National Power After AI

AUTHORS
Matthew Daniels
Ben Chang

JULY 2021

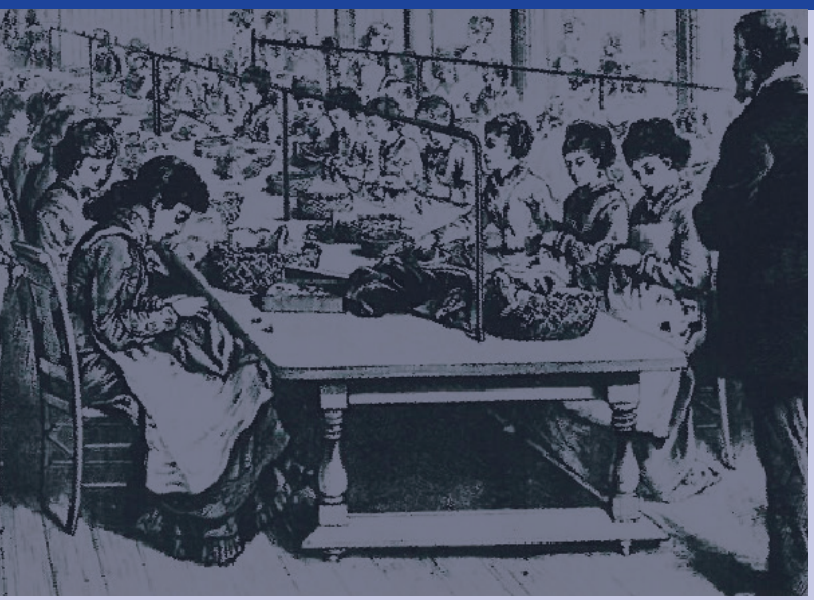

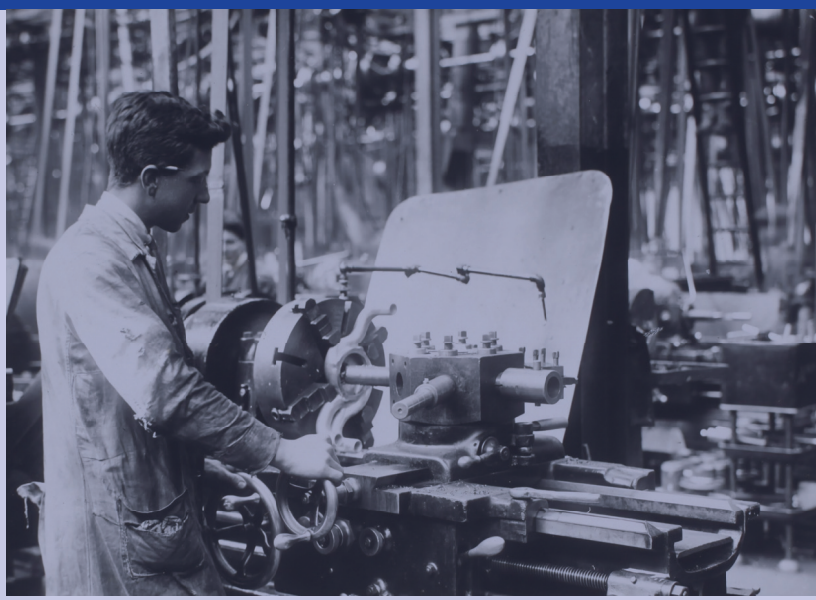

CSET CENTER for SECURITY and EMERGING TECHNOLOGY

CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

Established in January 2019, the Center for Security and Emerging Technology (CSET) at Georgetown's Walsh School of Foreign Service is a research organization focused on studying the security impacts of emerging technologies, supporting academic work in security and technology studies, and delivering nonpartisan analysis to the policy community. CSET aims to prepare a generation of policymakers, analysts, and diplomats to address the challenges and opportunities of emerging technologies. CSET focuses on the effects of progress in artificial intelligence, advanced computing, and biotechnology.

CSET.GEORGETOWN.EDU | CSET@GEORGETOWN.EDU

JULY 2021

National Power After AI

AUTHORS
Matthew Daniels
Ben Chang

ACKNOWLEDGMENTS

We would like to thank Igor Mikolic-Torreira, James Baker, Jack Clark, Remco Zwetsloot, Teddy Collins, Helen Toner, Jack Corrigan, Jeff Alstott, Maura McCarthy, Alex Friedland, Lynne Weil, David Lasker, Jared Dunnmon, Matt Mahoney, and Greg Allen for their comments on earlier drafts. We would also like to thank Andrew Imbrie for many thoughtful discussions and reading suggestions and Melissa Flagg for early input and ideas. This work benefitted directly from the early Office of Net Assessment summer study on AI in 2016. Of course, the authors are solely responsible for the views expressed in this publication and for any errors.

AUTHORS

Matthew Daniels was a Senior Fellow at CSET, where Ben Chang is an Andrew W. Marshall Fellow.

Since authoring this paper, Matthew Daniels has taken a position in the U.S. government. He completed his contributions to this paper prior to departing CSET. The views expressed herein are the authors' and do not necessarily reflect those of the U.S. government. Inquiries about this report should be directed to Ben Chang.

DOCUMENT IDENTIFIER

doi: 10.51593/20210016

Cover photos: (Top left) https://www.flickr.com/photos/ywcausa/8527933591 (Top right) https://unsplash.com/photos/-7Xb4TmVTN8 (Bottom) https://unsplash.com/photos/jHZ70nRk7Ns.

Contents

Introduction

How will artificial intelligence affect long-term U.S.-China competition?* Many analyses approach this question by focusing on how AI technologies may enhance each side's current capabilities–making aircraft, businesses, and nations, say, 10 percent faster and stronger. This perspective essentially suggests a broad race: the side that leverages modern AI technologies the most and soonest wins.

This view can mislead American strategists in two ways. First, it ignores the new vulnerabilities, costs, and accident risks associated with adopting modern AI technologies in certain settings.[1] Second, and most importantly, over the longer term, AI technologies will likely alter great power competitions in foundational ways, changing both how nations create power and their motives for wielding power against one another. In short, strategists are trying to create strategies for a game while the field, the players, the ball, and the rules could be changing.

Why? Major innovations themselves not only generate value for leading states, but also change the relative value of previously acquired assets, organizational practices, and defense strategies. Germany's development of *Blitzkrieg* during the interwar period, which represented a new way of using motorized vehicles and radios, is often cited as a military-specific example of important new organizational approaches. The German military effectively adapted its operational practices to suit new technologies. But the impact of major innovations can also be broader. For instance, the Industrial Revolution made new factors central to national power, includ-

*In this work, we use "artificial intelligence" to mean, as per the Defense Innovation Board, "a variety of information processing techniques and technologies used to perform a goal-oriented task and the means to reason in the pursuit of that task." See Appendix.

ing access to certain resources and the industrial capacity to leverage them. These broader effects take more time to appear, but their impact can be enormous: industrialization was not "revolutionary" because of the rapidity of change, as it unfolded in waves over decades, but because of its ultimate magnitude of change. With AI technologies, progressive substitution of machines for human cognitive labor may eventually have economic and social implications on a scale comparable to the Industrial Revolution. And like the Industrial Revolution, this AI revolution will change some fundamental elements of national power.

Of course, these foundational shifts can render some of the current processes and resources of a state obsolete, but they can also make what states are already doing, or already possess, *more* valuable. For example, the invention of railroads was a boon for those rich in steel precursors.[2] With AI, data-hungry algorithms may advantage authoritarian states, which already surveil and catalogue their own populations with little regard for human rights.[3]

We suggest an "evolutionary" view of technological change: major, widely diffused innovations are akin to environmental shifts, affecting the competitive capacity of states based on their existing trends in population, resources, institutions, character, and policies. Some previously "maladaptive" factors may become advantageous, and vice versa; states will adapt their institutions, organizations, and policies to the new environment in different ways and to varying degrees, and consequently gain or lose relative power as a result. Nations that primarily focus on AI technologies as offering marginal improvements in existing capabilities ("helping to build better mousetraps") will eventually miss larger opportunities to adapt. This paper is a first step into thinking more expansively about AI and national power. In what follows, we first explain this evolutionary view in greater detail before applying it to AI.

Ultimately, we seek pragmatic insights for long-term U.S. competition with authoritarian governments like that of China. For the foreseeable future, China's population and total economic size will very likely exceed those of the United States, even as its per capita GDP lags. This new challenge differs fundamentally from the United States' Cold War competition with the Soviet Union, and success will require thoughtful and timely diagnosis of modern environmental shifts in how states can produce power. These insights can guide our own investments as well as our approach to alliances. The United States has many sources of advantage and strength, and as Joseph Nye rightly observed, "Our greatest mistake in such a world would be to fall into one-dimensional analysis and to believe that investing in military power alone will ensure our strength."[4] This paper is a first step, intended to provoke new questions and provide a framework for assessing the relationship between AI and national power.*

*This work benefitted directly from the early Office of Net Assessment summer study on AI in 2016.

