2017 Award Winner
USAWC Student Awards Program

Strategic Robotpower: Artificial Intelligence and National Security

by

Lieutenant Colonel Patrick Sullivan
United States Army

Under the Direction of:
Dr. Andrew Hill

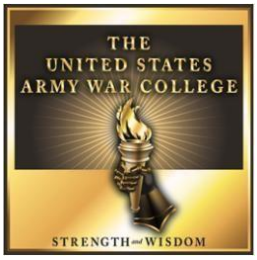

United States Army War College
Class of 2017

REPORT DOCUMENTATION PAGE

Form Approved--OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
01-04-2017	STRATEGY RESEARCH PROJECT	

4. TITLE AND SUBTITLE
Strategic Robotpower: Artificial Intelligence and National Security

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Lieutenant Colonel Patrick Sullivan
United States Army

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Dr. Andrew Hill

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army War College, 122 Forbes Avenue, Carlisle, PA 17013

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT Distribution A: Approved for Public Release. Distribution is Unlimited.
To the best of my knowledge this SRP accurately depicts USG and/or DoD policy & contains no classified information or aggregation of information that poses an operations security risk. **Author:** ☒ **PA:** ☒

13. SUPPLEMENTARY NOTES
Word Count: 6095

14. ABSTRACT
Artificial intelligence (AI) is on a developmental trajectory to be the central force in future society. There will be many benefits from this, but serious challenges as well, particularly for the national security setting. AI's development virtually guarantees that lethal autonomous weapons will someday be unleashed on the battlefield. Although these weapons could conceivably lower the human cost of war, they carry significant proliferation and collateral damage risk as well, and could make the decision to go to war easier. This would be inherently destabilizing to the Westphalian geopolitical order, which is already under strain due to AI's democratization of information. As the dominant artificial intelligence company, Google is best positioned to benefit from any decentralization and rebalancing of state power that occurs from AI-related disruption, with Silicon Valley as a whole becoming a political entity unto itself. Whatever the resultant decentralized/rebalanced power construct looks like, all stakeholders – transnational technology companies, nation-states, and what remains of the international system – will have a responsibility to provide collective good governance to ensure that AI's outcomes stay as positive as possible, especially as relates to privacy concerns, job displacement, and increased socio-economic inequality.

15. SUBJECT TERMS
Lethal Autonomous Weapons, Proliferation, Decline of Westphalian System, Rise of the Google State

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			**19b. TELEPHONE NUMBER** *(w/ area code)*
UU	UU	UU	UU	34	

Strategic Robotpower: Artificial Intelligence and National Security

(6095 words)

Abstract

Artificial intelligence (AI) is on a developmental trajectory to be the central force in future society. There will be many benefits from this, but serious challenges as well, particularly for the national security setting. AI's development virtually guarantees that lethal autonomous weapons will someday be unleashed on the battlefield. Although these weapons could conceivably lower the human cost of war, they carry significant proliferation and collateral damage risk as well, and could make the decision to go to war easier. This would be inherently destabilizing to the Westphalian geopolitical order, which is already under strain due to AI's democratization of information. As the dominant artificial intelligence company, Google is best positioned to benefit from any decentralization and rebalancing of state power that occurs from AI-related disruption, with Silicon Valley as a whole becoming a political entity unto itself. Whatever the resultant decentralized/rebalanced power construct looks like, all stakeholders – transnational technology companies, nation-states, and what remains of the international system – will have a responsibility to provide collective good governance to ensure that AI's outcomes stay as positive as possible, especially as relates to privacy concerns, job displacement, and increased socio-economic inequality.

Strategic Robotpower: Artificial Intelligence and National Security

Artificial intelligence (AI) promises to fundamentally change the way Americans live, work, and interact. In general, as technology grows more useful, demand grows in turn. This principle will perhaps never be more true than with AI, given the rapid and dramatic progress it has ushered in disparate areas such as medicine, entertainment, finance, and defense.[1] Although AI and related technologies carry the promise of great and broad societal benefit, they also carry significant challenges for national security and policymaking, given the close historical correlation between security and technological innovation.[2] The strategic impact of new technologies – usually measured in terms of disruption or destabilization – is difficult to predict.[3] Predicting the strategic impact of artificial intelligence is not just an issue of scale, however, but also of precedent: the very nature of AI demands that humanity consider its relationship with technology in entirely new ways.[4] Regardless of the difficulties and unknowns, this paper seeks to inform a strategic prediction by examining the implications of AI for the national security setting.

Our methodology will be to analyze artificial intelligence within the four main themes of the 2015 *National Security Strategy of the United States* (NSS) – defense of the United States and its allies, international order, values, and prosperity.[5] Although the Trump Administration will eventually publish its own NSS that outlines a new strategic vision for the United States, the four themes listed have remained fairly consistent since the first NSS in 1987, and should thus be considered a valid analytical framework.[6]

There are two core conclusions about artificial intelligence and national security that become apparent from this analysis. First, it is a near-certainty that future battlefields will prominently feature lethal autonomous weapons. Not only does this risk

proliferation and catastrophic collateral damage, but it may also increase the frequency and intensity of armed conflict. Second, the manner in which AI technologies are developing and currently being used will likely cause state power to be rebalanced, with transnational technology companies such as Google best positioned to most benefit. Within this emergent paradigm, social inequality and related destabilization could conceivably increase such that existing institutions and policies would be exceedingly hard-pressed to mitigate the negative effects. In order for the resultant geopolitical, economic, and social orders to persist, a more agile governance structure would need to arise as a replacement. Placing these findings and their supporting analyses in proper context, however, requires a lexis for artificial intelligence as a scientific discipline, as well as a short history of its development.

Artificial Intelligence in Perspective

The term “artificial intelligence” was coined in 1956 by a summer research project at Dartmouth College, with foundations in Alan Turing’s 1950 paper, “Computing Machinery and Intelligence,” which proposed the famous Turing Test (AKA “The Imitation Game”) to determine whether a machine could think.[7] Although there is no universally accepted contemporary definition for artificial intelligence, one of the discipline’s founders, American computer scientist Dr. Nils J. Nilsson, provides a useful one: “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.”[8]

This definition correlates to modern conceptions of human intelligence, which is generally understood to be efficient problem-solving ability that leverages past experience in a heuristic manner.[9] Comparing the definition of AI with this concept

suggests that no machine is close to human-level intelligence, yet AI-enabled machines have outperformed master-level human competitors in chess, the TV game show *Jeopardy!*, and the ancient Chinese strategy board game "Go," all of which are abstract and highly complex.[10] To account for this apparent paradox, it is necessary to distinguish between narrow (or weak) AI and general (or strong) AI. The supercomputer victors of the games listed – IBM's Deep Blue, IBM's Watson, and Google DeepMind's AlphaGo – are each examples of narrow AI, which is specially designed to handle a very limited number of problem domains (usually one) through various forms of pattern recognition. A general AI on the other hand is not limited, and can ostensibly perform any intellectual task as well (but not necessarily in the same way) as a human.[11] In light of this distinction, all current AI applications such as targeted e-commerce and commercial services, natural language processing, image recognition and labeling, personal digital assistants (e.g., Apple's Siri or Amazon's Alexa), augmented medical diagnostic tools, and self-driving cars are narrow. However, these applications also represent an important bridge to a potential future general AI, as they have each resulted from a new wave of AI-related research that centers on the concept of machine learning.

