REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
24-04-2018	Master's of Military Studies	SEP 2017 - APR 2018

4. TITLE AND SUBTITLE
Machine Learning and its Application in Marine Corps Operations

5a. CONTRACT NUMBER N/A
5b. GRANT NUMBER N/A
5c. PROGRAM ELEMENT NUMBER N/A

6. AUTHOR(S)
McCaffrey, Joseph, J, Major, USMC

5d. PROJECT NUMBER N/A
5e. TASK NUMBER N/A
5f. WORK UNIT NUMBER N/A

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
USMC Command and Staff College
Marine Corps University
2076 South Street
Quantico, VA 22134-5068

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)
Dr. John Gordon

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

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

13. SUPPLEMENTARY NOTES

14. ABSTRACT
A review of how machine learning, a subset of artificial intelligence, stands to fundamentally alter the way wars are fought and how the Marine Corps is presently postured to address this revolution in military affairs (RMA). Through research and documentation, evidence is provided demonstrating the Marine Corps is poorly postured to adapt to the RMA that machine learning is inducing. Recommendations are given for pragmatic solutions to alter structure, funding and partnerships to better align the Marine Corps for the near-term future.

15. SUBJECT TERMS
Machine Learning, Artificial Intelligence, Project Maven, Supervised Learning, Unsupervised Learning, Reinforcement Learning

16. SECURITY CLASSIFICATION OF: Unclass			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			USMC Command and Staff College
Unclass	Unclass	Unclass	UU	29	**19b. TELEPHONE NUMBER** *(include area code)* (703) 784-3330 (Admin Office)

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

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

MASTER OF MILITARY STUDIES

TITLE:

Machine Learning and Its Application in Marine Corps Operations

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

AUTHOR:

Joseph J. McCaffrey

AY 17-18

Mentor and Oral Defense Committee Member: J. W. Gordon
Approved:
Date:

Oral Defense Committee Member: MATTHEW FLYNN
Approved:
Date:

Executive Summary

Title: Machine Learning and its Application in Marine Corps Operations

Author: Major Joseph J. McCaffrey, United States Marine Corps

Thesis: Machine learning will significantly change the way future wars are fought. In order to leverage the inherent capabilities machine learning brings, the Marine Corps must make substantial changes.

Discussion: Machine learning, a subset of artificial intelligence, is a revolution in military affairs (RMA). It stands to fundamentally alter the way wars are fought. From image classification to speech recognition, robotics and autonomous vehicles, the possibilities are endless. However there remain real hurdles to this progress. Data acquisition and formatting is critical to success, and both are inherently difficult in government. Additionally, machine learning is not a panacea. There are problems machine learning can and cannot solve and it is important to identify the difference between the two. Accordingly, to leverage these new trends, the Marine Corps must understand the technology and be able and willing to adapt to it where required.

Conclusion: The Marine Corps is unprepared to adapt to this current RMA and needs to make substantial changes now in order to reverse the trend. The Marine Corps should adopt the current Project Maven process and create a unit under the Deputy Commandant of Information (DCI). This unit must begin to collate the Marine Corps' disparate data and create meaningful relationships with industry leaders using this data as incentive for private companies to participate.

DISCLAIMER

Table of Contents

Page

Introduction

Within the context of military operations, machine learning has the potential to radically change the United States' approach to war. We find ourselves in the midst of this revolution in military affairs (RMA) during a time of a larger, societal explosion in artificial intelligence. For the first time in history, there stands the potential to create an entity that is smarter than humans. Most experts agree, a computer will likely pass the Turing test (a widely accepted test of human level cognitive ability) before 2040.[1] The implications of such a landmark would be startling for society and revolutionary for the militaries of the world. Given the inevitability of this outcome, it is imperative that the Marine Corps prepare now for its realization.

Preparation for this revolution is no simple feat. There is significant work required to implement machine learning in a way that could provide a tangible benefit to the warfighter. There are poorly defined rules for many of the problems that Marines face and situations in war are infinitely unique. This uniqueness makes valuable extraction from machine learning models inherently difficult. As such, it is valuable to analyze where machine learning stands today and how, if applied appropriately, it can support Marine Corps operations. These efforts however, will not be successful if an appropriate framework is not in place. From organization to funding, the structure to leverage the power of machine learning must be realized in order to capitalize on its advantages. This work will not be completed in a vacuum. Political adversaries and others are greatly intensifying their investment in machine learning and many of them are unbound by regulations. The time to adapt and lead efforts in this new space is now. Machine learning will revolutionize the way future wars are fought, but in order to leverage the inherent capabilities machine learning brings, the Marine Corps must make substantial changes.

Machine Learning Overview

In March 2016 AlphaGo, a computer taught by machine learning only, defeated the human world champion of the ancient board game Go.[2] This was not merely an iteration of Deep Blue, the computer that beat chess champion Garry Kasporov in the mid-1990s. In reality, it was far different. Go is an incredibly complex game with more possible moves than there are atoms in the universe. Unlike Deep Blue, AlphaGo programmers could not simply teach the system to calculate the best possible move; it would take far too long to process and be of little use. Instead they had to show it what good Go playing looked like.[3] This, in and of itself, was a difficult task. Go is a game that requires critical thought, reasoning and intuition. Ask a Go player why he made a certain move and chances are they cannot tell you. As Michael Polanyi wrote in his book *The Tacit Dimension*, "…we know more than we can tell."[4] The Polanyi paradox, as it is known, provides that just as it is difficult to expound upon how one rides a bike, it is difficult for a Go player to explain how they play. As such, teaching a computer to play a game even humans have trouble explaining was a monumental challenge.