1 An Evolutionary Theory of Technological Competition

WHY STATIC, UNIVERSAL MEASURES OF POWER (THAT ARE USEFUL) DO NOT EXIST

Power is simply the relative capability of a state to achieve what it wants in international affairs. Power depends on one state's favorable factors relative to another. In one of the founding works of international relations, Hans Morgenthau proposed distinguishing nine elements of national power: geography, resources, industrial capacity, military preparedness, population, national character, national morale, quality of diplomacy, and quality of government.[5]

Since Morgenthau's writing, generations of analysts have sought a definitive way to measure national power that would, finally, allow accurate judgment of relative strength without fighting a war.[6] The search has included dozens of books and hundreds of journal articles offering competing methodologies and metrics.[7] For example: Should measures of useful access to resources include both steel and oil, or only steel? How should "soft power" be measured? What about the "latent" power that a state could theoretically draw from its population?[8]

Were such a universal, "objective" measure obtainable, the benefits would be enormous. We could easily answer questions such as, "who's ahead?" and "if it's not us, what should we do about it?" This quest, however, has not borne fruit. Proposed measures have tended to perform poorly when generalized.[9] History is full of surprises where states have

achieved victory even when "objective" metrics would predict their defeat: the United States owes its national existence to victory in the Revolutionary War over the British Empire, the superpower of the time.

Why? First and foremost, power is *always* contextual.[10] This is especially clear in military matters. A large military's skill at high-intensity conflict may not translate to skill at counterinsurgency; and factors that provide one country advantage relative to another can change. The world offers no "power particle" to measure objectively alongside other characteristics of nature—what we intuitively mean by "power" is mostly a generalization from particular observations.[11] Elements of power can also combine in surprising ways. Andrew Marshall offered the reminder that countries with relatively smaller populations and GNPs can pose substantial challenges to larger competitors: in 1938, Japan had roughly half the population and one-tenth the GNP of the United States,[12] but it built a navy that challenged the United States in wartime.[13] In part because of these issues, history is rife with leaders who have had a large gap between their beliefs and the reality of military capabilities.[14] Each competition should be analyzed carefully on its own, distinguishing elements of power, identifying key areas of competition, and working to diagnose the most important problems and opportunities in each area of competition.

MAJOR INNOVATIONS CHANGE THE SECURITY ENVIRONMENT, CHANGING WHAT GIVES RISE TO POWER

Military leaders throughout history are sometimes faulted for preparing to fight the previous war instead of the next one. We should likewise avoid strategizing to "win the previous competition." Just as changing adversaries from the Soviet Union to insurgents in Iraq and Afghanistan represented a new security environment and revealed the non-fungibility of power, major innovations also change the security environment as they are widely adopted. Such innovations do this in part by changing what assets, practices, and strategies give rise to power.

Differential impacts of emerging technologies often bring shifts in relative capabilities of individual countries.[15] Thinking about long-term competition in periods of rapid technological change therefore requires assessing how innovations change factors related to military and national power.

Major innovations can change the estimations of power in three ways:

- First, innovations ***introduce new elements of power.*** Major innovations, in changing how states generate power, can create new factors that must be considered in characterizing power. For example, the advent of railroads, internal combustion engines, and nuclear weapons dramatically increased

Perhaps least obviously, major innovations sometimes broadly alter what policies states pursue, by making certain kinds of behavior more valuable or less costly.

the importance of a state's access to steel, oil, and uranium, respectively.[16] New factors, however, are not only limited to materials. They may also encompass characteristics of a society's culture, organizations, or economic activities.[17]

- Second, innovations ***change the importance of existing elements of power.*** Major innovations also change the "coefficients" of existing elements of power, causing them to matter more or less than before. For example, Mongol light cavalry, modern navies, and ballistic missiles all changed how geographic barriers affected one's balance of power with geographic neighbors, eroding the effectiveness of simple remoteness, oceans, and armies still in the field, respectively, as shields against coercive power.[18] Industrialization meant the inventiveness of a nation's scientists and engineers became more important.

- Finally, innovations ***alter states' intermediate goals***. Perhaps least obviously, major innovations sometimes broadly alter what policies states pursue, by making certain kinds of behavior more valuable or less costly. While states retain the same ultimate ends, such as securing survival and prosperity, the intermediate, instrumental goals they pursue to reach those ends may shift. This can drive dramatic changes in state goals and policies. For example, before the Industrial Revolution, potential productivity gains in areas like agriculture and manufacturing were small and stable; this made conquering territory a primary means by which one group could increase its wealth and security.[19] During and after the Industrial Revolution, modern states could also pursue substantial military and economic growth by applying new technologies to increase productivity.

The next section discusses how these three changes manifest in the context of AI.

2 Power After AI: New Elements, Changed Factors, and Altered Goals

We offer early thinking about potential changes caused by AI: new elements of power, shifting importance for existing elements of power, and shifting intermediate goals for states. These are not definitive or complete results, but a starting place for broader thinking.

NEW ELEMENTS

One of the most familiar examples of new elements of power is associated with the Industrial Revolution, when machines began to help humans with physical labor in new and organized ways. The Industrial Revolution led to dramatic changes in the character of war and military power. A simple approximation is that, before the Industrial Revolution, any group's military power correlated most closely with its quantity of fieldable humans under arms, a measure of both taxable population and military potential. After the Industrial Revolution, any estimate of military power had to include a society's industrial capacity and access to resources to enable that capacity, which are measures of a society's ability to produce useful military hardware, such as ships, tanks, planes, and submarines.

It is useful to see AI technologies today as part of another large-scale transition: machines are increasingly helping humans with certain kinds of *cognitive* labor in new and organized ways.[20] This transition will span decades, with potential economic and social implications on a scale comparable to those of the Industrial Revolution. Today, as then, there are

U.S. defense leaders believe the rapidly growing military applications of AI technologies will be critical for the years ahead. State power will increasingly hinge on the new factors required to effectively adopt AI.

large questions about the future of economic production, human labor, and military capabilities. These future trends will define new elements of power.

U.S. defense leaders believe the rapidly growing military applications of AI technologies will be critical for the years ahead.[21] State power will increasingly hinge on the new factors required to effectively adopt AI. Four such factors often identified by existing literature include data, AI scientists and engineers ("AI talent"), computational power ("compute"), and AI-adapted organizations. Below, we explore the latter two in greater detail.

Ability to Access and Leverage Compute

The United States has historically used large-scale compute capabilities for analysis of nuclear weapons detonations and cryptanalysis.[22] More recently the U.S. government's uses have grown to include climate modeling and a variety of scientific applications. In the years ahead, the United States may also use large compute resources for creating and countering new AI capabilities.

For decades, cutting-edge AI systems have used steadily increasing quantities of compute resources, making improvements in compute capabilities a key driver of AI progress. This usage appears to have accelerated across the last decade: the compute used in the largest AI training runs has doubled every 3.4 months since 2012, growing more than 300,000 times from AlexNet in 2012 to AlphaGo Zero in 2018.[23] OpenAI researchers have shown that the 2010s appear to be the beginning of a new computing era for AI technologies, distinct from the preceding 40-50 years.[24]

For military applications where limited real-world data is available, techniques leveraging computer simulations instead of large quantities of data may further increase demand for compute.[25] Cloud compute may become vital for rapidly processing and fusing intelligence across platforms, while edge compute will be necessary for autonomous systems deployed in the field tasked with assessing and outthinking adversaries' equivalent systems.

As such, a nation's ability to leverage large quantities of computational power could become a new primary term feeding into its ability to influence international

affairs. For example, the key technical precursors required to manufacture cutting-edge AI chips are currently concentrated in the United States and allied countries–though semiconductor manufacturing capabilities more broadly, beyond just the most cutting-edge chips, may further grow the importance of Taiwan and South Korea as international trading partners.[26]

Importantly, compute resources must be configured in ways useful for modern AI capabilities. High-performance computing (HPC) systems currently maintained within the U.S. Government, such as in the Department of Energy, tend to be both specialized for non-AI functions and subject to system-specific security measures, posing challenges for broad, standardized utilization by other organizations. Consequently, commercial cloud compute resources may better serve the U.S. Government in deploying certain kinds of AI technologies, although potentially promising efforts to improve the use of U.S. HPC assets for AI are also underway.[27] Effective use will depend, too, on accessible software tools for using cloud compute systems–which may prove to be comparable to process and tooling approaches developed to make factories effective during industrialization in the United States.[28]

Compute resources can flow more easily than many traded goods. As computing infrastructure continues to grow, new ways of sharing access to large, regionally-concentrated quantities of compute, including through space internet constellations, may create new opportunities and incentives for international partnerships.

Ability to Manage Data and AI Safety & Security

Even when states possess the raw resources required to adopt some major innovation, they still must undertake the often-difficult process of institutional and organizational adaptation. Bureaucratic factors in organizations matter greatly: in militaries, competing civilian, interservice, and intra-service actors may promote or resist adoption of new technologies.[29] Resistance can include parochial forces that attempt to stymie adoption: for example, only direct pressure from Eisenhower moved the Air Force to adopt ICBMs instead of focusing solely on less survivable crewed bombers.[30] Organizational culture also has significant impacts: because mass armies threatened the pre-existing hierarchical power structure within many European militaries, many states failed to adopt Napoleon's innovation even after his dramatic string of victories.[31] During periods of rapid change, medium-sized powers may have opportunities to adopt innovations more speedily than larger powers.[32]

With AI, demands for organizational adaptations will be significant. Two factors are especially important: effective data pipelines and the effective management of security issues associated with modern AI technologies.

