Strategy Research Project

Implications of U.S. Ban on Artificially Intelligent Automated Killing Machines

by

Lieutenant Colonel John Gerard Lehane
United States Marine Corps

Under the Direction of:
Professor Samuel White

United States Army War College
Class of 2017

REPORT DOCUMENTATION PAGE	*Form Approved--OMB No. 0704-0188*

1. REPORT DATE *(DD-MM-YYYY)* 01-04-2017	**2. REPORT TYPE** STRATEGY RESEARCH PROJECT	**3. DATES COVERED** *(From - To)*
4. TITLE AND SUBTITLE Implications of U.S. Ban on Artificially Intelligent Automated Killing Machines		**5a. CONTRACT NUMBER** **5b. GRANT NUMBER** **5c. PROGRAM ELEMENT NUMBER**
6. AUTHOR(S) Lieutenant Colonel John Gerard Lehane United States Marine Corps		**5d. PROJECT NUMBER** **5e. TASK NUMBER** **5f. WORK UNIT NUMBER**
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Professor Samuel White		**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:** X

13. SUPPLEMENTARY NOTES
Word Count: 7,586

14. ABSTRACT
In November 2012, then Secretary of Defense Carter signed Department of Defense Directive 3000.09 Autonomy in Weapons Systems, which stymies the development and fielding of artificially intelligent autonomous lethal weapons systems. While that policy may be appropriate for weapons systems that presently exist, inhibiting their future development will place the United States at a significant tactical, operational, and strategic disadvantage in the future. This paper will examine the current pace of technological development, forecasts on how those trends may continue, the nascent development of automated killing machines, projected development of those machines, objections of governments and non-governmental organizations, and implications of either DoD Directive 3000.09 or future directives which may be even more restrictive.

15. SUBJECT TERMS
Robotics, Singularity, Future, Kurzweil, Autonomy

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**		39	**19b. TELEPHONE NUMBER** *(w/ area code)*
UU	UU	UU	UU		

Implications of U.S. Ban on Artificially Intelligent Automated Killing Machines

(7,586 words)

Abstract

In November 2012, then Secretary of Defense Carter signed Department of Defense Directive 3000.09 Autonomy in Weapons Systems, which stymies the development and fielding of artificially intelligent autonomous lethal weapons systems. While that policy may be appropriate for weapons systems that presently exist, inhibiting their future development will place the United States at a significant tactical, operational, and strategic disadvantage in the future. This paper will examine the current pace of technological development, forecasts on how those trends may continue, the nascent development of automated killing machines, projected development of those machines, objections of governments and non-governmental organizations, and implications of either DoD Directive 3000.09 or future directives which may be even more restrictive.

Implications of U.S. Ban on Artificially Intelligent Automated Killing Machines

> Once a new technology rolls over you, if you're not part of the steamroller, you're part of the road.
>
> —Stewart Brand[1]

In November 2012, then Secretary of Defense Ashton Carter signed Department of Defense (DoD) Directive 3000.09 *Autonomy in Weapons Systems*, which stymied the development and fielding of artificially intelligent autonomous lethal weapons systems.[2] *While that policy may be appropriate for weapons systems that presently exist, inhibiting their future development will place the United States (US) at a significant tactical, operational and strategic disadvantage in the future.* While the directive is not a complete ban, based upon public comments made by the former Secretary of Defense, such a ban is possible in follow-on directives. This paper will examine: the current pace of technological development, forecasts on how those trends may continue, the nascent development of Artificially Intelligent Automated Killing Machines (AIAKMs), projected development of those machines, objections of governments and non-governmental organizations, and implications of a more restrictive DoD Directive 3000.09 (or future directives). Based upon these implications, recommendations will be proposed on how the US may better take advantage of the forthcoming change. It is critical, however, to start with an understanding of how quickly things are changing.

In 1965, Intel engineer Gordon Moore observed a trend with respect to increased speed and capacity of silicon microchips, with a corresponding decrease in cost, energy use, and size. This observation, now known as Moore's Law, has been the "golden rule" for the electronics industry ever since.[3] Moore's Law states that "the number of transistors on a [micro]chip roughly doubles every two years."[4] The increase in computational power is matched by a reduction in production cost, and size.[5]

Moore's Law has held for more than 50 years, but there are some who believe that this pace can only be sustained for another five years.[6] There are three principal factors contributing to the end of Moore's Law. First, developers noted that once a chip's features were smaller than 90 nanometers there was significantly increased heat, which could not be dissipated effectively to maintain the circuits. Second, as circuit features shrink below 14 nanometers (less than 10 atoms wide), quantum principles begin to make them unreliable. Third, the machinery necessary to make continually smaller circuits is cost prohibitive to produce. This has curtailed the growth of individual circuit clock speed since 2004. Developers have since compensated by creating parallel processors, but those too are reaching design limits due to a phenomena known as electron leakage (meaning that electrons jump across circuits).[7] There are, however, some who believe that when Moore's Law reaches its end, something new will come along to replace it.

Dr. Raymond Kurzweil is among a growing group who believe that there will be a successor to Moore's Law, and he points out that Moore's Law had its own predecessors. Dr. Kurzweil has posited that Moore's Law is the fifth in a series of paradigms related to the development of computational power. "Kurzweil's Law" expands upon Moore's Law and posits that there is actually "an exponential growth in the rate of exponential growth" of computing power, and when the fifth paradigm concludes, it will be replaced by the next. Kurzweil's Law, when applied to the period from 1900 to 2000 demonstrates that through the previous paradigms, calculations per second per $1,000 has increased from approximately 0.000001 (or 1×10^{-6}) to approximately 10,000,000 (or 1×10^{8}), or a staggering gain of 10,000,000,000,000 fold (or 1×10^{14}) (see Figure 1).[8]

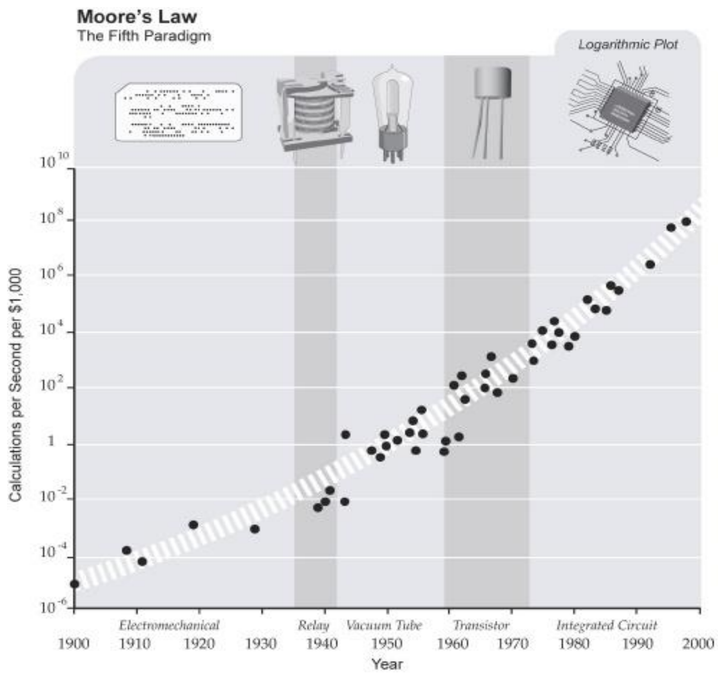

Figure 1. Moore's Law as the "Fifth Paradigm"[9]

The "sixth paradigm" of computing that Kurzweil's Law proposes has yet to take on a definitive form. In his 2005 book *The Singularity Is Near*, Dr. Kurzweil describes forms including nanotube circuits, and biological computers as possible contenders.[10] Others have pointed towards the promise of quantum computing as the silicon chip's successor. Quantum computers can theoretically perform at several orders of magnitude beyond the silicon chip by processing quantum bits, or qubits, vice the bits that current silicon-based computers use. Whereas a bit can only have a value of 1 or 0, a qubit may have values of 0, 1, neither or both due to the quantum mechanics phenomena known as superposition.[11] Determining the form of the sixth paradigm is beyond the scope of this paper, however, the acceptance of a sixth paradigm is a fundamental assumption supporting this paper. The sixth paradigm is a necessary enabler to develop an artificial intelligence (AI) that can perform better than the human brain.

