outperforming humans in those professions as well. Jobs particularly at risk of replacement are those with tasks that are optimization-based, have a structured environment, require low dexterity, and need little social interaction, e.g., telemarketers, radiologists, assembly-line inspectors, and truck drivers. Jobs that are in the safe zone for now are those that do require creativity, dexterity, and social skills, e.g., hairstylists, CEOs, and physical therapists.<sup>43</sup> Estimates of what percentage of the current workforce is potentially at risk in the next decades vary widely. A 2013 Oxford University study estimated that 47 percent of jobs in the US is at risk.<sup>44</sup> In contrast, in 2019 the Organization for Economic Cooperation and Development (OECD) estimated that, although 32 percent of jobs are likely to change radically, only 14 percent of jobs could disappear as a result of AI.<sup>45</sup> The 2019 McKinsey Group Institute (MGI) report is more positive, as it shows that only around 5 percent of occupations can be fully automated. However, within 60 percent of jobs, more than 30 percent of the tasks can be automated by today's technology.<sup>46</sup> Not all AI experts agree with this low percentage and predict that—as the steep progress of deep learning continues—these numbers will be much higher, around 40 percent of potential job replacement in 15 years.<sup>47</sup> Whether all jobs that potentially *can* be replaced *will* be replaced, and in what timeframe, remains to be seen as government regulations and social friction will most likely have a mitigating effect. Furthermore, cost is an important factor too. An AI doubly outperforming a human at ten times the cost is practically



useless. Nonetheless, the AI-induced risk of unemployment and disruption of the economic system requires great attention and precaution of all—most notably strategic policymakers.

### *Rapid AI Advancement*

As is often the case throughout history, mankind both underestimates and overestimates technological advancement.<sup>48</sup> AI progress experienced similar views, with periods of enthusiasm followed by “winters.” Sometimes, progress is counter intuitively slower or faster in certain areas. Known as Moravec’s paradox since the 1980s, it proved easier to train AI high-level intellectual or computational human abilities, than to give it the motor skills or perception of a child. One key requirement of AI, however, has been consistent and exponentially growing over the last 50 years: computing power. In 1965, Gordon Moore, co-founder of semiconductor chip manufacturer Intel, postulated that the number of transistors per square inch would double roughly every two years. Some argue that the trend, known as Moore’s Law, is near its end, with transistors now as small as 5 nm. While new barriers may slow down Moore’s Law, with new technologies—including quantum computing—there is no clear indication that progress will halt.

A second driver for AI advancement is engineering talent, with scientists causing breakthroughs with developments such as neural networks and machine learning. Such groundbreaking discoveries are then followed by a period of implementation, where theory becomes practice and entrepreneurs create real-world applications. Today, with the invention of deep learning in 2012, application is done though the third ingredient for AI advancement: data. Especially when computing power and engineering talent reach thresholds, data is the decisive factor in deep learning and AI progress. As Robert Mercer, founder of Cambridge Analytica

reportedly said, “there is no data like more data,” and weaker algorithms often outperform stronger well-designed ones, if exposed to significantly more data.

How rapid narrow AI will evolve, and whether it will evolve into AGI—or even a Superintelligence—is debated by experts. In a survey amongst the world’s leading experts in 2015, to guess the year we will have human-level AGI with at least 50 percent probability, the average answer was 2055, while some guessed over a hundred years or never. When asked again in 2017, the estimate had dropped to 2047.<sup>49</sup> Tegmark classifies the leading thoughts into three camps, centered around two questions: when will AGI happen, and will it be good for humanity?<sup>50</sup> Almost no expert thinks we will have AGI within the next few years, and few envision a dystopia where AGI will definitely be bad.

The first group of the three are what he calls digital utopians. They argue that AGI will most likely happen this century and the effect on humanity is mostly good. Amongst them are co-founder of Google Larry Page, roboticist Hans Moravec, AI pioneer Richard Sutton, and inventor Ray Kurzweil. In his book *The Singularity is Near*, Kurzweil sets the date for which human and machine intelligence merge—which he calls the Singularity—at 2045.<sup>51</sup> The second group are what Tegmark calls techno-skeptics. These thinkers also do not worry about AGI, because they think it is so difficult that it will either never happen, or at least not for a hundred years. Former chief scientist at Baidu, China’s equivalent of Google, Andrew Ng even fears that worrying too much would slow AI-progress down. Lastly, the group that envisions an AGI this century but thinks a good outcome is not guaranteed, is the beneficial-AI movement. Supporters of this mainstream view, including famous AI researcher Stuart Russel, entrepreneur Elon Musk, and Max Tegmark himself, argue that there are crucial questions that need answers first. And research needs to start today, in order to have answers ready when we need them.