The ability to deploy cutting-edge AI applications will increasingly depend on the quality of each organization's data pipeline. Modern machine learning methods are notoriously data-hungry, but simply possessing large quantities of data-collecting sensing platforms will be insufficient—for supervised learning applications, data must be structured, labeled, and cleaned; fusing data from many platforms, sources, and formats will represent its own herculean challenge for many militaries. Finally, these data pipelines must also be dynamic: data management itself must be monitored, in part to detect attacks, because "data poisoning" attacks can manipulate AI behavior by changing what lessons it learns.[33] Consequently, it will be increasingly important for military leaders to successfully implement organizational reforms to create and maintain effective data pipelines.

Military leaders must also learn to effectively manage the novel security issues associated with AI technologies. Relying on modern AI systems for safety- or mission-critical tasks carries challenges because many deep learning models are exceptionally hard to interpret.[34] Michael Jordan at UC Berkeley has analogized the creation of early large-scale AI models to building bridges before civil engineering was a rigorous discipline: "While the building blocks are in place, the principles for putting these blocks together are not, and so the blocks are currently being put together in ad-hoc ways. ... Just as early buildings and bridges sometimes fell to the ground—in unforeseen ways and with tragic consequences—many of our early societal-scale inference-and-decision-making systems are already exposing serious conceptual flaws."[35] A more developed engineering discipline for AI is needed to manage the risk of accidents from relying on opaque machines in the field.[36] In near-term military settings, effectively integrating new AI technologies will require special investment in test, evaluation, validation and verification (TEVV) processes by competent organizational leaders.[37]

More widely, many modern AI systems are not designed to work in the presence of malevolent actors. Potential security issues for deep learning systems include adversarial examples and model inversion, in addition to data poisoning and more traditional computer network and software attacks.[38] Adversarial examples refer to "inputs" (such as visual or audio patterns) to an AI system that cause the system to malfunction; model inversion refers to an ability to reverse-engineer the data used to train an AI system, which may include private or classified information. Despite these challenges, modern machine learning capabilities will be increasingly woven into G20 societies, economies, and military systems.* The U.S. position with

*For example, AI technologies will intersect with 5G and networking trends in cities as autonomous systems (like vehicles) in urban areas begin to have large quantities of interactions with other intelligent agents—working on everything from traffic coordination to utilities management and financial investments. The ability for intelligent systems to interact on large scales, safely and securely, will be critical.

The U.S. position with AI technologies for the next two or three decades appears analogous to the future that faced IT technologies in the 1990s: AI technologies are so valuable that they will be used despite substantial design and security issues.

AI technologies for the next two or three decades appears analogous to the future that faced IT technologies in the 1990s: AI technologies are so valuable that they will be used despite substantial design and security issues.

What might the future look like given these vulnerabilities? We can only speculate: in direct military settings, there may be new sub-competitions that resemble the emergence of electronic warfare after the invention of radar.[39] In economic systems, in addition to the potential for the novel security risks discussed previously, there is risk of physical manifestations of the kinds of problems currently seen in high frequency trading systems, such as rapid, unanticipated interactions among automated agents managing services in cities.[40] These issues may open new vulnerabilities to both individual rogue actors and state adversaries. Organizations that are able to adapt early to manage these new security issues will be advantaged.

Since states vary in their access to compute, data, AI talent, and useful organizational adaptations, they will also vary in their ability to benefit from modern AI technologies. Any national rankings based on these factors will be debatable, but the nations that generally lead in these metrics is unsurprising, and include: the United States, China, Japan, South Korea, the UK, Canada, Taiwan, Israel, France, Germany, and Russia. Advanced economies should be increasingly expected to focus their own investments and policies on improving their positions in these areas.

CHANGED FACTORS

Industrialization meant that a nation's stock of productive scientists and engineers counted more than it had in the past. With the arrival of AI, various previously recognized elements of national power will become more important, while others may become gradually less so. For illustrative purposes below, we discuss population size and scientific talent as contrasting examples: population size becoming less important, scientific talent becoming more important.

Population Size

As AI technologies increasingly substitute for human labor, total population size may become less important for national military and economic capacity. [41] Just as machines took over rote physical labor during industrialization, AI technologies will automate rote cognitive labor, from diagnosing maintenance needs to exploiting imagery intelligence. This may reduce the total quantity of human labor needed to maintain a military's operational capacity. In major wars, partially or fully autonomous AI platforms may further reduce a country's need to field humans in combat. As militaries rely more on autonomous systems for military operations, defense planners may come to count autonomous systems and their available domestic supply of AI chips the way they once counted soldiers and the available domestic recruiting pool of military-age adults.[42] Downstream, this could help technologically advanced nations compensate for demographic challenges, such as aging populations and low birth rates, a situation the United States, China, Japan, Western Europe, and Russia all face to varying degrees.[43]

Population trends continue to matter for national power—but AI technologies, like many other technologies of the past century, may further erode this importance.

Industrious Scientists and Engineers

Harnessing new technologies, both by developing technologies and accessing innovations created elsewhere, is an important means of growing power. Applications of AI can help in both areas, serving as a force multiplier on, and therefore increasing the importance of, productive scientists and engineers.

Recently, for example, DeepMind's AlphaFold achieved breakthrough rates of accuracy comparable to experimental methods in the protein-structure prediction challenge known as CASP.[44] By obviating the need for experimental protein structure assessment, a skill-demanding and time-intensive procedure, AlphaFold represents a large augmentation of human scientists' biosciences research. In a different domain of research, modern AI applications are able to help with chip design.[45] Researchers have demonstrated a deep learning system capable of designing the physical layout of computer chips more effectively than human engineers.[46] Google has used this system to design its next generation of Tensor Processing Units (TPUs), the company's specialized AI chips.

Likewise, rapid progress in machine translation, automatic literature review, and related tools means a given scientific discipline's state-of-the-art will become increasingly accessible and useful to well-organized groups of human scientists and engineers. Just as the printing press alleviated the need to travel from country to country to accumulate knowledge from different libraries, AI applications can lower the costs for researchers to access state-of-the-art knowledge in any field.

There are three ways that modern AI applications will contribute on a large scale to scientific discovery and engineering invention: they will contribute directly to new discoveries and engineered systems, especially in areas that involve searches over large spaces in data or design;[47] automate the physical work of science and engineering, such as "self-driving laboratories" that robotically automate experimental laboratory work;[48] and make global scientific knowledge more accessible to humans, such as by extracting knowledge from millions of articles as well as from articles in many different languages.[49]

Finally, there is an old debate about whether science advances most because of new ideas or new tools;[50] AI technologies appear able to contribute both. In the longer-term, AI may enable new and more creative forms of knowledge-generation that function as "pathfinders" for human brains, unlocking otherwise difficult-to-reach innovations. When AlphaGo beat Lee Sedol, its 37th move in the second game surprised human professionals. In the words of Go master Fan Hui, "It's not a human move. I've never seen a human play this move. So beautiful."[51] When AI behavior surprises us, we learn something new. Looking ahead, modern and future AI systems may be able to solve scientific puzzles that have thus far stumped humanity's best minds.[52]

Just as railways advantaged nations with access to steel, it appears that AI tools capable of augmenting science and engineering work will favor nations with the best existing "resources" of industrious scientists and engineers. This trend appears likely to deepen the advantages of nations that host, or can attract, a disproportionate fraction of the world's best in those fields.[53]

ALTERED GOALS

Finally, major innovations can alter state strategies, as different instrumental goals become more appealing for achieving a state's ultimate ends.

The Industrial Revolution again provides a clear example. Before industrialization, conquering territory was a primary way that one group could increase its wealth and security relative to others.[54] During and after the Industrial Revolution, in contrast, states have been able to pursue these ends effectively by increasing productivity—as well as by gaining access to international trading networks and new technologies to enable further military and economic growth. Territorial conquest by states in the modern era is rarer for many reasons—but not simply because states have become more beneficent, instead because changes in technology have reshaped how they can best achieve their goals.[55] In short, major innovations can alter what long-term competitions in each era are fundamentally *about*. In the standard "ends, ways, means" trichotomy, this corresponds to *ways*. States have the

same ends (security, wealth, prestige, influence, sovereign action), but the ways that competition is best pursued can change, such as through participation in globalized production chains instead of territorial conquest.

With AI technologies, there are two worrying possibilities: a broad movement toward authoritarianism and the greater use of advanced forms of population- and economy-targeting information warfare.

Social Control Temptations

A technological innovation rarely tilts intrinsically toward "freedom" or "authoritarianism." It is possible, however, to try to discern how new technologies may affect current social and economic systems in the future. Especially in authoritarian states like China, AI technologies may provide elites with tools that reduce contradictions between maintaining power and promoting economic growth through free markets. By making authoritarianism appear more feasible, this may generate an "authoritarian temptation" for the many states with malleable governance systems.