The processing power of the human brain is not an easy thing to measure, and there is no consensus on a precise figure. An internet search for “how many computations per second can the human brain do” will yield figures between 1 x 10^{14} and 1 x 10^{16}. In 2011, the *Scientific American* published an article that placed the computational power of the human brain at 2.2 billion megaflops (or 2.2 x 10^{15} calculations per second).[12] That figure is roughly consistent with Dr. Kurzweil’s hypothesis that the required computational power to equal a human brain would be 1 x 10^{16} calculations per second, which takes into account the human brain’s parallel processing efficiency and other advantages not present in a computer.[13] This paper will use the assumption that 1 x 10^{16} calculations per second is the requirement, as it appears to be the most conservative. If the trends noted earlier with respect to Moore’s Law and subsequently Kurzweil’s Law are extrapolated, then it is clear that a computer will achieve the 1 x 10^{16} calculations per second required to have the same calculation power as the human brain. This is not an original thought – in 2005, Dr. Kurzweil posited that this would take place around 2025.[14] Subsequent to 2025, the expansion in computing power will continue.

Dr. Kurzweil has stated that by 2025 a computer with capacity equivalent of the human brain will cost $1,000, and that by 2035 the human brain equivalent will cost 1 cent.[15] The trend continues through the 2040s when it is predicted that a single computer costing $1,000 exceeds the computational ability of the sum total of the human race (see figure 2).

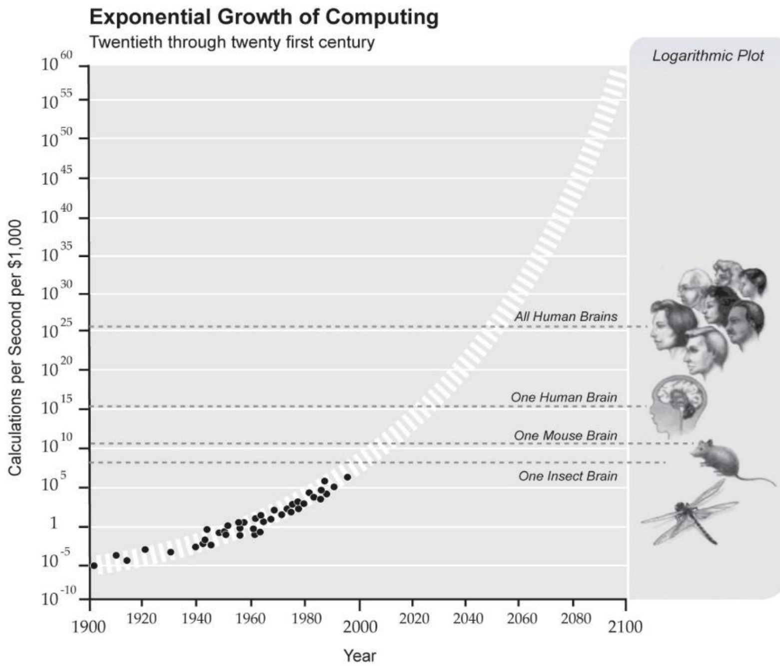

Figure 2. Exponential Growth of Computing Power into the 21st Century[16]

Even if the dates are affected by a mild slowing during the shift from silicon-based computers to the sixth paradigm that Dr. Kurzweil described, the trend is clear - it is reasonable to believe that there will be computers that exceed the computational capabilities of the human brain within a relatively short time. Calculations per second, however, are not the only measure of what it will take to have true AI. To gain an appreciation for true AI, it is necessary to consider the nature of human intelligence.

When viewed through an evolutionary perspective, human intelligence supported survival of the species by ensuring that human beings had the wherewithal to find food, avoid predators, mate, and protect offspring. As humans evolved, and language transformed societies, the brain too evolved to add capabilities. In this regard, the brain could be considered an “ultra-sophisticated – but special-purpose computer.”[17] If considered in this context, then there are some parallels to be drawn between human brains and computers with respect to how each system runs and processes information.

The makeup of how a computer runs can generally be broken down into two categories: information and rules that are a part of the computer, and the software that is loaded onto the computer. The information and rules are the Basic Input/Output System or BIOS. The BIOS sets the conditions for how a computer operates by ensuring that devices are connected and working properly, that the memory and other systems are running, and then turns over operation of the computer to software.[18] From there, the installed operating system, such as Microsoft Windows or Apple iOS, takes control of the computers operations. Thus, a computer operates on a mixture of hard-wired instructions and software. Similarly, human intelligence operates on a mixture of hard-wired instructions and software; in very simple terms, the BIOS is akin to human instincts and automatic nervous system behaviors, and the software is akin to learned behaviors. Drawing this parallel further, learned behaviors are analogous to “IF-THEN” statements in computer software. Within this analogy, it is possible to examine how many “IF-THEN” statements a human being processes per day and propose what it might take to create an equivalent AI.

A precise number of “IF-THEN” decisions made by human beings per day would be impossible to calculate, but there are many scholarly articles that estimate that human beings make 35,000 conscious decisions per day.[19] Some of these decisions are important, and some are trivial. One Cornell University study suggests that adults use over 200 of those decisions simply on what to eat.[20] For argument’s sake, 35,000 decisions per day will be used as an assumption for the number of decisions needed to reach human-equivalent AI. However, there is clearly more to comparative intelligence than simply numbers of decisions….for if that were the threshold, computer software would have reached human-equivalent AI decades ago. There is a multiplier that empowers those 35,000 human decisions to be more effective: pattern recognition.

Human beings use pattern recognition, or heuristics, to speed decision-making. This is learned behavior. The more times a human being sees something, the faster he or she will resort to the heuristic, thus multiplying the value of one of the 35,000 decisions made each day. The process is effective, but has limits. When human beings are confronted with a wholly unknown set of conditions, decision making slows. Moreover, when a familiar set of conditions are presented but a key element is overlooked or ignored, the heuristic may yield an incorrect decision. These limitations aside, heuristics have multiplied the value of the 35,000 daily decisions to a great extent. AI researchers have recognized this in several fields of application, but this paper will focus on the fields of neural networks and machine learning.

Neural networks have been developed based on models of how a human might approach a task through pattern recognition, with the first attempts beginning in the 1950s.[21] Modern day neural networks are enhancing or replacing human decision-making, and some researchers believe that this advancement will soon permeate everyday life.[22] Neural networks are already employed or under development to power self-driving automobiles, run facial recognition in surveillance systems, enhance multiplayer video games, and translate languages.[23] The field of Deep Learning combines neural networks together to allow for a parallel structure that can tackle abstract data. Contemporary examples of deep learning AI include Google's DeepMind, which defeated a world-class Go champion, and IBM's Watson, which won the game show Jeopardy.[24] While those systems aren't on par with a human mind, they are nonetheless impressive and offer a glimpse of the possible.

When the capabilities of current deep learning AI systems are projected into the future using Kurzweil's Law, it is clear that not only will an AI equal the decision-making capacity of the human mind, but will soon thereafter greatly exceed it. Combine this with the fact that computers can transfer information, or learned behavior, between each other in a manner that human beings cannot replicate. The result is that pattern recognition, which takes a human being a complete childhood and adolescence to build, could be rapidly shared across tens of thousands of systems. An added benefit is that this learning would not be lost when the human dies. This could result in a continually building pattern recognition database shared across neural networks, and that could rapidly exceed the sum total of any one human being's observed patterns.

Despite the astonishing cognitive potential of future computing systems, some might argue that this advancement would never result in true intelligence. While defining (or defending) what is or isn't truly intelligence exceeds the scope of this paper – it is also immaterial to the conclusions. In the end it will come down to a machine being able to process – by several orders of magnitude – more "IF-THEN" statements than humans. Compounding this fact, systems will be making the right connections, pattern recognition, from the increased amount of information observed. At this point, machines will be able to make faster and better decisions than human beings. The decisions may be potentially better because the machines will be able to consider a wider number of options in a shorter period of time, and not be limited by bias. The likely discussion about whether or not an AIAKM has genuine intelligence will be immaterial. AI will have access to vast information that it can process at incredible speeds, such that a human being will be unable to compete. While AI is the key enabler, an AIAKM still needs the weapons portion. Current weapon systems are already fielded, or under development to fulfill that requirement, and other forms that haven't even been considered yet are sure to come. Further, while an AIAKM has yet to be introduced to the battlefield, the effort to decrease human presence in direct combat, and replace them with machines has been underway for several years.

The US 2001 National Defense Authorization Act (NDAA) set goals for the DoD to acquire unmanned air and ground vehicles. Specifically, it mandated that "by 2010, one-third of the aircraft in the operational deep strike force air fleet are unmanned," and "by 2015, one-third of the operational ground combat vehicles are unmanned."[25] Of note, the ground combat vehicles described were associated with the now defunct Future Combat Systems (FCS) program. Nonetheless, the 2001 NDAA represented a significant step forward towards AIAKM development.