Regardless of beliefs in a Superintelligence, the potential effects of ever advancing narrow AI with human capabilities justifies timely action and attention. If some of the more recent estimates become reality, then over 40 percent of jobs are at risk in the next two decades in the United States. Importance of rising AI also outweighs that of climate change—which is getting a lot of media attention—in its urgency and impact. AI might even provide the solution for global warming. Furthermore, as Tegmark points out, AI might dominate what happens with other global issues, including terrorism, poverty, migration, unemployment, and wars.<sup>52</sup> An AI arms race between the United States, China, and to a lesser extent Russia must be avoided.

#### *AI Leadership: China and the United States*

For the last two decades, the United States has dominated the world's technology markets with Silicon Valley borne companies such as Google, Microsoft, Amazon, and Facebook. This still holds today but one nation is catching up rapidly: China. A turning point for China's approach to AI occurred in May 2017, when Google's AlphaGo beat Chinese Go player Ke Jie, the number one ranked player in the world. In ancient Chinese culture, Go is not just a game but an artform Chinese scholars had to master. The loss of Ke Jie was a Sputnik moment for China. Less than two months later, in July 2017, the Chinese government released its *New Generation Artificial Intelligence Plan*, detailing how it plans to lead the AI frontier by 2030.<sup>53</sup> An ambitious strategic plan that outlines how in 2020 China's AI technology will be on par with the United States, with government funding of over 150 billion USD. By 2025, with increased funding surpassing 400 billion USD, China envisions AI to be their main driver for economic transformation. By 2030, China's goal is to be the world's primary AI innovation center, with a funding of one trillion USD; this is in line with its strategy to become an economic superpower, leading by innovation.

The White House has a similar plan, highlighted on its new website launched in March 2019.<sup>54</sup> The plan outlines the government's strategy and efforts in maintaining AI leadership, as described in the United States' AI initiative announced in Executive Order 13859. In summary, the initiative aims to pursue five pillars for advancing AI: promoting sustained R&D investment, unleashing Federal AI resources, removing barriers to AI innovation, empowering American workers with AI-focused education and training opportunities, and promoting an international environment that is supportive of American AI innovation and its responsible use. Furthermore, it leverages AI to improve the government's own internal processes. Section 8 of the Executive Order directs the NSA to "[develop] an action plan to protect the United States advantage in AI and AI technology critical to United States economic and national security interests against strategic competitors and adversarial nations."<sup>55</sup> Today, the only relevant strategic competitor of relevance for the United States in AI leadership is China.

In his book *AI Superpowers: China, Silicon Valley, and the New World Order*, Kai-Fu Lee, one of the most prominent figures in the Chinese internet sector, describes how China is rapidly increasing its position in the AI market, and may very well surpass the United States. Lee distinguishes the balance of current and future AI capabilities between the two nations across four waves: internet AI, business AI, perception AI, and autonomous AI. Where the United States currently holds a significant advantage in all four segments, he predicts that already in five years, China either closes the gap or surpasses the United States.<sup>56</sup>

The reason for this potential shift in leadership can be explained by looking at the four building blocks for becoming an AI superpower in an AI driven economy: abundant data, persistent entrepreneurs, progressive government support, and expert AI scientists. For the first three, China is already in a more favorable position.<sup>57</sup> Very rapidly, through what Lee calls

China's alternate internet universe—dominated by mobile communication—China now has the world's richest data ecosystem. This data is used by tenacious entrepreneurs, supported by China's mass entrepreneurship and innovation policy, and the ever-growing technology hub in Beijing's Zhongguancun neighborhood—the beating heart of China's AI movement and equivalent of Silicon Valley. According to Lee, the third requirement, government support, is also in China's favor due to its techno-utilitarian political culture. Where the United States aggressively punishes missteps in funding technology, China's political will rewards proactive investment and adoption;<sup>58</sup> a culture that—together with the pressure of China's ambitious 2030 roadmap—might also lead to the acceptance of more risk, creating less safe AI applications. Lastly, Lee concludes that elite AI expertise remains mostly in the United States. However, although Silicon Valley's edge is slightly overestimated and the gap is closing, expertise is less relevant in the age of AI implementation and deployment. While discoveries and breakthroughs such as deep learning are eventually required for progress, they are few and far between, and followed by an age of implementation, where elite expertise is needed less and entrepreneurship more so.

As less relevant expertise remains the only advantage of the United States, China is quickly steaming towards a dominant position in AI technology. There is an important factor, however, that could alter this course: large corporations. Although numerous companies pour resources into AI research and application, the seven key ones are Google, Facebook, Microsoft, Amazon, and Chinese Baidu, Alibaba, and Tencent. Academics cannot compete with these large companies and therefore focus mainly on research in finding the next breakthrough—ideally to publish it. If, however, such a breakthrough occurs within a sealed cooperate environment, this could lead to a return to an age of discovery, where AI expertise is again the dominant factor.