First, AI technologies are likely to reduce the costs of controlling populations under authoritarian rule. Automating mass collection, processing, and analysis of data is likely to decrease the marginal cost of controlling additional citizens, thus reducing the resources required to indefinitely sustain totalitarianism. With access to hundreds of millions of cameras, social media postings, bank accounts, automated analysis of emotions and sentiment, and other data streams, AI-empowered algorithms can perform much of the work previously done by secret police in pre-AI authoritarian states.[56] Automated surveillance methods are likely to scale more effectively than manual surveillance, which requires some amount of human labor per citizen to be controlled. For example, Lichter et al. analyzed official Stasi records from East Germany, finding that more than 1.5 percent of the population was either officially employed or unofficially used as informers by the secret police.[57] Beyond the quantity of people involved in human surveillance operations, automated surveillance may impose lower economic costs on a society than human surveillance.[58]

On this matter, China appears poised to benefit from feedback cycles between AI deployment and data aggregation—the Chinese government is already using AI technologies to enhance population control, as well as to profile and control its ethnic minorities.[59] In these early efforts, the Chinese government is collecting large quantities of data, from facial scans to DNA; COVID-19 has only deepened PRC data collection on its citizens.[60] This data will help fuel new AI development for social control in Chinese firms. Future AI applications could, in turn, help China manage its data and drive more expansive collection, continuing the cycle.

China will likely export versions of these capabilities to authoritarian governments globally in the 2020s and 2030s, as it has already begun to do. According to recent CSET research, since 2008, over 80 countries have adopted Chinese surveillance technologies.[61] These tools will help authoritarian governments worldwide deepen their holds on power.[62]

Second, and more speculatively, AI progress may benefit authoritarian states by reducing the costs and consequences of state interventions into internal markets. The classic critique of centrally planning complex economies is that attempting to do so poses intractable optimization problems.[63] For many practical reasons, from human organizational factors to corruption, AI technologies are unlikely to change this. However, AI technologies could reduce, to some degree, the negative consequences of state interventions in markets.

For example, AI applications may help gather and interpret the volumes of information necessary for more effective economic controls. An analogous effect is visible inside large firms in both China and the United States today: companies like eBay, Taobao, Amazon, and Uber apply machine learning to mine large volumes of sales data to better match demand and supply. Modern machine learning tools enable automatic pattern analysis, improved forecasting, and natural language processing for predicting demand and performing sentiment analysis. Google's "Smart Bidding," for example, uses machine learning to optimize conversions for ads; California uses AI to predict electricity demand, more effectively controlling the power grid and reducing blackouts.[64] Walmart's internal logistical management has analogs to a centrally planned micro-economy.[65] There are many challenges to using analogous tools effectively for state economic policy, perhaps most of all the variable goals of planners themselves. But these trends suggest national-level strategic planning may be able to benefit from better information by applying modern machine learning tools to data accessible by states.

Leaders of authoritarian states like China may find themselves facing lower costs for sustaining domestic political and economic control; leaders of authoritarian-leaning states may find themselves handed these tools by China.

The effects of AI on population control and state interventions in markets are not certain. In the near term, however, it seems likely that Chinese elites at least *believe* that AI may help them better control their society, and so too may elites in other states.

Information Warfare

Besides increasing the fitness of authoritarian governments more generally, AI-enhanced information warfare may lower the costs of both influencing foreign populations and pursuing economic warfare policies at scale. If mass opinion can

be decisively influenced by the clash between AI influence systems, for example, China may determine its best bet for reabsorbing Taiwan is heavy investment in AI-empowered propaganda.

Information attacks can also target economic systems and financial markets, especially AI systems associated with managing equities investments. An unintentional, early demonstration of this possibility occurred in 2013, when U.S. trading algorithms responded to disinformation posted by the AP's Twitter account after it was hacked.[66] Information warfare may be increasingly linked to economic warfare, not just political disruptions.

Higher-end, AI-empowered information warfare is a more speculative, longer-term capability. Chris Wiggins has characterized current technical trends as enabling "reality jamming": the potential for synthetic, targeted, and optimized disinformation at web-scale.[67] In this future, current computational propaganda concerns are just the tip of the iceberg. The bigger issue is the potential for large-scale machine-generated information that is highly targeted at particular individuals or subpopulations, evolved to maximally shape particular behaviors, and potentially able to affect anyone with web access.[68]

Leveraging these developments, governments may attempt to shape perceptions of other populations more frequently than in the past.[69] OpenAI self-censored full publication of its GPT-2 language-generation model in 2019, for example, because it was concerned that generating close-to-human text would enable nefarious actors to proliferate disinformation. It is easy to imagine states pursuing similar capabilities for their own ends.[70] According to recent CSET research, GPT-2's successor, GPT-3, may be especially potent at generating disinformation at scale when steered by a skilled human operator and editor, opening up the possibility of highly effective human-machine teaming.[71]

These trends may pose challenges for democratic societies, though it is still too early to make clear judgments. Three unresolved questions exist today: First, if a long-term risk in authoritarian systems is intellectual conformity, an analogous effect in democracies may be mob majoritarianism.[72] This inherent challenge in democratic societies could turn out to be exacerbated by modern information technologies and make organizational reforms even more difficult. Second, more research is needed to understand the balance between democracies' ability to use disagreements and diverse information to advance new explanations and solutions, and the potential for information attacks to undermine political stability.[73] And third, most fundamentally, Western democracies, and particularly the U.S. system of government, are based on a foundation of individual freedom where individuals are the

best judges of their own interests. It is not yet obvious how Western institutions will adapt to machines that can anticipate—or shape—individuals' own preferences, states, and choices better than the individuals themselves can.[74]

In the context of international competition, leveraging AI technologies to alter target states' national priorities or political stability through information warfare would represent "winning without fighting" *par excellence*.

3 Conclusions and Key Points

In this evolutionary theory of technological competition, AI's effects on national power fall into three categories: new elements of power, changed factors, and altered goals. Exploring *new elements* required for successful AI adoption, such as compute and organizational adaptations, helps us understand when, how, and why some societies may be better positioned than others to benefit from major innovations. Similarly, the idea of changed factors helps focus on how existing elements of national power may have changing importance, such as population size and industrious researchers. Finally, thinking about altered goals of states in competition shows how major innovations can reshape the ways that states engage in competition, such as enacting new domestic political and economic controls and leveraging AI-enabled information attacks on other states' social and economic systems. This research offers a way to start thinking about these issues together, and hopes to spur new, wider thinking and work.

Creating new conceptual tools for U.S. decision-makers and analysts to make sense of AI technologies' effects is vital to American prosperity. Over the long term, these technologies will create significant changes in U.S.-China competition.

Creating new conceptual tools for U.S. decision-makers and analysts to make sense of AI technologies' effects is vital to American prosperity. Over the long term, these technologies will create significant changes in U.S.-China competition.

From this research, we see three early sets of insights into opportunities for U.S. leaders:

- Thinking of long-term competitions in an evolutionary framework makes large, broadly-diffused technology changes akin to environmental shifts. Like a volcanic eruption or the start of an ice age,

broad adaptations are valuable and some states will be better at adapting than others. It is useful to begin thinking about how AI technologies can *create new elements of power, change the importance of existing elements of power, and alter the goals of states in competition*. Getting a better sense of AI's effects in each of these factors will be critical for major powers. The United States has a number of opportunities: studying the approaches of other countries, especially U.S. competitors and medium-sized, quickly-changing countries;[75] developing strategies for global leadership in producing, using, and sharing compute resources; supporting development of AI engineering as a rigorous discipline in the United States and leveraging humans trained in it; continuing to push DOD and IC organizational reforms for how data is managed and leveraged; and leveraging AI tools, cross-training between AI and other disciplines, and high-skilled STEM immigration to access new breakthroughs in science and engineering more widely.

- AI technologies may change not only what states can do, but also what they want. Major innovations can broadly alter intermediate, instrumental objectives that states pursue by making certain kinds of behaviors more valuable or less costly. This can drive dramatic changes in state goals and policies. The United States may look for new opportunities in technology-related democracy promotion; shaping AI technologies themselves to favor democracies, such as by supporting development of AI technologies with less dependence on centralized data; [76] and developing approaches to more rapidly adapt social and economic institutions to "information attacks" by AI systems.

- Finally, effects of technological change can be highly asymmetric: *new elements, changed factors, and altered goals* may have very different manifestations in different countries. For the United States, this means learning from its competitors without mirror imaging them and sharing insights with allies before assuming they should symmetrically match U.S. policies. Perhaps most significantly, it may also mean looking ahead to how AI technologies may affect the aims and interests of U.S. allies and partners.

The scale of possible impacts from major technologies is obvious: the United States benefitted greatly from growth connected to technological and economic changes in the 40 years from 1880 through 1920; and China has also already benefitted from a mix of technological and economic changes in its resurgence from 1980 through 2020.[77] Recent history demonstrates that getting technology right is

critical for long-term national flourishing—and determining trajectories for the United States and China over the next 20 to 30 years.

Can we sketch the longer-term future? Only speculation is possible today:

Broad historical examinations tend to suggest that more successful societies present fewer obstacles to long-term change and, especially, limit the costs of intellectual conformity. They seek to maximize the benefits of pluralism, competition, and mechanisms to share, challenge, and supplement new knowledge.[78]

A key challenge for China will be limiting the long-term costs of intellectual conformity induced by an authoritarian government. A favorable factor for China will be the dynamic organizations it has built over the last 20 years, which may remain able to adapt and benefit from organizational learning as the world continues to change over the next 10 to 20 years. In the longer term, however, continued evolution seems increasingly challenging for China under the CCP and absent substantial pluralism; many of its main challenges for net economic-technological growth are likely to persist, while the benefits of its dynamic organizations are likely to decline over time.

