Artificial Intelligence and Relative Combat Power

To what extent can Artificial Intelligence be used to shift the balance of relative combat power in mid- to high-intensity ground combat?

Name: Joanna van der Merwe
Student number: s1426664
Program: MA International Relations, Global Conflict in the Modern Era
Supervisor: Dr. Lukas Milevski
Date: 30 August 2019
Word Count: 15 215 words

For my godfather, Capt. V. Martinelli SAN (Retd).

Acknowledgements

This thesis would not have been possible without my father, Sapper Danie van der Merwe, especially his patience and persistence in teaching me how to use the correct military ranks and titles and helping me understand how certain weapons function. Lt. Col. Björn de Heer for providing me the opportunity to write this thesis as part of my internship at the Land Warfare Centre of the Royal Netherlands Army, as well as for his constant support. Dr. Lukas Milevski for his understanding and continued belief in the unluckiest student. Arvid Halma, who helped translate my research and thoughts on Artificial Intelligence into a coherent chapter. And everybody else (you know who you are) who put up with me throughout the writing process.

Table of Contents

Introduction

“10 years from now, if the first person through a breach isn’t a friggin robot, shame on us” – former US Deputy Secretary of Defence Bob Work

It is not hard to imagine why the former US Deputy Secretary of Defence would expect that by 2025 robots will be on the frontline of wars (Roff & Danks, 2018). As Professor Klaus Schwab argues “we are at the beginning of a revolution” – the fourth industrial revolution – characterised by “the convergence of digital technologies with breakthroughs in material science and biology” leading to entirely new ways in which society will function (Schwab, 2017; Benioff, 2017). One of the fundamental technological developments that underpins this revolution is the advancement of Artificial Intelligence (AI), upon which many future technologies will be reliant in the same way current technology is reliant on electricity (Scharre, 2018). Growth in the abundance of data, increased computing power and breakthroughs such as Deep Learning have led to a renewed drive to create weapons systems that enhance military capabilities, one of which is robots on the front line rather than humans.

These new technological developments and the growth of AI have led to many arguing that militaries are on the verge of a Revolution in Military Affairs: an RMA (Schousboe, 2019; Brose, 2019; Goure, 2017)). These technologies pose both a great opportunity and threat for the militaries who will actively deploy them, but what both of these look like is hard to determine. Some of the opportunities include ‘better’ wars in which International Humanitarian Law (IHL) is upheld, fewer casualties are reported (both civilian and military), the use of force is more efficient and deterrence more effective. However, the threats look almost similar to the opportunities: The use of especially AI, may lead to difficulties in upholding IHL, and may prove to be incapable of handling the complexities of combat (Scharre, 2018; Schousboe, 2019; Brose, 2019). Furthermore, the democratisation of these technologies due to open-source movements, cyber espionage, cost reduction etc. and their ubiquity means that these developments may become more widely available to opponents, reducing the Western technological superiority (Microsoft News Center, 2016; Scharre, 2018; The National Interest Staff, 2019; FitzGereld, et al., 2016).

As the debate around the effects of technological developments such as AI progresses, a deep understanding of the technology and of warfighting is needed in order to assess the role that technology will have and what role we, as humans, want it to have. By focussing too much on the revolutionary aspects of AI, researchers run the risk of ignoring those elements of warfighting capability thus far determining the West's higher relative combat power, that cannot be replicated by a machine (Schousboe, 2019). At the moment, a clear image of how these technologies will (or not) revolutionise the military does not exist. It is the job of researchers, policy-makers, and those at strategic level, to ensure that sufficient understanding of the emerging trends is developed so that soldiers tasked with using these tech systems are given tools that enhance their capabilities, rather than become a risk and burden in the heat of fighting. This paper aims to contribute to this understanding of a specific technological development by answering the question *to what extent can Artificial Intelligence be used to shift relative combat power in mid- to high-intensity ground combat?*

Chapter 1 looks at the existing knowledge in the area of Revolution in Military Affairs (RMA), providing a working definition of what constitutes an RMA. Based on this definition it explores the arguments around recent developments that have been thought to be RMA's. Finally, it explores the modern system of force employment as the determinant factor of relative combat power for ground forces, having emerged as a result of the RMA that took place during World War One. Following this, chapter 2 provides an overview of the research design, including the case study selection process and the limitations and delimitations that were applied to this paper. It also discusses the availability of primary and secondary sources for the type of analysis used in this paper. Chapter 3 provides an overview of AI and the most recent development of Deep Neural Networks (DNNs) as one of the most promising areas for tasks such as automated image recognition. Chapter 4 provides a necessary overview of the first day of Operation Anaconda, the selected case study. Chapter 5 brings together chapter 3 and chapter 4 by discussing the requirements for DNNs to enable functioning in the context of the Operation, and what the implications are for challenges and opportunities that would have appeared on the battlefield. Finally, the conclusion discusses what the findings mean for the broader framework of a new RMA, as well as for the future of military developments around AI.