According to Kai-Fu Lee, the most likely candidate for such a breakthrough is—by far—the American company Google.<sup>59</sup>

### **Future Scenarios**

With the real possibility that China surpasses the United States in AI leadership, and having analyzed some of the AI trends in the military, economy, and information domain, it is fair to conclude that, while the impact of AI in the future is uncertain, not all possible scenarios are positive. To visualize potential future worlds, this paper identifies and focuses on three trends that influence and inform the focal decision on whether decision-makers need to increase focus and resources on AI development and policy. These three trends will act as scenario logics, or axes, along which the future scenarios will differ. The first trend is that China is catching up with the United States and is likely to be the dominant AI player in the next decade. The second trend is that AI technology will significantly increase military capabilities and power. The third trend is that AI will increasingly impact the economy and potentially cause massive global unemployment. All three trends are impactful but also uncertain in outcome. Approaching these trends in a binary way, i.e., they either become true or false, leads to eight possible scenarios. These eight scenarios can be grouped in three potential worlds: a mostly positive future, a future with risks for international security, and a dangerous future. Annex A provides an overview.

#### *Mostly Positive Future*

A future scenario where the United States dominates the field of AI, and where AI predominantly has positive effects on the economy, not leading to massive unemployment, can be categorized as mostly positive with the least amount of risk. Whether AI will significantly increase military capabilities (scenario A) or not (scenario B) is less relevant here; it will benefit

the United States if it does but does not significantly change or hinder United States leadership if it does not.

In these two scenarios, the coming decade is dominated by an age of discovery, where American engineering talent and AI experts rapidly advance AI technology through breakthroughs as deep learning did in the past. While global entrepreneurs, including Chinese, implement and apply new technologies, leadership is held by those who make the new discoveries while maintaining the advantage in AI expertise—i.e., the United States. Furthermore, the age of discovery also negates the advantages of data abundance in China. Moreover, most new AI discoveries are made in-house by large American Silicon Valley based companies, predominantly Google. These companies will increase their power at the expense of their Chinese counterparts. Other AI discoveries originate in academics, with better opportunities for entrepreneurs—including Chinese—to implement and take advantage. Because the United States is leading the AI market, it is in a good position to introduce and enforce laws, ethics, and processes that lead to a safer application of AI. Cyber Warfare is used more often by adversaries, including not-state actors. However, as an offensive tool, cyber is losing its advantage against the United States.

In summary, the two scenarios where the United States retains AI leadership, and unemployment due to AI advancement is mitigated, paint a mostly positive AI future. While China and Russia will continue to challenge United States AI dominance, this future world has the least amount of risk for international security.

*A Future with Risks*

A future world where the United States dominates the field of AI, but one wherein AI evolution is so fast and impactful that it leads to massive unemployment, is full of risk (scenario C and D). In scenario C, AI alters military capabilities that benefit American military power but, as this enables a mostly autonomous force countering Russian ambitions, it also further amplifies the risk of unemployment.

In both scenarios, the coming decade is one of AI discovery, where mostly American experts rapidly advance AI technology. AI systems are now capable of replacing most jobs—both manual and cognitive labor—conducting them faster, more efficient, and at a lower cost than their human predecessors. While the technology also creates new jobs, opportunities that are seized by entrepreneurs, AI advancement is so rapid that the job market cannot keep up—leading to dangerously high unemployment rates. Although the United States enjoys AI leadership, the swift shift in workforce is underestimated and preparations, laws, and regulations prove insufficient. In both scenarios, China and Russia will continue to challenge United States AI leadership.

A future where AI does not lead to massive unemployment and where the impact on military capabilities is marginal, but where China dominates the field of AI, is also one with risk (scenario E). Here, the next decade is one of AI implementation, with no new groundbreaking discoveries, negating the minor initial advantage in AI expertise by the United States. Entrepreneurs in China, with government support, take advantage of the ever-increasing amount of data generated, and Chinese companies grow at the expense of their American counterparts. China succeeds in executing its AI roadmap—leading the frontier in 2030. While the impact on



employment is small, China is in a far better position to take advantage of the positive benefits that AI adds to the economy—further aiding China’s strategic goal to remain the world’s leading economic power.

In summary, although all three scenarios have different dynamics, they describe a world where AI advancement creates risk. Whether the United States enjoys AI leadership but has difficulties mitigating the negative effects on the job market, or China leads the AI frontier taking advantage of the minor benefits in the economic and military domain, both paths lead to a world with increased risk to international security and stability.

### *A Dangerous Future*

A more dangerous future is one in which China dominates the field of AI, and the impact of AI advancement in one of the two domains—military or economy—is large (scenario F and G). The most dangerous and uncertain future unfolds, however, when China gains AI leadership, providing them significant advantage in military capability, while increased unemployment adds to uncertainty, instability, and insecurity (scenario H).

If the impact on Chinese military power is minor, but AI leads to massive unemployment (scenario F), the United States is now in a suboptimal position to influence the future global policy for improvement. As AI is now capable of replacing a large percentage of human workers, China and its entrepreneurs lead the frontier in the inevitable economic reform. Chinese preparations and government policy mitigate most of the negative effects on the economy but focus far less on the wellbeing of the population—leading to increased inequality.