A likely challenge for the United States will be institutional and organizational sclerosis, which will make organizational learning and adaptation challenging over the next decade. Interactions between AI technologies and democratic institutions increase uncertainty and may exacerbate these challenges. Weighing against these factors is Samuel Huntington's reminder of the United States' multidimensional sources of power and ability for self-renewal.[79] The most favorable factors for U.S. vitality and competition with authoritarian governments coincide with its enduring strengths: areas such as its cultural values and pluralism, overall approach to governance, and access to global talent.[80] In the longer term, the United States' central challenges appear more temporary, and its greatest advantages more enduring—a favorable outlook achievable with thinking and work today.

Appendix

In 1948, after John von Neumann gave a talk on computing machines in Princeton, a member of the audience asked the canonical question: *Of course, machines can't really think, can they?* Von Neumann replied, "You insist that there is something a machine cannot do. If you will tell me precisely what it is that a machine cannot do, then I can always make a machine which will do just that!"[81] Part of the challenge of defining AI has been that defining intelligence and thinking in humans continues to be difficult.

This paper uses the definition of AI from the Defense Innovation Board: *a variety of information processing techniques and technologies used to perform a goal-oriented task and the means to reason in the pursuit of that task*.[82] More colloquially, AI can be thought of as a broad discipline and set of technologies centered on creating machines that can make decisions relatively well under uncertainty.[83]

It is useful to distinguish AI from *autonomy*. The former is defined above; the latter is best thought of as some degree of delegation of decision-making agency to another entity, which could be a human or a machine.[84] Systems can have neither, both, or one of these two things. For example, an autonomous military system can be unintelligent, as in the case of a landmine, or an intelligent system can support humans without autonomy, as in the case of an information system for a pilot.

The 2010s were the third period of global excitement about AI. The first period occurred in the 1960s, centered in the United States and the UK, and the second period occurred in the 1980s, centered in the United States and Japan. Both periods were associated with significant investment and optimism for cascading breakthroughs in machine intelligence. Both periods were followed by "AI winters": periods of widespread divestment from AI R&D and the belief that earlier expectations had far exceeded reality.[85] The current period will probably be remembered as being centered in the United States and China, though with substantial activity in the UK, Europe, Canada, Japan, Israel, and South Korea.

Since the 2010s, most excitement about AI has focused on *machine learning* (ML), and, within ML, mostly on applications of neural networks (*deep learning*). ML is a broad subfield of AI that centers on inference from data and overlaps substantially with statistics and optimization. "Neural networks" refers to a family of statistical models for extracting patterns from large quantities of data, originally inspired by the behavior of biological neurons.

While the rediscovery and improvement of neural nets started the current AI wave in the late 2000s, specific trends over the last 20 to 30 years enabled the success of recent applications: global growth and diffusion of compute resources; large quantities of digital data globally; and the connection of these two by the global internet. For this reason, the foundation of modern AI advancements is often called the "triad" of new algorithms, compute resources, and data.[86]

Endnotes

1. On AI-specific vulnerabilities, see Andrew Lohn, "Hacking AI: A Primer for Policymakers on Machine Learning Cybersecurity," CSET, December 2020, https://cset.georgetown.edu/research/hacking-ai/. On AI-specific accidents, see Tim G. J. Rudner and Helen Toner, "Key Concepts in AI Safety: An Overview," *CSET*, March 2021, https://cset.georgetown.edu/research/key-concepts-in-ai-safety-an-overview/.
2. Emily O. Goldman and Richard B. Andres, "Systemic effects of military innovation and diffusion," *Security Studies* 8 (1999), 116.
3. Dahlia Peterson, "Designing Alternatives to China's Repressive Surveillance State," *CSET*, October 2020, https://cset.georgetown.edu/research/designing-alternatives-to-chinas-repressive-surveillance-state/; Tim Hwang, "Shaping the Terrain of AI Competition," *CSET*, June 2020, https://cset.georgetown.edu/research/shaping-the-terrain-of-ai-competition/; and Andrew Imbrie, Ryan Fedasiuk, Catherine Aiken, Tarun Chhabra, and Husanjot Chahal, "Agile Alliances: How the United States and Its Allies Can Deliver a Democratic Way of AI," CSET, February 2020, https://cset.georgetown.edu/research/agile-alliances/.
4. Joseph Nye, *The Paradox of American Power: Why the World's Only Superpower Can't Go It Alone*, (New York: Oxford University Press, 2002), p. 12.
5. Morgenthau carefully distinguished that some of these are relatively stable over time, whereas others are more subject to frequent change. Some also have useful sub-elements to consider, making the complete set: geography, access to resources (including food and raw materials), industrial capacity, military preparedness (including technology, leadership, and quantity and quality of armed forces), population (including distribution and trends), national character, national morale, quality of diplomacy, and quality of government. See: Hans Morgenthau, *Politics Among Nations: The Struggle for Power and Peace* (New York: Alfred A. Knopf, 1948), pp. 102-152.
6. Indeed, for some political scientists, the lack of such a consensus measure in fact is the cause of all wars, as if states were to agree who would win in advance, there would be no reason to pay the costs of war at all, instead of simply bargaining. See James D. Fearon, "Rationalist Explanations for War," *International Organization* 49, no. 3 (1995), 381, 390-401; and Bernard Brodie, *War and Politics* (New York: MacMillan Publishing Co., Inc., 1973), 35-6, 63.
7. One 2011 survey of the literature catalogued 69 different proposals of equations for measuring national power. See Karl Hohn, "Geopolitics and the Measurement of National Power," PhD diss., (University of Hamburg, 2011). Some selected examples:

   Joseph Nye distinguished military power, economic power, and soft power. He noted that "power resources" for the United States in the 20th Century included economic scale, scientific and technological leadership, location, military forces and alliances, universalistic culture and liberal international regimes; for the 21st century he suggested the corresponding elements as technological leadership, military and economic scale, soft power, and being a hub of transnational communications. See: Joseph Nye, *The Paradox of American Power: Why the World's Only Superpower Can't Go It Alone*, (New York: Oxford University Press, 2002), pp. 4-12.

   Robert Gilpin focuses power simply as the "military, economic, and technological capabilities of states," and notes that he intentionally leaves out "important and intangible elements that affect the outcomes of political actions, such as public morale [and] qualities of leadership." See: Robert Gilpin, *War & Change in World Politics*, (New York: Cambridge University Press, 1981), p. 13-14.

   The Correlates of War project, a widely used quantitative database for studying warfare, defines a "Composite Index of National Capability" (CINC) in terms of a country's share of world population, urban

population, iron and steel production, energy consumption, military expenditure, and military personnel. (See: Singer, J. David, Stuart Bremer, and John Stuckey, "Capability Distribution, Uncertainty, and Major Power War, 1820-1965," in Bruce Russett (ed.) *Peace, War, and Numbers*, (Beverly Hills: Sage, 1972), pp. 1948, as well as. https://correlatesofwar.org/data-sets/national-material-capabilities.)

More recently, Michael Beckley has argued that traditional measures of power conflate gross resources with net resources, and thus fail to account for a country's burdens in addition to its assets. Thus, he proposes the use of "GDP * GDP per capita." See "The Power of Nations: Measuring What Matters," *International Security* 43.2 (2018): 7-44.

8. The best overview of this quest is Ashley J. Tellis, Janice Bially, Christopher Layne, and Melissa McPherson, "Measuring National Power in the Postindustrial Age," *RAND Corporation*, 2000.
9. This statement includes whether such measures are used quantitatively to predict who wins a war, or whether war will occur, or whether settlement terms will favor one side or another. See: Ibid., 17.
10. For various discussions of this, see: Stephen Biddle, *Military Power: Explaining Victory and Defeat in Modern Battle* (Princeton: Princeton University Press, 2004); David A. Baldwin, "Power Analysis and World Politics: New Trends versus Old Tendencies," World Politics 161 (1979): 161-94; Jeffrey Hart, "Three Approaches to the Measurement of Power in International Relations," *International Organization* 30 (1976), 289-305.
11. Almond and Genco (1977) most famously made this point about how to think about political phenomena in general. See Gabriel A. Almond and Stephen J. Genco, "Clouds, Clocks, and the Study of Politics," *World Politics* 29.4 (1977): 489-522.
12. Paul Kennedy, *The Rise and Fall of the Great Powers*, (New York: Random House, 1987), p. 199.
13. Andrew Marshall, "RMA Update," Memorandum for the Record, 2 May 1994.
14. For many examples collected in one place, see: Herbert Goldhamer, "Reality and Belief in Military Affairs," RAND Corporation, 1977.
15. For others who have offered ways of thinking about this, all with substantial detail, see: George Modelski and William R. Thompson, *Leading Sectors and World Powers* (Columbia: University of South Carolina Press, 1996); Paul Kennedy, *The Rise and Fall of the Great Powers* (New York: Random House, 1987); Robert Gilpin, *War and Change in World Politics* (UK: Cambridge University Press, 1981).
16. On steel, see Goldman and Andres, "Systemic effects of military innovation and diffusion," 116. On the internal combustion engine and oil, see W. G. Jensen, "The Importance of Energy in the First and Second World Wars," *The Historical Journal* 11 (1968): 538-54. On uranium, see R. Scott Kemp, "The Nonproliferation Emperor Has No Clothes: The Gas Centrifuge, Supply-Side Controls, and the Future of Nuclear Proliferation," *International Security* 38 (2014): 39-78, especially 41-4.
17. Michael C. Horowitz, *The Diffusion of Military Power: Causes and Consequences for International Politics* (Princeton: Princeton University Press, 2010).
18. On the Mongols, see Goldman and Andres, "Systemic effects of military innovation and diffusion," 102, 88-9. On modern power projection and the loss of American "free security," see C. Vann Woodward, "The Age of Reinterpretation," *The American Historical Review* 66 (1960): 1-19. On nuclear weapons, see Thomas C. Schelling, *Arms and Influence* (New Haven: Yale University Press, 2008), 30-1.
19. Gilpin, p. 23.
20. This is an old idea about artificial intelligence, but was brought to our attention in the modern context by Richard Danzig in July 2016, during discussion as part of the review board of a DOD Summer Study.
21. A voluminous literature discusses AI's military applications. See an overview at Daniel S. Hoadley and Kelley M. Sayler, "Artificial Intelligence and National Security," *Congressional Research Service*, November 10, 2020, https://fas.org/sgp/crs/natsec/R45178.pdf. For further reading, see Robert O. Work and Shawn Brimley, "20YY: Preparing for War in the Robotic Age," CNAS, January 2014, https://s3.amazonaws.com/files.cnas.org/documents/CNAS_20YY_WorkBrimley.pdf?; Luttwak (ONA, 2020); Christian Brose, *The Kill Chain: Defending America in the Future of High-Tech Warfare* (New York: Hachette Books, 2020).
22. National Research Council, "Getting Up to Speed: The Future of Supercomputing," (Washington, DC: The National Academies Press, 2005), https://www.nap.edu/catalog/11148/getting-up-to-speed-the-future-of-supercomputing.

