

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE (DD-MM-YYYY) 01-05-2018		2. REPORT TYPE Future War Paper		3. DATES COVERED (From - To) Jul 2017 - May 2018	
4. TITLE AND SUBTITLE Synthetic Cognition: The Implications of Artificial Intelligence on Marine Corps Combat Effectiveness at the Tactical and Operational Levels of War				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Gruba, Matthew, J				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) School of Advanced Warfighting 2076 South Street Quantico, VA 22134				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Marine Corps University 2076 South Street Quantico, VA 22134				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Artificial intelligence (AI), or intelligent behavior by machines, is a decades-old concept that is now rapidly maturing as advances in computing power catch up to theoretical applications. AI has the potential to dramatically improve efficiencies and lethality over the next decade for organizations throughout the Department of Defense (DoD). The Marine Corps, as the smallest service in the DoD, is positioned to significantly benefit from the cost savings and operational enhancements that AI will soon offer at the tactical and operational levels of war. AI could benefit the fleet in many areas, including scheduling optimization, target selection and image processing, and wargaming simulation. Technologically, these AI enhancements rely only on reasonable extensions of existing theory and application, and require no major new technological breakthroughs to implement over the next ten to fifteen years. Ongoing national AI policy development and continued emphasis by senior Marine Corps leaders will both be essential to AI effective implementation relative to global adversaries in the years to come.					
15. SUBJECT TERMS Artificial Intelligence, United States Marine Corps, Future Warfare, Innovation, Neural Network, Computer Science, Decision Support, Wargaming, Scheduling, Targeting					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 22	19a. NAME OF RESPONSIBLE PERSON Lynn, Charles, LtCol, USMC
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (include area code) 703-784-6800

*United States Marine Corps
School of Advanced Warfighting
Marine Corps University
2076 South Street
Marine Corps Combat Development Command
Quantico VA 22134*

FUTURE WAR PAPER

Synthetic Cognition: The Implications of Artificial Intelligence on Marine Corps Combat Effectiveness at the Tactical and Operational Levels of War

**SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF OPERATIONAL STUDIES**

Major Matthew J. Gruba

AY 2017-18

Mentor: Bradley Meyer, PhD

Approved: *Bradley J. Meyer Ph.D*

Date: *15 May 2018*

DISCLAIMER

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

Artificial intelligence (AI), or intelligent behavior by machines, is a decades-old concept that is now rapidly maturing as advances in computing power catch up to its theoretical applications. It comprises a broad range of capabilities, but may generally be described as the derivation, via automated processing, of operational or computational models by detecting patterns in large amounts of data or in the interactions of large numbers of variables. Current applications in the civilian world are myriad, and include predicting consumer preferences based on web-surfing behavior, directing autonomous vehicles, identifying faces, and trading in securities.

The private sector, motivated by competitive market forces, continues to aggressively leverage advances in AI to create new products (both evolutionary and revolutionary) and generate operational process efficiencies. Both large and small organizations can access AI's benefits; because AI is software-based, implementation costs are low. Within the Department of Defense (DoD), innovation in AI is certainly a persistent research and development area, but AI has yet to become a significant daily presence at the operational or tactical levels of most warfighting components.

AI development takes place in the cyber realm, but its eventual full application has the potential to benefit, in one way or another, every organization in the DoD. Leaders must recognize this to properly conceptualize of AI as a force multiplier across all warfighting functions, and not merely as an esoteric tool for cyber specialists. The Marine Corps, as the smallest service in the DoD, is ideally positioned to exploit the cost savings and operational enhancements that AI will offer in coming decades.

AI is perhaps the ultimate technological manifestation of the Marine Corps' hallmark of doing more with less. As Moore's Law continues to increase the computing power of local, end-user devices and workstations, smaller operational and tactical units are increasingly able to access AI functionality. Because AI processes typically involve the manipulation of very large data sets and multiple variables, they require a substantial amount of computing power (and, in the field, battery power) to complete useful tasks within a useful amount of time. Early AI users had to employ large mainframe computers, minicomputers, or servers to accomplish this. As recently as a decade ago, the Marine Corps had little use for AI at the tactical level because contemporary tactical systems lacked the necessary processing power and memory to support it.