The most dangerous future is where AI significantly increases Chinese military capabilities (scenario G and H). A decade long age of implementation prevents the United States

to exploit its initial AI expertise, and subsequently increases chances that the next AI breakthrough is from within China. Now being the dominant player in the economy, with or without massive impact on the job market, China enjoys a significant advantage in military power. As AI power continues to grow, exponentially or linearly, the power gap with the United States will increase and is difficult to overcome. Furthermore, as China possesses AI leadership, there is an increased risk of criminal organizations and hostile non-state actors taking advantage of autonomous AI possibilities. A risk that is now far more difficult to mitigate by the United States.

### **The Way Ahead**

The above described eight scenarios, categorized into three worlds, show that only in two cases an AI future is mostly positive (scenario A and B). Both require United States AI leadership and marginal impact on the unemployment rate. All other scenarios describe a future with varying levels of risk to international stability and security. To best prepare for a future with minimal risk, maintaining AI leadership is key, followed by planning to mitigate the prospect of unemployment as a result of AI technology. Computing power and AI breakthroughs are two leading indicators that help predict which scenario is most likely to unfold.

While a detailed roadmap on maintaining AI-leadership is beyond the scope of this paper, three broad lines-of-effort (LOE) are key—all of which are in the direct sphere of influence of strategic decision-makers. The first LOE is retaining AI expertise, the only real AI advantage the United States currently has over China. This effort includes increased government funding for education and research, in part to mitigate most AI advancements originating in the commercial sector. The second LOE is increased government support for entrepreneurs. Leading companies

in Silicon Valley require the right conditions to advance and stay ahead of their Chinese counterparts, with minimal limiting policies. As complete corporate freedom creates risk, increased focus on AI policy is required to establish the right balance. One of the most important and controversial topics that need continued government attention and policy is that of data collection—the third LOE. To mitigate the disadvantage of ever-growing data abundance in China, United States policy needs to allow for big data collection by companies but take privacy concerns into account. Several CEOs, including Facebook’s Mark Zuckerberg have called for more government regulations.<sup>60</sup>

Mitigating a rapid increase in unemployment as a result of AI is difficult and requires years of preparation. Whether AI will be capable of replacing a large percentage of the human workforce after the next decade is debated amongst experts. However, the United States cannot afford to be unprepared and must prioritize AI policy and strategy, and support think tanks such as the FLI.

To help predict which scenario is likely to unfold, two leading indicators are of note: computing power and AI breakthroughs. If in the next decade, computing power continues to grow exponentially, or even experiences boosts with implementations of new technology such as quantum computing, chances of AI leadership are in favor of the United States. If, however, computing power growth experiences hurdles, or even slows down to linear growth, China has time to catch up and increases its chances of AI leadership, albeit in a more distant future.

Other indicators are AI breakthroughs, which throughout history are few and far between. If none occur in the next decade, this is to the advantage of China as the current age of AI implementation prolongs. In this event, AI impact on the military and economy are likely marginal. The opposite is true when AI breakthroughs follow in rapid succession. Then, chances

of AI leadership are in favor of the United States, albeit with larger impact on the military and employment.









Attaching likelihoods to each scenario is beyond the scope of this paper, but also beside the point. None of the specific scenarios or worlds are extremely likely or unlikely. The point is that the uncertainty and high stakes ask for full and immediate attention of decision-makers to best address and account for the scale and severity of the impact.

### **Conclusion**













The future of AI and its impact on the world is uncertain. While recent developments such as deep learning fielded interesting, useful, and even revolutionary applications, the anticipated true impact of AI in the next decades is likely far greater. While some argue AI is merely a revolution comparable to the industrial revolutions in the past, many argue that this one is different. While the effects of the industrial revolution were great, it was the creation of tools for mankind to utilize or partially replace their jobs. Now, the tool we create has the potential to not only replace our jobs but replace our thought process and decision-making itself—especially the closer we get to AGI. While the United States currently enjoys AI leadership, China is likely to catch up and take over this role in the next decade. With potential significant impact of AI on the military and economic domain, this puts the United States in an uncomfortable position—even without losing their leadership role. Staying ahead of China requires a solid AI strategy and focus from policymakers and mitigating the possible negative effects of AI on employment requires years of preparation. The United States cannot afford to bet on a mostly positive future.

## Annex A: Three Future 2030 AI Worlds













### A Mostly Positive Future

				<p><b>Scenario A</b> The United States retains AI leadership, AI significantly effects Military power, and AI does not cause massive unemployment.</p>
				<p><b>Scenario B</b> The United States retains AI leadership, AI only marginally effects Military power, and AI does not cause massive unemployment.</p>

### A Future With Risks

				<p><b>Scenario C</b> The United States retains AI leadership, AI significantly effects Military power, and AI leads to massive unemployment.</p>
				<p><b>Scenario D</b> The United States retains AI leadership, AI only marginally effects Military power, and AI leads to massive unemployment.</p>
				<p><b>Scenario E</b> China exercises AI leadership, AI only marginally effects Military power, and AI does not cause massive unemployment.</p>