Eight years later, in 2009, the United States Air Force (USAF) published the *Unmanned Aircraft System (UAS) Flight Plan 2009-2047*. This flight plan detailed the way ahead for USAF UASs for over three decades, and contained actionable information affecting Doctrine, Organization, Training, Materiel, Leadership, Education, Personnel, Facilities, and Policy (DOTMLPF-P).[26] The flight plan stated that the use of "UAS are compelling where human physiology limits mission execution" due to factors such as "persistence, speed of reaction, [and] contaminated environment."[27] The document was more than just a future concept; when it was published it reflected nascent practice. By August of 2009, just three months after the flight plan was published, it was reported that the USAF annually trained more pilots to fly UASs than to fly manned aircraft.[28] While the flight plan did not provide an in-depth discussion on AIAKMs, it did describe "Auto [Target] Engagement" capabilities emerging between fiscal years '25 and '47 (see Figure 3), and stated that "[a]ssuming legal and policy decisions allow, technological advances in artificial intelligence will enable UAS to make and execute complex decisions."[29]

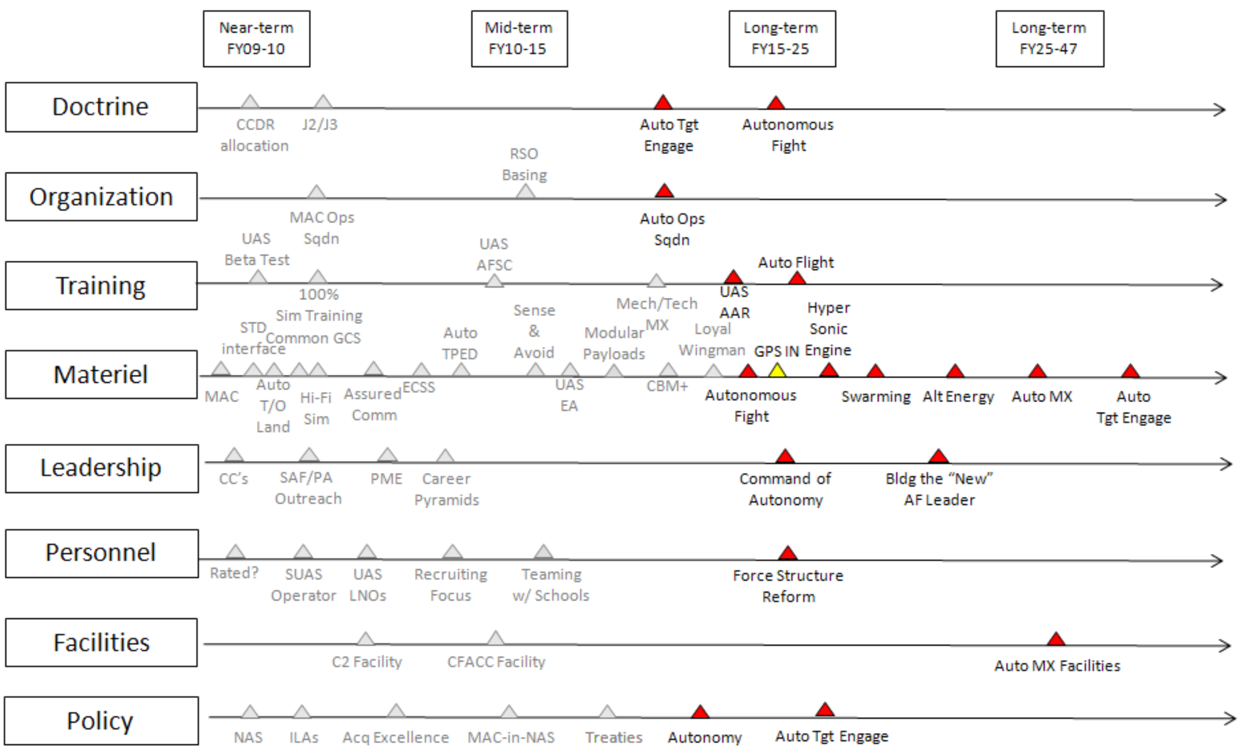

Figure 3. The USAF's *Unmanned Aircraft System (UAS) Flight Plan 2009-2047*[30]

UASs got a head start, and have remained ahead in both development and fielding when compared to Unmanned Ground Vehicles (UGVs). The nearest ground equivalent of the USAF's UAS Flight Plan was the Future Combat Systems (FCS) program, however, in the same year that the USAF published its UAS Flight Plan, the Army's FCS program was terminated.[31] The FCS program featured an extensive line of

UGVs including the Armed Robotic Vehicle (ARV) and the Multifunctional Utility/Logistics and Equipment (MULE). Prior to FCS's cancellation, the ARV underwent several promising tactical demonstrations with various models between 2004 and 2008.[32] The MULE (see Figure 4) came in transport, countermine and assault support models, and survived the 2009 demise of the FCS program, only to be terminated in 2011 citing cost and technology immaturity.[33] The termination of FCS was, by no means, the end of UGVs, but does highlight the principal challenge they must overcome – operating on the ground has more variables that operating in the air. Efforts to overcome those challenges are still going on.

Figure 4. Lockheed Martin's Multifunctional Utility/Logistics and Equipment (MULE)[34]

As recently as 2016, the United States Marine Corps (USMC) experimented with Qinetic's Modular Advance Armed Robotic System (MAARS) (see Figure 5) as a part of a future infantry battalion.[35] MAARS can be armed with a machine gun, and a grenade launcher, and is equipped with various sensors; all of this is mounted on a small tracked chassis.[36] MAARS could conceivably take the place of a machine gunner or other ground combatant, but requires a remote operator. In this regard, MAARS, like the MULE, the ARV and existing UASs are not AIAKMs. Existing UASs and UGVs are "man in the loop" systems, meaning that they require a remote operator and cannot make their own decisions. There are, however, recent developments towards integrating AI with unmanned systems.

Figure 5. Qinetiq's Modular Advanced Armed Robotic System (MAARS)[37]

In November of 2016, Russian weapons developers claimed in open press to have a robot, dubbed 'Flight' that could automatically detect and engage human-sized targets at ranges in excess of four miles.[38] Not to be outdone, the Defense Advanced Research Projects Agency (DARPA) has claimed that it is actively pursuing projects advancing AI with the ultimate goal of creating a sentient machine.[39] Plans are even reportedly underway to develop drone wingmen to accompany manned F-35 and F-22 aircraft, and will be controlled in part through AI built into the aircraft.[40] This development is somewhat ironic, given that the frail human beings piloting advanced aircraft are a limiting factor with respect to flight performance. Regardless, it is clear that many are working on taking advantage of the capabilities AI can bring to weapons systems. This progress is, however, not without its detractors.

In April of 2016, Human Rights Watch and Harvard Law School's International Human Rights Clinic jointly issued a call on all states to ban weapons that did not have "meaningful human control" (meaning control over the selection and engagement of individual targets, not simply offering criteria to the weapons system and allowing it to make its own decision). Both organizations suggested that nations should become signatories to international law prohibiting fully autonomous weapons, and adopt national laws prohibiting the same.[41] Some nations preceded this call with action.

In June of 2013, the Parliament of the United Kingdom (UK) held hearings on "Lethal Autonomous Robots", and although the members of Parliament recognized that nations possessing such capabilities would have a significant advantage over those who didn't, they elected to affirm that the nation would not seek acquisition or development of such capabilities.[42] The UK's position was reinforced by Foreign Office Minister Alistair Burt in a subsequent interview where he stated "The UK has unilaterally decided to put in place a restrictive policy whereby we have no plans at present to develop lethal autonomous robotics" and adding that "we do not intend to [formalize] that in a national moratorium."[43] The UK had this debate a year after US DoD policy was established.

In November of 2012, the US Secretary of Defense signed a directive entitled *Autonomy in Weapons Systems.* That directive stated that "[a]utonomous and semi-autonomous weapon systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force."[44] While the directive doesn't ban the weapons systems outright, it does place significant restrictions on them that all but rule out the use of AI to independently make life or death decisions. The policy specifically states that semi-autonomous systems must not attack targets that have not been previously selected. If the system should have lost communications, the system can be used for: the defense of static installations and "manned platforms", and, they may use "non-lethal, non-kinetic force" applied "against material targets." All systems not meeting these criteria must be individually approved by the Undersecretary of Defense for Policy, the Under Secretary of Defense for Acquisition, Technology and Logistics, and the Chairman of the Joint Chiefs of Staff.[45] This policy, has been surrounded in controversy within the DoD.

In August of 2016, Frank Kendall, Undersecretary of Defense for Acquisition, Technology and Logistics raised concerns that the US may be hobbling itself in future conflicts by insisting on human control over weapons systems. Mr. Kendall's concerns are at odds with Robert Work, US Deputy Secretary of Defense, who stated that in the future humans must be working with machines, but Mr. Kendall expressed sincere doubt that "human-machine teaming is superior to machines untethered to human judgment."[46] This public display of difference in professional opinions was likely the cause of follow-on commentary by the Secretary of Defense.