Programmers started by showing the system 30 million moves played by human experts. By analyzing these moves with a general-purpose algorithm, the machine learned how to execute these human strategies with startling precision, essentially mimicking the world's best Go players. This initial system played human champion Lee Sedol and won four games to one. That alone was an impressive feat but the AlphaGo team was not satisfied with their sole loss. They set about creating a new version of the system which they called AlphaGo Zero. This time instead of showing the system human expert moves, they simply told it the rules of Go and implemented a self-play algorithm, allowing the system to incrementally learn by playing itself.[5] As expected, at first it performed awfully. Much like a child, its moves were random with seemingly no strategy to back its play. However, that quickly changed. Within 36 hours it played almost 4.3 million games and was demonstrating expert human strategies, all self-learned.[6] After only 72 hours, it was able to beat the first version of itself 100 games to zero. Then it played some of the best human players in the world and chalked up a record of 60-0. To date, no human has beaten it. Human Go players, including former world champions, describe AlphaGo Zero's strategies as conservative yet superhuman, making unexplainable moves that somehow pay off dozens of moves later. Humans are now using the strategies AlphaGo Zero created to improve their own game. Amazingly a game once thought impossible to teach to a computer was not only taught to a computer but surpassed all human knowledge of the ancient game in only three days.

If machine learning can facilitate AlphaGo Zero, so the argument goes, then what more practical impacts can it have? A discussion of how the Marine Corps can benefit from machine learning must begin with an in-depth look at what machine learning is and what problems it can and cannot solve.

Machine learning is a subset of the often-discussed artificial intelligence (AI) (figure 1); in other words, it is one way to achieve artificial intelligence.[7] The terms are often used interchangeably but it is important to understand the difference. As defined by expert Luca Scagliarini, "Machine learning is an application of AI that provides systems the ability to

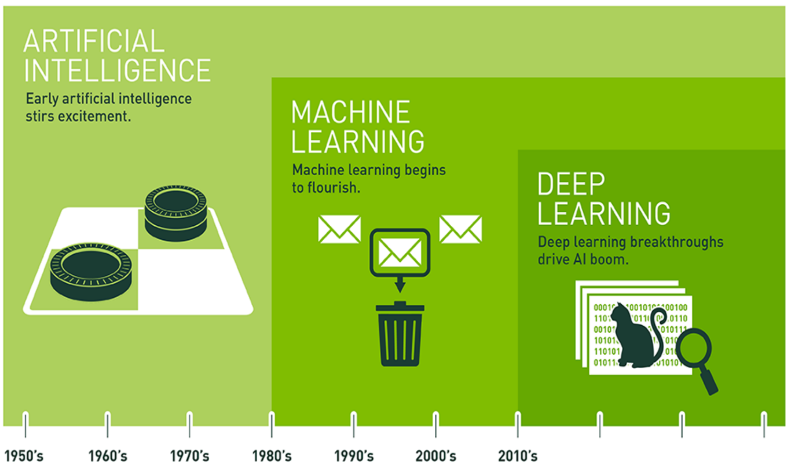

Figure 1: Artificial intelligence evolution (NVIDIA)

automatically learn and improve from experience without being explicitly programmed."[8]

Machine learning is hardly new. In 1957 Frank Rosenblatt, a Cornell psychologist developed an algorithm he called Perceptron in support of a Navy funded program.[9] The goal was for the computer powered algorithm to act as a binary classifier, that is, to look at an image and classify it into one of two categories. However, computing power at the time was far too slow to effectively process the algorithm and the experiment ultimately failed. Unfortunately, many at the time mistook the failure of Perceptron as proof machine learning was inherently impossible. It would not be until the 1980s that a true appreciation for the Perceptron algorithm would be realized. At that time multiple papers were published demonstrating how it could be successful if the algorithm was layered together and applied with new computing technology. This so called multi-layered Perceptron has become what is today known as deep neural networks or deep learning, a further subset of machine learning (figure 1).

Since 1980 the field of machine learning has boomed. As computing power has exponentially increased, interest within the mathematics and computer science communities has flourished, creating a positive feedback loop resulting in further improvements of the algorithms that drive machine learning. Today, machine learning is largely broken up into three different types: supervised, unsupervised, and reinforcement.[10] In order to understand how the Marine Corps stands to benefit from this resurgence of machine learning it is critical to understand these types.

Supervised Learning

Children learn how to identify objects not by explicit rules but through example. A parent wishing to teach a child what a dog looks like does not say, “a dog has four legs, a black nose, floppy ears and is furry.” They simply point to a picture of a dog and say, “dog”. They do this enough and eventually the child learns what a dog is. It is in this way that supervised learning works but on a much larger scale. A computer is shown what the programmer wants it to learn and is told what it is. It is then shown thousands, if not millions, of examples and it eventually learns to distinguish the object from others.[11]

Supervised learning is great for classification problems such as classifying images into groups (cats, boats, cars, etc.).[12] More practically, it is used for optical character recognition, turning handwriting into typed font. Again, this is done by training the computer on thousands of labeled images of human handwriting, applying machine learning algorithms and testing it to ensure accuracy. When conducted under ideal circumstances one should expect accuracy to exceed 99% on a simple classification problem.[13] If it is less, the model likely needs to be fine-tuned or further trained. It is in this way that hospitals are beginning to train computers to diagnose disease through medical imagining, for instance, showing the computer thousands of images labeled by human experts (doctors in this case) and allowing it to learn what cancer looks like on an x-ray.[14]

There are inherent weaknesses of supervised learning. For one, as Frank Rosenblatt found out, you need massive amounts of labeled data to train the model with. For general image classification this is relatively easy. Though tedious, any human can sort through thousands of images and label general objects in them (e.g. dog, car, grass). When the image classification is more complex, such as identifying cancer in an x-ray, it requires not only obtaining the images but having an expert in the field (radiologist) label them. There are a number of generic labeled training sets available for public use but creating labeled training sets for specific uses is often half the battle in implementing an effective supervised learning model. Additionally, though increasingly less of a problem, is computing power. The process of training, or physically showing the model the training set, is extremely computing intensive. For Frank Rosenblatt this was prohibitive, but today cloud computing allows even individuals to leverage vast amounts of computing power at affordable prices. Modern computer games have further assisted in this effort as graphic processing units (GPU), once used to ensure the modern gamer could play stutter free, are now being leveraged for their unique ability to simultaneously compute at levels 100 times that of the standard computer processing unit (CPU).