As this paper has strict limitations, the research can be broadened at a later stage. However, the author has chosen the following elements in order to stay within these limitations:

- Conventional forces engaged in ground combat
- Tactical level of warfare

- The theory of Modern System Force Employment
    - Suppressive Fire – one of the five characteristics included in this theory
- Based on this theory, mid- to high-intensity combat
- Deep Neural Networks and automated image recognition

Chapter 1: Literature Review

1.1. The Revolution in Military Affairs

Current discussions around the effects of technological development, including weapon development and outcomes of the broader information revolution, concern whether the technology is or will cause a Revolution in Military Affairs (RMA). An RMA can be defined as translation of major military innovations into increased military effectiveness, causing sudden increase in relative combat power application, and resulting in fundamental changes in tactical warfare patterns. This fundamentally alters the "character and execution of operations" (Benjamin Lambeth as cited by Rose, 2012), leading to new ways in which current and future wars are fought. This is coupled with overwhelming victory for one side, sparking a wave of reforms in militaries as the new tactical pattern is adopted. Although RMA as a field of study only appeared during the 1990s, in response to the United States (US) victory in the 1991 Gulf War scholars have applied the theories retrospectively, attempting to identify causal relations between certain factors (including technological developments and socio-cultural changes in society or state) and the related RMA. However, despite an existing definition of RMA, the number of revolutions – including when and why they occurred – varies widely and remains a topic of scholarly debate (Rose, 2012; Adamsky & Bjerga, 2010; Franck, 2004; Brun, 2010; Shapiro, 1999).

The Information-Technology Revolution

The importance of understanding RMA dynamics in the current climate of technological development, stems from arguments that the 1991 Gulf War represented an RMA resulting from significant improvements in Information Technologies (IT). Academics such as Dima Adamsky, Kjell Inge Bjerga and Elizabeth Stanley-Mitchell argue that the culmination of these technological improvements resulted in a *change in tactical patterns* intertwined with a new level of military effectiveness (Adamsky & Bjerga, 2010; Stanley-Mitchell, 2001). The military innovations that played a role in the IT-RMA resulted from interaction between various push and pull factors. In this context, push factors make changes possible whilst pull factors create needs for change in executing military operations. Theories regarding the revolutionary power of IT vary according to the understanding of interactions between push and pull factors. Shapiro identifies three main groups at the root of the revolution (where the pull will come from), namely broader society, politics, and the military (Kott, 2008; Shapiro, 1999)[1].

Firstly, theories concerning the revolution beginning in a general societal change encompass the idea that the Information Revolution challenges the Westphalian system, which will lead away from state-centric warfare and towards warfare that in current terminology involves 'non-state actors' as key players of the future (non-Westphalian) power structures. This will pull militaries into wars that differ from previous wars, in terms of the means by which they are fought and ends which they aim to achieve. They also challenge the notion of what would constitute a military, and who or what these militaries would fight for. Another facet of overarching societal change is the idea that IT developments lead to emergence of new organisational structures. It is argued that thriving structures in future will be smaller and more flexible units of a dispersed network, rather than the static, hierarchical structures that emerged during the Industrial Revolution. Thus, supporters of these theories hypothesise the RMA will be characterised by opposing networks, and military organisations will adopt these new organisational structures. Furthermore, those engaged in warfare may not necessarily represent a state (Arquilla & Ronfeldt, 1993; Shapiro, 1999; Brun, 2010).

The second group of theories argues that changed warfare will result from IT influencing politics. IT affects the political aspect of war, as it allows for increased flow of and accessibility to information. This influences how political actors believe they should or are permitted to act, by the groups they represent. The 'free flow' of information that these technologies facilitate alters the relative importance of different types of operations and tactics, thus altering the patterns of warfare (Shapiro, 1999).

These first two theory groups were formulated just after the 1991 Gulf War, attempting to predict how the IT-RMA would manifest in those years. Although certain elements of the theories can be seen twenty-eight years later with the rise of actors such as Al Qaeda and the Islamic State, the characteristics of warfare prior to the IT-RMA persist.

[1] Many theories attempted to predict how the IT-RMA would manifest in the years after the Gulf War, and are written in the future tense. This paper will maintain the tense used by the original authors, as it can be argued that the broader information revolution is still ongoing and therefore cannot necessarily be analysed as a historical event.

The third group focusses on a revolution rooted solely within the military sector. This type of revolution concerns changes regarding "the weapons used to fight battles; the targets they attack; the systems that provide command and control, logistical, and intelligence support for weapons; and the organizations that use the weapons systems" resulting from IT developments and thus, a change in tactical patterns of operations (Shapiro, 1999). Despite the civilian nature of major IT developments, the theory of an IT-RMA as defined in this paper, falls within this third group. Therefore, the dynamics of an RMA rooted in the military sector is further explored in this section.

The IT-RMA

IT-RMA pioneers included the US, the Soviet Union and Israel. However, each differed in whether their technological developments came before or after their intellectual conceptual developments. The Soviet Union began with conceptual and intellectual exploration of changes that could come about from trends in technological development, which were yet to materialise but which they predicted would occur. On the other hand, the US and Israel only began to explore the implications of the Information Revolution for execution of their military campaigns, once these technologies were already integrating into their armed forces (Adamsky & Bjerga, 2010).