23. Dario Amodei and Danny Hernandez, "AI and Compute," *OpenAI*, 16 May 2018, https://openai.com/blog/ai-and-compute/.
24. Girish Sastry, Jack Clark, Greg Brockman, and Ilya Sutskever, "Addendum: AI and Compute," OpenAI, 7 November 2019, https://openai.com/blog/ai-and-compute/.
25. Xue Bin Peng, Lerrel Pinto, Alex Ray, Bob McGrew, Jonas Schneider, Josh Tobin, Marcin Andrychowicz, Peter Welinder, Pieter Abbeel, and Wojciech Zaremba, "Generalizing from Simulation," *OpenAI*, October 19, 2017, https://openai.com/blog/generalizing-from-simulation/. For a recent study substituting simulated for real-world data in a military context, see Li Ang Zhang, Jia Xu, Dara Gold, Jeff Hagen, Ajay K. Kochhar, Andrew J. Lohn, and Osonde A. Osoba, "Air Dominance Through Machine Learning – A Preliminary Exploration of Artificial Intelligence–Assisted Mission Planning," RAND, 2020, https://www.rand.org/pubs/research_reports/RR4311.html.
26. CSET has a line of research both explaining and advising on how to maintain this state of affairs. See Saif M. Khan, "Securing Semiconductor Supply Chains" (Washington, DC: Center for Security and Emerging Technology, January 2021); Saif M. Khan, "The Semiconductor Supply Chain: Assessing National Competitiveness," *CSET*, January 2021, https://cset.georgetown.edu/research/the-semiconductor-supply-chain/; Saif M. Khan, "U.S. Semiconductor Exports to China: Current Policies and Trends" CSET, October 2020, https://cset.georgetown.edu/wp-content/uploads/U.S.-Semiconductor-Exports-to-China-Current-Policies-and-Trends.pdf; Saif M. Khan and Carrick Flynn, "Maintaining China's dependence on democracies for advanced computer chips," *Brookings*, April 2020, https://www.brookings.edu/research/maintaining-chinas-dependence-on-democracies-for-advanced-computer-chips/; Saif M. Khan, "AI Chips: What They Are and Why They Matter," *CSET*, April 2020, https://cset.georgetown.edu/research/ai-chips-what-they-are-and-why-they-matter/.
27. E. A. Huerta, Asad Khan, Edward Davis, Colleen Bushell, William D. Gropp, Daniel S. Katz, Volodymyr Kindratenko, Seid Koric, William T. C. Kramer, Brendan McGinty, Kenton McHenry, and Aaron Saxton, "Convergence of artificial intelligence and high performance computing on NSF-supported cyberinfrastructure," *Journal of Big Data* 88 (2020).
28. This suggestion offered by Jack Clark in early comments on this paper.
29. The innovation literature is large. Seminally, see Barry R. Posen, *The Sources of Military Doctrine: France, Britain, and Germany Between the World Wars* (New York: Cornell University Press, 1984); Stephen P. Rosen, *Winning the Next War: Innovation and the Modern Military* (New York: Cornell University Press, 1991); Theo G. Farrell and Terry Terriff, *The Sources of Military Change: Culture, Politics, Technology* (Colorado: Lynne Rienner Publishers, 2002).
30. Edmund Beard, *Developing the ICBM: A Study in Bureaucratic Politics* (New York: Columbia University Press, 1976).
31. Goldman and Andres, "Systemic effects of military innovation and diffusion."
32. Andrew Marshall, "RMA Update," Memorandum for the Record, 2 May 1994; Horowitz, *The Diffusion of Military Power*.
33. Marcus Comiter, "Attacking Artificial Intelligence: AI's Security Vulnerability and What Policymakers Can Do About It," *Belfer Center for Science and International Affairs*, August 2019, https://www.belfercenter.org/sites/default/files/2019-08/AttackingAI/AttackingAI.pdf.
34. As Alan Turing wrote in 1950, "An important feature of a learning machine is that its teacher will often be very largely ignorant of quite what is going on inside." (Alan Turing, "Computing Machinery and Intelligence," *Mind*, Volume LIX, Issue 236, October 1950, p. 458.) With deep learning, this problem is especially acute due to the scale of statistical models involved – for example, ResNet, a commonly used image classification architecture, uses around 5*107 parameters. What is layer 27 of a hundred-layer neural network doing? (See, for example: Leilani H. Gilpin, David Bau, Ben Z. Yuan, Ayesha Bajwa, Michael Specter, and Lalana Kagal, "Explaining Explanations: An Overview of Interpretability of Machine Learning," arXiv, February 3, 2019, https://arxiv.org/pdf/1806.00069.pdf.) Intuitively, it is difficult for a human being to understand the inner workings of the model with any precision. For a very good effort see: Chris Olah, Arvind Satyanarayan, Ian Johnson, Shan Carter, Ludwig Schubert, Katherine

Ye, and Alexander Mordvintsev, "The Building Blocks of Interpretability," *Distill*, 2018, https://distill.pub/2018/building-blocks/.

35. Michael Jordan, "Artificial Intelligence—The Revolution Hasn't Happened Yet," *Harvard Data Science Review*, July 2019, https://hdsr.mitpress.mit.edu/pub/wot7mkc1.
36. Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané, "Concrete Problems in AI Safety," *arXiv*, July 25, 2016, https://arxiv.org/pdf/1606.06565.pdf.
37. Michèle A. Flournoy, Avril Haines, and Gabrielle Chefitz, "Building Trust through Testing: Adapting DOD's Test & Evaluation, Validation & Verification (TEVV) Enterprise for Machine Learning Systems, including Deep Learning Systems," *CSET / WestExec Advisors*, 2020, https://cset.georgetown.edu/wp-content/uploads/Building-Trust-Through-Testing.pdf.
38. For modern machine learning systems specifically, these issues are introduced in places like: Nicolas Papernot et al., "SoK: Towards the Science of Security and Privacy in Machine Learning," Nov 2016, arXiv:1611.03814v1; Gamaleldin Elsayed, Ian Goodfellow, and Jascha Sohl-Dickstein, "Adversarial Reprogramming of Neural Networks," June 2018, arXiv:1806.11146v1; and Nicholas Carlini et al., "On Evaluating Adversarial Robustness," February 2019, arXiv:1902.06705v2.
39. A companion reading for thinking about this might be: R. V. Jones, *The Wizard War: British Scientific Intelligence*, 1939-1945, (New York: Coward, McCann & Geoghegan: 1978).
40. Perhaps the best introduction to this was provided by Tim Hwang et al.: "For a heart-stopping few minutes on May 6, 2010, the Dow Jones Industrial Average dropped a staggering 1,000 points—and inexplicably proceeded to recover all of those losses within the following few minutes. The Flash Crash, as it was later dubbed, remains the biggest one-day point decline in Dow Jones history. After a five-month investigation, the SEC reported that the sudden loss and gain that day was the result of an unusually large number of contracts being sold by a mutual fund, which triggered a wave of aggressive sell-off activity from untold numbers of firms running automated high frequency trading programs. No human agency was at the heart of the momentary crash. Instead, it appears that unanticipated interactions among multiple automated scripts designed to buy and sell stock produced the precipitous fall and rise in prices. Financial robots may also be behind the otherwise inexplicable correlations between mentions of the actor Anne Hathaway in the news and increases in the stock price of Warren Buffet's Berkshire Hathaway fund." See: Tim Hwang, Ian Pearce, and Max Nanis, "Socialbots: Voices from the Fronts," *Interactions*, March-April 2012. More recently, and closer to a direct example, in 2013 trading systems responded to information from the AP's twitter feed after it had been hacked by (apparently) Syrian dissidents, causing a temporary drop of $130B. See: Max Fisher, "Syrian hackers claim AP hack that tipped stock market by $136 billion. Is it terrorism?" *Washington Post*, 23 April 2013, https://www.washingtonpost.com/news/worldviews/wp/2013/04/23/syrian-hackers-claim-ap-hack-that-tipped-stock-market-by-136-billion-is-it-terrorism/.
41. Daron Acemoglu and Pascual Restrepo, "Demographics and Automation," *NBER*, March 2018, https://www.nber.org/papers/w24421.
42. Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (New York: W. W. Norton & Company, 2018).
43. Todd Schneider, Gee Hee Hong, and Anh Van Le, "Land of the Rising Robots," *IMF*, June 2018, https://www.imf.org/external/pubs/ft/fandd/2018/06/japan-labor-force-artificial-intelligence-and-robots/schneider.pdf.
44. Ewen Callaway, "'It will change everything': DeepMind's AI makes gigantic leap in solving protein structures," *Nature*, November 30, 2020, https://www.nature.com/articles/d41586-020-03348-4.
45. For a useful overview: Jeffrey Dean, "The Deep Learning Revolution and Its Implications for Computer Architecture and Chip Design," *arXiv*, 13 Nov 2019, https://arxiv.org/ftp/arxiv/papers/1911/1911.05289.pdf.
46. Azalia Mirhoseini, Anna Goldie, Mustafa Yazgan, et al. "A graph placement methodology for fast chip design," Nature 594, 207–212 (2021). https://doi.org/10.1038/s41586-021-03544-w.
47. For example, suggesting valuable hypotheses to test or engineering design configurations. This goal has had a resurgence in the 2010s: in 2016 Hiroaki Kitano, creator of Robocup, proposed a grand challenge for AI systems capable of making Nobel-worthy scientific discoveries: Hiroaki Kitano, "Artificial intelligence