Modern desktop workstations, which are ubiquitous across deployed and garrisoned Marine forces, are now much closer to possessing both the onboard processing capacity to execute AI functions and the network connectivity to interface with cloud AI services. This means that the cost of implementing tactical AI over the next one to two decades would be relatively low, with little, if any, additional hardware required beyond normal, ongoing upgrades. Costs would be limited primarily to two areas: software development or licensing, and training and rollout. These would accrue over several years leading up to full implementation. Tablets and other field-expedient warfighter systems already in use could both receive AI output and provide data-point input in near- or real-time via current and near-next-generation tactical networks. Operational-level AI applications would be similarly inexpensive, and could easily leverage the substantial command and control assets of a typical divisional (or higher) command operations center and G-6 department.

A multitude of problems exist across the Marine Corps that AI could address. They typically involve situations in which there is a complicated rule-based or constraint-based production or evaluation process that is a part of a unit's ongoing battle rhythm. At the tactical level, such problems span the administrative, combat, and troop welfare realms. In the administration of flying squadrons, for instance, pilots assigned to the S-3 department devote thousands of hours each year to creating flight schedules, which require the synthesis of countless variables, including aircraft availability, crewing rules, medical statuses, flight training syllabi, and flight qualifications. This requires a large S-3 staff to ensure schedule legality, safety, and flight-hour efficiency.

In the combat operations realm, continuous pattern analysis of the massive amounts of data that modern sensor systems generate currently overwhelms the capacity of human processors, and commands often cannot fully leverage organic "big data" collection capabilities and rapidly relay useful interpretations back to the warfighter. Previous generations of warfighters bemoaned the dearth of timely battlefield information; today, they are drowning in it.

At the operational level of war, the Marine Corps' wargaming of critical operational plans and orders is limited by the knowledge, expertise, placement, and schema of the planning team and associated role-players. Current wargaming techniques are, in effect, manual simulations. They take time, and are necessarily simplified representations of reality. Wargames help to reveal holes in a plan, but with a more robust injection of real-world variables (a level of complexity too ungainly for current methods), they could become accurate enough to also be used as predictive tools.

The Obama Administration published a 40-page national research and development strategic plan for AI in 2016. However, it is only a framework, providing broad recommendations about identifying AI opportunities, coordinating efforts, and creating an AI workforce.¹ It has yet to be substantially developed. In contrast, China has identified AI as a national strategic priority, and appears to be spending billions more on it than is the U.S. For several years, Chinese researchers have produced more deep learning research papers than their American counterparts.² AI is, in China's view, the "next big thing," and presents a substantial emergent opportunity to gain a strategic advantage against global competitors. The implications for U.S. national defense are self-evident, and there is much that the DoD (and the Marine Corps) can do in response.

Although AI has yet to reach the fleet in any broad, comprehensive way, it could easily do so over the next ten to fifteen years based on current and near-future Marine Corps assets and infrastructure. It would not require additional scientific breakthroughs or new technologies. Basic elements of the AI capabilities that could benefit the Marine Corps already exist. The chief change would be organizational: existing standard operating procedures carry a good degree of inertia, and augmenting them with AI, while dramatically increasing efficiency and lethality in the long term, would require an investment of training time during the initial rollout years. However, the relatively low cost of AI implementation, coupled with its potential to dramatically improve mission effectiveness and troop welfare, would generate a return-on-investment that few other innovations could match in the years to come. The Marine Corps stands on the cusp of an AI revolution in military affairs, and in order to maintain its future lethality must avail itself, fiscally and philosophically, of the array of new capabilities that recent and emerging

technologies have made feasible in the field of AI and that have manifold applications across the tactical and operational levels of war.