### A Dangerous Future

				<p><b>Scenario F</b> China exercises AI leadership, AI only marginally effects Military power, and AI leads to massive unemployment.</p>
				<p><b>Scenario G</b> China exercises AI leadership, AI significantly effects Military power, and AI does not cause massive unemployment.</p>
				<p><b>Scenario H</b> China exercises AI leadership, AI significantly effects Military power, and AI leads to massive unemployment.</p>

## Endnotes

- 
- <sup>1</sup> <https://www.hansonrobotics.com>.
- <sup>2</sup> Peter Schwartz, *The Art of the Long View* (New York: Doubleday, 1996), 241-48.
- <sup>3</sup> Alan Turing, "Computing Machinery and Intelligence," *Mind*, Volume 49, issue 236, October 1950, 433–60, <https://academic.oup.com/mind/article/lix/236/433/986238>.
- <sup>4</sup> Graham Oppy and David Dowe, "The Turing Test," *The Stanford Encyclopedia of Philosophy*, February 8, 2016, <https://plato.stanford.edu/entries/turing-test/>.
- <sup>5</sup> Jef Akst, "Machine, Learning, 1951," *The Scientist*, May 1, 2019, <https://www.the-scientist.com/foundations/machine--learning--1951-65792>.
- <sup>6</sup> Arthur Samuel, "Some Studies in Machine Learning Using the Game of Checkers," *IBM Journal of Research and Development* 3, issue 3 (July 1959), 210-29, <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5392560>.
- <sup>7</sup> <http://jmc.stanford.edu>.
- <sup>8</sup> John McCarthy et al., *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, August 31, 1955, <http://www-formal.stanford.edu/jmc/history/dartmouth/dartmouth.html>.
- <sup>9</sup> Introduced in 1957 by American psychologist Frank Rosenblatt, the perceptron network is an artificial model of communicating neurons; pioneering work for what later became more complex modern neural networks.
- <sup>10</sup> The paper *Artificial Intelligence: A General Survey*, also known as the Lighthill Report, was compiled by James Lighthill in 1973 and provided a pessimistic overview of the current state and prognosis of AI progress. The report was one the key reasons for a decrease in government AI funding at that time.
- <sup>11</sup> Initiated by Japan's Ministry of International Trade and Industry (MITI) in 1982, the FGCS was to use parallel computing with the aim to create a new generation of supercomputers, providing a basis for future AI developments.
- <sup>12</sup> Alex Roland and Philip Shiman, *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983-1993* (Cambridge, MA: MIT Press, 2002), 1.
- <sup>13</sup> Chris Smith et al., "The History of Artificial Intelligence," University of Washington, December 2006, 18, <https://courses.cs.washington.edu/courses/csep590/06au/projects/history-ai.pdf>.
- <sup>14</sup> <https://stanfordhealthcare.org/stanford-health-care-now/2014/cyberknife-technology-20th-anniversary.html>.
- <sup>15</sup> Roger Parloff, "The Deep-Learning Revolution," *Time: Artificial Intelligence: The Future of Humankind*, special edition, 2017.
- <sup>16</sup> [https://ec.europa.eu/research/fp7/understanding/fp7inbrief/what-is\\_en.html](https://ec.europa.eu/research/fp7/understanding/fp7inbrief/what-is_en.html).
- <sup>17</sup> The White House, *The Networking & Information Technology Research & Development Program: Supplement to the President's FY2020 Budget*, September 2019, 4, <https://www.whitehouse.gov/wp-content/uploads/2019/09/fy2020-nitrd-ai-rd-budget-september-2019.pdf>.
- <sup>18</sup> *Lexico: Powered by Oxford*, s.v. "artificial intelligence," accessed December 21, 2019, [https://www.lexico.com/definition/artificial\\_intelligence](https://www.lexico.com/definition/artificial_intelligence).
- <sup>19</sup> *Merriam-Webster*, s.v. "artificial intelligence," accessed November 8, 2019, <https://www.merriam-webster.com/dictionary/artificial-intelligence>.
- <sup>20</sup> Max Tegmark, *Life 3.0: Being Human in the Age of Artificial Intelligence* (New York: Alfred A. Knopf, 2017), 39.
- <sup>21</sup> <https://goertzel.org/who-coined-the-term-agi/>.
- <sup>22</sup> Tegmark, *Life 3.0*, 81.
- <sup>23</sup> In his 1950 paper *Programming a Computer for Playing Chess*, Claude Shannon calculated that the lower bound for the total number of games of chess is  $10^{120}$ . This number is larger than the atoms in the known observable universe. With the boardgame Go, the lower bound is far greater.
- <sup>24</sup> <http://www.aaai.org/>.
- <sup>25</sup> Tegmark, *Life 3.0*, 22.
- <sup>26</sup> <https://futureoflife.org/>.
- <sup>27</sup> Michael Horowitz et al., "Artificial Intelligence and International Security," *Centre for New American Security*, July 10, 2018, 4, <https://www.cnas.org/publications/reports/artificial-intelligence-and-international-security>.
- <sup>28</sup> Chris Bing, "The tech behind the DARPA Grand Challenge winner will now be used by the Pentagon," *Cyberscoop*, August 11, 2017, <https://www.cyberscoop.com/mayhem-darpa-cyber-grand-challenge-dod-voltron/>.
- <sup>29</sup> Jessica Cussins, "AI Researchers Create Video to Call for Autonomous Weapons Ban at UN," *The Future of Life Institute*, November 14, 2017, <https://futureoflife.org/2017/11/14/ai-researchers-create-video-call-autonomous-weapons-ban-un/>.