The Information Revolution provided a number of push factors, allowing the military to change the characteristics of operations in the 1991 Gulf War. According to Stanley-Mitchell, the IT-RMA was enabled by ability to combine military innovations such as precision-strike weapons, and information technologies developed in the civilian sector (Stanley-Mitchell, 2001). This does not say the two types of technologies were not present in conflicts prior to the 1991 Gulf War, or that information did not play an important role in determining outcomes of prior campaigns. However, during this war the first instances of ability to network together different technological developments allowed for an increase in military effectiveness, which lead to an overwhelming victory for the US military. During the Vietnam War, precision-strike bombing was already utilised, but the lack of ability to effectively communicate the necessary information resulted in reduced decisiveness of the weapon effects. Therefore, during the 1991 Gulf War, when thanks to the Information Revolution the information could be communicated effectively and in a timely manner, these precision-strike capabilities turned into an overwhelming weapon against President Saddam Hussein's army. This war was characterised by increased reliance on airpower to overpower the opponent with little need to engage ground troops, especially in mid- to high-intensity close ground combat. Instead of engaging conventional infantry forces, special forces were used to identify and communicate targets for airpower. US engagements such as in Kosovo followed that same warfare pattern with similarly overwhelming results, leading many to believe that they had in fact experienced an RMA resulting from IT developments (Kott, 2008; Stanley-Mitchell, 2001; Nye & Owens, 1996).

The Other Revolution in Military Affairs

Despite the distinct change in character of the 1991 Gulf War and following engagements, criticisms of it as IT-RMA, and related theories, have developed. In his paper, "While You're Making Other Plans – The Other RMA", Brigadier General Itai Brun argued that whilst Western militaries have been focused on an RMA based on technological developments, those without access to these technologies have been experiencing their own revolution: The Other Revolution in Military Affairs (O-RMA) (Brun, 2010). The push factor of the O-RMA was lack of access to technology, whilst the pull factor was the gap between themselves and the technologically superior militaries they faced. This O-RMA has characteristics in direct opposition to the technology they try to counter, and is an acknowledgement of persistent technological inferiority.

Firstly, militaries implementing characteristics from the O-RMA focus on ensuring survivability of their forces by taking steps such as using low-signature systems and forces, for instance commando and guerrilla forces and tactics, and suicide bombers. Secondly, the rockets and missiles used by these forces are technically simple and low- cost, yet still penetrate far behind enemy lines with little chance of interception. Furthermore, these rockets are hidden in civilian facilities, as it is known that technologically superior armies who up until now have tended to be Western, are governed by laws that often render retaliation against such facilities unacceptable (e.g. by International Humanitarian Law). Additionally, these rockets and missiles aim to incur a large number of civilian and military casualties, in order to exploit aversion to casualties as a pressure point for ending the war. This is a result of the media effect[2], as the information influences how politicians respond. Those following the O-RMA also attempt to draw opponents into close combat, as this undermines the ability to make use of

[2] The media influence public opinion through their reporting and, in turn, what the public in democratic countries will allow their government to do (United States Marine Corps, 1998).

technological superiority. In close combat, actions such as aerial strikes become harder to undertake without endangering friendly troops (Brun, 2010).

Brun argues that the characteristics described above are developing into a broad enough area to constitute an RMA. However, the parallel existence of two opposing RMAs, the O-RMA and the IT-RMA, questions whether either of them are actual revolutions. As highlighted in the definition of an RMA, a new tactical pattern that leads to overwhelming victory results in a wave of reforms as militaries attempt to adopt this new pattern in order to maintain relative combat power. Furthermore, the characteristics of the O-RMA are argued to have emerged from Iraq's defeat during the 1991 Gulf War, demonstrating that the IT-RMA did not lead to universal adoption of similar tactical patterns (Brun, 2010).

By applying the aforementioned definition of an RMA to historical revolutions, the adoption of new tactical patterns of warfare is clearly demonstrated. During the Napoleonic RMA, France implemented various reforms such as the 'citizens' army' and a radically new, smaller unit structure. This allowed for implementation of new tactical patterns, such as strategic surprise and operational disruption. After Prussia experienced an overwhelming defeat, it adopted reforms allowing it to withstand and counter the radically new form of war brought about by Napoleon's military. These reforms created a structure that allowed Prussia to copy Napoleon's tactical fighting patterns. This sparked a wave of reforms as other militaries adopted the tactical patterns still used today as the basis for some aspects of military thought throughout the world (Franck, 2004).

Both the Afghan War that began in 2001 and the 2003 Iraq War demonstrated clashes between militaries focussed on IT-RMA, and actors who did not follow this pattern. These wars showed that technology alone did not create overwhelming victory. Although the latest military technological innovations were present, they did not translate into military effectiveness. The IT-RMA has not led to a universal attempt to mimic US actions in the first Gulf War, but rather to development of doctrines using weaknesses in a technology- dependent military and drawing a technologically superior army into a type of combat, such as close ground combat, where technology-based doctrine no longer provides major advantage. Therefore, it is key to assess the tactical patterns in these types of combat, to assess what the determinant in relative combat power is and gain insights into what could cause a shift and a possible RMA.