AI comprises a large constellation of varied technologies and sub-fields. Broadly defined, it describes automated processes that enable human-like behavior in man-made platforms. An example is machine learning, which involves completing operations or tasks without specific programming direction on precisely how to do so. In many respects, AI is an attempt to mimic the operational patterns of the human brain, which is perhaps the most complex network known to man.

The brain comprises about 86 billion neurons, each with hundreds of thousands of synapses, creating trillions of functional contacts.³ No two human brains are entirely alike; they all possess similar components, but there is no universal blueprint for all of these components down to the cellular or neuronal level. In short, the neural pathways that enable someone to, say, navigate their way to work every day using their normal route will be arranged entirely differently than those of another person with the identical commute. People of similar cognitive capability will manage this task with correspondingly similar success, despite independently evolving differing neural architectures with which to get there.

Next, compare the human brain's organically-evolved architecture to that of the formally-engineered modern microprocessor. Continuing manufacturing advances in microchip wafer photolithography are pushing miniaturization down to the five-to-seven-nanometer range. This means that the paths that electrical signals travel on a chip are as narrow as 5 billionths of a

meter. By comparison, a strand of human hair is 80,000-100,000 nanometers thick, and a brain synapse is 20-40 nanometers wide.⁴ A schematic of a microchip the size of a small fingernail would need to cover an entire wall to be of a visually useful scale.

Unlike the brain, a microchip's structure does not grow and evolve. Developers must pre-engineer and then manufacture the chip in whole. They must lay out every gate, node, pathway, and semiconductor in microscopically precise fashion, standardized for every chip in that line, for the chip and its supported processor to function. A small design error or manufacturing inconsistency could render the entire processor inoperative. The chip possesses no organic workaround or evolutionary learning capacity when this happens.

Much like the hardware of the microprocessor, the programming that drives it is typically very rigid. The programmer must clearly define the operation with discrete inputs and outputs. The algorithm is mathematical and proceeds in a conceptually linear, iterative, or stepwise fashion, via devices like conditional functions (*if this, then that*) and loop processing (*do this until that*). A program doesn't change unless a programmer manually updates it. This means that the program can do one thing or one set of things very effectively—and generally with remarkable reliability compared to a human—but has virtually no capacity to adjust to changing circumstances or inputs. For example, a simple program to distinguish the shapes of two types of objects from each other, say, red pyramids and a blue blocks, might use color as a determinant (*if the object is red, then it must be a pyramid*). But what if the program encounters a blue pyramid? As written, it would be unable to generate a correct response. The program would fatally fail. But a 5-year-old human child who had never seen a blue pyramid (only red pyramids and blue

blocks) would likely still be able to correctly state that it is a pyramid. (A skill which, as an adult, she might use to visually distinguish between a journalist shouldering a camera and an attacker shouldering a missile, with varying success.)

The child's brain, in the absence of an expected input, possesses the neurological dexterity to draw from other sources of information to arrive at the correct interpretation. By age 5, the child would likely know the difference between a shape and a color *without necessarily being explicitly taught that difference*. The brain achieves these logical and computational jumps as it interprets something new by using its network of neurons and synapses.

The neural net concept underpins AI theory and is a simplified version of the human brain's network. It is the reason behind AI's considerable potential impact on the future warfighter. Neural nets are not new. Developers began work on the earliest one, called the Perceptron, in the 1950s. But it had only two layers of neurons (an input layer and an output layer), which mathematically limited the complexity of operations it could perform. Theorists knew that adding additional intermediate neural layers would dramatically increase the neural net's effectiveness, but didn't know how to train those layers.⁵ As a result, throughout the 1960s and 1970s, most researchers turned their focus away from neural net development.

Two developments reversed this trend and drove what has become the AI revolution. First, in 1986, researchers achieved a breakthrough in thinking about how to incorporate intermediate layers into a neural net. Second, advances in computational power provided the raw processing power that a neural net requires during execution for it to be of sufficient value.