- <sup>30</sup> James Vincent, “Putin says the nation that leads in AI will be the ruler of the world,” *The Verge*, September 2017, <https://www.theverge.com/2017/9/4/16251226/russia-ai-putin-rule-the-world>.
- <sup>31</sup> Tom Simonite, “For Superpowers, Artificial Intelligence Fuels New Global Arms Race,” *Wired*, August 9, 2017, <https://www.wired.com/story/for-superpowers-artificial-intelligence-fuels-new-global-arms-race/>.
- <sup>32</sup> Samuel Bendett, “Get Ready, NATO: Russia’s New Killer Robots are Nearly Ready for War,” *the National Interest*, March 7, 2017, <https://nationalinterest.org/blog/the-buzz/russias-new-killer-robots-are-nearly-ready-war-19698>.
- <sup>33</sup> Daniel S. Hoadley, “Artificial Intelligence and National Security,” *Congressional Research Service*, January 30, 2019, 22, <https://crsreports.congress.gov/product/pdf/R/R45178>.
- <sup>34</sup> Distil Research Lab, “2019 Bad Bot Report,” *Distil Networks*, 2019, <https://resources.distilnetworks.com/i/1110664-2019-bot-report-executive-summary/>.
- <sup>35</sup> Onur Varol et al., “Online Human-Bot Interactions: Detection, Estimation, and Characterization,” *arXiv.org*, March 9, 2017, <https://arxiv.org/abs/1703.03107>.
- <sup>36</sup> Erica Anderson, “Building trust online by partnering with the International Fact Checking Network,” *Google’s The Keyword blog*, October 26, 2017, <https://www.blog.google/topics/journalism-news/building-trust-online-partnering-international-fact-checking-network/>.
- <sup>37</sup> <https://deepfakedetectionchallenge.ai/>.
- <sup>38</sup> Andy Patrizio, “IDC: Expect 175 zettabytes of data worldwide by 2025,” *Data Center Explorer*, accessed November 10, 2019, <https://www.networkworld.com/article/3325397/idc-expect-175-zettabytes-of-data-worldwide-by-2025.html>.
- <sup>39</sup> United Nations Office of Drugs and Crime, “Estimating Illicit Financial Flows Resulting from Drug Trafficking and Other Transnational Organized Crimes,” Research report (October 2011), 11, [https://www.unodc.org/documents/data-and-analysis/studies/illicit\\_financial\\_flows\\_2011\\_web.pdf](https://www.unodc.org/documents/data-and-analysis/studies/illicit_financial_flows_2011_web.pdf).
- <sup>40</sup> *PricewaterhouseCoopers*, accessed December 28, 2019, <https://www.pwc.com/us/en/services/consulting/library/artificial-intelligence-predictions-2019.html>.
- <sup>41</sup> Irving Wladawsky-Berger, “The Impact of Artificial Intelligence on the World Economy,” *The Wall Street Journal*, November 16, 2018, <https://blogs.wsj.com/cio/2018/11/16/the-impact-of-artificial-intelligence-on-the-world-economy/>.
- <sup>42</sup> Erik Brynjolfsson and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies* (New York: W. W. Norton & Company, 2014).
- <sup>43</sup> Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order* (Boston: Houghton Mifflin Harcourt, 2018), 347.
- <sup>44</sup> Carl Benedikt Frey, and Michael A. Osborne, “The Future of Employment: How Susceptible are Jobs to Computerization?” *Oxford University*, September 17, 2013, [https://www.oxfordmartin.ox.ac.uk/downloads/academic/The\\_Future\\_of\\_Employment.pdf](https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf).
- <sup>45</sup> <https://www.oecd.org/employment/employment-outlook-2019-highlight-en.pdf>.
- <sup>46</sup> Susan Lund et al., “The future of work in America: People and places, today and tomorrow,” *McKinsey & Company*, July 2019, <https://www.mckinsey.com/featured-insights/future-of-work/the-future-of-work-in-america-people-and-places-today-and-tomorrow>.
- <sup>47</sup> Don Reisinger, “A.I. Expert Says Automation Could Replace 40% of Jobs in 15 Years,” *Fortune*, January 10, 2019, <https://fortune.com/2019/01/10/automation-replace-jobs/>.
- <sup>48</sup> In 1901, Wilbur Wright famously said that “Man will not fly for 50 years.” Just two years later the Wright Brothers successfully conducted their first manned flight.
- <sup>49</sup> Tegmark, *Life 3.0*, 157.
- <sup>50</sup> Tegmark, *Life 3.0*, 31.
- <sup>51</sup> Ray Kurzweil, *The Singularity is Near* (New York, NY: Viking, 2005).
- <sup>52</sup> Tegmark, *Life 3.0*, 48.
- <sup>53</sup> Graham Webster et al., “Full Translation: China’s New Generation Artificial Intelligence Development Plan,” *New America*, 2017, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.
- <sup>54</sup> The White House, “Artificial Intelligence for the American People,” accessed January 2, 2020, <https://www.whitehouse.gov/ai/>.
- <sup>55</sup> The White House, “Maintaining American Leadership in Artificial Intelligence,” *National Archives*, February 11, 2019, 3971, <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>.