Strategic, Operational and Tactical Levels of Combat

There are 3 levels of military planning namely strategic, operational and tactical. As the definition of an RMA used for this paper highlights the importance in a change of tactical patterns of operation however, changes at this level have direct effects on the other two. Therefore, radical changes at the tactical level could lead to changes at all levels.

The strategic level is concerned with achieving national or coalition goals (Headquarters of the Army, 1993). Exactly what this entails has evolved since its first definition by Clausewitz as the “use of battle for the purpose of war” (Strachan, 2012). This definition focused solely on use of the military, making war the business of generals and not politicians. However, it came to represent the link between military and political decision-making bodies of states. This can be attributed to the realisation by those studying military victories, that at this level diplomacy and economics affect the chances of a state achieving goals which it previously would only have used the military to achieve. Thus, the strategic level has come to focus on war as the business of both politicians and generals (Strachan, 2012). The operations level provides the vital link between strategic objectives set out by generals and politicians at state level, and forces deployed on the ground. It converts strategic objectives into methods for using force. It ensures that the sum of smaller tactical movements achieves the larger goals of the war. This includes “design, organization, and conduct of campaigns and major operations” (Headquarters of the Army, 1993). The tactical level of war is focused on military engagements in which military units are manoeuvred by the operational level. The successes and failures at this level, set the conditions under which operational manoeuvres and decisions can be made (Headquarters of the Army, 1993). As the definition of RMA indicates, it is at the tactical level that the emerging patterns may or may not constitute an RMA.

Note that the three levels are inextricably linked. Therefore, significant changes at the tactical level will have significant effects on the operational level and then the strategic level, and vice versa. The decisions at each of the levels are not just defined by the actors themselves but also by actions and reactions of the enemy. The relationship between actions taken by actors involved in the different levels is not strictly linear but rather brings together the different parts (Luttwak, 2003). Although Stephen Biddle’s theory of force employment, discussed in the next section, focusses on the tactical and operational levels and this paper only covers the tactical level, the implications of findings at this level also have operational and strategic effects.

1.2. The Modern System of Force Employment

When proponents of the IT-RMA were drawn into contemporary wars where their technological superiority was reduced, such as Afghanistan 2001 and Iraq 2003, the battles maintained a tactical pattern that pre-dates the 1991 Gulf War. This brought into question whether the IT-RMA truly constitutes a revolution or simply the emergence of a new type of operation based on new types of technology. Therefore, if the IT-RMA is not seen as the latest revolution, then the last RMA to produce a tactical pattern adopted by all actors engaged in combat occurred during the First World War (WW1). It is accepted that during WW1 a revolution took place, though, the causes are debated. Some authors attribute the end of static force entrenchment to the breakthrough of the latest technological development at the time, the tank. Stephen Biddle, on the other hand, attributes the breakthrough to emergence of the modern system of force employment (Biddle, 2004). Until 1918, lethality of weapons systems had been increasing and forces faced by this firepower had begun to converge on the same ideas how to survive and make use of this. The history of how countries like United Kingdom developed these ideas will not be covered here, but successful implementation of these ideas of the modern system of forces employment broke the entrenchment. Despite increased importance of airstrikes and increased reliance on special forces during and following the 1991 Gulf War, modern warfare has shown that insertion of ground forces in mid- to high- intensity combat is still necessary. The patterns that emerged during WW1 are still determinant of relative combat power between forces (Biddle, 2004; Biddle, 2005; Cohen, 2004). Even those countries and actors which Brun argues are experiencing the O-RMA – including the 1991 Iraq military, Taliban forces in Afghanistan and Al Qaeda in Iraq –make use of the modern system of force employment elements as far as they are able. When engaged in this level of mid- to high-intensity ground combat, both militaries in the IT-RMA and actors in the O-RMA display tactical patterns that emerged during WW1 (Biddle, 2004; Brun, 2010). The key element of all these characteristics described below, is that that they are the most persistent post-WW1 tactical pattern of mid- to high- intensity combat – despite all technological developments since.

Characteristics

The term ‘force employment’ refers to decisions made by commanders in war, as well as doctrines that provide prescriptions and instructions used by commanders and soldiers. It is the non-material factors that an army relies on. Eliot A. Cohen refers to it as the "professional ability and temper of the organization" resulting in a specific pattern of "how forces are disposed and used" (Cohen, 2004; Biddle, 2005). According to Biddle, force employment is the determining factor between more and less powerful militaries, rather than preponderance or arguments based solely on technological superiority or development. It is force employment that enables a military to make use of its material capabilities, converting these capabilities into military effectiveness.