Figure 1: Neural net layer diagram from seminal 1986 backwards-propagation research by David Rumelhart, Geoffrey Hinton, and Ronald Williams.⁶

In 1986, David Rumelhart, Geoffrey Hinton, and Ronald Williams developed an approach called backwards propagation, or backprop, to enable the deep learning that is the basis for AI.⁷ It is employed during a neural net's training phase and works best with very large data sets. Backprop starts at the output layer by looking at whether the output was correct. For example, does an image show a person carrying a shoulder-fired missile, or not? (If not, the person might be shouldering a video camera, carrying nothing at all, etc.) If the output was incorrect, the backprop process works in reverse, down through the intermediate neuronal layers, identifying the degree to which each neuronal layer contributed to the error and, based on this, applying a correction by adjusting the weights of the connections between those neuronal layers.⁸ Using

large volumes of training data, the backprop process eventually trains the neural net to recognize the desired pattern. This process describes the essence of machine learning.

Backprop works quite well on applications like image processing. But given the tremendous pixel-level variability of multiple images showing the same object, and given the number of (and relationship between) objects that a user might seek to drive a decision cycle, the process demands considerable computing power, far in excess of what was available even five to ten years ago. Today, advances in photolithography, cooling technology, photonics, and even new progress in quantum computing all contribute to making AI computational processes much more executable than they were a generation ago. Furthermore, progress in bandwidth availability and optimization allow platforms in austere locations to upload data via the cloud for processing.

AI continues to evolve apace. But the technology is already at a sufficient stage of development to benefit warfighters in measurable, concrete ways. Processing power has finally caught up to theory, thereby enabling execution. We now have the programming knowledge and computational ability to bridge the gap between the organic learning/evolving nature of the human brain and the rigid architecture of the engineered microchip. The Department of Defense (DoD) is positioned to benefit considerably by implementing AI in the tactical and operational realms, given AI's relatively low cost (because it is software-based) and its dramatic processing capability. What follows are some possible applications of AI within the Marine Corps that the organization could reasonably implement across the fleet over the next ten to fifteen years. They illustrate the remarkable advances that are possible in the relatively near term; advances that potential adversaries are also aggressively pursuing.

While the most visible (and feted) Marine Corps achievements tend to occur during kinetic combat operations, such achievements represent only a small portion of our ongoing activities as a fighting force. Success in combat relies heavily on readiness, which is why Marine leaders have always emphasized the importance of training. Developing proficiency in mission-essential tasks and warfighting skills requires realistic, repetitive training, which units must accomplish within the constraints of time and money. Marine aviation offers an excellent example of these considerations, and of how AI could dramatically improve the efficiency with which aviators gain and maintain mission proficiency in today's resource-constrained environment.

Like all Marine units, fleet Marine flying squadrons develop annual training plans, from which they derive quarterly and monthly training plans. Squadron operations departments then generate weekly and daily flight schedules to execute training and support flights. These schedules are based on a squadron's annual allocated flight hours, which are in turn based on aircraft and crews assigned, expected maintenance readiness, and other factors, all of which effectively serve as constraints on the total amount of flight training time available to develop and maintain flight crew proficiency in assigned combat mission sets.

Within each squadron, operations personnel devote thousands of hours each year to generating and updating flight schedules. This occupies the time of several aviators throughout each day and requires considerable training and experience due to the liability involved: the daily flight schedule is a signed order by the commanding officer, who certifies that every line of the flight schedule is legal and complies with all flight regulations, higher orders, and standard operating

procedures (SOPs). Daily schedule writers are usually junior copilots, who are most in need of study and preparation time for their syllabus flights.