In layperson's terms, machine learning is the ability of computers to learn from data, as opposed to being explicitly programmed.[12] It does not represent a specific algorithm for singular problem-solving, but rather a general approach to solve many different problems.[13] Machine learning stems from cloud computing resources and Internet-based data gathering, both of which are recent innovations and whose commercial potential has radically exceeded initial expectations.[14] Coupled with new

computing models such as artificial neural networks that replicate how the human brain functions, extreme data processing speeds enabled by advances in quantum mechanics, and recently achieved evolutionary benchmarks such as machine vision and basic linguistic prediction, machine learning portends a near-term paradigm in which machines are capable of operating and adapting to changing real-world circumstances without human control.[15] At a minimum, such machine autonomy would serve as a necessary condition for general artificial intelligence, which has been a core objective of AI research since the beginning.[16]

An additional concept commonly associated with general artificial intelligence is the technological singularity, whereby a sufficiently "smart" and self-aware machine could theoretically modify its own code recursively, producing new and more intelligent versions of itself in a chain reaction until a superintelligence emerged.[17] This superintelligence, by its very nature and method of creation, would far surpass the ability of humans to understand and control, thus potentially posing an existential threat. Although the technological singularity has not occurred, and a malevolent-seeming superintelligence remains science fiction, the risk associated with their development is increasingly the subject of serious scientific and philosophical inquiry.[18] These inquiries are beyond the scope of this paper, however. In order to remain aligned with the analysis that follows, the reader should assume that any negative implications of general artificial intelligence are manageable within the human condition, and fall short of the existential threat threshold.

Despite the association with a potential superintelligence and the resultant "robot apocalypse," artificial intelligence is not about building a mind *per se*, rather providing a

gateway to expanded human potential through improved problem-solving tools.[19] In this context, the evolution from narrow to general AI is really just a shift from systems that have additive intelligent capabilities to ones that are sufficiently intelligent in total so that humans can better integrate with them.[20] Also, it bears remembering that artificial intelligence is still a nascent discipline and frontier science in comparison to technological history as a whole.[21] Thus, while the various fears of AI's unknowns are certainly valid, as related technologies become commonplace, they will no longer be seen as AI and even newer and more enigmatic technology will emerge.[22]

Through this continuous process of improvement and demonstrated utility, it is reasonable to expect that artificial intelligence will follow its current trajectory to become the central force in society with impacts that are mostly positive.[23] This is not guaranteed, however; the direction and effects of previous major technological shifts have not been consistently positive, which is an indictment of the technology itself as well as the supporting economic and policy landscape for a particular society.[24] This reality is disconcerting, as is the unprecedented pace of underlying change being driven by AI. The extraordinary complexity of artificial intelligence systems, which can make their performance unpredictable, and their development outside of the government's ability to regulate are factors that aggravate the growing concern about AI's potentially disruptive and destabilizing effects. As the remainder of this paper will demonstrate, these effects are inherent to national security and implicitly demand strong policy and strategy responses.

Killer Robots – AI and Defense

The history of military applications of artificial intelligence tacks closely to the history of the discipline itself. The Defense Advanced Research Projects Agency

(DARPA) has conducted its own AI research for the better part of four decades, in addition to providing grant monies to academia and private industry through the Strategic Computing Initiative.[25] This research spawned AI-based navigation and sensing for explosive ordnance disposal robots and other unmanned ground systems, which built upon earlier efforts to use ground robotics for remote control of vehicles and mounted weapons.[26] RAND Corporation's Rule-Oriented System for Implementing Expertise (ROSIE) was an early AI-based targeting tool that served as a precursor to AI systems now commonly used for target identification, discrimination, and recommendation in remotely-piloted aircraft (RPA), C-RAM (counter-rocket, artillery, and mortar) platforms, and the Aegis missile defense system.[27] Taken together with the increased importance of artificial intelligence in computerized war-gaming and related decision aids, these applications show that military AI is moving away from exclusively "dull, dirty, and dangerous" tasks to ones that demand greater autonomy and complexity.[28]

The appeal of AI-enabled autonomous weapons is practically self-evident. Autonomous weapons could reduce "boots on the ground" requirements, enable greater precision in targeting, and increase speed in decision-making, thereby lowering the human cost of war.[29] Additionally, autonomous weapons will likely be much cheaper to produce over time, since removing the operator allows miniaturization and simplifies systems integration.[30] This could help break the vicious cycle of defense acquisition, in which the U.S. military seems to get less capability at greater expense generation-over-generation.[31] Moreover, Great Power competitors such as China and Russia are rapidly closing the technological and doctrinal gaps that ensure the United States' current

advantage in precision strike and power projection, which makes reinvestment in these technologies unaffordable and illogical in the long run.[32] The effectiveness of these technologies in a future operating environment increasingly characterized by hybrid warfare is debatable regardless, whereas artificial intelligence is, in the words of Deputy Secretary of Defense Robert O. Work, "The one thing that has the widest application to the widest number of Department of Defense (DOD) missions" moving forward.[33] Russia's Chief of General Staff and one of the intellectual fathers of hybrid warfare, General Valery Gerasimov, endorses this view, predicting a future battlefield dominated by learning machines.[34] Additionally, numerous AI researchers have forecasted autonomous weapons as the third military-technical revolution, after gunpowder and nuclear weapons.[35]

Although autonomous weapons tend to get associated with the more fantastical elements of science fiction, such as the "killer robots" of the Terminator franchise, it is important to note that mobile general artificial intelligence (a more precise name for a robot) is not necessary to develop an autonomous weapon; much of the foundational technology exists today and there are several precursor systems already in use.[36] These include the Samsung SGR-A1 Sentry Gun in the Demilitarized Zone on the Korean Peninsula, and the Israeli Defense Force's SentryTech system, both of which have settings that allow for lethal engagement without human intervention.[37] Even the C-RAM systems used by the U.S. military are *de facto* autonomous with lethal potential, at least collaterally; although a human operator supervises the system, he or she would have to react near-instantaneously to override an engagement.[38]

Although no nations have fully committed to the development of AI-enabled autonomous weapons, none have disavowed them either.[39] For the United States, the current trend is towards development. The U.S. Army Strategy for Robotics and Autonomous Systems envisions autonomous systems incorporated into combined arms maneuver by 2035, with dynamic force and mission autonomy to follow in the 2040s.[40] DARPA initiated a program in 2013 to integrate machine learning in a wide variety of weapon systems, and the Office of Naval Research is funding several studies in support of its broader vision to "develop autonomous control that intelligently understands and reasons about its environment … and independently takes appropriate action."[41] From a U.S. policy perspective, DOD Directive 3000.09, *Autonomy in Weapon Systems*, calls for the identification of operational needs that can be satisfied by autonomous weapon systems while neither encouraging nor prohibiting these systems having lethal capabilities.[42] Exploiting the space created by this opacity, the U.S. Army's last research and development budget submission outright describes lethal ground autonomous weapons.[43]