Most importantly, supervised learning models require that humans understand the underlying pattern. In order for the model to identify cancer in an x-ray it requires a doctor who knows what cancer looks like in an x-ray. This inherently limits supervised machine learning models to executing tasks that humans can already accomplish. The machine, however, has the advantage of being able to accomplish these tasks with much greater speed, more accuracy, and less bias than their human counterparts, making supervised learning models an essential method of the machine learning paradigm.

Unsupervised learning

While we explicitly attempt to teach our children many things about the world around us, most of what humans learn is in an unsupervised manner. A child learns about gravity not by being taught Newton’s law of universal gravitation, but rather by tripping and falling or spilling milk. They realize, through experience, there is a natural tendency for everything to be drawn downwards, not up. This is an example of unsupervised learning, detecting patterns through the observation of unlabeled data. Unsupervised machine learning models learn the same way. These models are fed unlabeled data in the hope that they will be able to discern some useful pattern that we humans cannot. The algorithms that support unsupervised learning models have grown in popularity and power over the last few years. They are particularly good at what is called *clustering*. Clustering involves taking unlabeled data and clustering it into groups that humans may be able to discern helpful information from.[15] For instance, an unsupervised model may be fed unlabeled purchase data from Amazon and subsequently identify that individuals who buy one product are more likely to buy another, seemingly unrelated product. This information can help companies make better informed decisions in regard to marketing, recommendations, and distribution.

However, just as with human learning, these patterns do not always match reality. Human assumptions based on previous observations of unlabeled data make us susceptible to bias. We tend to overfit data. There is good evolutional purpose for this as we need to be able to extract a lot of information from a limited sample size but at times, it can lead us to erroneous conclusions. Computers fed with inaccurate or incomplete data, can make the same mistakes. Indeed, the challenge of effectively implementing an unsupervised machine learning model often revolves around confirming whether detected patterns are, in fact, a reflection of reality or merely an over or underfitting of the original data.

Reinforcement Learning

The final primary type of machine learning model in use today is reinforcement learning. Again, attempting to mimic the way humans learn, reinforcement learning models are designed to give the computer (known as an agent in this model) a reward for completing an appropriate task and a penalty for failing. For example, humans may give a child a treat for cleaning up after themselves and a time out for not doing so. Eventually the child learns if he wants to avoid a penalty and get a treat, it is best to clean up after himself. So too, do reinforcement learning models seek to learn the best strategy (known as policy). AlphaGo Zero is an example of a model trained using reinforcement learning. When it won a game, it would get a reward (in the form a point) and when it lost a game it would get a penalty (losing a point).[16] It then became a matter of instructing the model to maximize the points it received. After 4.3 million games of playing itself and being rewarded or penalized after every game, it had figured out an incredibly effective strategy to win, a strategy that in fact exceeded all human strategies.

As with the other learning models, reinforcement learning has its drawbacks. In order to effectively train such a model, its human teacher must know what to reward it for and what to penalize it for. For the game of Go, this is relatively easy: reward the model for winning and penalize it for losing. Yet for more complicated problems this can become increasingly difficult. If teaching a self-driving car, do you reward it or penalize it for swerving to avoid a collision with another car? What if that swerving results in it hitting an animal, or worse yet, a pedestrian? In order to teach models to solve these complex problems, it requires humans to deeply consider what is most important to us. It is this very consideration that concerns many machine learning critics who fear that human teachers will reward machines for precisely the wrong things.

Algorithms

As the individuals that use these systems will likely not understand the math that underlies them, there becomes an inherent distrust of the system as whole. It is important to remember that machine learning is not magic. Underlying these models are tried and tested mathematical functions. Algorithms such as the softmax function or naïve Bayes classifier are derived from relatively ancient mathematics.[17] Machine learning gave these old functions new light by applying computing power to create models based on millions of multivariable samples. Just a few decades ago this was impossible. This is not to say new mathematics are not being realized in the machine learning revolution, they in fact are at an incredible rate, but the basis for these new algorithms are rooted in sound theories.

In order for the Marine Corps to move forward successfully with machine learning, it must be able to overcome this natural distrust. This will come by demonstrating on a small scale the powers of the models to supplement and even better our own human decision making. The next section will discuss some ways in which this is already being done. It will also explore how to improve upon these successes to demonstrate that machine learning works and that it can be viably applied to the warfighting functions.

Current DoD Machine Learning Projects

The DoD has already begun to explore the possibilities of machine learning, albeit in a limited manner. Project Maven is the first large scale effort by the DoD to implement such a solution.[18] Leveraging a supervised learning model, the Algorithmic Warfare Cross-Functional Team (AWCFT), Project Maven's official title, is working to provide a classification solution to the abundance of drone footage currently handled by the Department. Currently, intelligence analysts are required to view all video from a drone and document everything they see. Across the force, this equates to hundreds of thousands of hours of footage, most of it benign. Analysts must document everything from number of people seen, to cars, weapons and all other matter of things. Additionally, there is vast amounts of additional sensor data that is often discarded as there is simply not enough time or people to review and catalogue it. Project Maven seeks to solve this problem by using machine learning, a problem that is well within its current capability. Open source classification tools are available today from a number of well-known companies to include Microsoft, Google and Amazon. These tools are available for on-site instantiations, meaning Project Maven and others can leverage the power of the tools without being required to have their data hosted on commercial companies' equipment. At present, Project Maven is only using unclassified information to train its models, but the on-site option gives them the ability to port the capability to the classified realm. This system will likely be available to others within the DoD to conduct their own model training on.

While Project Maven represents a step forward for the DoD in regard to applying machine learning, its current scope is not trend setting in the larger community. Classification problems, even over video as opposed to static images, is well explored and documented. YouTube, owned by Google, maintains likely the largest training video set, with over 450,000 hours of labeled video.[19] This set can be leveraged by anyone to classify video of their own (Figure 2).