Generating a flight schedule is, essentially, a complex optimization problem. It involves multiple inputs (allocated flight hours, aircraft available, crews available, daily flight windows, moon phase, etc.) that are constrained by both practical and regulatory factors (maintenance cycles, specific designations of each crew member for a given mission type, medical readiness, crew completion of prerequisite flights the day before, crew day and crew rest rules, crew currency, time of year, etc.). Remarkably, aside from the adoption of database platforms, the schedule writing process has remained largely unchanged for decades. A pilot (usually the squadron's pilot training officer) fills out the weekly flight schedule on a spreadsheet and plugs in flights and crews to try to maximize the utility of the coming week's available flight hours. He bases this on the training needs of the pilots and aircrew (dozens of people, nearly all of whom are progressing through at least one training syllabus), while not violating hundreds of flight scheduling regulations promulgated by upwards of ten sources. Mistakenly scheduling a crew member who is not current or qualified for a given mission could endanger the crew and would make the daily flight schedule (derived from the weekly schedule) an illegal order from the commanding officer. Also, writing a schedule that "satisfices"—rather than fully optimizes—the squadron's training time will, by the end of the year, result in less flight time and slower syllabus progression for aviators, who are already frequently shortchanged by low fleet-wide aircraft readiness rates.

Marine Corps flying squadrons have been using an automated flight scheduling database platform for years, but it provides only a limited error-checking validator and does not perform what takes most of a schedule writer's time: optimizing the scheduling of every crew member on every flight to fully maximize the training benefit for the squadron's crewmember population over the training year. That is a far more complex task than simple database management. It is the type of processes that, in previous generations, only humans could do (and do imperfectly). But this sort of multivariate optimization problem could be easily tackled by an AI-based software application. A common AI programming language like Python, with help from existing commercial-off-the-shelf applications like IBM's CPLEX Optimizer, would be well-suited to this task. Just as neural nets accurately identify an image after analyzing the pixels of many such images, AI programming can accurately create an optimized schedule based on a large constellation of learned constraints and desired outcomes.

The enhanced capabilities of such an AI-based optimization tool beyond what human schedule writers currently provide would serve as force multipliers at squadrons throughout the fleet. Currently, every time an event changes on a daily flight schedule due to maintenance, weather, or crew issues, the upcoming daily and weekly schedules need to be manually readjusted by schedule writers to reflect the change, often on short-notice. For example, if a flight is cancelled that is a prerequisite for a subsequent syllabus flight the following day that is part of a multiple-aircraft event, it could cause major shifts in that event, or possibly its cancellation, which could, in turn, impact scheduled syllabus events in subsequent days. With such changes happening on the fly (not planned on the weekly schedule), the potential for errors is higher in the revised daily schedules. If an AI system managed the schedule optimization, it could automatically and

continuously run re-optimizations of subsequent flight schedules as the duty officer entered daily flight completion information into the system. Further, the optimizer would not only update and re-optimize the next day's and current week's schedule, but all such schedules going as far into the future as the user desired. An entire year's worth of daily flight schedules could already exist, all of them continuously updated as the squadron completed each day's flight operations.

An AI scheduler would add enhanced capabilities while simultaneously reducing administrative schedule-writing time by over 70,000 hours per year across the fleet, based on a very conservative estimate of 20 hours of schedule work per week per squadron. That is the equivalent of about 35 full-time positions. This would allow more time for schedule writers to fly and to focus on flight syllabus studies and tactics development, improving the Marine Corps' return on investment on its aviators.

Turning to the application of machine learning in combat operations, a major focus of both public and DoD attention continues to be the role of AI in analyzing the wealth of data supplied by imagery and sensor systems, chiefly those aboard unmanned aerial vehicles (UAVs). In 2011 alone, the U.S. Air Force collected 325,000 hours (about 37 years) of video using its drones. The limiting factor in imagery analysis is no longer the availability of imagery. Rather, it is the availability of human processing capacity necessary to analyze it. "Information overload" caused the deaths of 23 Afghan civilians from a drone strike in 2010, during which decision-makers could not adequately process the volume of incoming video.⁹