<sup>56</sup> Lee, *AI Superpowers*, 306.

<sup>57</sup> Lee, *AI Superpowers*, 189.

<sup>58</sup> Lee, *AI Superpowers*, 193.

<sup>59</sup> Lee, *AI Superpowers*, 211.

<sup>60</sup> Mark Zuckerberg, “The Internet needs new rules. Let’s start in these four areas,” *The Washington Post*, March 30, 2019, [https://www.washingtonpost.com/opinions/mark-zuckerberg-the-internet-needs-new-rules-lets-start-in-these-four-areas/2019/03/29/9e6f0504-521a-11e9-a3f7-78b7525a8d5f\\_story.html](https://www.washingtonpost.com/opinions/mark-zuckerberg-the-internet-needs-new-rules-lets-start-in-these-four-areas/2019/03/29/9e6f0504-521a-11e9-a3f7-78b7525a8d5f_story.html).



## Bibliography

- Allen, Greg. "China's Artificial Intelligence Strategy Poses a Credible Threat to U.S. Tech Leadership." *Council on Foreign Relations*. December 4, 2017, <https://www.cfr.org/blog/chinas-artificial-intelligence-strategy-poses-credible-threat-us-tech-leadership>.
- Allen, Greg, and Taniel Chan. "Artificial Intelligence and National Security." *Belfer Center for Science and International Affairs*. July 2017, <https://statewatch.org/news/2017/jul/usa-belfer-center-national-security-and-ai-report.pdf>.
- Auslin, Michael. "Can the Pentagon Win the AI Arms Race?" *Foreign Affairs*, October 19, 2018, <https://www.foreignaffairs.com/articles/united-states/2018-10-19/can-pentagon-win-ai-arms-race>.
- Bing, Chris. "The tech behind the DARPA Grand Challenge winner will now be used by the Pentagon." *Cyberscoop*, August 11, 2017, <https://www.cyberscoop.com/mayhem-darpa-cyber-grand-challenge-dod-voltron/>.
- Bostrom, Nick. *Superintelligence: Paths, Dangers, Strategies*. New York, NY: Oxford University Press, 2014.
- Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: Norton & Company, 2014.
- Cummings, M.L. "Artificial Intelligence and the Future of Warfare." *Chatham House*, January 2017, <https://www.chathamhouse.org/sites/default/files/publications/research/2017-01-26-artificial-intelligence-future-warfare-cummings-final.pdf>.

- Cummings, M. L., Heather M. Roff, Kenneth Cuckier, Jacob Parakilas, and Hannah Bryce. "Artificial Intelligence and International Affairs." *Chatham House*, June 14, 2018, <https://www.chathamhouse.org/sites/default/files/publications/research/2018-06-14-artificial-intelligence-international-affairs-cummings-roff-cukier-parakilas-bryce.pdf>.
- Cussins, Jessica. "AI Researchers Create Video to Call for Autonomous Weapons Ban at UN." *The Future of Life Institute*, November 14, 2017.
- Frey, Carl Benedikt, and Michael A. Osborne. "The Future of Employment: How Susceptible Are Jobs to Computerization?" *Oxford University*, September 17, 2013, [https://www.oxfordmartin.ox.ac.uk/downloads/academic/The\\_Future\\_of\\_Employment.pdf](https://www.oxfordmartin.ox.ac.uk/downloads/academic/The_Future_of_Employment.pdf).
- Hass, Ryan, and Zach Balin. "US-China relations in the age of artificial intelligence." *Brookings*, January 10, 2019, <https://www.brookings.edu/research/us-china-relations-in-the-age-of-artificial-intelligence>.
- Horowitz, Michael C., Gregory C. Allen, Edoardo Saravalle, Anthony Cho, Kara Frederick, and Paul Scharre. "Artificial Intelligence and International Security." *Centre for New American Security*, July 10, 2018, <https://www.cnas.org/publications/reports/artificial-intelligence-and-international-security>.
- Jensen, Benjamin M., Christopher Whyte, and Scott Cuomo. "Algorithms at War: The Promise, Peril, and Limits of Artificial Intelligence." *International Studies Review*, June 24, 2019, <https://doi.org/10.1093/isr/viz025>.
- Kurzweil, Ray. *The Singularity is Near*. New York, NY: Viking, 2005.
- Lee, Kai-Fu. *AI Superpowers: China, Silicon Valley, and the New World Order*. Boston: Houghton Mifflin Harcourt, 2018.