The DOD Directive 3000.09 is a bit clearer on the control of autonomous weapon systems, saying they must employ "appropriate levels of human judgement."[44] In the near term at least, this is understood to mean that a human operator will always be in-the-loop (the human controls the weapon) or on-the-loop (the human supervises the weapon and can override it if necessary).[45] This includes manned-unmanned teaming concepts favored by "centaur warfighting" and the DOD's "Third Offset" acquisition strategy, which relies heavily on artificial intelligence technologies.[46] Additionally, as autonomous weapons development gains momentum, concomitant concerns over

independence and lethality are assuaged somewhat by the belief that the systems will be defensive in nature, similar to the quasi-autonomous sentries already in use. This is false comfort, however. As Paul Scharre, Director of the Future of Warfare Initiative at the Center of a New American Security, has observed, “if there was an easy way to delineate between offensive and defensive weapons, nations would have agreed long ago to only build ‘defensive’ weapons.”[47]

Commitments to keeping humans in or on-the-loop are not sufficient to allay concerns about ceding control of lethal decisions to a machine, or to avoid creating potentially uncontrolled killer robots. On the contrary, there are numerous incentives that make the development of lethal autonomous weapons borderline inevitable as long as artificial intelligence continues to deliver on its technological promise. Although AI-enabled machines are not yet as “smart” as humans, they are far superior at solving multiple control problems very quickly, given their ability to process massive amounts of information to detect patterns without suffering fatigue, recognition error, bias, or emotional interference.[48] At the tactical level, military operations are basically just a series of control problems, and decision-making in competitive environments tends to accelerate; this is what Sun Tzu was getting at when he described speed as “the essence of war.”[49] Thus, a human operator in or on-the-loop detracts from the very advantage that autonomous weapons and other military AI applications provide.[50] Add to this the fragility of communication links in a hybrid operating environment with cyber and anti-access elements, and the operational imperative to delegate actions – including lethal ones – directly to machines becomes clear.[51]

There is also the “first mover” principle at play. Militaries have an intrinsic motivation to develop superior capabilities than their adversaries, and the first competitor to maximize AI’s potential to fundamentally change the character of future warfare would enjoy a significant tactical and operational advantage.[52] The United States is already at the forefront for AI technologies, has a defense industrial base that leads the world in complex systems engineering and integration, and has tremendous practical experience with RPAs and unmanned ground vehicles to scale up from.[53] Accordingly, the U.S. is in the best position to be the first mover for lethal autonomous weapons and to gain the attendant tactical and operational advantages.

Regardless of how the United States sees its first mover advantage, the development of lethal autonomous weapons is highly probable due to factors beyond U.S. control. Most artificial intelligence and machine learning research is openly occurring in the private sector and academia, untethered to military contracts and generally without an eye towards military applications (no matter how obvious these applications may be). Even technologies developed at DARPA typically do not remain classified.[54] Moreover, a formal ban under the auspices of the Convention on Certain Conventional Weapons or similar agreement is unlikely, since restrictions of this form usually have disproportionate impact on states that most rely on the technologies related to the weapon, i.e., the United States and China in this case.[55] Thus, Great Power competitors will have access to the foundational AI technologies for lethal autonomous weapons, and given the incentives already described, will likely seek to develop new and dangerous concepts of operation that leverage them.[56] Accordingly, using its first mover advantage to define the probable lethal autonomous weapons

paradigm is arguably the most responsible and stabilizing choice the United States can make for the evolving national security setting, as opposed to having the paradigm thrust upon it.

Once this paradigm gets defined, irrespective of origin, there will be a military imperative to adapt to it. This includes proactively dealing with the foreseeable consequences, of which proliferation and system unpredictability are the most alarming. Regarding proliferation, it almost certain that rogue states and/or violent extremist organizations would either design their own non-discriminatory lethal autonomous weapons, or remove safeguards from a system already developed by a more responsible actor.[57] History is replete with unsuccessful attempts to control technology once it is loosed; gunpowder and submarines are both prominent examples.[58] There are many artificial intelligence researchers who fear the technological trajectory of autonomous weapons is such that they "will become the Kalashnikovs of tomorrow," with particular utility in assassinations, ethnic cleansing, destabilizing governments, and population control.[59] Absent an unlikely ban on lethal autonomous weapons, or even more unlikely suppression of the foundational AI technologies, proliferation risk simply becomes system tension that has to be managed.[60]

Unpredictable performance of lethal autonomous weapons can create accidental and collateral damage risk that would also have to be managed. It is human nature to anthropomorphize machines, but artificial intelligence and cascading technologies are functionally different than any form of human cognition, and will thus act in ways not anticipated by their developers.[61] In fact, this is part and parcel with the whole concept of autonomy.[62] For machine learning, an AI system is trained on inputs and outputs,

often unsupervised, until *voilà!* – it just works. In this fashion, machine learning is akin to a "black box," with developers sacrificing understanding of AI system behavior and shortchanging control in favor of performance.[63] Artificial neural networks in particular can sometimes yield odd and unpredictable results, and if an artificial intelligence is based on quantum computer modeling, then the AI itself will embody quantum indeterminacy.[64] Additionally, an autonomous AI in a physical environment is subject to the "open world" conundrum, in which the system is bound to encounter conditions that were not anticipated when it was designed and built.[65]

Thus, lethal autonomous weapons will inevitably produce errors, and not necessarily ones a human operator would produce if they were in or on-the-loop.[66] These errors will additionally be difficult to correct or prevent from reoccurring; not only could the sheer complexity of the weapon system prevent an error's cause from being auditable, it is difficult to take corrective action if you don't understand how the weapon system is behaving and why.[67] Also, automation bias – through which humans demonstrate uncritical trust in automation and its outputs – could create denial that an error had even occurred.[68]

The negative outcomes of the accidental and collateral damage risk accrued with a lethal autonomous weapon's error production are fratricide and civilian casualties. Although AI-related error can reasonably be expected to occur much less frequently than human error, lethal autonomous weapons have higher damage potential over possibly orders of magnitude more social-technological interactions, some of which will have not been anticipated by the system's designers.[69] Moreover, one has to take any artificial intelligence on interface value, so a lethal autonomous weapon's error would

likely repeat with a consistent level of force until some external agent intervened.[70] Human error, on the other hand, tends to be idiosyncratic and one-off given a human operator's (presumed) common sense, moral agency, and capacity for near-real time consequence management.[71]

Unjust War – AI and International Order

In addition to altering the tactical and operational environments, the lethal autonomous weapons paradigm and other military artificial intelligence applications pose strategic risk. Not only will artificial intelligence potentially change the criteria for war and how it is conducted, but it might also fundamentally disrupt the geopolitical landscape in which war is waged.