The neural net architecture of AI, along with backprop procedures for machine learning, are well suited to imagery analysis. Bringing AI into the analysis process is the only realistic hope that the DoD has of comprehensively examining massive, continuous flow of video and pictures from globally distributed, persistent sensors. There is simply an insufficient number of human analysts—by multiple orders of magnitude—on hand to analyze the data, and the problem will only get worse as our data sources get increasingly robust. The Marine Corps, as the smallest service, has the potential to feel this shortfall the most acutely, and is best positioned to benefit from the enhanced capabilities that AI could provide in this area over the next ten to fifteen years.

A given neural net learns to determine whether what it sees in an image is a specific thing that it is designed to look for (Is this a shoulder-fired missile, or isn't it?). That task alone is complex enough, but Marines are interested in far more than the presence of a single type of weapon system. Decision-makers want to know whether the movement pattern of such a system and adjacent personnel indicate its level of readiness, where it might be going next, and (for rules-of-engagement decisions) the likelihood of its imminent use. An AI system would need to collect and process this information about hundreds of types of weapons and enemy activities and then present human monitors with extracted targeting information, which the system could collate based on things like a target's priority, the confidence level of the AI's assessment of a target, the imminency of a targeted weapon's use against friendly personnel, etc. These variables would then inform the human decision-maker about whether and how to engage the target. The size, complexity, and number of neural nets that warfighters require are considerable if the nets are to provide human-like (or better) analysis; recall that the human brain contains trillions of

functional contacts. AI nets seeking to mimic or surpass human computational abilities are staggeringly extensive.

Raw computational power to execute these neural nets is not lacking, but it doesn't reside within the sensor system. The sensor must transmit its data to the AI system for analysis. In the case of live video, this requires substantial and persistent bandwidth. The data stream must also provide sufficient throughput speed to deliver real-time or near-real-time analysis and facilitate time-sensitive decisions, whether those decisions are automated or made by a human in the decision loop. Consequently, alongside sensor and AI net development, the DoD needs to continue to improve onboard data compression capabilities, which are rapidly becoming the new limiting factor in timely AI imagery analysis.¹⁰ Without them, the Marine Corps must augment its organic data connections with extremely costly commercial satellite bandwidth.

An additional challenge to full implementation of AI for imagery analysis is the fact that neural nets, despite their significant advances in recent years, are still relatively easy to fool. Recently, a research group at MIT called Labsix was able to trick Google's Cloud Vision API image classifier into misclassifying a group of skiers as a dog. Interestingly, the team had no proprietary knowledge of exactly how Google's classifier worked. By gradually altering pixels in each image fed into the classifier, the team caused the classifier to follow a machine-learning process that led to an incorrect answer, which the classifier produced well before the image was sufficiently changed to render the skiers unrecognizable to a human observer.¹¹ A similar misidentification on the battlefield of a non-weapon as a weapon could have grave consequences.

Despite these ongoing challenges, the DoD is making rudimentary progress on a viable AI imagery analysis system. Project Maven, a pilot project aimed at furthering AI implementation across the DoD, began in late-April 2017. Initially comprising a team of only six members (including Marine Corps representation), the project had a working system analyzing live video data from a number of UAV platforms within six months of inception. It benefitted from a narrowly-defined target set (counter-ISIS activities) and extensive interaction with civilian AI advisors; indeed, none of the team members had technical experience in AI.¹² The follow-on activity within the DoD's primary development communities that Maven could spur will likely shorten rollout time for a fleet-ready systems in the years to come, but only if the DoD retains AI (and technological development in general) in its strategic planning focus.