to win the nobel prize and beyond: Creating the engine for scientific discovery." *AI magazine* 37, no. 1 (2016): 39-49. A recent data brief by CSET also surveys how AI technologies have been accelerating growth of new science and engineering research clusters across a broad span of disciplines, see: Matthew Daniels, Autumn Toney, Melissa Flagg, and Charles Yang, "Machine Intelligence for Scientific Discovery and Engineering Invention," CSET, May 2021, https://cset.georgetown.edu/publication/machine-intelligence-for-scientific-discovery-and-engineering-invention/.

48. For example: Kevin Williams, Elizabeth Bilsland, Andrew Sparkes, Wayne Aubrey, Michael Young, Larisa N. Soldatova, Kurt De Grave et al. "Cheaper faster drug development validated by the repositioning of drugs against neglected tropical diseases." *Journal of the Royal Society Interface* 12, no. 104 (2015): 20141289.
49. For example, see: Volodymyr Kuleshov, Jialin Ding, Christopher Vo, Braden Hancock, Alexander Ratner, Yang Li, Christopher Ré, Serafim Batzoglou, and Michael Snyder, "A machine-compiled database of genome-wide association studies," *Nature Communications* 10, 3341 (2019), https://doi.org/10.1038/s41467-019-11026-x.
50. See, for example: Freeman Dyson, "Is Science Mostly Driven by Ideas or by Tools?" *Science*, Vol. 338 (December 2012): 1426-1427.
51. Cade Metz, "In Two Moves, AlphaGo and Lee Sedol Redefined the Future," *Wired*, March 16, 2016, https://www.wired.com/2016/03/two-moves-alphago-lee-sedol-redefined-future/.
52. Iain M. Cockburn, Rebecca Henderson, and Scott Stern, "The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis," in *The Economics of Artificial Intelligence: An Agenda*, eds. Ajay Agrawal, Joshua Gans, and Avi Goldfarb (Chicago: University of Chicago Press, 2017), 115-46.
53. Other work has made analogous arguments, though from a different perspective, see: Remco Zwetsloot and Zachary Arnold, "Foreign Brains Help America Compete," *Wall Street Journal*, January 30, 2020; Remco Zwetsloot and Dahlia Peterson, "The US-China Tech Wars: China's Immigration Disadvantage," *The Diplomat*, December 31, 2019; Remco Zwetsloot, Roxanne Heston, and Zachary Arnold "Strengthening the U.S. AI Workforce," CSET, September 2019, https://cset.georgetown.edu/publication/strengthening-the-u-s-ai-workforce/.
54. Gilpin, p. 23.
55. Stephen Brooks, *Producing Security: Multinational Corporations, Globalization, and the Changing Calculus of Conflict* (Princeton: Princeton University Press, 2005).
56. Ross Andersen, "The Panopticon Is Already Here," *The Atlantic*, September 2020, https://www.theatlantic.com/magazine/archive/2020/09/china-ai-surveillance/614197/; Ben Angel Chang, "AI and US-China Relations," in Nicholas D. Wright, ed., *AI, China, Russia, and the Global Order: Technological, Political, Global, and Creative Perspectives* (DOD SMA: December 2018).
57. Andreas Lichter, Max Löffler, and Sebastian Siegloch, "The long-term costs of government surveillance: Insights from Stasi spying in East Germany," *SOEPpapers on Multidisciplinary Panel Data Research* 865 (2016): 1-60. Available online: https://www.econstor.eu/bitstream/10419/146890/1/869045423.pdf.
58. This remains speculative. For example, some evidence suggests surveillance itself depresses economic activity by eroding social trust, causing individuals to reduce their productive activity. (See: Lichter, Löffler, and Siegloch, "The long-term costs of government surveillance: Insights from Stasi spying in East Germany," p. 22). Other studies have found interpersonal trust to correlate with entrepreneurship and innovation. (See: Stephen Knack and Philip Keefer, "Does Social Capital Have an Economic Payoff? A Cross-Country Investigation," *The Quarterly Journal of Economics* 112 (1997): 1251-88.) This effect was particularly acute because Stasi informants retained their normal roles as colleagues, family, and friends, and so the knowledge of Stasi presence caused widespread doubt and fear. Automated technological surveillance would plausibly avoid these effects. (See: Lichter et al., "The long-term costs of government surveillance," p. 22.)
59. Paul Mozur, "One Month, 500,000 Face Scans: How China Is Using A.I. to Profile a Minority," *New York Times*, April 2019; Josh Chin and Liza Lin, "China's All-Seeing Surveillance State Is Reading Its Citizens' Faces," *Wall Street Journal*, June 2017.
60. Liza Lin and Shan Li, "Chinese Citizens Must Scan Their Faces to Register for New Mobile-Phone Service," *Wall Street Journal*, December 2019; Sui-Lee Wee, "China Uses DNA to Track Its People, With the Help of American Expertise," *New York Times*, February 2019. Shan Li, "Made-in-China Censorship for Sale," *Wall*

Street Journal, March 2020; Paul Mozur, Raymond Zhong, and Aaron Krolik, "In Coronavirus Fight, China Gives Citizens a Color Code, With Red Flags," *New York Times*, March 2020. CSET has also considered possible future trends based on COVID-driven increases in surveillance: https://www.cset-foretell.com/blog/surveillance-creep.