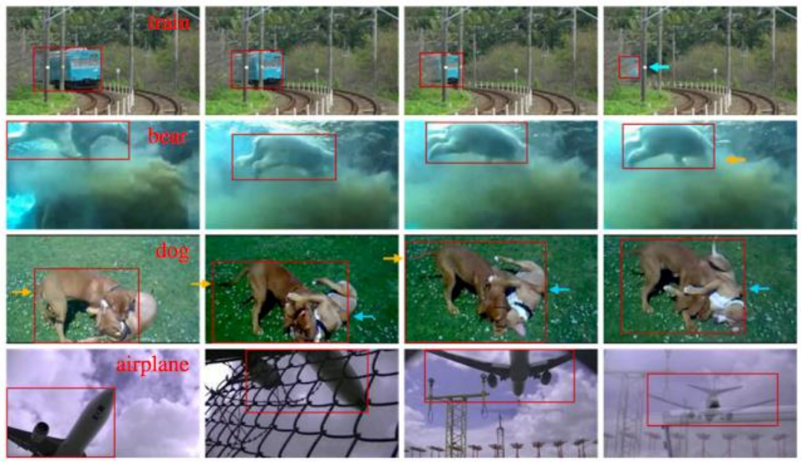

Figure 2: Google's Video Classifier (Google)

Project Maven's unique requirements, identifying weapons types, individuals digging and other drone unique video, requires its own training set (YouTube's training set likely would not be effective for training Maven's model). Therein lies the larger, though not insurmountable problem with using machine learning for classification. As previously discussed, it requires a massive labeled training set. This is manpower intensive as it requires humans to go through hundreds of thousands of hours of video and label everything they see. While this takes a long time, it will ultimately pay off as the system will be able to automatically identify these objects once it is fully trained.

<u>Potential Future Projects</u>

Project Maven's approach is intentionally limited, in order to get a product to the warfighter quickly and prove machine learning's viability. However, to compete on a more global scale, Project Maven will need to increase substantially. As classification is a great problem for machine learning to tackle, it is worth exploring what other problems can be solved with classification. Erik Brynjolfsson and Andrew McAfee of MIT in their book *Machine, Platform, Crowd: Harnessing Our Digital Future* argue that machine learning is best applied towards problems that are dull, dirty or dangerous.[20] Many operations within the Marine Corps fall into at least one if not all three of these categories. As such, there are many areas in which the Corps can benefit by investing in machine learning systems today.

Imagery Classification

One problem set that fits into the dull category of Brynjolfsson and McAfee's task breakdown, is that of classifying satellite imagery. With government and commercial satellites constantly gathering images from space, the work of identifying objects within that imagery lands squarely on analysts. Just like drone footage, this is manpower intensive and often a poor allocation of analysts' time. Additionally, as humans, analysts are susceptible to making mistakes, particularly when tasked with dull work. There is overwhelming evidence that human error rates skyrocket when tasked with monotonous work over long periods of time.[21] As such, automatic classification powered by a machine trained model could be used today to reduce errors and greatly increase the value of satellite imagery to decision makers. Models implemented could automatically identify changes to military bases, to include counting aircraft, vehicles, ships and other equipment. This information could be fed into a searchable database, allowing analysts to spend more time using all source intelligence to develop a clear profile of a particular adversary. This data could further be combined with additional sensor data to provide extremely clear pictures of civilian and military activity around the globe, nearly all autonomously.

All the necessary tools exist today to allow the Marine Corps to implement such a solution for a relatively low cost. For example, companies such as Digital Globe, from which the

US Government buys satellite imagery, use Google's Tensorflow application to implement their change detection algorithm.[22] This algorithm automatically detects changes between imagery dates, quickly identifying new buildings and roads. They have taken it a step further by identifying changes in parking lot traffic (Figure 3) and airplanes at airports.[23] While today this data is provided to investors in an attempt to gain a trading advantage it could, relatively easily, be ported to support military operations. The National Geospatial-Intelligence Agency (NGA) is already thinking along these lines. In an RFI published by NGA in September 2017 they said, "[we are] conducting a market research study on the demand for a public-private partnership (PPP) to produce artificial intelligence algorithms trained on geospatial data to help defend national security."[24] In other words, NGA is exploring providing its imagery to private companies who in turn would use it to train machine learning models and subsequently sell those models. NGA would receive these trained models in return for providing the imagery. This type of innovative partnership with industry may be the key to Marine Corps success in the future AI environment.

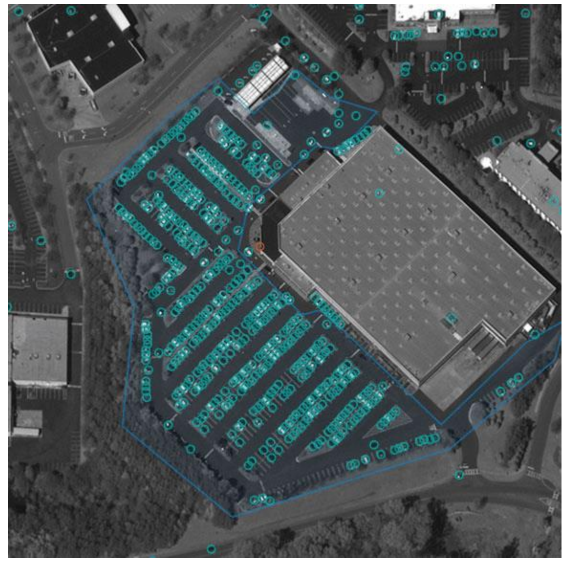

Figure 3: DigitalGlobe's Change Algorithm (Digitalglobe.com)

<u>Speech Recognition</u>

Speech recognition is another category that has already impacted consumer electronics and stands poised to greatly improve military applications as well. When broken down into audio spectrum, a speech recognition model merely matches spoken words with the profile of words it has been trained on. As such, it is a form of classification, albeit a complicated one. Speech recognition has improved tremendously over the last five years in large part due to machine learning. Google's speech recognition now recognizes 110 languages and variants and can operate in noisy environments, returning results in real-time.[25] Google then takes the users speech and further fine tunes its model, creating an incredibly accurate system.