Although AI can potentially reduce the human cost of war within individual conflicts, this reduction could make conflicts themselves occur more frequently and with greater intensity.[72] The human cost of war is an important consideration in the decision to use violent force; in fact, the principle of *jus ad bellum* (literally "justice of war," i.e., the conditions under which a state can rightly or justly resort to the use of force) is often predicated on this calculation.[73] If there is less potential for casualties, however, then governments could potentially operate with less restraint in using the military instrument to secure national interests.[74] The way in which the United States brazenly challenges other nations' sovereignty through the use of armed RPAs is an indication of this lessening of restraint. Also, the 2015 National Security Strategy commits to "avoiding costly large-scale ground wars," but underwrites the use of proxies and asymmetric methods to combat violent extremist organizations, therein lowering the threshold for war.[75] This latter point and recent history suggest that the United States and its partners have a troubling predilection for the "Jupiter Complex" – using force to exact righteous

retribution against perceived evil adversaries. Artificial intelligence could potentially create a positive feedback loop for this behavior, thereby encouraging it .[76]

The human cost of war would likely normalize over time with the proliferation of lethal autonomous weapons, but if certain populations had developed a sense of invulnerability in the interim, then they would be less prepared for the burdens of war moving forward.[77] Accordingly, policy-makers would need to be careful that military AI applications do not become a detriment to the population's sense of national identity and collective purpose. Policy-makers should be further cautioned that, despite the appeal and apparent low cost of AI-enabled warfare, any military engagement is inherently destabilizing.[78] For the lethal autonomous weapons paradigm, at least at the outset, the destabilizing effects center on a potential availability gap for foundational AI technologies. An availability gap could reinforce and exacerbate global inequalities, as well as incentivize a "first strike" or new forms of extremism to close the gap.[79] In response, the AI-advantaged nation's basic assumptions of deterrence and compellence would be severely challenged, thus changing the tenets of their defense strategy.[80]

Lethal autonomous weapons in particular also represent a responsibility gap, which goes beyond the accidental and collateral risk concerns associated with errors and system failures.[81] It is doubtful that an artificial intelligence with lethal capacity could uphold the two central elements of *jus in bello* ("justice in war," or the law that governs how force is to be used), namely discrimination and proportionality. Regarding discrimination, it is difficult to restrict the use of lethal autonomous weapons to a self-regulated set of narrowly constructed scenarios since, again, the system's developers

would not be able to anticipate every interaction the system will encounter.[82] The system's inherent "framing problem," in which it would inevitably have incomplete understanding of its external environment, just adds to its intractability.[83] The developer could attempt to install a sense of compassion, empathy, and mercy in the system, but installed ethics are liable to become obsolete due to similar issues of framing and unanticipated interactions.[84] Thus, since the lethal autonomous weapon will likely not possess the equivalence of moral agency or social norming – to reiterate, there is nothing inherently human about a machine – it will ruthlessly and relentlessly carry out its assignments, with none of Clausewitz's fog or friction to temper its actions.[85] The lack of moral agency also precludes a lethal autonomous weapon from being held accountable, since it cannot fear punishment or learn from it.[86]

A machine's capacity for proportionality fails along a similar line of argument. Although a lethal autonomous weapon's lack of emotion checks the passions of war, which can lead to atrocities and revenge killings, it also constrains prediction of an adversary's emotions and actions.[87] Proportionality requires judgement of this form and depends greatly on context which, given its framing problem, a lethal autonomous weapon is also unlikely to be able to process.[88]

Artificial intelligence and lethal autonomous weapons could also complicate efforts to stabilize crises.[89] Crisis settings often demand quick decisions with incomplete information; if an ill-considered decision creates an unanticipated lethal autonomous weapon interaction with an error or failure outcome, then an unintended "flash war" could result.[90] Although this type of conflict could start quickly, there is no guarantee that it would end quickly. On the contrary, given the rapid and high volume interactions of

complex autonomous systems, the flash war would conceivably spiral out of human control and be difficult to stop.[91] The ensuing chaos would support the interests of rogue states, violent extremist organizations, and practitioners of hybrid warfare. Accordingly, these groups could be expected to create conditions for flash wars through their own use of lethal autonomous weapons, or by hacking fail-safes in someone else's weapon if there was a need to avoid attribution.[92]

The flash war scenario suggests a forthcoming geopolitical environment in which artificial intelligence allows military power to be decoupled from traditional indices such as population size/growth and gross domestic product.[93] This decoupling could fundamentally change the character of alliances and security cooperation agreements, since apparently weak states would no longer need the protection of ostensibly strong ones. Alliances would also potentially no longer be influenced by forward basing and access considerations, since AI-enabled additive manufacturing, small high-density power generation, and miniaturization will likely change power projection modalities.[94] Adam Elkus, a Cybersecurity Fellow at the public policy think tank New America, describes this course as, "[artificial] intelligence creating a new form of meta-geopolitics that will reshape notions of national power."[95] Within this new form meta-geopolitics, proliferation of lethal autonomous weapons could democratize violence, providing individuals and groups with state-level instruments of military power.[96] Different AI technologies will similarly democratize information, as they are already, thus also providing these smaller entities with state-level instruments of softer types of power.

Artificial intelligence is promoting technological growth and diffusion of knowledge at unprecedented rates, and its transformative effects on society seem to be

speeding up.[97] With this diffusion of knowledge, human and strategic interests are merging, sometimes forcefully so.[98] Evidence abounds for this rise in human agency, from the hacktivist collective Anonymous' use of AI programs to support the Arab Spring, to AI-enabled social media fueling demands for social justice across the world, to social impact investing and entrepreneurship spawning what Nicholas Kristof of the *New York Times* has called "DIY foreign aid."[99] Previously dormant socio-economic, cultural, and ethnic fault lines are fracturing, and state power is proving increasingly inadequate to tamp down the conflicts that get unleashed, assuming that the state is even paying attention to the forces of disruption.

The inability of state power to manage AI-related disruption has invited questions about the future suitability of the Westphalian system of nation-state sovereignty. Although predictions about the return of the city-state or the rise of a "new medievalism" are probably oversold, it is clear at a minimum that human agency is forcing the redistribution and decentralization of power to non-state actors.[100] Given the deep interconnection of commerce and geopolitics that already exists with globalization, foremost among these non-state actors are transnational corporations, some of which are as powerful as nation-states yet beholden to none.[101] And given the outsized influenced of artificial intelligence on the approaching social and economic orders, the most consequential transnational corporations and, by extension, the most powerful non-state actors will be technology companies.

Rise of the Google State – AI, Prosperity, and Values

The proverbial alpha in the technology company pecking order will likely be Google. Google is all-in with machine learning, employing the largest number of PhDs in the field – outstripping even academia – and their "AI First" growth strategy envisions

the widening industrial applications of self-programming computers.[102] These applications could potentially transform the basis of economic growth for countries throughout the world, and Google is well-positioned to dominate this new economic order in ways that go far beyond its talent advantage.