61. Dahlia Peterson, "Designing Alternatives to China's Repressive Surveillance State," *CSET*, October 2020, https://cset.georgetown.edu/research/designing-alternatives-to-chinas-repressive-surveillance-state/.
62. Jessica Chen Weiss, "Understanding and Rolling Back Digital Authoritarianism," *War on the Rocks*, February 17, 2020, https://warontherocks.com/2020/02/understanding-and-rolling-back-digital-authoritarianism/; Jessica Chen Weiss, "An Ideological Contest in U.S.-China Relations? Assessing China's Defense of Autocracy," forthcoming in *Security and US-China Relations: Differences, Dangers, and Dilemmas*, eds. Avery Goldstein and Jacques deLisle. Available online: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3427181.
63. Ludwig von Mises, *Human Action* (Chicago: Contemporary Books, Inc., 1963), 678-80; More colloquially, see Cosma Shalizi, "In Soviet Union, Optimization Problem Solves You," *Crooked Timber*, May 30, 2012, http://crookedtimber.org/2012/05/30/insoviet-union-optimization-problem-solves-you/.
64. Paul R. Milgrom and Steve Tadelis, "How Artificial Intelligence and Machine Learning Can Impact Market Design," forthcoming in *The Economics of Artificial Intelligence*, eds. Ajay K. Agrawal, Joshua Gans, and Avi Goldfarb (Chicago: University of Chicago Press, 2019), 1-24. Available online: https://www.nber.org/books/agra-1.
65. Leigh Phillips and Michal Rozworski, "The People's Republic of Walmart: How the World's Biggest Corporations are Laying the Foundation for Socialism" (New York: Verso, 2019).
66. Max Fisher, "Syrian hackers claim AP hack that tipped stock market by $136 billion. Is it terrorism?" *Washington Post,* 23 April 2013, https://www.washingtonpost.com/news/worldviews/wp/2013/04/23/syrian-hackers-claim-ap-hack-that-tipped-stock-market-by-136-billion-is-it-terrorism/.
67. Chris Wiggins, "Reality Jamming: technology-enabled misinformation at scale," independent abstract, 2017; Susan McGregor, Chris Wiggins, Joan Donovan, Matt Jones, Jonathan Albright, and Sam Thielman, "Reality Jamming: The Future of Information Online," *Tow Center*, December 11, 2017, https://medium.com/tow-center/reality-jamming-the-future-of-information-online-3ad5cb0d932e.
68. See also: Matt Chessen, "The MADCOM Future," *The Atlantic Council*, 2017, https://www.atlanticcouncil.org/wp-content/uploads/2017/09/The_MADCOM_Future_RW_0926.pdf.
69. The Office of Net Assessment sponsored early work on this. For example, see: Michael J. Mazarr, Abigail Casey, Alyssa A. Demus, Scott W. Harold, Luke J. Matthews, Nathan Beauchamp-Mustafaga, James Sladden, "Hostile Social Manipulation: Present Realities and Emerging Trends," *RAND Corporation*, 2019, https://www.rand.org/pubs/research_reports/RR2713.html.
70. Alec Radford, Jeffrey Wu, Dario Amodei, Daniela Amodei, Jack Clark, Miles Brundage, and Ilya Sutskever, "Better Language Models and Their Implications," *OpenAI*, February 14, 2019, https://openai.com/blog/betterlanguage-models/; Dipayan Ghosh and Ben Scott, "Digital Deceit: The Technologies Behind Precision Propaganda on the Internet," *New America*, January 23, 2018, https://www.newamerica.org/public-interest-technology/policypapers/digitaldeceit/, 26-8; Sarah Kreps and Miles McCain, "Not Your Father's Bots: AI Is Making Fake News Look Real," *Foreign Affairs*, https://www.foreignaffairs.com/articles/2019-08-02/not-your-fathers-bots.
71. Ben Buchanan, Andrew Lohn, Micah Musser, and Katerina Sedova, "Truth, Lies, and Automation: How Language Models Could Change Disinformation," *CSET,* May 2021, https://cset.georgetown.edu/publication/truth-lies-and-automation/.
72. For example, Tocqueville's admonition: "If, in place of all the diverse powers that hindered or slowed beyond measure the rapid development of individual reason, democratic peoples substituted the absolute power of a majority, the evil would only have changed character. Men would not have found the means to live independently; they would only have discovered, a difficult thing, a new face of servitude. I cannot say it enough: for those who see liberty of the mind as a holy thing, and who hate not only the despot but also despotism, there is in that something to make them reflect deeply. For me, when I feel the hand of power pressing on my head, knowing who is oppressing me matters little to me, and I am no more inclined to put

my head in the yoke, because a million arms present it to me." (Alexis de Tocqueville, *Democracy in America*, Edited by Eduardo Nolla, Translated by James T. Schleifer (Indianapolis: Liberty Fund, 2012), Vol. II, Section 1, Chapter 2.)

73. Henry John Farrell and Bruce Schneier, "Common-Knowledge Attacks on Democracy," *Berkman Klein Center* Research Publication No. 2018-7, Available online: https://ssrn.com/abstract=3273111
74. Richard Danzig, "An Irresistible Force Meets a Moveable Object: The Technology Tsunami and the Liberal Order," *Lawfare Research Paper Series* 5.1 (2017), https://assets.documentcloud.org/documents/3982439/Danzig-LRPS1.pdf, 4-7.
75. For learning from U.S. competitors, see: Peter Westwick, "Lessons from Stealth for Emerging Technologies," *CSET*, March 2021, https://cset.georgetown.edu/publication/lessons-from-stealth-for-emerging-technologies/, pp. 25-26.
76. Tim Hwang, "Shaping the Terrain of AI Competition," *CSET*, June 2020, https://cset.georgetown.edu/publication/shaping-the-terrain-of-ai-competition/.
77. The United States emerged as a major power following industrialization of its economy and society over the half-century from 1875 to 1925. This period was particularly unstable for the international system, with rapid technological change and uneven growth—the U.S. fraction of global manufacturing output more than doubled from 14.7 percent in 1880 to 39.3 percent in 1928. (See: Paul Kennedy, *The Rise and Fall of the Great Powers*, (New York: Random House, 1987), p. 202.) As industrialization transformed the U.S. economy and society, population growth allowed the United States to harness these changes into national power. U.S. population increased from 44M in 1874 to 114M in 1924. (See: Hans Morgenthau, *Politics Among Nations*, (New York: Knopf, 1956), p. 114.) This was well above the populations of Germany, Japan, France, Britain, and Italy. (Paul Kennedy, *The Rise and Fall of the Great Powers*, (New York: Random House, 1987), p. 199.) The United States in this period became both industrialized and populous relative to other countries and, by 1920, was the strongest power on the planet. A century later, China is undergoing an analogous shift, but with still-uncertain results. China began instituting major economic reforms in 1979. In the period from approximately 1980-2020, China transformed its economy, society, and military, partly by harnessing modern information technologies. In this period, China's fraction of global GDP (by PPP) increased from 2.3 percent in 1980 to 18.3 percent in 2017, while the United States declined from 24.3 percent to 15.3 percent in the same period. (Wayne Morrison, "China's Economic Rise: History, Trends, Challenges, and Implications for the United States," *Congressional Research Service*, June 2019, p. 10. https://www.everycrsreport.com/files/20190625_RL33534_088c5467dd11365dd4ab5f72133db289fa10030f.pdf) China already had the population needed to harness economic reforms for growth. If U.S. growth in economic power came from industrialization and population growth, China's could be described as coming from large-scale capital investment and productivity growth—the latter due to both resource reallocations and imported technologies and processes. (Morrison, "China's Economic Rise", p. 6-7) China's trajectory for the next 20-30 years, however, remains highly uncertain. Whether China can continue to sustain substantial economic growth depends in significant part on the degree to which it can make new technology and innovation a source of future growth. (Morrison, "China's Economic Rise", p. 7-8.) China also faces enormous demographic, environmental, public health, and peripheral security challenges that will impose large costs on its government. (See, for example: Michael Beckley, *Unrivaled: Why America Will Remain the World's Sole Superpower*, (New York: Cornell University Press, 2018), pp. 120-134.)
78. See, for example: Joel Mokyr, *A Culture of Growth: The Origins of the Modern Economy*, (Princeton: Princeton University Press, 2017). Paul Kennedy, *Rise and Fall of the Great Powers* (New York: Random House, 1989).
79. Samuel Huntington, "The US-decline or renewal." *Foreign Affairs* 67 (1988): 76.
80. A modern reflection on these strengths is presented in Richard Danzig et al., "A Preface to Strategy: The Foundations of American National Security," *JHU Applied Physics Laboratory*, 2018.
81. E.T. Jaynes was in the audience and noted the exchange. E.T. Jaynes, *Probability Theory: The Logic of Science*, (St. Louis, MO: Washington University, 1996), p. 4.

82. Defense Innovation Board, "AI Principles: Recommendations on the Ethical Use of Artificial Intelligence by the Department of Defense," https://media.defense.gov/2019/Oct/31/2002204458/-1/-1/0/DIB_AI_PRINCIPLES_PRIMARY_DOCUMENT.PDF.
83. For example, Nils Nilsson, a pioneer of AI research, writes: "artificial intelligence is that activity devoted to making machines intelligent, and intelligence is that quality that enables an entity to function appropriately and with foresight in its environment. According to that definition, lots of things — humans, animals, and some machines — are intelligent. Machines, such as 'smart cameras,' and many animals are at the primitive end of the extended continuum along which entities with various degrees of intelligence are arrayed. At the other end are humans, who are able to reason, achieve goals, understand and generate language, perceive and respond to sensory inputs, prove mathematical theorems, play challenging games, synthesize and summarize information, create art and music, and even write histories." See: Nils Nilsson, *The Quest for Artificial Intelligence* (New York: Cambridge University Press, 2010).
84. This was most recently reiterated by the Defense Science Board. See: Defense Science Board, "Summer Study on Autonomy," *U.S. Department of Defense*, June 2016.
85. The first AI winter was actually initiated by a government report in the UK, commonly referred to as the *Lighthill Report*.
86. Ben Buchanan, "The AI Triad and What It Means for National Security," *CSET*, August 2020: https://cset.georgetown.edu/wp-content/uploads/CSET-AI-Triad-Report.pdf. Further, many primers on modern AI and ML systems now exist and are not reproduced here. See, for example: Mark Halpern, "No Ghost in the Machine," *The American Scholar*, Spring 2020, https://theamericanscholar.org/no-ghost-in-the-machine/#.Xnq96G4pCu6; John Launchbury, "A DARPA Perspective on Artificial Intelligence," *Defense Advanced Research Projects Agency (DARPA)*, March 2017, https://www.darpa.mil/about-us/darpa-perspective-on-ai; Ben Buchanan and Taylor Miller, "Machine Learning for Policymakers: What it is and why it matters," *Harvard Belfer Center for Science and International Affairs*, June 2017, https://www.belfercenter.org/sites/default/files/files/publication/MachineLearningforPolicymakers.pdf; Andrej Karpathy, "AlphaGo, in context," https://medium.com/@karpathy/alphago-in-context-c47718cb95a5; Michael Jordan, "Artificial Intelligence—The Revolution Hasn't Happened Yet," *Harvard Data Science Review*, https://hdsr.mitpress.mit.edu/pub/wot7mkc1.

CSET CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

CSET.GEORGETOWN.EDU | CSET@GEORGETOWN.EDU