Systems such as these have many potential applications in military operations, most obviously application in signals intelligence. As in imagery intelligence, analysts bear the brunt of the burden in reviewing and cataloguing troves of audio information. Machine learning can apply speech recognition to not only transcribe what is being said in audio but also automatically translate it from any language into English or vice versa. This could greatly reduce the need for specialized, cleared translators, which are often difficult to hire and retain in the numbers needed. Additionally, with the implosion of digital devices comes the need to interact and manipulate them. This is traditionally a friction point for warfighters. With simple speech recognition available today, a warfighter could toggle through his tablet to call for fire using voice commands only. This would allow him to stay heads up in a kinetic situation, improving situational awareness and reducing human error.

Computer Vision

The British Machine Vision Association defines computer vision as being, "…concerned with the automatic extraction, analysis and understanding of useful information from a single image or a sequence of images." Classification problems involving images indeed must rely on computer vision. For a simple image, this is often attained by flattening the image into what is known as feature vectors.[26] These feature vectors are an abstraction of the image into numbers that a computer can understand. For video, this is done frame by frame. Combined with technology such as Light Detection and Ranging, better known as LIDAR, one can develop a computer that quite literally sees. Computer vision is integral to a number of problems that machine learning can help solve. Autonomous vehicles, as will be discussed in detail later, rely heavily on computer vision as a self-driving car must constantly observe the external environment in order to successfully navigate. Robotics will need to have some sort of vision to operate at a level where they can be considered useful. Combining vision, as humans think of it, with data, creates a wealth of possibilities for robotic systems. In the near term these systems will likely be limited to simple tasks, but eventually could execute complex tasks to include moving just like humans. Boston Dynamics has created a fully functioning "humanoid." It believes the world is designed for humans and any useful robot will need to move like a human. Currently, its humanoid can walk, do back flips and pick up small boxes. It has also developed a robot that emulates a mule and could be used by militaries to transport heavy equipment over rugged terrain.

Figure 4: Boston Dynamic's "Humanoid"

Autonomous Vehicles

Perhaps one of the most exciting byproducts of the machine learning boom is self-driving cars. Loup Ventures, a large venture capital firm, predicts by 2040, 95% of all vehicles sold will be self-driving capable, potentially a 3.6 trillion-dollar market.[27] Rio Tinto, an international mining company, has been using autonomous dump trucks at its iron ore operations for nearly two years and just announced a purchase of $5 billion in additional autonomous vehicles.[28] These vehicles, powered by machine learning, use pre-defined GPS routes to maneuver vehicles through the rugged terrain surrounding their mines. Such technology could be easily ported to military operations, allowing autonomous supply vehicles to traverse main supply routes (MSRs) without exposing human escorts to the inherent dangers of the battlefield. Autonomous vehicles paired with computer vision cameras could additionally detect disturbances in ground providing indicators of a potential improvised explosive device (IED). This would improve safety of manned and unmanned vehicles alike. Autonomous vehicles could significantly reduce risk and give commanders additional flexibility in their operations.

Machine Learning Challenges

Having reviewed just a few of the many capabilities that machine learning can bring to the warfighting functions it is clear that wars will be fought differently in the near future. However, the Marine Corps is poorly positioned to take advantage of this RMA. This poor position is largely a result of not addressing the challenges that inherently lie in building effective machine learning models. These challenges are relatively simple but require significant effort in overcoming. Appropriately managing data is of paramount importance but it is closely related to the element of randomness. Additionally, it is vital to benchmark the Marine Corps' current position in the AI space with that of its adversaries.

Data

At present, data drives machine learning. Without a large amount of data, the machine learning models of today cannot effectively learn. While there has been some progress with a technique called transfer learning, in which a system extracts data from one problem and applies it to solve problems of an entirely different nature, the data required is still large. The process of gathering appropriate data and labeling it (in the case of supervised learning), is at best a very laborious process. At worst it is impossible, because a number of problems humans would like to apply machine learning to, there simply is not enough data to train a system on. For example, we cannot train a system to develop cancer drugs because we know so little about how and why cancer forms. John Launchbury, a Director at the Defense Advanced Research Projects Agency (DARPA) argues the next wave of AI will require models to abstract tons of information from limited amounts of data like humans do.[29] Additionally, even in cases where we do have the data, there may be problems in converting the data into a form that a computer can understand. For example, how do we digitize the biological processes of the human body? In order for machine learning to help us develop new live saving drugs, it will need to be able to view and understand these processes. Breaking such processes down into 1's and 0's is a feat humans have yet to accomplish. For now, those who have access to large troves of discernable data will continue to lead the way in AI development.

This information owner advantage could be leveraged by the Marine Corps, but it is not. The aforementioned NGA RFP is an example of an information owner using the leverage of the data they have, to gain a decisive advantage. The Marine Corps gathers and retains streams of information but not in aggregate, not in any standardized format and not shared. For example, the Marine Corps' Global Combat Support System (GCSS-MC) contains records of every piece of hardware in the Marine Corps inventory, but it remains largely a stand-alone system, operating on a proprietary Oracle framework.[30] It does what it is supposed to do but it does not talk to other systems and its raw data is not consumable in an open environment. Therefore, there is currently no ability to train machine learning models with GCSS-MC data. This is critical if the Marine Corps wants be able to implement capabilities such as a predictive maintenance model. Such a system would tell the unit when equipment needs maintenance as opposed to adhering to arbitrary preventative maintenance timelines that ultimately waste time and money.

The Marine Corps faces an uphill battle in ensuring that data such as that generated in GCSS-MC is formatted appropriately. These requirements need to be written into contracts but that will not happen until the need is clearly articulated up and down the chain of command. Formatting can be adjusted after data creation, but it takes time and thus, on a large scale, is expensive. Getting the formatting right at data creation will save time and money.