An additional way for the Marine Corps to leverage AI's potential is by capitalizing on AI's most fundamental aim: to make a machine behave like a human. At the tactical and operational levels of war, the Marine Corps Planning Process (MCP) serves as the primary tool to develop courses of action (COAs) and vet those COAs based upon how the enemy is likely to respond to them. The wargaming step of MCP facilitates this by assigning one or more members of the planning team to play the adversary and react based on the team's analysis of the enemy during the problem framing step, which requires assumptions about the enemy's doctrine (and degree of adherence to it), creativity, and mindset. The inherent limitation to this approach is that the adversary role-players come from the same organization as the friendly forces, and despite their best efforts will still be influenced to some degree by their own norms, biases, and assumptions, which may differ considerably from those of the actual enemy and produce wargaming movements inconsistent with enemy behavior in execution.

AI provides a method to dramatically increase independent “thought” on the part of the wargaming adversary by either augmenting the adversary role players or replacing them completely with a computer-generated enemy. An AI neural net, just like an actual enemy on the battlefield, learns from previous experiences and adapts its assessments and subsequent behaviors without specific direction from planning team members or computer programmers. Deep learning machines have been demonstrating their effectiveness against chess grand masters for several years; similar algorithms would be useful in the far less structured (and more complex) realm of wargaming and could reasonably become viable over the next ten to fifteen years.

An excellent primer for approaching AI use in wargaming already exists. In the early 2000s, a joint French and Brazilian computer science research team collaborated on a short study to determine how to develop AI heuristics for gameplay behavior in a popular Napoleonic turn-based battle simulation called Battleground. The aim of the study was to find ways to apply more abstract (human-like) decision-making concepts in gameplay to make it more realistic. The researchers sought to augment existing AI planning and learning concepts with *parallel* processing (i.e., accurately portraying behaviors and decisions and each level of command concurrently within the game simulation). This would simulate how a commander’s orders and intent would be carried out by subordinate units down to the small-unit level. The research demonstrated a conceptual basis for applying abstraction within the action space to convert high-level orders into tactically concise instructions for subordinates, making the simulated opponent far more human-like in behavior.¹³

In the coming years, the Marine Corps could easily build from this existing research to develop a human-like artificial adversary for more objective enemy actions. In their study, the researchers developed different heuristics for different tactical priorities (based on Napoleonic-era tactics): speed, stealth, safety, or a combination of these. In a similar fashion, the Marine Corps could develop a simulation platform that would broadly direct enemy behavior consistent with the enemy's assessed most likely or most dangerous course of action, and then allow the AI to play out its moves independently and organically, based on that initial guidance. The development costs of such a platform would be relatively small, and training commands could eventually incorporate it as a software module into the Marine Corps' existing simulation architecture.

These examples of future AI integration across the Marine Corps are by no means exhaustive, and represent only a sample of the numerous means through which AI could benefit the fleet. Technologically, they rely only on reasonable extensions of existing theory and application, and require no major new breakthroughs to implement over the next fifteen years. Like all innovations, one of AI's greatest barriers is likely to be institutional inertia, and the unwillingness of some to abandon familiar, tested methods in favor of new ones. For this reason, the design of the rollout and training plan for new AI applications will be nearly as important as the technologies themselves. The aggressive adoption of AI for use in espionage and warfare by our rivals and potential adversaries, including Russia and China, highlights the criticality of our continued focus on these emerging technologies.¹⁴ AI, along with all technological development, must remain a strategic and budgetary priority for the DoD in the years to come, and the Marine Corps must strenuously lobby for investments in AI that could provide the smallest of the

services with a significant asymmetrical advantage—at a fairly low cost—relative to its larger interservice competitors.

¹ The White House, *The National Artificial Intelligence Research and Development Strategic Plan*, (Washington: Government Printing Office, 2016), 37.

² Cade Metz, “As China Marches Forward on A.I., the White House Is Silent,” *The New York Times*, last modified February 12, 2018, <https://www.nytimes.com/2018/02/12/technology/china-trump-artificial-intelligence.html>.

³ Kayt Sukel, “The Synapse—A Primer,” *The Dana Foundation*, last modified March 15, 2011, <http://www.dana.org/News/Details.aspx?id=43512>.