"AI First" represents what some observers have described as "institution building and consolidation of power on a scale and at a pace unprecedented in human history."[103] Google has unmatched reserves of data in a massive cloud and supercomputing architecture that spans 13 countries distributed over four continents, which provides them a distinct strategic advantage over technology company competitors.[104] This is because the company that possesses the data controls the algorithms for continued machine learning research and physical system integration. Eric Schmidt, CEO of Alphabet, Google's parent company, has predicted that big data will be the "new oil" – a commodity so consequential in the global economic order that nations will fight over it.[105] And given Google's existing computing architecture, their data reserves are likely to increase outside of their core business model, since smaller companies will be more inclined to tap into Google's cloud resources as opposed to investing the billions of dollars to build their own secure architecture.[106]

With Google's incumbent control of big data, it will be able to use artificial intelligence to refine its products better than any potential competitor or market entrant, thereby ensuring its dominance until the market fundamentally changes or some successor paradigm forces a new global economic order.[107] Such a change is unlikely in the foreseeable future. On the contrary, it is more likely that Google and its transnational technology company brethren in Silicon Valley dominate much more than the economic

order. On September 12, 2016, for the first time in history, the five largest public corporations by market capitalization were technology companies, each of which are heavily invested in AI.[108] Greater awareness of their collective power will increasingly force Silicon Valley companies to act in their own self-interest (as any transnational company should be expected to), and the resources and influence that these companies command will ensure that their actions are politically significant, regardless of intent. In this manner, technology will continue to become a sort of political entity unto itself, with technology companies not realistically having the option to stay neutral in the public policy space, with either their products or their stated positions.[109] Within this vein, it bears recognizing that President Trump was elected over the near-universal objection of Silicon Valley, and that several prominent Silicon Valley executives and investors have led the call for California's secession from the United States.[110]

If AI-enabled social media platforms and search engines do not remain politically neutral, then the democratic process can suffer greatly as a result, therein comprising yet another disruption vector for AI into the national security setting. Predictive algorithms could greatly improve the reach and effectiveness of robo-calls, social media bots, and gerrymandering of voting districts, which would suppress democratic participation and increase fractiousness in the current highly polarized political environment.[111] Artificial intelligence already acts as an "invisible authority" on the Internet that only reflects back its image of consumers, and the pervasiveness of AI applications to help organize a highly complex world make society amenable to relinquishing control.[112] Hence consumers' willingness to give up so much of their privacy to Google and others (to include governmental entities, tacitly or transitively) to

enable passive collection of the data that is necessary for machine learning AI applications to work and further develop.[113] Thus, the stage is set for technology companies to serve as "Big Brother," controlling ideas to potentially nefarious ends, and along the way reducing the sociability and consciousness that make us all "human."[114]

As it stands, AI-enabled social media is not used as tool to bring people together but rather to sow division, particularly among groups who are already feeling left behind by the emerging global economic order. Prominent among these groups are workers most susceptible to job displacement that is being caused by AI-driven automation, and for whom Silicon Valley has become an agent of destruction.[115]

The have been more jobs lost to technology in the United States over the last decade than any other sector, and this development is likely to accelerate as automation costs come down and AI technologies gain even greater primacy in all facets of the economy.[116] Due to the nature of automation *vis-à-vis* current narrow AI capabilities, the jobs lost and threatened are highly concentrated among lower skilled and less educated workers. According to a study completed by former President Obama's National Science and Technology Council, between nine and 47 percent of all American jobs are at risk of displacement over the next two decades, with 83 percent of these jobs being concentrated in the lower middle class.[117] This continues a trend since the latter half of the 20th Century of increasing system bias towards skilled labor, as well as production's increasing reliance on capital at the expense of labor as a whole.[118] For workers who have already been displaced, their demonstrated difficulty in matching extant skills to the AI economy's in-demand jobs indicates the long-term disruptive

potential of job displacement, with certain parts of the workforce moving into permanent unemployment and (presumably) poverty.[119]

These in-demand jobs are at the very high-end or very low-end of the pay scale. The resultant "hollowing out" of the middle class will increasingly bifurcate the job market between the two extremes of low-skill/low-pay and high-skill/high-pay, thereby increasing social tensions and promoting a belief that capitalism's "winner take all" ethos is not working for most.[120] Since many people derive a significant amount of meaning, identity, and self-worth from their employment, the potential for malaise and its negative effects--crime, social dereliction, etc.–-become clear.[121]

The inequalities created by the AI economy are potentially far greater and more disruptive than those related to mere job displacement. There is additional segregation that could occur exclusively in the high end of the skill and pay scales. In an economy where machines are doing most of the work, virtually all of a company's returns would go to investment as opposed to labor.[122] This in turn would create a premium for intellectual capital – those select few from the uppermost stratum who can direct the ever increasing resources in the most profitable and visionary ways. In short, the same Silicon Valley titans who are currently shaping and setting the AI economy. Additionally, it is much easier to forecast the jobs that will be lost to or threatened by artificial intelligence than it is to predict what jobs will be in demand and how much production will actually rely on them.[123] Thus, the AI economy could actually represent "superstar biased" technological change, in which labor becomes virtually non-existent (or *de facto* irrelevant from a policy perspective), thereby fundamentally changing the nature of production and work.[124] Accordingly, economies would have to be reorganized to enable

a new form of resource allocation other than compensation for labor.[125] And as Silicon Valley is best positioned to benefit from the AI economy, its political weight will likely prove decisive in dictating what that reorganization will be.

Conclusion – Governance is Destiny

Social upheaval, political turmoil, privacy concerns, and good old class warfare are each reasons for many prognosticators to suggest that societies have nothing to fear from killer robots or the like, as people will rise up well before machines do.[126] No matter when and in what context AI-induced disruption occurs, however, it is disruption just the same and thus inherently relevant to national security. Since disruption is already occurring despite artificial intelligence technologies still being relatively immature, disruption will likely get worse before it gets better. In turn, the national security setting will likely be much less stable in the interim.

And yet, artificial intelligence is here to stay, and will continue to gain influence and utility. It will one day become the central force in society, for good and for ill. Although the national security implications of AI are consequential and potentially severe, one can nonetheless take it as an article of faith that the benefits of AI will outweigh its costs. The recognition that technology is not destiny helps with this; throughout history, various economies have had access to similar levels of technology, but saw qualitatively different outcomes because of different policies and institutions.[127] Indeed, the critical factor that will shape the AI future is the same factor that shapes the AI present: governance. As this paper has shown, what constitutes governance will change with the decentralization and rebalancing of state power. Regardless of where the power lies--within nation-states, within Silicon Valley, and/or within what remains of the international system--there will be a set of strategic leaders to wield it. May they

demonstrate the intelligence, artificial and human, to do so responsibly and with an eye towards a best possible future for all.

Endnotes

[1] Raffi Khatchadourian, "The Doomsday Invention," *The New Yorker Online*, November 23, 2015, http://www.newyorker.com/magazine/2015/11/23/doomsday-invention-artificial-intelligence-nick-bostrom (accessed February 25, 2017).

[2] Klaus Schwab, "The Fourth Industrial Revolution: What it Means and How to Respond," *World Economic Forum: Global Agenda Online*, January 14, 2016, https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/ (accessed February 25, 2017).