Even more important than formatting is sharing. If the data is not shared, it cannot be used. For the Marine Corps to truly leverage the data it retains, it needs to be able to see it all in one location. Such a repository does not exist but it should. This challenge is not unique to the Marine Corps, but it appears the Marine Corps has not even begun to address it while others, even some in government, have. The City of San Diego for instance, uses a tool called Airflow to share multiple streams of data across platforms with their machine learning models.[31] Data will continue to be a linchpin for machine learning. Any organization that wants to take its efforts in the space seriously must address this problem first.

Randomness

Machine learning still does a poor job at solving problems involving prediction of human behavior. There is a fine line between complex and stochastic problems. While machine learning may help identify patterns in complex problems, stochastic problems by definition, have no patterns. As such, machine learning at present does not apply well to these sorts of problems. Unfortunately, it is inherently difficult for humans to identify the line between complex and random. There is risk in attempting to apply machine learning to a seemingly complex problem set only to find out later that it is in fact stochastic. For example, high frequency traders use machine learning in an attempt to predict stock market movements but they do this minute by minute. These predictions become useless when the system is asked to predict periods out to days or weeks, as the number of affecting variables grows to infinity and it becomes all but impossible to predict accurately. The very short-term prediction of stock market movement over the next minute is arguably a complex problem, as in all but the rarest of circumstances a finite number of variables will affect its movement. However, as the prediction time increases, the accuracy of high frequency trading models drops to zero.

In general, the Marine Corps must exercise extreme caution when thinking that machine learning may be able to accurately predict situations which involve human interaction or behavior as the involvement of humans adds an element of randomness that is difficult, if not impossible, to pinpoint. As such, one who seeks to employ a machine learning system should carefully consider the type of problem one is attempting to solve.

Adversarial Implementation

The machine learning discussion does not occur in a vacuum. While the United States determines how to move forward with machine learning, so too does its adversaries.

Tech giant Yandex, often considered Russia's Google equivalent, offers many of the same open source machine learning systems as Google does. Unlike the US and Google however, the Russian government imparts heavy control over Yandex, often times influencing which search results appear first.[32] It is likely that Russia will continue to lean heavily on its commercial tech arm in a way that the US government cannot. This will enable them to stay at the forefront of AI technology without necessarily having to invest in research and development. A competitive edge in AI can clearly be seen as a military advantage.

China is perhaps even more threatening. China has publicly announced plans to become the world leaders in AI by 2030 and has begun investing to ensure that happens. Its unique government structure enables them to leverage commercial industry to progress government technology. Private companies in China also benefit from a regulatory environment that is lax in regard to personal privacy. Successful machine learning implementation requires extremely large data sets. In the US, there are strict rules for how companies acquire and maintain their customer's private information. In China, there are not. This allows Chinese companies such as Baidu (Chinese Google equivalent) to maintain massive data sets of private information that would be much more difficult for a US company to accomplish.[33] You have not explained how this is a threat. This inherently enables the Chinese government to train models on massive amounts of data not available to anyone else, giving them a decisive advantage in the space.

Historical Context

While it is outside the scope of this project to explore historical cases of military technological advantages in depth, it is important to look at the many factors that drive technical advantages and the shortfalls of leaning too heavily on them. While machine learning stands to revolutionize modern warfare, the focus must remain, as Max Boot noted, "[on] the soldier struggling to kill or avoid being killed, and to his commander struggling to master the remorseless logic of carnage."[34] Technology alone does not win wars.

We are in the midst of the fourth technological revolution of modern times, the Information Revolution, of which machine learning and its parent, artificial intelligence, are just a part. As in all technological revolutions, a military will not gain an advantage by merely possessing a technology. Training, strategy, leadership, employment and other factors all play critical roles in deciding a technical advantage. Furthermore, a technical advantage does not necessarily equal a battlefield advantage. Since 1945, every failed US military effort has been against a technologically inferior enemy.[35] As a superpower, maintaining a technological edge is not enough. Even perfect implementation of machine learning will not equate to a panacea or an uncontested technological advantage, especially as it pertains to unconventional threats. However, not maintaining a decisive advantage in the machine learning space (and technology in general) invites conventional threats.

Recommendations

Artificial intelligence and its machine learning subset stand to redefine technology. At present, the Marine Corps is poorly positioned to take advantage of this new wave of technology. While ventures like Project Maven are absolutely crucial it is hardly enough. Andrew Ng, former Chief Scientist at AI giant Baidu and Stanford professor, says, “AI is the new electricity.”[36] If he is right, much more will be needed than a small project with limited goals. Therefore, the Marine Corps should take a three-pronged approach in addressing this new realm and positioning itself for the future. Addressing structure, data and partnerships is the key to success in the machine learning space.

Structure

The Marine Corps should immediately create an AWCFT (Project Maven) like unit. This unit should fall under the new Deputy Commandant for Information (DCI) and their authorities. This new unit should be dedicated to artificial intelligence and matrixed across the Corps. Centralizing structure and effort is not always the optimal approach but AI requires specialized skills and technology, all of which are still relatively new. Centralizing this effort will allow for more effective recruiting and sharing of information in addition to quicker allocation of resources. Making this organization subordinate to DCI will give it the required clout to effectively operate. Organizations would then be able to submit requests for the development and training of AI powered programs, similar to how Cyber Command operates today.

As the repositories of developed models grow, organizations may find that what they are trying to accomplish has already been done and can leverage this previous work.

Data

While the Marine Corps may not be postured today to leverage machine learning, it is the owner of perhaps the most valuable resource in applying the technology: data. The amounts of data that the Marine Corps collects and maintains is staggering. However, most of this data is disparate and unlabeled, analogous to having millions of books in a warehouse unsorted. Yet in reality there is not actually one warehouse, there are hundreds. The information one needs may very well be in these warehouses but it is unlikely one could find it. Collecting, cataloguing and labeling this data is critical to any effort that seeks to employ machine learning. The aforementioned AI unit must have the authority to collect data from organizations across the Marine Corps and information across classification levels. This would be a monstrous task and while it is one that ultimately seeks to reduce manpower across the force, the process itself would ironically require a herculean, human effort.