In September 2016, Secretary of Defense Ashton Carter stated to reporters that while the US was actively pursuing AI, that it would never unleash "truly autonomous killing machines," emphasizing that "whenever it comes to the application of force, there will never be true autonomy, because there'll be human beings (in the loop)."[47] The public back-and-forth on autonomous weapons leaves room for ambiguity and debate. It could be argued that DoD Directive 3000.09 was never intended to ban the development of AIAKMs. One could argue instead, that the directive was intended to: provide guidance for development; provide political top cover while developing controversial weapons; and, require explicit approval to move forward with the development of specific weapons systems. Letter and intent of DoD Directive 3000.09 aside, it is conceivable that public opinion will generate sentiment supporting a more restrictive follow-on policy in the US. As such, this paper will interpret the former Secretary of Defense's words as the basis for an assumption that there is a long-term intent and potential to prohibit the development of AIAKMs. It is also worth noting that the public back-and-forth within the DoD is also echoed by influential individuals outside of the department.

In 2015, Elon Musk and Stephen Hawking called for "a ban on offensive autonomous weapons beyond meaningful human control" citing fears that the weapons might fall into the hands of terrorists, be used to subdue populations, or even be tasked with committing genocide.[48] While those calling for a ban seem more numerous, AIAKMs also have proponents. Georgia Tech Professor, Ron Arkin, in an April 2016 interview with National Public Radio pointed out that while the on-going debate was important, the benefits or the potential benefits of AIAKMs should not be overlooked. Arkin stated that "they can assume far more risk on behalf of a noncombatant than any human being in their right mind would" and that they "can potentially have better sensors to cut through the fog of war. They can be designed without emotions such as anger, fear, frustration – which causes human beings, unfortunately, to err."[49] Professor Arkin is a proponent, and strikes at the heart of the matter – human beings are fallible.

To demonstrate just how fallible human beings are, it's worthwhile to examine an experience common to most adults – driving. Over the two previous decades in the US, the average number of people killed annually in automobile accidents is approximately 33,000.[50] 94% of those deaths have been attributed to human error.[51] During the three year period from 2013 to 2015, the National Highway Transportation Safety Administration (NHTSA) recorded an average of over 6 million police-reported crashes annually resulting in over 2.3 million injuries in the US.[52] Despite these statistics, nearly everyone has likely heard someone say, "There's no way a computer will ever be a better driver than me." And yet, Google's self-driving vehicles "have driven more than 500,000 miles without a crash attributable to the automation."[53] That is an accident rate of 0.0%, and accounts for 0 fatalities and 0 injuries. While fatal motor vehicle accidents in the US have been in steady decline since the mid-1960s where they peaked at over 50 fatalities for every billion miles traveled, as of 2013 there were still approximately 10 fatalities for every billion miles traveled.[54] That steady decline, while positive, is poor when compared to the NHTSA's report which noted that in Tesla automobiles, the crash rate dropped 40% after the automatic steering feature was introduced.[55] A RAND corporation study published in 2016 concluded that self-driving vehicles would lead to "substantial reductions in crashes and the resulting human toll."[56] It appears as though systems operating AI well below the threshold of human intelligence are already making better decisions than human beings, at least when it comes to driving. In addition to making better decisions, artificially intelligent vehicles are not subject to the key contributors to distracted driving identified by the NHTSA, namely texting, cell phone use, eating, grooming, reading, watching videos or adjusting the radio.[57]

Many of the same human beings who drive cars in the US also serve in its armed forces, and carry with them the same fallibilities that make them risky drivers. While some might argue that it is a bit of a stretch to compare combat with driving, it could be argued that both are executed largely through human pattern recognition and there is room for error when human beings encounter previously unrecognized patterns, draw bad conclusions from incomplete observations, are scared, are fatigued, or are distracted. It is difficult to determine exactly what percentage of battlefield errors can be attributed to each of these categories, but there is some data that offers at least anecdotal evidence that the parallels are valid. During the 1991 Gulf War, 35 of the 148 Americans killed lost their lives to friendly fire. The Gulf War was not an isolated incident – in 1758 Colonel George Washington killed 13 friendly British troops in a firefight during the French and Indian War; Stonewall Jackson died from wounds after being hit by friendly fire; during the Vietnam War, American helicopters killed American soldiers on Hamburger Hill.[58] Other undesirable outcomes such as panic, fear, commitment of atrocities, and more can be tied to human fallibility. The US will continue to face these outcomes as long it avoids proceeding forward in deploying AI in its killing machines in the same manner it appears to be headed with self-driving automobiles.

In the near future it is likely the machine will be able to best the human being on the battlefield. Regardless, arguments are still made that manned-unmanned teaming is the optimal way ahead. This is thinking is flawed and sub-optimal. It presumes that the human will always provide added value to the machine, irrespective of the time required to add this value. That may be true in the short term, but it is logical to assume that the human being will eventually be an impediment, and the manned-unmanned team of human and AI will be bested by a pure AI – the human being will take too long to offer meaningful input. Keep in mind, Kurzweil's Law predicts that if you are year behind in development, your adversary might have AI orders of magnitude more powerful. Given this, it is appropriate to anticipate that there are significant implications of banning the development of AIAKMs.

The first implication of is that the US is will likely lose battles (or perhaps even a war) against a peer competitor who elects to field AIAKMs in significant numbers. As discussed earlier, the hardware to kill is already deployed, and if coupled with AI, it is difficult to conceive how an American soldier, sailor, airman, or Marine could possibly overcome a force so equipped. Moreover swarming drones, nano-machines, or other emerging technology, coupled with AI make the prospects for a human on a future battlefield even more daunting. The argument that human creativity and intuition will somehow always be able to overcome the machine has been proven wrong time and again. The machine has bested chess masters, Jeopardy contestants, and world-class Go players. It has already demonstrated a statistically significant ability to drive better than its human counterpart. Therefore, the US must pay close attention to near peer competitors, such as China and Russia, that possess a significant industrial base, have demonstrated proficiency with AI, and have national interests and goals that are at odds with those of the US. In such a contest, the best equipped American Soldiers will find themselves as incapable, irrelevant, and outclassed by AIAKMs in a fashion similar to the battleship during the Second World War when attacked by aviation.

The US will find that "man-in-the-loop" systems remain susceptible to hacking and would operate at a distinct handicap against a "closed loop" AIAKM. The "closed loop" system could operate with far fewer, if any, signals inputs because onboard AI would be making its own decisions, and storing necessary information on board. The world has already witnessed the susceptibility of "man-in-the-loop" systems to hacking. One reported instance of hacking was reported in 2009 when the Iranian military successfully jammed the communications of an RQ-170 Sentinel. The Iranians hacked into the drone's Global Positioning System (GPS) software and hijacked the aircraft, ultimately landing it in Iran.[59] This incident prompted Professor Todd Humphreys at the University of Texas to explore whether or not Iran's claims were feasible. "Using equipment costing less than $2,000", he was successful at replicating Iran's claims, and did so in the presence of representatives from the US Department of Homeland Security.[60] One internet source quotes the unit cost of the RQ-170 at $6,000,000.[61] Assuming that the unit costs are even remotely close to that figure, a $2,000 investment in countermeasures provides a very good return against a "man in the loop" system. The US military is likely to recognize that it cannot sustain this cost differential and must address this vulnerability.

With respect to cost, the US may find that it can no longer afford the personnel costs associated with the contemporary all-volunteer military. AIAKMs will likely be cheaper to field and sustain than human combatants. Former Secretary of the Army, Eric Fanning, estimated that 10,000 soldiers costs roughly 1.6 billion dollars annually (or $160,000 per soldier) when considering recruitment, construction, equipment, modernization, and other indirect costs.[62] A 2014 article (which appears to be consistent with many other articles) by George Mason University's Mercatus Center points out that ballooning personnel, healthcare, and retirement costs will continue to consume an increasing and unsustainable portion of the DoD budget.[63] Military personnel costs in fiscal year 2014 consumed nearly a third of the DoD budget (see Figure 6). In contrast, AIAKMs will not need to be recruited, housed in barracks or government housing, receive medical care, or collect retirement after twenty years of service.