[3] Adam Elkus, "The AI Wars?" *Slate: Future Tense – Citizen's Guide to the Future*, blog entry posted January 20, 2015, http://www.slate.com/articles/technology/future_tense/2015/01/what_artificial_intelligence_does_and_does_not_mean_for_security_and_geopolitics.html (accessed January 23, 2017)

[4] Nayef Al-Rodhan, "The Security Implications and Existential Crossroads of Artificial Intelligence," *Georgetown Journal of International Affairs Online*, April 2, 2015, http://journal.georgetown.edu/the-security-implications-and-existential-crossroads-of-artificial-intelligence/ (accessed February 25, 2017).

[5] Barack Obama, *National Security Strategy* (Washington, DC: The White House, February 2015), 7-23. Also note that the first theme is actually "security" vice "defense," but I substituted defense to avoid confusion.

[6] All sixteen National Security Strategies (1987-2015) can be accessed at http://nssarchive.us/.

[7] Executive Office of the President, Committee on Technology, *Preparing for the Future of Artificial Intelligence* (Washington, DC: National Science and Technology Council), October 12, 2016, https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf (accessed January 27, 2017).

[8] Standing Committee of the One Hundred Year Study of Artificial Intelligence, *Artificial Intelligence and Life in 2030: Report of the 2015 Study Panel* (Stanford, CA: Stanford University, September 2016), https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf (accessed February 25, 2017). See also Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (Cambridge, UK: Cambridge University Press, 2010).

[9] Paul J. Springer, *Military Robots and Drones* (Santa Barbara, CA: ABC-CLIO, 2013), 33. Also consider "human-like" intelligence, to account for exceptionally "smart" animals such as dolphins, octopuses, and chimpanzees.

[10] Brent D. Sadler, "Fast Followers, Learning Machines, and the Third Offset Strategy," *Joint Force Quarterly Online 83* (October 1, 2016): http://ndupress.ndu.edu/JFQ/Joint-Force-Quarterly-83/Article/969644/fast-followers-learning-machines-and-the-third-offset-strategy/ (accessed February 25, 2017). See also "Google's AlphaGo Beats Go Master Lee Se-Dol," *BBC: Technology News Online*, March 12, 2016, http://www.bbc.com/news/technology-35785875 (accessed March 6, 2017).

[11] Edward M. Lerner, "A Mind of Its Own, Part I: Artificial Intelligence," *Analog Science Fiction & Fact,* September 2016, 38-49.

[12] Chris Griffith, "Killer Robots," *The Australian Online*, November 3, 2016, http://www.theaustralian.com.au/life/personal-technology/killer-robots-are-closer-than-ever/news-story/74c20a6a186d9df3d5993da59c6475ad (accessed February 25, 2017).

[13] Executive Office of the President, *Preparing for the Future of Artificial Intelligence.*

[14] Standing Committee, *Artificial Intelligence and Life in 2030.*

[15] Ibid. See also Randy Eshelman and Douglas Derrick, "Relying on the Kindness of Machines? The Security Threat of Artificial Agents," *Joint Force Quarterly Online* 77 (April 1, 2015): http://ndupress.ndu.edu/Media/News/News-Article-View/Article/581874/jfq-77-relying-on-the-kindness-of-machines-the-security-threat-of-artificial-ag/ (accessed February 25, 2017). See also John Markoff, "Microsoft Spends Big to Build a Computer out of Science Fiction," *New York Times Online*, November 20, 2016, https://www.nytimes.com/2016/11/21/technology/microsoft-spends-big-to-build-quantum-computer.html (accessed November 21, 2016). See also Gideon Lewis-Kraus, "The Great A.I. Awakening," *The New York Times Magazine Online*, December 14, 2016, https://www.nytimes.com/2016/12/14/magazine/the-great-ai-awakening.html (accessed February 25, 2017).

[16] Standing Committee, *Artificial Intelligence and Life in 2030*. See also Executive Office of the President, *Preparing for the Future of Artificial Intelligence.*

[17] Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford, UK: Oxford University Press, 2014).

[18] In addition to Bostrom's work, see also James Barrat, *Our Final Invention: Artificial Intelligence and the End of the Human Era* (New York: St. Martin's Press, 2013); Nikola Danaylov, *Conversations with the Future: 21 Visions for the 21st Century* (Toronto: Singularity Media, Inc., 2016); George Zarkadakis, *In Our Own Image: Savior or Destroyer? The History and Future of Artificial Intelligence* (New York: Pegasus Books, 2016).

[19] Khatchadourian, "The Doomsday Invention." See also Lewis-Kraus, "The Great A.I. Awakening."

[20] Lewis-Kraus, "The Great A.I. Awakening."

[21] Elkus, "The AI Wars?"

[22] Standing Committee, *Artificial Intelligence and Life in 2030.*

[23] Ibid.

[24] Executive Office of the President, *Artificial Intelligence, Automation, and the Economy*, December 20, 2016, https://obamawhitehouse.archives.gov/sites/whitehouse.gov/files/documents/Artificial-Intelligence-Automation-Economy.PDF (accessed January 31, 2017).

[25] Springer, *Military Robots*, 34.

[26] P.W. Singer, *Wired for War: The Robotics Revolution and Conflict in the 21st Century* (London: Penguin Press, 2009), 24-30.

[27] Jonas Stewart, *Strong Artificial Intelligence and National Security,* Final Report (Newport, RI: Naval War College Joint Military Operations Department, May 18, 2015), http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=ADA622591 (accessed February 13, 2017). See also Jeffrey L. Caton, *Autonomous Weapon Systems: A Brief Survey of Developmental, Operational, Legal, and Ethical Issues,* The LeTort Papers (Carlisle Barracks, PA: U.S. Army War College, Strategic Studies Institute), http://www.strategicstudiesinstitute.army.mil/pubs/display.cfm?pubID=1309 (accessed February 25, 2017). See also Human Rights Watch, *Losing Humanity: The Case against Killer Robots,* (Washington, DC: Human Rights Watch, November 19, 2012), https://www.hrw.org/report/2012/11/19/losing-humanity/case-against-killer-robots (accessed February 25, 2017).

[28] Paul Scharre, *Robotics on the Battlefield Part I: Range, Persistence, and Daring*, 20YY Series (Washington, DC: Center for a New American Security, May 2014), https://s3.amazonaws.com/files.cnas.org/documents/CNAS_RoboticsOnTheBattlefield_Scharre.pdf (accessed March 2, 2017). See also United States Army, "U.S. Army Strategy for Robotic and Autonomous Systems (RAS)."

[29] Al-Rodhan, "The Security Implications and Existential Crossroads."

[30] Arthur Herman, "The Pentagon's 'Smart' Revolution," *Commentary Magazine Online*, June 26, 2016, https://www.commentarymagazine.com/articles/pentagons-smart-revolution/ (accessed February 25, 2017).

[31] Ibid.