Partnership

Lastly, machine learning is not an effort the Marine Corps can endeavor alone. The level of expertise needed to make breakthroughs in this arena is extremely high. No military organization will be able to recruit and retain the top PhD level experts that will ultimately lead the machine learning evolution. These experts are certain to reside in private industry and while the likes of China and Russia will compel experts in their country to assist in government efforts, the United States cannot. As such, the Marine Corps must continue to engage in partnership with industry. However, in order for these partnerships to be effective, they must be mutually beneficial. The troves of data the Marine Corps retains is highly lucrative to private industry for a variety of reasons and any worthwhile partnership will include a two-way transfer of data. Currently, the Marine Corps is structured to inherently protect all its data, regardless of relevance. This approach must be readdressed in the 21st century if the Marine Corps wishes to build and maintain meaningful relationships with industry partners.

Conclusion

AlphaGo demonstrated that machine learning, as a subset of artificial intelligence, is poised to change the world. The human condition is being tested and while many fear the possibilities, those who are unable to adapt to the changes are sure to fail. The Marine Corps is behind in adapting to this current RMA and needs to make substantial changes now in order to reverse the trend. The Marine Corps should adopt the current Project Maven approach and create a like unit under the DCI. This unit must begin to collate disparate data and create meaningful relationships with industry leaders using their data as incentive for private companies to participate. These actions would help the Marine Corps get on the right path to leveraging the benefits of machine learning and ensuring they retain a technological edge over their adversaries well into the future.

As Max Boot notes in his book *War Made New*, "…the only test worth anything is the test of battle, and it is here that we must look for the impact of technology upon military affairs".[37] If the Marine Corps is unable to translate the uses of machine learning into a decisive battlefield advantage, all will be for naught. The best capabilities in the world mean nothing if one is unable to apply them in meaningful ways or otherwise waste them pursuing impossible dreams.

Notes

[1] Aleksandar. Todorovic, "Has The Turing Test Been Passed?" last modified 2017, http://isturingtestpassed.github.io//.

2 Cade Metz, "Why the Final Game Between AlphaGo and Lee Sedol Is Such a Big Deal for Humanity," WIRED, last modified June 3, 2017, https://www.wired.com/2016/03/final-game-alphago-lee-sedol-big-deal-humanity/.

[3] DeepMind, "AlphaGo Zero: Learning from Scratch," DeepMind, accessed December 28, 2017, https://deepmind.com/blog/alphago-zero-learning-scratch/.

[4] Michael Polanyi and Amartya Kumar. The tacit dimension (Chicago, Ill: The University of Chicago Press, 2013).

[5] DeepMind.

[6] Ibid.

[7] Michael Copeland, "The Difference Between AI, Machine Learning, and Deep Learning? | NVIDIA Blog," The Official NVIDIA Blog, last modified December 15, 2017, https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/.

[8] Luca Scagliarini, "What is Machine Learning? A Definition - Expert System," Expert System: Artificial Intelligence: Cognitive Computing Company, last modified October 5, 2017, http://www.expertsystem.com/machine-learning-definition.

[9] Charles Tappert, "Rosenblatt's Contributions," Seidenberg School of CSIS, accessed January 20, 2018, http://csis.pace.edu/~ctappert/srd2011/rosenblatt-contributions.htm.

[10] Andrew Ng, "Introduction to Machine Learning" (Lecture, Machine Learning, Stanford University, Coursera.com, 2017).

[11] Ibid.

[12] Ibid.

[13] Ibid.

[14] Ibid.

[15] Ibid.

[16] DeepMind.

[17] Ibid.

[18] Gregory Allen, "Bringing AI to the Fight Against ISIS," Bulletin of the Atomic Scientists, last modified December 23, 2017, https://thebulletin.org/project-maven-brings-ai-fight-against-isis11374.

[19] "YouTube-8M: A Large and Diverse Labeled Video Dataset for Video Understanding Research." Research at Google. Accessed January 22, 2018. https://research.google.com/youtube8m/.

[20] Andrew McAfee and Erik Brynjolfsson, Machine Platform Crowd: Harnessing Our Digital Future (New York: W. W. Norton & Company, 2017), 74.

[21] M.G Stewart, "Modelling Human Error Rates for Human Reliability Analysis of a Structural Design Task - ScienceDirect," ScienceDirect.com, accessed January 22, 2018, http://www.sciencedirect.com/science/article/pii/0951832092900975?via%3Dihub.

[22] Valery Komissarov, "AI Applications for Satellite Imagery and Satellite Data," TechEmergence, last modified September 10, 2017, https://www.techemergence.com/ai-applications-for-satellite-imagery-and-data/.

[23] Ibid.

[24] "RFI: DEVELOPING ALGORITHMS ON GEOSPATIAL DATA - Federal Business Opportunities: Opportunities," Home - Federal Business Opportunities. Accessed January 20, 2018, https://www.fbo.gov/index?s=opportunity&mode=form&id=692a228f86775dab6cd36dbf7b1494f1&tab=core&_cview=0.

[25] "Cloud Speech-to-Text - Speech Recognition," Google Cloud, accessed April 16, 2018, https://cloud.google.com/speech/.

[26] Adrian Rosebrock, "How To Describe and Quantify an Image Using Feature Vectors," PyImageSearch, last modified April 28, 2014, https://www.pyimagesearch.com/2014/03/03/charizard-explains-describe-quantify-image-using-feature-vectors/.

[27] Gene Munster, "Auto Outlook 2040: The Rise of Fully Autonomous Vehicles," Loup Ventures, accessed January 22, 2018, http://loupventures.com/auto-outlook-2040-the-rise-of-fully-autonomous-vehicles/.