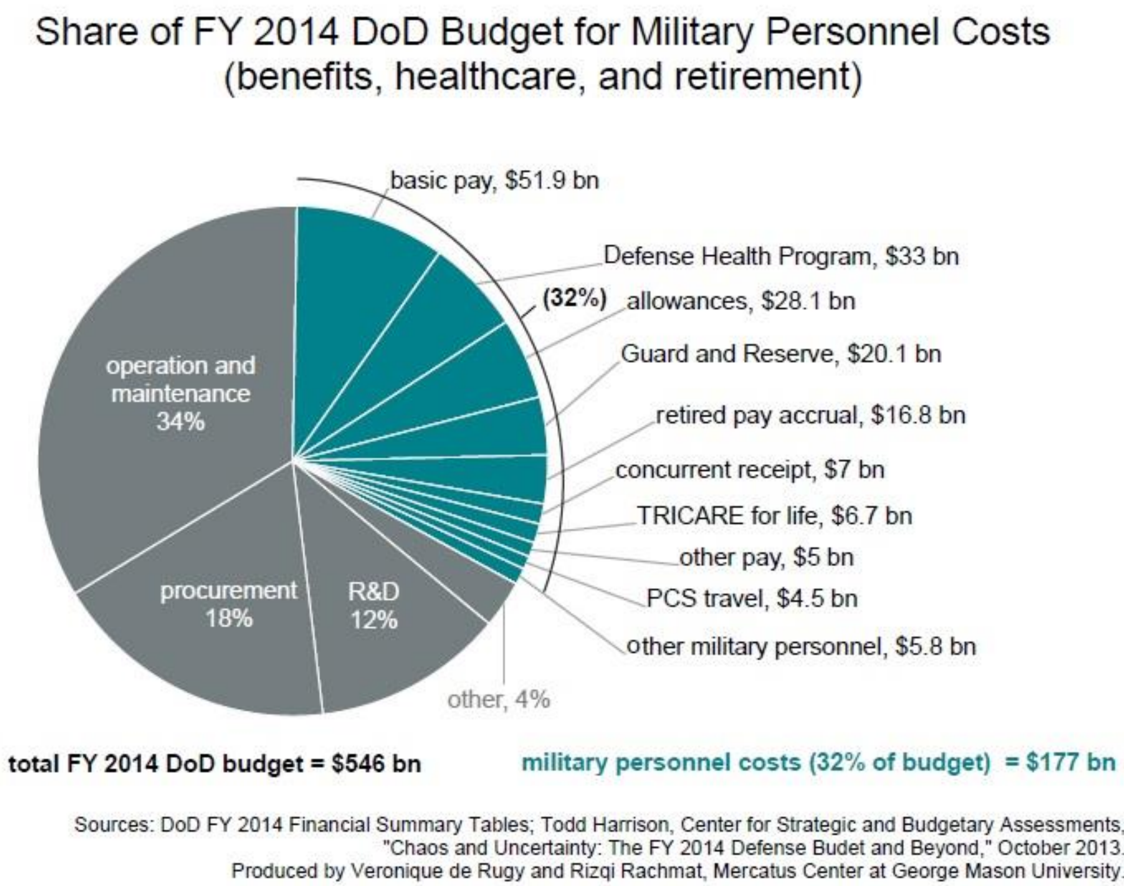

Figure 6. Mercatus Center's Analysis of the 2014 Department of Defense Budget.

Soldiers, Sailors, Airmen, and Marines are also expensive if they are killed in combat. Serviceman Group Life Insurance (SGLI) death benefits are $400,000.[64] Additionally, a death gratuity payment is $100,000.[65] There are also a myriad of other costs involving: funeral honors; programs to support surviving family members; and, recruiting and training a replacement. *USA Today* recently ran a story citing the cost of the average Improvised Explosive Device (IED) at $416.[66] A quick internet search will show that an AK-47 can be purchased (in the US) for $529, and a box of 640 rounds of 7.62mm x 39 ammunition can be purchased for $149.[67] Adversaries of the US may be acquiring these weapons for even less, but this cost is illustrative of the financial differential the US is trading when it loses service members in combat. If the 16 years of sustained combat operations that the US has engaged in since September 11th is indicative of years to come, the financial costs of human combatants is not sustainable. By contrast, there are already robotic human analogs that are cheaper than the costs of the service member described above, although they may not be as capable as a soldier yet. Kawada Industries produces the Humanoid Robot HRP-4 (see Figure 7), which can perform many of the tasks that a human being can, and costs $300,000.[68] Given Kurzweil's Law, it is reasonable to believe that the cost will come down, and the future versions of HRP-4's capabilities will increase. If developed into an AIAKM, this is a viable means to offset the financial costs associated with the death of human combatants. Human costs, however, are not just associated with loss of life.

Figure 7. Humanoid Robot Series HRP-4[69]

The US may also potentially forgo the cost savings that AIAKMs could provide with respect to sustainment. AIAKMs don't need to be given drink, fed, immunized, or exercised while waiting for a conflict to start – they can be kept in storage. Moreover, organizations supporting an AIAKM-equipped force could be smaller and more efficient that those supporting contemporary human formations. The AIAKM-equipped force would consume significantly less food and water, assuming that AIAKMs replace the majority of combat personnel. The AIAKM-equipped force would logically experience far fewer casualties or diseases, therefore it would require far less medical personnel and materiel support. If equipped with contemporary weapons, the AIAKM-equipped force could be more accurate than human soldiers, thus requiring less ammunition and fewer weapon system repair parts. The AIAKM-equipped force may likely require more repair parts, batteries, and other unique items to keep the machines running; but, given likely efficiencies, it seems reasonable that this would represent an overall decrease in the logistical requirements of the force. An argument could be made that an AIAKM force would require significant maintainer support, but this forgoes the likely development of self-repairing systems, other machines that might conduct repairs, or disposable AIAKMs. The opportunity for cost reductions is compelling.

Adversaries of the US are increasingly aware of what it costs to project power, and some have already adopted cost imposing campaigns against the US. In 2004, Al Qaeda (AQ) leader Osama Bin Laden clearly articulated this awareness when he stated that his organization was pursuing a “policy in bleeding America to the point of bankruptcy.”[70] The AQ leader went on to cite examples including the September 11 attacks costing AQ less than $500,000 to conduct, as well as previous success against the Russians in Afghanistan with limited expenditure on equipment.[71] If an AQ analog, or a nation state were to attempt a similar cost imposing campaign in the future, the US economy might not be able to withstand the impact. However, the financial impact of such a campaign would only be part of the story.

The US may find that it could no longer find an appropriate quantity or quality of personnel to field the all-volunteer military. Without the AIAKM to compensate for this shortage, the US may have to lower accession standards or even resort to a draft to mitigate the risks. Since the 1980s, there has been a decline in propensity of Americans to want to join the military. A 2009 Strategic Analysis Incorporated study on recruiting the all-volunteer force showed that youth propensity towards military service has declined (see Figure 8). The inability to find sufficient qualified personnel for enlistment was also expressed as a concern in 2015 by a DoD official responsible for military accessions, noting that propensity to enlist had fallen by 27 percentage points, and attributed the decline to a “lot of media coverage of the military over the last 10 years [that] has highlighted wounded warriors, sexual assault and some of the negative aspects of military service."[72] If the US is forced to resort to a draft to fill its military ranks, the sentiment of the domestic population will weigh more heavily on its decisions to go to war or conduct military interventions, especially against AIAKM-equipped adversaries.

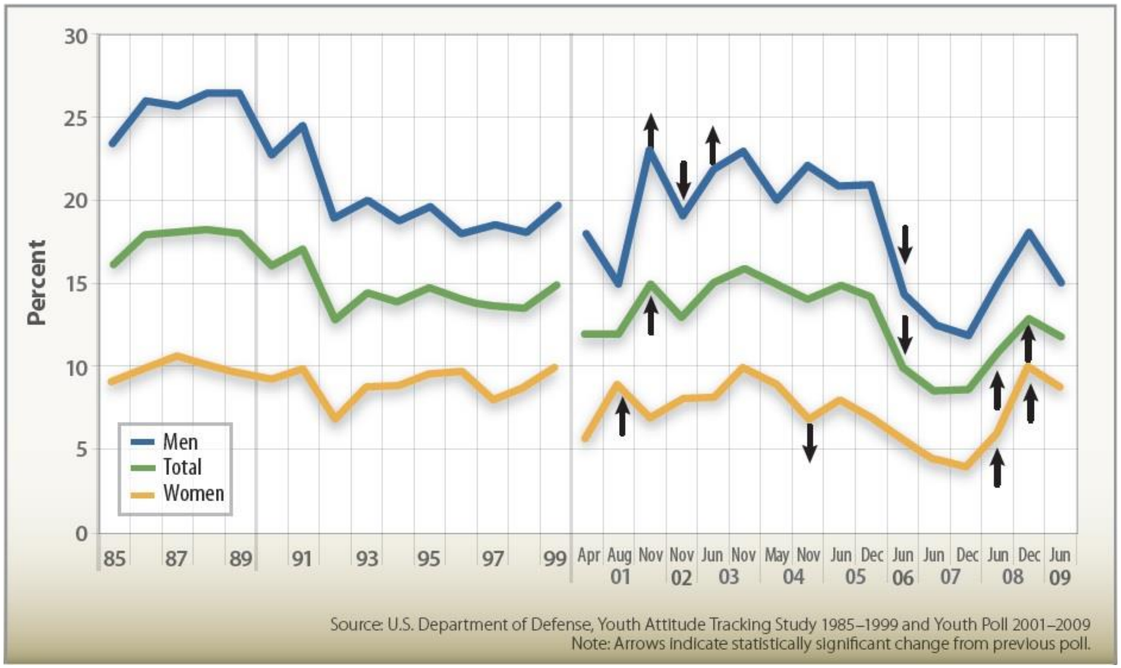

Figure 8. Youth Propensity for Military Service has Decreased Significantly[73]