[32] Robert Martinage, *Toward a New Offset: Exploiting U.S. Long-Term Advantages to Restore U.S. Global Power Projection Capability* (Washington, DC: Center for Strategic and Budgetary Assessments, 2014), http://csbaonline.org/uploads/documents/Offset-Strategy-Web.pdf (accessed February 25, 2017).

[33] Matthew Rosenberg and John Markoff, "At Heart of US Strategy, Weapons That Can Think," *New York Times Online*, October 26, 2016, http://www.nytimes.com/images/2016/10/26/nytfrontpage/scannat.pdf (accessed February 25, 2017).

[34] Sadler, "Fast Followers."

[35] Future of Life Institute, "Autonomous Weapons: An Open Letter from AI & Robotics Researchers," https://futureoflife.org/open-letter-autonomous-weapons/ (accessed February 25, 2017).

[36] Scharre, *Robotics on the Battlefield Part I.*

[37] Brad Allenby, "What is 'Military Artificial Intelligence'?" *Slate: Future Tense – Citizen's Guide to the Future*, blog entry posted December 2, 2016, http://www.slate.com/articles/technology/future_tense/2016/12/the_difficulty_of_defining_military_artificial_intelligence.html (accessed February 25, 2017). See also Human Rights Watch, *Losing Humanity.*

[38] Human Rights Watch, *Losing Humanity.*

[39] Paul Scharre, *Autonomous Weapons and Operational Risk,* Ethical Autonomy Project (Washington, DC: Center for a New American Security: February 2016), https://s3.amazonaws.com/files.cnas.org/documents/CNAS_Autonomous-weapons-operational-risk.pdf (accessed February 25, 2017).

[40] U.S. Department of the Army, "U.S. Army Strategy for Robotic and Autonomous Systems (RAS)."

[41] Al-Rodhan, "The Security Implications and Existential Crossroads." See also Caton, *Autonomous Weapon Systems: A Brief Survey.*

[42] Caton, *Autonomous Weapon Systems: A Brief Survey.* See also Executive Office of the President, *Preparing for the Future of Artificial Intelligence.*

[43] Program Element 0603827A, Soldier Systems-Advanced Development, Project S54, Small Arms Improvement. See Caton, *Autonomous Weapon Systems: A Brief Survey.*

[44] Ashton Carter, *Autonomy in Weapon Systems,* Department of Defense Directive 3000.09 (Washington, DC: U.S. Department of Defense, November 21, 2012), http://www.dtic.mil/whs/directives/corres/pdf/300009p.pdf (accessed March 12, 2017).

[45] Ibid.

[46] Martinage, *Toward a New Offset.* See also Scharre, *Robotics on the Battlefield Part I*. See also Paul Scharre, "Robotics on the Battlefield Part II: The Coming Swarm," *Center for a New American Security: 20YY Series*, October 2014, https://s3.amazonaws.com/files.cnas.org/documents/CNAS_TheComingSwarm_Scharre.pdf (accessed March 2, 2017).

[47] John Markoff, "Pentagon Turns to Silicon Valley for Edge in Artificial Intelligence," *New York Times Online*, May 11, 2016, https://www.nytimes.com/2016/05/12/technology/artificial-intelligence-as-the-pentagons-latest-weapon.html (accessed February 25, 2017).

[48] Kareem Ayoub and Kenneth Payne, "Strategy in the Age of Artificial Intelligence," *Journal of Strategic Studies Online 39*, November 2015, http://www.tandfonline.com/doi/pdf/10.1080/01402390.2015.1088838?needAccess=true (accessed February 25, 2017).

[49] Scharre, *Robotics on the Battlefield Part I.* See also Scharre, "Autonomous Weapons and Operational Risk." See also Sun Tzu, *The Art of War* (Sweden: Chiron Academic Press, 2015).

[50] Stewart, *Strong Artificial Intelligence and National Security.*

[51] Robert O. Work and Shawn Brimley, "20YY: Preparing for War in the Robotic Age," *Center for a New American Security: 20YY Series*, January 2014, https://s3.amazonaws.com/files.cnas.org/documents/CNAS_20YY_WorkBrimley.pdf (accessed February 25, 2017).

[52] Sadler, "Fast Followers." See also Herman, "The Pentagon's 'Smart' Revolution."

[53] Martinage, *Toward a New Offset*.

[54] Rosenberg and Markoff, "At Heart of US Strategy." See also Christopher Mims, "We're Fighting Killer Robots the Wrong Way," *The Wall Street Journal Online*, August 17, 2015, https://www.wsj.com/articles/were-fighting-killer-robots-the-wrong-way-1439784149 (accessed February 25, 2017).

[55] Allenby, "What is 'Military Artificial Intelligence'?"

[56] Mims, "We're Fighting Killer Robots the Wrong Way."

[57] Herman, "The Pentagon's 'Smart' Revolution." See also Griffith, "Killer Robots."

[58] Eshelman and Derrick, "Relying on the Kindness of Machines?"

[59] Future of Life Institute, "Autonomous Weapons: An Open Letter."

[60] Edward M. Lerner, "A Mind of Its Own, Part II: Superintelligence," *Analog Science Fiction & Fact,* October 2016, 25-35.

[61] Ayoub and Payne, "Strategy in the Age of Artificial Intelligence."

[62] Human Rights Watch, *Losing Humanity*. See also Eshelman and Derrick, "Relying on the Kindness of Machines?"

[63] Luke Dormehl, "Algorithms: AI's Creepy Control Must Be Open to Inspection," *The Guardian Online*, January 1, 2017, https://www.theguardian.com/commentisfree/2017/jan/01/algorithms-ai-artificial-intelligence-facebook-accountability (accessed January 4, 2017). See also Lerner, "A Mind of Its Own, Part II," 25-35.

[64] Scharre, *Autonomous Weapons and Operational Risk*. See also Lerner, "A Mind of Its Own, Part I," 38-49.

[65] Executive Office of the President, *Preparing for the Future of Artificial Intelligence.*

[66] The AI Now Report, "The Social and Economic Implications of Artificial Intelligence Technologies in the Near-Term," Summary of the AI Now Public Symposium, Hosted By the White House and New York University's Information Law Institute, July 7, 2016, https://artificialintelligencenow.com/media/documents/AINowSummaryReport_3.pdf (accessed February 25, 2017).

[67] Executive Office of the President, National Science and Technology Council, *The National Artificial Intelligence Research and Development Strategic Plan,* Report of the Network and Information Technology Research and Development Subcommittee (Washington, DC: The White House, October 13, 2016), https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/national_ai_rd_strategic_plan.pdf (accessed November 22, 2016).

[68] Scharre, *Autonomous Weapons and Operational Risk*.

[69] Ibid.

[70] Dormehl, "Algorithms: AI's Creepy Control."

[71] Scharre, *Autonomous Weapons and Operational Risk*.

[72] Stewart, *Strong Artificial Intelligence and National Security*. See also Schwab, "The Fourth Industrial Revolution."

[73] Human Rights Watch, *Losing Humanity*. See also Sadler, "Fast Followers."

[74] Al-Rodhan, "The Security Implications and Existential Crossroads." See also Caton, "Autonomous Weapon Systems: A Brief Survey."