At the core of the modern system of force employment is the aim of reducing the soldier's exposure to enemy fire, whilst enabling friendly force movement and slowing down the movement of enemy forces (Biddle, 2004). In his book, *Military Power: Explaining victory and defeat in modern battle*, Biddle provides a description of this system's characteristics. It is these characteristics that marked the radical departure from the pre-WW1 system. Adoption of the modern system required changes at all levels of the military and changes in the organisational structures of the army, that led to this change being classed as an RMA during WW1. The five tactical characteristics are as follows:

- cover and concealment
- dispersion
- small-unit independent manoeuvre
- suppression
- combined arms integration

To achieve cover and concealment, soldiers need to use terrain features in order to deny a visible target for weapons to be aimed at. Due to technological developments that increase the lethality of weapons, finding adequate cover is increasingly difficult. The possibility of achieving cover and concealment requires the dispersed use of small independent units. Prior to WW1, forces were made up of relatively large units that moved in relatively larger formations (Franck, 2004; Biddle, 2004). Apart from allowing for adequate use of cover, the shift towards smaller, dispersed units places fewer troops within the blast radius of increasingly lethal weapons. Suppressive fire forces opponents to take cover, allowing own forces to move forward between different areas of cover. Prior to the emergence of the modern system, firepower attempted to destroy dug-out positions rather than enabling ground forces to move forward and destroy these positions. As every weapon or weapon system has strengths and weaknesses, the ‘combined arms’ approach reduces the overall vulnerability of a single-weapon by complementing the strengths and weaknesses of different weapons.

The interaction between these characteristics is fluid, and each can be implemented to a different extent at different points of engagement. This results in variance not only between different actors but also variance within a single military at different points in time, and between different parts of the military. When opposing sides both attempt to use the modern system, the varying extents to which they succeed will significantly affect the relative combat power in mid- to high-intensity ground combat (Biddle, 2004).

These five characteristics are applicable to actors in both offensive and defensive postures. Acting defensively, the modern system requires the same characteristics, but adapted to withstand the more fluid attacks by forces using the modern system. For example, cover and concealment is still necessary and carried out in a similar way. The aim of this characteristic in a defensive posture is to prevent the attackers from concentrating their firepower on an easily identifiable defensive position, whilst making it harder for attacking units to know where to move towards. Similarly, dispersed units in defensive positions reduce the number of losses caused by artillery strikes, and are harder to suppress. In order to successfully use dispersion, one needs a larger number of smaller units with higher degrees of independence. The implementation of this system is not achievable by all armed groups, and requires a high level of training and refinement (Biddle, 2005; Biddle, 2004).

Implementing Modern System Force Employment

Pre-WW1 doctrine prescribed tactical manoeuvres that involved the advance of infantry troops in mass formation. These tactical patterns had emerged during the Napoleonic Wars. However, as the lethality, range and accuracy of weaponry increased, this type of manoeuvre became costly. By the end of the 1800’s, strategists came to believe that “frontal assaults would in the future be not simply prohibitively expensive, but statistically impossible” (Howard, 1986). Despite this, troops still used these doctrinal prescriptions, and Germany’s Schlieffen Plan for WW1 invasion of France and Belgium can be seen as implementation of these tactical patterns at strategic level. However, these troops could only move forward a certain distance before reaching the point where lethality of small arms fire was too costly, creating a “no-man’s land” characteristic of WW1 entrenched warfare (Ellis, 1976).

The offensive and defensive tactical characteristics that emerged to overcome no-man's land were more complex than before, as they required loosening of the rigid structures that previously governed military manoeuvres. The increased independence of units and the decreased size of units require a larger number of commanders able to create their own unique plans, enabling them to make use of characteristics in the section above. They also need these capabilities much earlier in their careers. Furthermore, to achieve effective suppressive fire and combined arms, tight coordination between suppressive fire and forward movement is required. Because these widely spread units move at different speeds, coordination is not easy. As the modern system applies to mid- and high-intensity close combat, artillery is reliant on effectively coordinated and communicated plans in order to avoid fratricide. Tactical commanders become key in creating these plans, as they move further away from operational command. Therefore, not only is the environment of high-intensity combat complex in itself, the modern system force employment is complex and difficult to achieve, placing commanders at all levels under a high level of pressure (Biddle, 2004).

Chapter 2: Research Design

2.1. Research Question

Based on literature and theories exploring balances of power and the causes of revolutions in military affairs, the theory of modern system force employment is an exciting framework that can be used to analyse potential effects of AI on military balances of power through shifts in relative combat power at the tactical level. In order therefore to provide insight into the possibilities of AI as a revolutionary technology, this paper will explore the following question:

To what extent can Artificial Intelligence be used to shift relative combat power in mid- to high- intensity ground combat?

By answering this question using the theory of modern system force employment, the paper explores a small part of the larger possibility that AI may act as catalyst for the next Revolution in Military Affairs and/or shifts in military balances of power.

2.2. Methodology

The remainder of this paper contains chapters on Artificial Intelligence, the case study - a key day in the USA setback of Operation Anaconda, and a qualitative analysis of the insights that this provides into whether AI would have a positive or negative effect on the higher relative combat power held by the USA during the Operation. The case study is established using standard process tracing methodology.

2.3. Case Study Selection

Selection Criteria

As the theory of modern system force employment is limited to the use of conventional land forces, at the time of writing the number of wars in which this has occurred in modern history is limited. Therefore, the first selection criterion established for the case study was that at least one side involved in the operation used conventional ground forces[3]. The second criterion following from this, was that these ground forces used the modern system of force employment as established by Stephen Biddle. The third criterion was that the combat engagement reached the level of mid- to high-intensity[4]. This includes close engagement between the opposing sides.