The US may find decreased domestic acceptance for loss of American lives, especially if the population recognizes that other nations are sending AIAKMs and not sending their sons and daughters to war. Moreover, if adversaries are equipped with AIAKMs, there may be loss of confidence in the US military's ability to win, which may likely in-turn decrease the willingness of the population to support a conflict. In AIAKM-equipped nations, warfare may become further removed from the population as they have fewer visible reminders of the cost of war. In contrast, the US media coverage is likely to increase on Americans wounded or killed by AIAKMs. This increased media coverage is likely to increase visibility of the cost of war, which may further undermine domestic support for conflict. At some point, Americans will ask why the nation should risk their sons and daughters when other nations do not. Popular support may similarly decline if it appears that the fight against AIAKMs is unwinnable. Loss of popular support will limit US freedom of action, especially when a conflict is not related to a vital or existential national interest. Second order effects of this phenomena may place existing US mutual defense treaties at risk, which in turn may jeopardize the existing international system and world order. If a NATO partner is attacked by, and defends itself with, AIAKMs, will the US be willing to commit human soldiers to fulfill treaty obligations? Failing to fulfil treaty obligations could jeopardize US military leadership worldwide.

The US may find itself in the unenviable position of being the junior partner in an alliance when an ally fields forces consisting of more capable AIAKMs. Since World War II, the US has enjoyed the position of being the senior partner in treaties and alliances. This leadership may be: explicit, such as in the North Atlantic Treaty Organization (NATO) where the US still holds the senior command billet; or implicit, such as the series of alliances throughout the Pacific where the US is the *de facto* lead. Fortunately, this has been a natural fit since the US has had the more powerful military in these relationships. If the US, however, bans the development of AIAKMs and its allies do not, it may find itself as the less capable partner or unable to fulfill its treaty obligations. Some alliance candidates in the Pacific that may be able to assume the leading role as a more capable military include the Republic of Korea and Japan.

Both the Republic of Korea and Japan have robust economies, aging populations, and a cultural affinity for robotics. In 2014, Samsung Electronics produced the SGR-1 robotic machine gun for service on the Korean De-Militarized Zone. The unit costs $200,000 and was immediately put into service.[74] While fielding this weapon may not be representative of the future AIAKM, it does demonstrate a willingness to take nascent steps. Today, Japan is devoting great sums of money to researching robotics for geriatric care, and that investment is projected to increase substantially over the next twenty years.[75] While geriatric care is not a weapons system, much of the technology may have some applications that can be easily militarized once perfected. Given current investments, it is conceivable that both of these countries would continue to develop these types of weapons systems, and do not appear to have any prohibitions on creating AIAKMs. Keeping in mind the mutual defense treaties that the US has with both countries, it is conceivable that the US could find itself obliged to commit its young men and women whilst either Japan, or the Republic of Korea, could commit machines. Ironically, the roles and missions of the machines may be of a lesser value. Additionally, the US soldiers may not be able to effectively work alongside AIAKMs.

The US may experience interoperability challenges as allied forces modernize their organizations with AIAKMs. Contemporary US command and control systems are designed to interface within a human being's effective span of control and organizational paradigm. AIAKMs may have an increased span of control, and radically different organizational paradigms. There is no current necessity for individual AIAKMs to hold rank or positions of authority based upon seniority or time in service; rank or position of authority may be a foundational argument when allied military forces interact. Additionally, given the likely processing capabilities of AI, it may not be effective or efficient for allied AIAKMs to interact with US forces composed of human beings or manned/un-manned teams.

Another significant implication of a US ban on AIAKMs is the potential that adversaries employing AIAKMs capability may be able to place the US nuclear deterrent at risk without fear of reprisal. Artificially intelligent swarms of nano-machines might approach undetected and pierce the ventilation system of Intercontinental Ballistic Missile (ICBM) silos, and then, either kill the silo crew or damage critical portions of the missile. Those systems could inflict similar damage on American bombers. Artificially intelligent nuclear powered attack submarines could conceivable follow American ballistic missile submarines indefinitely, as they wouldn't need to surface for food replenishment, crew swap or other constraints that manned submarines have. Given enough time, the artificially intelligent attack submarines could potentially conduct a coordinated attack to destroy this final arm of America's nuclear trident. If a coordinated attack was conducted on all three portions of the trident, this could leave the US vulnerable to a nuclear strike without a second strike capability and the Mutual Assured Destruction (MAD) of an adversary. In this regard, a US ban on AIAKMs could represent an existential threat.

The implications of a US ban on autonomous weapons are nearly all negative. Based solely on the nuclear scenario, the risk of not developing these systems is too high. Efforts to reach international consensus to ban AIAKMs may prove far more difficult than preventing nuclear proliferation; nuclear materials are traceable and can be detected. Further, the technologies involved with AIAKMs also have commercial applications that may make detection and compliance more difficult to define and gain international agreement. Even if a treaty can be developed, the barrios to entry are much lower for AIAKMs than they are for nuclear weapons. Non-state actors may be able to develop AIAKMs below the radar of international detection. With these implications in mind, a national strategy must be developed to effectively align the ends, ways, and means related to AIAKMs.

The Secretary of Defense should consider immediately rescinding DoD Directive 3000.09 *Autonomy in Weapons Systems*, which prevents the development, testing and employment of autonomous and semi-autonomous systems "that can independently select and discriminate targets."[76] Alternatively, the Secretary of Defense should encourage the development of these systems and enact policies that support their adoption. Further the US should gain and maintain the lead on development of AI. If it is not politically feasible to rescind the policy due to public sentiment, then the DoD should focus on supporting dual-use applications of AI that could be rapidly ported over to support future killing machines. This approach, however, is sub-optimal because it risks not having the AI capability on hand when it is needed. An argument could be made that the nation could develop this capability once it enters a protracted war. This argument is flawed because an AIAKM-equipped adversary may win the war before the US has the opportunity to fight a protracted war. Setting policy in place now, could lay the groundwork for what should be a massive, but precedent, undertaking to support the future defense of the nation in a new kind of warfare.

In 1939, German physicists learned to split a uranium atom. Two years later scientists convinced the US President that the destructive potential of the atomic bomb was so incredible that $2 billion dollars and 120,000 Americans were devoted to developing the first atomic bomb by 1945.[77] Had the Axis powers developed the atomic bomb first, the Allies may have lost the war. The destructive and revolutionary potential of AI is possibly more impactful than the atomic bomb. Therefore, it is recommended that the Nation invest heavily in research and development to ensure that the US is the first nation to gain and maintain the lead on true AI – this could be a project on the same scope and scale as the Manhattan Project. The parallels between the atomic bomb and AI are ominous. Whichever country achieves deep AI first will achieve a marked, and possibly unrecoverable, advantage over everyone else on the planet.

The previous Secretary of Defense, Ashton Carter, initiated proposals for public-private partnerships with Silicon Valley technology firms.[78] It is critical that these nascent efforts continue if the US is going to create a national effort for AI because the talent to create AI is largely in the private sector. Google is already working on what some have termed the "Manhattan Project of AI."[79] It is recommended that the DoD partner with the companies like Google on these types of projects, if Google, *et al* are willing. If commercial firms are not willing, then the US government should develop a "go it alone" plan to develop the AI Manhattan Project. The imperative behind this is made clear by an understanding of Kurzweil's Law – if the US does not have the lead, then a competitor with a year's more research could develop an AI overmatch. Investment in this type of project will be financially daunting, however, the risks or not doing so far more discouraging. The DoD could also assume risk on modernization of other programs that would become duplicative or obsolete when AIAKMs come to fruition and offset some of the financial burden.