⁴ Sukel, “The Synapse—A Primer,” <http://www.dana.org/News/Details.aspx?id=43512>.

⁵ James Somers, “Is AI Riding a One-Trick Pony?,” *MIT Technology Review* 120, no. 6 (2017): 30.

⁶ Somers, “Is AI Riding a One-Trick Pony?,” 31.

⁷ David Rumelhart, Geoffrey Hinton, and Ronald Williams, “Learning Representations by Back-Propagating Errors,” *Nature*, 323 (1986): 533-536.

⁸ Somers, “Is AI Riding a One-Trick Pony?,” 33.

⁹ Ted Johnson and Charles Wald, “The Military Should Teach AI to Watch Drone Footage,” *Wired*, last modified November 26, 2017, <https://www.wired.com/story/the-military-should-teach-ai-to-watch-drone-footage/>.

¹⁰ Johnson and Wald, “The Military Should Teach AI to Watch Drone Footage,” <https://www.wired.com/story/the-military-should-teach-ai-to-watch-drone-footage/>.

¹¹ Jackie Snow, “Computer Vision Algorithms Are Still Way Too Easy to Trick,” *MIT Technology Review*, last modified December 20, 2017, <https://www.technologyreview.com/the-download/609827/computer-vision-algorithms-are-still-way-too-easy-to-trick/>.

¹² Gregory Allen, “Project Maven Brings AI to the Fight Against ISIS,” *Bulletin of the Atomic Scientists*, last modified December 21, 2017, <https://thebulletin.org/project-maven-brings-ai-fight-against-isis11374>.

¹³ Vincent Corruble, Charles Madeira, and Geber Ramalho, *Steps Toward Building a Good AI for Complex Wargame-Type Simulation Games* (Paris: Université Pierre et Marie Curie, 2002), 1-5.

¹⁴ Allen, “Project Maven Brings AI to the Fight Against ISIS,” <https://thebulletin.org/project-maven-brings-ai-fight-against-isis11374>.

Bibliography

- Allen, Gregory. "Project Maven brings AI to the fight against ISIS." *Bulletin of the Atomic Scientists*. Last modified December 21, 2017. <https://thebulletin.org/project-maven-brings-ai-fight-against-isis11374>.
- Corruble, Vincent, Charles Madeira, and Geber Ramalho. *Steps Toward Building a Good AI for Complex Wargame-Type Simulation Games*. Paris: Université Pierre et Marie Curie, 2002.
- Johnson, Ted and Charles Wald. "The Military Should Teach AI to Watch Drone Footage." *Wired*. Last modified November 26, 2017. <https://www.wired.com/story/the-military-should-teach-ai-to-watch-drone-footage/>.
- Metz, Cade. "As China Marches Forward on A.I., the White House Is Silent." *The New York Times*. Last modified February 12, 2018. <https://www.nytimes.com/2018/02/12/technology/china-trump-artificial-intelligence.html>.
- The National Artificial Intelligence Research and Development Strategic Plan*. (Washington: Government Printing Office, 2016).
- Rumelhart, David, Geoffrey Hinton, and Ronald Williams. "Learning Representations by Back-Propagating Errors." *Nature*, 323 (1986): 533-536.
- Snow, Jackie. "Computer Vision Algorithms Are Still Way Too Easy to Trick." *MIT Technology Review*. Last modified December 20, 2017. <https://www.technologyreview.com/the-download/609827/computer-vision-algorithms-are-still-way-too-easy-to-trick/>.
- Somers, James. "Is AI Riding a One-Trick Pony?" *MIT Technology Review* 120, no. 6 (2017): 30.
- Sukel, Kayt. "The Synapse—A Primer." *The Dana Foundation*. Last modified March 15, 2011. <http://www.dana.org/News/Details.aspx?id=43512>.
- Unmanned Systems Integrated Roadmap, FY2013-2038*. (Washington: Department of Defense, 2013).