Operation Anaconda

Operation Anaconda as described in chapter 5, was the first combat engagement since Vietnam that required deployment of conventional US forces against the non-conventional forces of Al Qaeda and the Taliban. In the first days of the operation the US forces were engaged in mid- to high-intensity combat for a continued period of approximately 4 days, using characteristics of modern system force employment.

2.4. Limitations and Delimitations

Delimitations

Suppression

Due to restrictions and subject complexity, this paper can only research one of the five characteristics that make up the modern system of force employment. The characteristic Suppression was selected as characteristic, on the basis of the determining role it played within the case study: Operation Anaconda. This characteristic is not implied as necessarily the highest operational importance in all cases. The following chapter will provide an overview of the Operation and the role of suppressive fire in determining the outcome. As highlighted in previous chapters, no characteristic stands alone but forms a part of an interactive whole. Thus, changes in ability to implement one characteristic bring about significant changes in the others.

In Biddle's book *Military Power*, he distinguishes between modern system usage in offensive or defensive postures. As established in the literature review, this paper focusses on the tactical level of offensive posture.

---

[3] Conventional forces fighting conventional wars include "combat fought between military forces on or over major land masses [excluding] war at sea, and [excluding] strategic bombing against civilian targets" (Biddle, 2004). Therefore, conventional forces in this case include those involved in fighting close range mid - to high intensity traditional combat.
[4] Mid- to high- intensity combat can be defined as battlefields that are "chaotic, intense, and highly destructive" (Department of the Army & United States Marine Corps, 1996)

Moral and Ethical Dilemmas

Another factor in the selection of the characteristic focused on in this paper is the moral and ethical dilemmas that arise in the use of technologies that may remove humans from decision-making cycles. This is especially important within a violent setting that directly or indirectly may result in taking human life. The author acknowledges that without answers to the moral and ethical challenges the question regarding implementation of AI in the context of the selected case study cannot be fully answered. The paper assumes that the AI would be morally and ethically acceptable as a technology to provide suppressive fire. This assumption is purely to allow for exploration of the results of this research. Additionally, the paper views suppressive fire as a task that does not primarily aim to take human life further than assisting human soldiers in their task of overcoming the enemy.

Limitations

Theory

The use of theory to explore the use of AI is limited to tactics by conventional land forces in mid- to high-intensity close combat. Therefore, due to the specific nature of the theoretical framework the paper does not purport to provide generalizable conclusions applicable across all branches of the military and all their tactical patterns of combat. Furthermore, it does not attempt to justify the use of modern system force employment nor argue for or against the use of AI in combat. The paper provides a methodological approach, that can be used for further such research to determine whether the use of AI at certain levels of military engagement will have an effect on the tactical patterns that determine relative balances of power.

2.4. Sources

Primary Sources

Access to primary sources about this operation was challenging for a number of reasons. When studying the tactical level of warfare, one has to interview and analyse the experience of a large number of soldiers who were involved at both the tactical and operational level. Firstly, gaining access to soldiers who were present at the tactical level was an obstacle that could not be overcome within the author’s current position, time and resources limitations. This operation took place over ten years ago, so many participants may no longer be serving in the military.

Locating and reaching out is a major challenge. Furthermore, even with access to the soldiers, the intensity of this form of combat can cause biases in what is remembered about the operation. This is due to the psychological effects of trauma as well as the fog-of-war[5]. The Royal Netherlands Land Warfare Centre with whom this research was conducted, was not represented during the operation, and primary sources could not be located within the Netherlands. Secondly, in order to gain an objective and wholistic view of the combat operation, the number of interviews needed is large as each soldier only sees an assigned piece of the battlefield.

Another challenge is generally limited declassified information availability about operations and tactical patterns. Sean Naylor explains this difficulty in his book *Not a Good Day to Die: The Untold Story of Operation Anaconda*. Naylor explains that Operation Anaconda was covered by gag orders until 2004. Furthermore, even after Naylor had submitted applications for release of information through the US Freedom of Information Act, he only received heavily redacted documents. This demonstrates just one of the challenges to gaining access and using primary sources for tactical level research (Naylor, 2006).

Due to the limited number of operations suitable for the purposes of this paper, the challenges facing primary source collection for these types of operation and analyses, and the limitations of a master thesis, this paper relies on secondary sources.

Secondary Sources

The need to rely on secondary sources played an important role in case study selection for this paper, as objective and well researched sources had to be readily available in a large enough quantity to allow a thorough tactical level understanding of the operation.

Operation Anaconda is a well-established and studied operation, with a high level of objectivity available in a number of sources. These sources are easily accessible in terms of time and clearance needs. These sources include books such as *Operation Anaconda America's First Major Battle in Afghanistan* by Laster W. Crau and Dodge Billingsley. This book is based on the authors' extensive interviews with over a hundred people, many of which were directly involved in the operation at tactical level. Additionally, one of the authors was present

[5] The fog-of-war refers the uncertainty that surrounds all warfare and distorts memory as Clausewitz states that "Each Commander can only fully know his own position; that of his opponent can only be known to him by reports, which are uncertain" (Clausewitz, 1968)

alongside the US forces during the operation. The second book is *Not a Good Day to Die: The Untold Story of Operation Anaconda*, in which a multi-page section can be found with names of those interviewed for the book. Furthermore, the author was "present at the rehearsals for and some of the combat during" the operation (Naylor, 2006). In addition to these two key books, a number of academic and non-academic articles and interviews were written about the operation. Therefore, in the absence of access to primary sources, Operation Anaconda was selected for the number and quality of secondary sources available about the tactical level of the engagement.