[28] "Rio Tinto to Expand Autonomous Fleet As Part of $5 Billion Productivity Drive," Rio Tinto - Rio Tinto, last modified December 18, 2017, http://www.riotinto.com/media/media-releases-237_23802.aspx

[29] John Launchbury, "DARPA Perspective on AI," Defense Advanced Research Projects Agency, n.d.https://www.darpa.mil/about-us/darpa-perspective-on-ai..
[30] "Global Combat Support System - Marine Corps (GCSS-MC)," U.S. Marine Corps, last modified August 4, 2017, https://marinecorpsconceptsandprograms.com/programs/command-and-controlsituational-awareness-c2sa/global-combat-support-system-marine-corps.
[31] Maksim Percherskiy, "Data Science On The Silicon Beach," Partially Derivative, last modified March 6, 2017, http://partiallyderivative.com/podcast/2017/03/06/maksim-percherskiy.
[32] Louise Matsakis, "How Russia Polices Yandex, Its Most Popular Search Engine," Motherboard, last modified November 24, 2017, https://motherboard.vice.com/en_us/article/8x5z54/how-russia-polices-yandex-its-most-popular-search-engine.
[33] Ibid.
[34] Max Boot, War Made New: Technology, Warfare, and the Course of Modern History (New York: Gotham Books, 2006), 165.
[35] Jeffery Record. Beating Goliath, Why Insurgencies Win.
[36] Shana Lynch, "Andrew Ng: Why AI is the New Electricity | The Dish," Stanford News, last modified March 14, 2017, https://news.stanford.edu/thedish/2017/03/14/andrew-ng-why-ai-is-the-new-electricity/.
[37] Boot, 65.

Bibliography

Acemoglu, Daron, and Pascual Restrepo. “Robot and Jobs: Evidence from US Labor Markets.” *MIT Scripts*, (March 2017). http://pascual.scripts.mit.edu/research/robots_jobs/robots_and_jobs_2017.pdf.

Agrawal, Ajay, Joshua Gans and Avi Goldfarb. “The Obama Administration’s Roadmap for AI Policy.” *Harvard Business Review, (*Dec 2016). https://hbr.org/2016/12/the-obama-administrations-roadmap-for-ai-policy

Allen, Gregory. "Bringing AI to the Fight Against ISIS." Bulletin of the Atomic Scientists. Last modified December 23, 2017. https://thebulletin.org/project-maven-brings-ai-fight-against-isis11374.

Boot, Max. *War Made New: Weapons, Warriors and the Making of the Modern World*. New York City: Gotham Books, 2006.

Copeland, Michael. "The Difference Between AI, Machine Learning, and Deep Learning? | NVIDIA Blog" The Official NVIDIA Blog. Last modified December 15, 2017. https://blogs.nvidia.com/blog/2016/07/29/whats-difference-artificial-intelligence-machine-learning-deep-learning-ai/.

DeepMind. "AlphaGo Zero: Learning from Scratch.” Accessed December 3, 2017. https://deepmind.com/blog/alphago-zero-learning-scratch/.

Developing Algorithms on Geospatial Data. Federal Business Opportunities. Accessed January 22, 2018. https://www.fbo.gov/index?s=opportunity&mode=form&id=692a228f86775dab6cd36dbf7b1494f1&tab=core&_cview=0.

Gray, Colin S. *Another Bloody Century: Future Warfare*. London, England: Weidenfeld and Nicolson, 2005.

Komissarov, Valery. "AI Applications for Satellite Imagery and Satellite Data." TechEmergence. Last modified September 10, 2017. https://www.techemergence.com/ai-applications-for-satellite-imagery-and-data/.

Launchbury, John. "DARPA Perspective on AI." *Defense Advanced Research Projects Agency*. n.d. https://www.darpa.mil/about-us/darpa-perspective-on-ai.

Lynch, Shana. "Andrew Ng: Why AI is the New Electricity | The Dish." Stanford News. Last modified March 14, 2017. https://news.stanford.edu/thedish/2017/03/14/andrew-ng-why-ai-is-the-new-electricity/.

Matsakis, Louise. "How Russia Polices Yandex, Its Most Popular Search Engine." Motherboard. Last modified November 24, 2017. https://motherboard.vice.com/en_us/article/8x5z54/how-russia-polices-yandex-its-most-popular-search-engine.

McAfee, Andrew, and Erik Brynjolfsson. *Machine Platform Crowd: Harnessing Our Digital Future*. New York: W. W. Norton & Company, 2017.

Metz, Cade. "Why the Final Game Between AlphaGo and Lee Sedol Is Such a Big Deal for Humanity." WIRED, June 3, 2017. https://www.wired.com/2016/03/final-game-alphago-lee-sedol-big-deal-humanity/.

Munster, Gene. "Auto Outlook 2040: The Rise of Fully Autonomous Vehicles." Loup Ventures. Accessed January 22, 2018. http://loupventures.com/auto-outlook-2040-the-rise-of-fully-autonomous-vehicles/.

Ng, Andrew. "Introduction to Machine Learning." Lecture, Machine Learning, Stanford University, Coursera.com, 2017.

Polanyi, Michael, and Amartya Kumar. *The Tacit Dimension.* Chicago, Ill: The University of Chicago Press, 2013.

Rosebrock, Adrian. "How To Describe and Quantify an Image Using Feature Vectors." PyImageSearch. Last modified April 28, 2014. https://www.pyimagesearch.com/2014/03/03/charizard-explains-describe-quantify-image-using-feature-vectors/.

Scagliarini, Luca. "What is Machine Learning? A Definition - Expert System." Expert System: Artificial Intelligence: Cognitive Computing Company. Last modified October 5, 2017. http://www.expertsystem.com/machine-learning-definition/.

Stewart, M.G. "Modelling Human Error Rates for Human Reliability Analysis of a Structural Design Task - ScienceDirect." ScienceDirect.com | Science, Health and Medical Journals, Full Text Articles and Books. Accessed January 22, 2018. http://www.sciencedirect.com/science/article/pii/0951832092900975?via%3Dihub.

Todorovic, Aleksandar. "Has The Turing Test Been Passed?" Last modified 2017. http://isturingtestpassed.github.io/.