[75] Obama, *National Security Strategy*, 7-10. See also Human Rights Watch, *Losing Humanity*.

[76] Caton, "Autonomous Weapon Systems: A Brief Survey."

[77] Scharre, *Robotics on the Battlefield Part II.*

[78] Al-Rodhan, "The Security Implications and Existential Crossroads."

[79] Ayoub and Payne, "Strategy in the Age of Artificial Intelligence." See also Al-Rodhan, "The Security Implications and Existential Crossroads."

[80] Work and Brimley, *20YY: Preparing for War in the Robotic Age*.

[81] Human Rights Watch, *Losing Humanity.*

[82] Human Rights Watch, *Mind the Gap: The Lack of Accountability for Killer Robots* (Washington, DC: Human Rights Watch, April 9, 2015), https://www.hrw.org/report/2015/04/09/mind-gap/lack-accountability-killer-robots (accessed February 25, 2017).

[83] Human Rights Watch, *Losing Humanity.*

[84] Lerner, "A Mind of Its Own, Part I," 38-49. See also Ayoub and Payne, "Strategy in the Age of Artificial Intelligence." See also Executive Office of the President, *The National Artificial Intelligence Research and Development Strategic Plan*.

[85] Ayoub and Payne, "Strategy in the Age of Artificial Intelligence." See also Al-Rodhan, "The Security Implications and Existential Crossroads."

[86] Human Rights Watch, *Mind the Gap*.

[87] Human Rights Watch, *Losing Humanity*. See also Sadler, "Fast Followers."

[88] Human Rights Watch, *Losing Humanity.*

[89] Ayoub and Payne, "Strategy in the Age of Artificial Intelligence." See also Work and Brimley, *20YY: Preparing for War in the Robotic Age.*

[90] Scharre, "Autonomous Weapons and Operational Risk."

[91] Scharre, *Robotics on the Battlefield Part I.*

[92] Scharre, "Autonomous Weapons and Operational Risk."

[93] Work and Brimley, *20YY: Preparing for War in the Robotic Age.* See also Herman, "The Pentagon's 'Smart' Revolution." See also Scharre, *Robotics on the Battlefield Part I.*

[94] Work and Brimley, *20YY: Preparing for War in the Robotic Age*. See also Herman, "The Pentagon's 'Smart' Revolution."

[95] Elkus, "The AI Wars?"

[96] Herman, "The Pentagon's 'Smart' Revolution." See also Schwab, "The Fourth Industrial Revolution." See also Scharre, *Robotics on the Battlefield Part I.*

[97] Eshelman and Derrick, "Relying on the Kindness of Machines?" See also Schwab, "The Fourth Industrial Revolution." See also Brad Stone, "Silicon Valley Reckons with its Political Power," *Bloomberg Online*, December 22, 2016, https://www.bloomberg.com/news/articles/2016-12-22/silicon-valley-reckons-with-its-political-power (accessed January 23, 2017).

[98] Al-Rodhan, "The Security Implications and Existential Crossroads."

[99] Anne-Marie Slaughter, "The Only Way Forward: Can the New World Order Be Saved By Humanism?" *Foreign Policy Online Global Thinkers 2016*, https://gt.foreignpolicy.com/2016/essay/the-only-way-forward?df8f7f5682 (accessed March 2, 2017). See also Al-Rodhan, "The Security Implications and Existential Crossroads."

[100] Michael Vaughan, "After Westphalia, Whither the Nation State, Its People, and Its Governmental Institutions?" September 29, 2011, https://espace.library.uq.edu.au/data/UQ_266787/AfterWestphalia.pdf (accessed February 25, 2017).

[101] "Is History's Most Ruthless Company the Future of Business?" *Slate Online*, December, 2016, http://www.slate.com/articles/arts/taboo/2016/12/taboo_east_india_company_circa_2016.html (accessed January 10, 2017). See also Vaughan, "After Westphalia."

[102] Khatchadourian, "The Doomsday Invention." See also Lewis-Kraus, "The Great A.I. Awakening."

[103] Lewis-Kraus, “The Great A.I. Awakening.”

[104] Ibid.

[105] Rob Price, “Alphabet’s Eric Schmidt: Big Data is So Powerful, Nation States Will Fight Over It,” *Business Insider Online*, March 9, 2017, http://www.businessinsider.com/google-eric-schmidt-countries-will-fight-over-big-data-alphabet-cloud-2017-3 (accessed March 9, 2017).

[106] Ibid.

[107] Executive Office of the President, *Artificial Intelligence, Automation, and the Economy.*

[108] Stone, “Silicon Valley Reckons With Its Political Power.”

[109] Vaughan, “After Westphalia.” See also Dormehl, “Algorithms: AI’s Creepy Control.”

[110] Stone, “Silicon Valley Reckons With Its Political Power.”

[111] Standing Committee, *Artificial Intelligence and Life in 2030*. See also Sydney Finkelstein, “Algorithms are Making Us Small-Minded,” *BBC Online*, December 12, 2016, http://www.bbc.com/capital/story/20161212-algorithms-are-making-us-small-minded (accessed December 13, 2016).

[112] Dormehl, “Algorithms: AI’s Creepy Control.”

[113] Stone, “Silicon Valley Reckons With Its Political Power.”

[114] Slaughter, “The Only Way Forward.” See also Standing Committee, *Artificial Intelligence and Life in 2030*. See also Schwab, “The Fourth Industrial Revolution.”

[115] Stone, “Silicon Valley Reckons With Its Political Power.”

[116] Martinne Geller and Ben Hirschler, “Impact of Job-Stealing a Growing Concern at Davos,” *Reuters Online*, January 20, 2017, http://www.reuters.com/article/us-davos-meeting-robots-idUSKBN1540H0 (accessed January 21, 2017).

[117] Executive Office of the President, *Artificial Intelligence, Automation, and the Economy*.

[118] Executive Office of the President, *Artificial Intelligence, Automation, and the Economy*. See also Standing Committee, *Artificial Intelligence and Life in 2030.* See also Schwab, “The Fourth Industrial Revolution.”

[119] Executive Office of the President, *Artificial Intelligence, Automation, and the Economy.* See also Cade Metz, “The AI Threat Isn’t Skynet. It’s the End of the Middle Class,” *Wired Business Online*, February 10, 2017, https://www.wired.com/2017/02/ai-threat-isnt-skynet-end-middle-class/ (accessed March 10, 2017). See also The AI Now Report, “The Social and Economic Implications of Artificial Intelligence.”

[120] Schwab, “The Fourth Industrial Revolution.” See also Metz, “The AI Threat Isn’t Skynet.”

[121] Ibid.

[122] Schwab, “The Fourth Industrial Revolution.”

[123] Standing Committee, *Artificial Intelligence and Life in 2030.*

[124] Executive Office of the President, *Artificial Intelligence, Automation, and the Economy.*

[125] Ibid.

[126] Metz, “The AI Threat Isn’t Skynet.”

[127] Executive Office of the President, *Artificial Intelligence, Automation, and the Economy.*