Levesque, Elizabeth Mann. “The role of AI in education and the changing US workforce.”

*Brookings*, October 18, 2018, <https://www.brookings.edu/research/the-role-of-ai-in-education-and-the-changing-u-s-workforce/>.

Lund, Susan, James Manyika, Liz Hilton Segel, André Dua, Scott Rutherford, and Brent Macon.

“The future of work in America: People and places, today and tomorrow.” *McKinsey & Company*, July 2019, <https://www.mckinsey.com/featured-insights/future-of-work/the-future-of-work-in-america-people-and-places-today-and-tomorrow>.

Mayer-Schonberger, Victor, and Kenneth Cukier. *Big Data: A Revolution That Will Transform*

*How We Live, Work, and Think*. New York: Houghton Mifflin Harcourt, 2018.

Mara, Karlin. “The implications of artificial intelligence for national security strategy.”

*Brookings*, November 1, 2018, <https://www.brookings.edu/research/the-implications-of-artificial-intelligence-for-national-security-strategy/>.

Meserole, Chris. “What is machine learning?” *Brookings*, October 4, 2018,

<https://www.brookings.edu/research/what-is-machine-learning/>.

O’Hanlon, Michael E. “The role of AI in future warfare.” *Brookings*, November 29, 2018,

<https://www.brookings.edu/research/ai-and-future-warfare/>.

Parloff, Roger. “The Deep-Learning Revolution.” *Time: Artificial Intelligence: The Future of*

*Humankind*, special edition, 2017.

Reisinger, Don. “A.I. Expert Says Automation Could Replace 40% of Jobs in 15 Years.”

*Fortune*, January 10, 2019, <https://fortune.com/2019/01/10/automation-replace-jobs/>.

Roland, Alex, and Philip Shiman. *Strategic Computing: DARPA and the Quest for Machine*

*Intelligence, 1983-1993*. Cambridge, MA: MIT Press, 2002.

- Samuel, Arthur. "Some Studies in Machine Learning Using the Game of Checkers." *IBM Journal of Research and Development* 3, issue 3 (July 1959), 210-29, <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=5392560>.
- Scharre, Paul. "A Million Mistakes a Second." *Foreign Policy*, September 12, 2018, <https://foreignpolicy.com/2018/09/12/a-million-mistakes-a-second-future-of-war/>.
- Schwartz, Peter. *The Art of the Long View*. New York: Doubleday, 1996.
- Shadbolt, Nigel, and Roger Hampson. *The Digital Ape: How to Live (in Peace) With Smart Machines*. New York: Oxford University Press, 2018.
- Smith, Chris, Brian McGuire, Ting Huang, and Gary Huang. "The History of Artificial Intelligence." University of Washington, December 2006, <https://courses.cs.washington.edu/courses/csep590/06au/projects/history-ai.pdf>.
- Tegmark, Max. *Life 3.0: Being Human in the Age of Artificial Intelligence*. New York: Alfred A. Knopf, 2017.
- Turing, Alan. "Computing Machinery and Intelligence." *Mind*, Volume 49, issue 236, October 1950, 433–60, <https://academic.oup.com/mind/article/lix/236/433/986238>.
- The White House. "Maintaining American Leadership in Artificial Intelligence." *National Archives*, February 11, 2019, 3971, <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>.
- The White House. "The Networking & Information Technology Research & Development Program: Supplement to the President's FY2020 Budget." *The White House*, September 2019, 4, <https://www.whitehouse.gov/wp-content/uploads/2019/09/fy2020-nitrd-ai-rd-budget-september-2019.pdf>.

Vincent, James. "Putin says the nation that leads in AI will be the ruler of the world." *The Verge*, September 2017, <https://www.theverge.com/2017/9/4/16251226/russia-ai-putin-rule-the-world>.

Webster, Graham, Rogier Creemers, Paul Triolo, and Elsa Kania. "Full Translation: China's New Generation Artificial Intelligence Development Plan." *New America*, August 1, 2017, <https://www.newamerica.org/cybersecurity-initiative/digichina/blog/full-translation-chinas-new-generation-artificial-intelligence-development-plan-2017/>.

Wladawsky-Berger, Irving. "The Impact of Artificial Intelligence on the World Economy." *The Wall Street Journal*, November 16, 2018, <https://blogs.wsj.com/cio/2018/11/16/the-impact-of-artificial-intelligence-on-the-world-economy/>.

Zilber, Neri. "The Rise of the Cyber-Mercenaries: What happens when private firms have cyberweapons as powerful as those owned by governments?" *Foreign Policy*, August 31, 2018, <https://foreignpolicy.com/2018/08/31/the-rise-of-the-cyber-mercenaries-israel-nso/>.