Chapter 3: Artificial Intelligence

Although the term 'Artificial Intelligence' (AI) was only coined in 1955 by researchers John McCarthy, Marvin Minky, Claude Shannon and Nathaniel Rochester at Dartmouth University, their research basis traces back to the 1940s. The work beginning in the 1940s explored logical functioning of neural networks in the brain. Alan Turing is considered the 'father of AI' due to his 1950 publication "Computing Machinery Intelligence". In this, he proposed the question "Can machines think?", establishing the fundamental question behind AI development (Negnevitsky, 2005; Anthes, 2017; McCulloch & Pitts, 1990; National Science and Technology Council , 2016). Findings from McCulloch and Pitts in exploring neural networks, Alan Turing and the research from Dartmouth provide the foundation for all AI research until now.

3.1. Defining Artificial Intelligence

There is no single accepted definition of AI. This is attributable to the interdisciplinary nature of the field, the broad applicability of technologies often referred to, and the preference of researchers to either employ a broad or narrow definition for their work. In their book, *Artificial Intelligence A Modern Approach*, Stuart Russel and Peter Norvig group definitions into four main categories based on whether the focus of the definition is on "thought process or reasoning", or "behaviour" (2010). A distinction is also made between groups that benchmark by intelligence measurement, human behaviour and abilities, or rationality (Russell & Norvig, 2010; De Spiegeleire, et al., 2017).

Because this paper explores cutting-edge developments in AI, a broad definition is taken. As highlighted by the National Science and Technology Council, the most important core aspects of AI research and development is "automation or replication of intelligent behaviour" (National Science and Technology Council , 2016). Therefore, the working definition for this paper is:

> *An algorithm with the ability to carry out tasks that require intelligence (in terms of ability to perceive, reason and act) when performed by a human, including activities such as learning, decision-making, and problem solving.*

Under this definition, the algorithm can be used to predict, identify, analyse and select targets, and to provide more accurate and detailed models. Furthermore, it can achieve cognitive engagement with a human, and without human intervention actively use its insights to decide on suitable actions (Russell & Norvig, 2010; National Science and Technology Council , 2016; Davenport & Ronanki, 2018).

A further classification necessary to understand the capability difference between types and levels of AI, focuses on comparison with human capabilities. Three groups are defined:

1) *Artificial Narrow Intelligence* – "machine intelligence that equals or exceeds human intelligence for specific tasks"
2) *Artificial General Intelligence* – "machine intelligence meeting the full range of human performance across any task"
3) *Artificial Super Intelligence* – "machine intelligence that exceeds human intelligence across any task" (De Spiegeleire, et al., 2017)

Armed forces using AI as defined by this paper is not new. There are many examples of systems that involve firing projectiles in a similar manner to what would be required if inserted into Operation Anaconda[6]. However, it is widely accepted that existing AI technologies all fall into the first category of Narrow Intelligence whilst the other two groups may be developed in the future (De Spiegeleire, et al., 2017; The MITRE Corporation, 2017).

3.2. Deep Neural Networks and Automated Image Recognition

A large number of subfields exist within AI research and development, such as "computer vision, natural language processing, robotics (including human-robot interactions), search and planning, multi-agent systems, social media analysis (including crowdsourcing), and knowledge representation and reason" (The MITRE Corporation, 2017). The field of computer vision is of particular interest as the basis of much commercial and defence sector self-driving vehicle development. Within the field of computer vision there are various subtasks that an AI, or combination of algorithms, can expect to undertake such as "detection, localization, recognition, understanding, classification, categorization, verification and identification"

[6] It must be noted that these systems are only allowed to function fully autonomously in defensive stances under the supervision of a human.

(Andreopoulos & Tsotsos, 2013). However, there is no universal agreement what each subfield means and what exactly that task would look like when carried out by AI. Furthermore, a number of technologies can be used to achieve computer vision, such as automatic image recognition or object detection through visual recognition of images, radar, sonar etc (The MITRE Corporation, 2017; Launchbury, 2017).

One of the latest AI developments is Deep Neural Networks (DNNs). As discussed earlier, machine learning and neural networks were early AI fields explored by researchers, with neural networks as foundation of machine and statistical inference learning (Scharre, 2018). Neural networks were a foundational area of research during the 1950s birth of the AI field. However, it took years for neural networks to develop present maturity levels, slowly becoming the "technology of choice" for various applications of AI. Currently developing Deep Neural Networks support deep learning as an AI function. DNN are essential for renowned recent developments such as IBM's Watson, Facebook algorithm, AlphaGO, and Uber's self-driving cars. For the purpose this paper, focus is on DNN abilities for automated image recognition such as in self-driving cars (The MITRE Corporation, 2017; McCulloch & Pitts, 1990).