While the future cannot be predicted with certainty, the trends described by futurists like Dr. Kurzweil clearly indicate that technology will continue to increase in computing power and decrease in cost. That trend may yield AI that, when combined with lethal weapons systems, may leave the human being unable to compete effectively against artificial intelligence automated killing machines on the battlefield. The potential of AI represents a revolutionary change to the character of warfare, with only one clear parallel – the atomic bomb. If the US elects to not participate, and more importantly develop it first, it will find itself a victim of those who do. This is a challenge the Nation cannot afford to ignore.

Endnotes

[1] Brainy Quote, "Stewart Brand Quotes," https://www.brainyquote.com/quotes/quotes/s/stewartbra172275.html (accessed March 16, 2017).

[2] Ashton Carter, *Autonomy in Weapon Systems,* Department of Defense Directive 3000.09 (Washington DC: U.S. Department of Defense, November 21, 2016), 2, http://www.dtic.mil/whs/directives/corres/pdf/300009p.pdf (accessed December 13, 2016).

[3] "50 Years of Moore's Law," http://www.intel.com/content/www/us/en/silicon-innovations/moores-law-technology.html (accessed December 6, 2016).

[4] Intel Corporation, "Moore's Law," https://docs.google.com/viewer?a=v&pid=sites&srcid=ZGVmYXVsdGRvbWFpbnxjd3VlZmM0MTZ8Z3g6MzI0MzYyMTkyZmZjM2I5Yw (accessed March 19, 2017).

[5] Ibid.

[6] Tony Simonite, "Moore's Law Is Dead. Now What?" *Technology Review,* May 13, 2016, https://www.technologyreview.com/s/601441/moores-law-is-dead-now-what/ (accessed December 11, 2016).

[7] M. Mitchell Waldrop, "The Chips Are Down For Moore's Law: The Semiconductor Industry Will Soon Abandon its Pursuit of Moore's Law," *Nature: International Weekly Journal of Science,* February 9, 2016, http://www.nature.com/news/the-chips-are-down-for-moore-s-law-1.19338 (accessed February 1, 2017).

[8] Raymond Kurzweil, "The Law of Accelerating Returns," http://www.kurzweilai.net /articles/art0134.html?printable=1 (accessed March 16, 2017).

[9] Raymond Kurzweil, "Moore's Law: The Fifth Paradigm," http://www.singularity.com/ charts/page67.html (accessed December 6, 2016).

[10] Raymond Kurzweil, *The Singularity is Near* (New York: Penguin Press, 2005), 112-113.

[11] University of Waterloo, "Quantum Computing 101," https://uwaterloo.ca/institute-for-quantum-computing/quantum-computing-101 (accessed February 1, 2016).

[12] Mark Fischetti, "Computers Versus Brains: Computers are Good at Storage and Speed, but Brains Maintain the Efficiency Lead," *Scientific American,* November 1, 2011, https://www.scientificamerican.com/article/computers-vs-brains/ (accessed February 1, 2017).

[13] Kurzweil, *The Singularity is Near*, 137-138.

[14] Ibid., 125.

[15] Kurzweil, "The Law of Accelerating Returns."

[16] Raymond Kurzweil, "Exponential Growth of Computing – Twentieth Through Twenty First Century," http://www.singularity.com/images/charts/ExponentialGrowthofComputing.jpg (accessed December 6, 2016).

[17] Hans Moravec, "The Future of Artificial Intelligence," *Scientific American*, March 23, 2009, https://www.scientificamerican.com/article/rise-of-the-robots/ (accessed December 11, 2016).

[18] Justin Phelps, "How to Enter Your PC's BIOS," *PC World,* October 5, 2011, http://www .pcworld.com/article/241032/how_to_enter_your_pcs_bios.html (accessed December 11, 2016).

[19] Neil Farber, "Decision-Making Made Ridiculously Simple! 8 Key Factors to Help You Make Those Tough Choices," *Psychology Today*, July 11, 2016, https://www.psychologytoday.com/blog/the-blame-game/201607/decision-making-made-ridiculously-simple (accessed December 12, 2016).

[20] Susan S. Lang, "Mindless Autopilot Drives People to Dramatically Underestimate How Many Daily Food Decisions They Make, Cornell Study Finds," *Cornell Chronicle*, December 22, 2006, http://news.cornell.edu/stories/2006/12/mindless-autopilot-drives-people-underestimate-food-decisions (accessed December 12, 2016).

[21] Brian D. Ripley, *Pattern Recognition and Neural Networks* (New York: Cambridge University Press, 1996), 1-3.

[22] Deepak Garg, "Are We Ready to Face the 'Deep Learning' Revolution?" *Times of India: Breaking Shackles Blog,* blog entry posted February 8, 2017, http://blogs.timesofindia.indiatimes.com/ breaking-shackles/are-we-ready-to-face-the-deep-learning-revolution/ (accessed February 9, 2017).

[23] Cade Metz, "AI Is About To Learn More Like Humans – With A Little Uncertainty," *Wired*, February 3, 2017, https://www.wired.com/2017/02/ai-learn-like-humans-little-uncertainty/ (accessed February 27, 2017).

[24] James Maguire, "Artificial Intelligence: When Will the Robots Rebel?" February 3, 2017, http://www.datamation.com/data-center/artificial-intelligence-when-will-the-robots-rebel.html (accessed February 9, 2017).

[25] *National Defense Authorization Act for Fiscal Year 2001,* Public Law 106-398, 106th Congress (October 30, 2000), Sections 220-221, http://www.dod.mil/dodgc/olc/docs/2001 NDAA.pdf (accessed March 19, 2017).

[26] U.S. Department of the Air Force, *United States Air Force Unmanned Aircraft Systems Flight Plan 2009-2047* (Washington, DC: U.S. Department of the Air Force, May 18, 2009), 3, https://fas.org/irp /program/collect/uas_2009.pdf (accessed March 28, 2017).

[27] Ibid., 14.

[28] Walter Pincus, "Air Force Training More Pilots for Drones Than for Manned Planes," *Small Wars Journal,* blog entry posted August 11, 2009, http://ebird.osd.mil/ebfiles/e200 90811695598.html (accessed March 19, 2017).

[29] U.S. Department of the Air Force, *United States Air Force Unmanned Aircraft Systems Flight Plan 2009-2047*, 50.

[30] Ibid.

[31] Christopher G. Pernin et al., *Lessons from the Army's Future Combat Systems Program* (Santa Monica, CA: RAND Corporation, 2012), iii, http://www.rand.org/pubs/monographs/MG1206.html (accessed March 28, 2017).

[32] Ibid., 285-286.

[33] Joe Pappalardo, "Why the Army Killed the Robotic Vehicle MULE," *Popular Mechanics,* August 3, 2011, http://www.popularmechanics.com/military/a6839/why-the-army-killed-the-robotic-vehicle-mule (accessed March 28, 2017).

[34] Ibid.

[35] Paul Szoldra, "The Marines Are Testing This Crazy Machine Gun-Wielding Death Robot," *Business Insider*, July 20, 2016, http://www.businessinsider.com/marines-testing-robot-2016-7 (accessed March 28, 2017).

[36] Mark Rutherford, "Robotics Rodeo Puts Umanned Tech Front and Center," September 3, 2009, http://news.cnet.com/8301-13639_3-10339238-42.html (accessed March 16, 2017).

[37] QintetiQ North America, "Modular Armed Advanced Robotic System," https://www.qinetiq-na.com/products/unmanned-systems/maars/ (accessed March 15, 2017).

[38] William Watkinson, "Russia Developing Killer Robot That 'Can Detect and Shoot a Human From 4 Miles Away'," *International Business Times,* November 12, 2016,

http://www.ibtimes.co.uk/russia-developing-killer-robot-that-can-detect-shoot-human-4-miles-away-1591224 (accessed December 12, 2016).

[39] Annie Jacobsen, "Inside the Pentagon's Efforts to Build a Killer Robot," *Time Magazine Online,* October 27, 2016, http://time.com/4078877/darpa-the-pentagons-brain/ (accessed December 7, 2016).

[40] Kris Osborn, "F-35 Will Include Artificial Intelligence," *RealClear Defense,* January 21, 2017, http://www.realcleardefense.com/2017/01/21/f-35will_include_artificial_intelligence_289640.html (accessed February 1, 2017).

[41] Human Rights Watch, "Killer Robots and the Concept of Meaningful Human Control: Memorandum to Convention on Conventional Weapons (CCW) Delegates," April 11, 2016, https://www.hrw.org/news/2016/04/11/killer-robots-and-concept-meaningful-human-control (accessed December 11, 2016).

[42] Parliament of the United Kingdom, "House of Commons Hansard Debates for 17 Jun 2013," June 17, 2013, http://www.publications.parliament.uk/pa/cm201314/ cmhansrd/ cm130617/debtext/130617-0004.htm (accessed December 12, 2016).

[43] Steve Ranger, "Robots of Death, Robots of Love: The Reality of Android Soldiers and Why Laws for Robots are Doomed to Failure," *The Tech Republic,* http://www.techrepublic.com/article/robots-of-death-robots-of-love-the-reality-of-android-soldiers-and-why-laws-for-robots-are-doomed-to-failure/ (accessed December 12, 2016).

[44] Carter, *Autonomy in Weapons Systems*, Department of Defense Directive 3000.09, 2.

[45] Ibid., 3.

[46] Sydney J. Freedberg Jr., "Should US Unleash War Robots? Frank Kendall Vs. Bob Work, Army," *Breaking Defense,* August 16, 2016, http://breakingdefense.com/2016/08/should-us-unleash-war-robots-frank-kendall-vs-bob-work-army/ (accessed December 4, 2016).

[47] Sydney J. Freedberg Jr. and Colin Clark, "Killer Robots? 'Never,' Defense Secretary Carter Says," *Breaking Defense,* September 15, 2016, http://breakingdefense.com/2016 /09/killer-robots-never-says-defense-secretary-carter/ (accessed November 22, 2016).

[48] Matt Novak, "Elon Musk and Stephen Hawking Call for Ban on Autonomous Military Robots," *Gizmodo,* July 27, 2015, http://paleofuture.gizmodo.com/elon-musk-and-stephen-hawking-call-for-ban-on-autonomou-1720383277 (accessed March 15, 2017).

[49] Eyder Peralta, "Weighing the Good and the Bad of Autonomous Killer Robots in Battle," *National Public Radio,* April 28, 2016, http://www.npr.org/sections/alltechconsidered /2016/04/28/476055707/weighing-the-good-and-the-bad-of-autonomous-killer-robots-in-battle (accessed December 12, 2016).

[50] National Highway Traffic Safety Administration (NHTSA), "Fatality Analysis Reporting System (FARS) Encyclopedia," https://www-fars.nhtsa.dot.gov/Main/index.aspx (accessed February 12, 2017).

[51] NPR Staff, “Google Makes the Case for A Hands-Off Approach to Self-Driving Cars,” *National Public Radio: All Tech Considered,* February 24, 2016, http://www.npr.org/sections/alltech considered/2016/02/24/467983440/google-makes-the-case-for-a-hands-off-approach-to-self-driving-cars (accessed February 12, 2017).

[52] National Highway Traffic Safety Administration (NHTSA), “Quick Facts 2015,” December 2016, https://crashstats.nhtsa.dot.gov/Api/Public/ViewPublication/812348 (accessed February 27, 2017).

[53] James M. Anderson et al., *Autonomous Vehicle Technology: A Guide for Policymakers* (Santa Monica, CA: RAND Corporation, 2016), xix, http://www.rand.org/content/dam/rand/pubs/research_reports/RR400/RR443-2/RAND_RR443-2.pdf (accessed February 11, 2017).

[54] Ibid., 14.

[55] Russ Mitchell, “U.S. Ends Investigation of Fatal Tesla Crash and Finds ‘No Safety Defects’ in Car’s Autopilot,” *The Los Angeles Times Online,* January 19, 2017, http://www.latimes. com/business/autos/la-fi-hy-tesla-autopilot-20170119-story.html (accessed February 11, 2017).

[56] Anderson et al., *Autonomous Vehicle Technology: A Guide for Policymakers*, 16.

[57] Official U.S. Government Website for Distracted Driving, “Facts and Statistics,” https :// www.distraction.gov/stats-research-laws/facts-and-statistics.html (accessed February 27, 2017).

[58] Nancy Benac, “The Long, Unfortunate History of Friendly Fire Accidents in U.S. Conflicts,” *PBS Newshour - The Rundown: A Blog of News and Insight*, June 11, 2014, http://www.pbs.org/newshour/rundown/long-unfortunate-history-friendly-fire-accidents-u-s-conflicts/ (accessed February 11, 2017).

[59] “How Iran Hacked Super-Secret CIA Stealth Drone,” *RT,* December 15, 2011, https://www.rt.com/usa/iran-drone-hack-stealth-943/ (accessed March 28, 2017).

[60] Katia Moskvitch, “Are Drones the Next Target for Hackers?” *BBC Future*, February 6, 2014, https://www.rt.com/usa/iran-drone-hack-stealth-943/ (accessed March 28, 2017).

[61] Visually, “One of Our Drones is Missing,” http://visual.ly/rq-170-sentinel-spy-drone (accessed March 28, 2017).

[62] Anthony Capaccio, “Trump’s Bigger Army Could Cost $12 Billion by Fanning’s Math,” *Bloomberg Politics*, January 19, 2017, https://www.bloomberg.com/politics/articles /2017-01-19 /trump-s-bigger-army-could-cost-12-billion-by-fanning-s-math (accessed February 27, 2017).

[63] Veronique de Rugy, “Personnel Costs May Overwhelm Department of Defense Budget,” February 4, 2014, https://www.mercatus .org/publication/personnel-costs-may-overwhelm-department-defense-budget (accessed March 28, 2017).

[64] U.S. Department Of Veterans Affairs, Life Insurance, “Service Members’ & Veterans’ Group Life Insurance,” http://benefits.va.gov/insurance/sgli.asp (accessed March 19, 2017).

[65] Department of Defense Military Compensation, "Death Gratuity," http://militarypay.defense.gov/Benefits/Death-Gratuity/ (accessed March 19, 2017).

[66] Tom Venden Brook, "Aghan IEDs: Warfare on the Cheap," *USA Today Online*, June 24, 2013, http://www.usatoday.com/story/nation/2013/06/24/price-list-for-ieds-shows-how-insurgents-fight-on-a-budget/2447471/ (accessed March 28, 2017).

[67] J&G Sales, "J&G Sales: Firearms / Ammunition / Accessories," http: //www.jgsales.com (accessed March 28, 2017).

[68] Smashing Robotics, "Thirteen Advanced Humanoid Robots for Sale Today," April 16, 2016, https://www.smashingrobotics.com/thirteen-advanced-humanoid-robots-for-sale-today/ (accessed February 27, 2017).

[69] Robotics Today, "HRP Series: HRP 4," http://www.roboticstoday.com/robots/hrp-4 (accessed February 27, 2017).

[70] "Bin-Laden: Goal is to Bankrupt U.S.," *CNN,* November 1, 2004, http://www.cnn.com/2004/WORLD/meast/11/01/binladen.tape/ (accessed March 28, 2017).

[71] Ibid.

[72] Tom Philpott, "Attractions of Service Obscured by Hard Costs of War," *Stars and Stripes Online*, August 27, 2015, http://www.stripes.com/news/us/attractions-of-service-obscured-by-hard-costs-of-war-1.364957 (accessed February 13, 2017).

[73] Barbara A. Bicksler and Lisa G. Nolan, *Recruiting an All-Volunteer Force: The Need for Sustained Investment in Recruiting Resources – An Update* (Arlington, VA: Strategic Analysis Incorporated, 2009), 9, http://www.people.mil/ Portals/56/Documents/MPP/AP/Bicksler_ Recruiting_Paper_2009.pdf (accessed March 28, 2017).

[74] Mark Prigg, "Who Goes There? Samsung Unveils Robot Sentry That Can Kill From Two Miles Away," *The Daily Mail*, September 15, 2014, http://www.dailymail.co.uk/sciencetech/ article-2756847/Who-goes-Samsung-reveals-robot-sentry-set-eye-North-Korea.html (accessed February 27, 2017).

[75] Danielle Muoio, "Japan is Running Out of People to Take Care of the Elderly, So It's Making Robots Instead," *Business Insider,* November 20, 2015, http://www.businessinsider.com /japan-developing-carebots-for-elderly-care-2015-11 (accessed February 27, 2017).

[76] Carter, *Autonomy in Weapons Systems,* Department of Defense Directive 3000.09, 1.

[77] U.S. History, "The Manhattan Project," http://www.ushistory.org/us/51f.asp (accessed February 27, 2017).

[78] Mark Pomerleau, "What Stands in the Way of DOD-Silicon Valley Partnership?" *Defense Systems*, May 13, 2015, https://defensesystems.com/articles/2015/05/13/obstacles-to-dod-sillicon-valley-partnership.aspx (accessed February 27, 2017).

[79] Dan Rowinski, "Google's Game of Money Ball in the Age of Artificial Intelligence," *Readwrite*, January 29, 2014, http://readwrite.com/2014/01/29/google-artificial-intelligence-robots-cognitive-computing-moneyball/ (accessed February 27, 2017).