The key difference between earlier neural networks and DNN is how they learn. For a neural network, a human builds the learning model. Building this model requires human identification of each feature of an object, picture, language etc. by which it classifies. For example, the human has to define a daisy has a green stem, five white petals, a yellow centre etc. A DNN removes the need for human identification of all features. A DNN can be fed the picture of the flower, and will identify each of the features present to build the model. The human then assists by associating that model with an object, 'labelling the data'. For example, the DNN uses pre-classified data consisting of different flower pictures. It identifies the features such as edges, texture, colour, and optical flow, all the way down to pixel level. As it analyses these flowers, a human approves how much it correctly classifies. The DNN is then provided with pictures of flowers not previously classified, and applies the model it has learned to these pictures to provide output of what classifications it has identified. This output may look as follows: 65,3% daisy, 23% rose, 11,7% unknown[7]. Its performance is measured by how much it classifies wrongly. Complexities in simple image recognition include foreshortening, aspect, occlusion

[7] For the purposes of this paper, explanation of this process is simplified.

and deformation. These lead to distorted images that may lead to misclassification or inability to classify (Launchbury, 2017; Russell & Norvig, 2010).

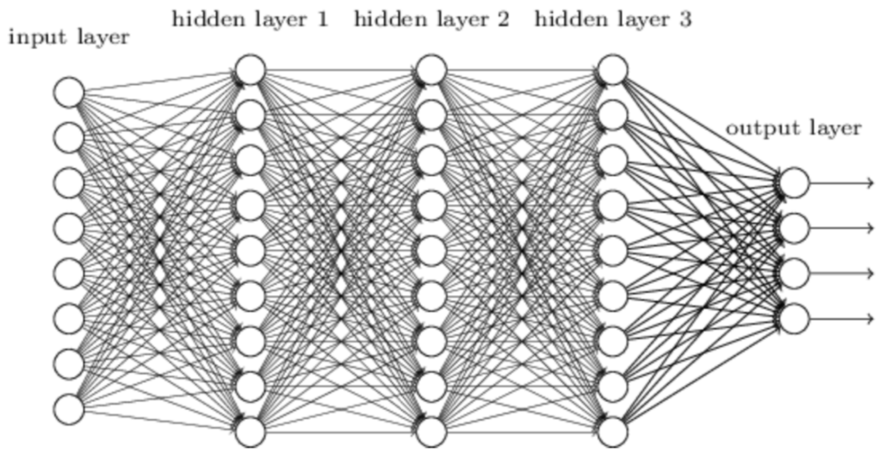

Figure 1: A simplified illustration of a Deep Neural Network. A modern DNN can have thousands to millions of nodes (the circles) and hundreds of hidden layers Source: (The MITRE Corporation, 2017)

Neural network functioning is based upon the strengths of connections that exist within the network. The network functions as follows – an input data is fed into the network and passed through many layers of these nodes, where each node calculates a weight and passes the data on to the next layer, eventually producing the classification. These weights are continuously adjusted until the classification matches the labelled input data, known as supervised learning of the DNN. The more layers, the more complex the network, the more complex tasks the network can handle. These technologies have outperformed humans on the same tasks, and it is known that DNN works efficiently in facial recognition, object detection and understanding (Scharre, 2018; The MITRE Corporation, 2017).

Rules

For functioning in the military, AI needs three levels of rules – higher level, middle level and lower level rules. The higher level includes legal obligations such as national laws of the country and International Humanitarian Law (IHL). They are the least granular and most flexible of the three levels, and are difficult to encode into a DNN via supervised learning due to their broad nature and their 'interpretability'. Middle level rules include documents such as operational objectives and rules of engagement. They are more granular and less flexible, designed for the context of operation in which the military is engaging. However, they are still a challenge to encode via the supervised learning techniques of DNN. These first two levels remain a functional area of DNN that has to be determined by a human. The lower level rules produce a set of classifications for what is detected, and (based on human understanding of the middle and higher-level rules) the way in which the DNN will handle certain classifications have to be decided and parameters set by a human. Therefore, the higher and middle classification features have to be written and programmed in a similarly to machine learning and neural networks (as opposed to deep learning and deep neural networks).

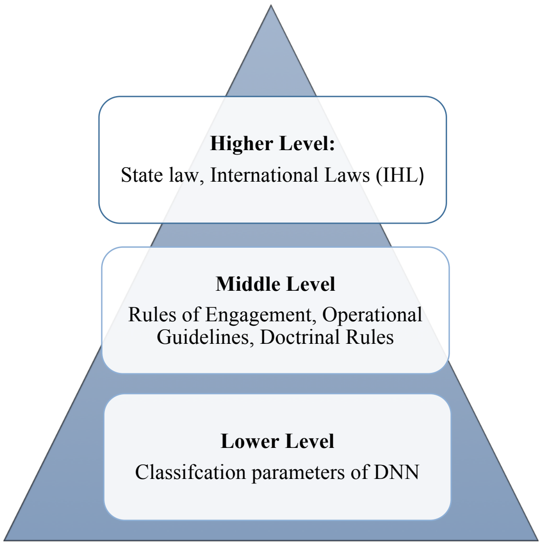

Figure 2: The parameters within which a DNN functioning in military settings can be divided into three main levels, with each level requiring a different type of intervention by human and different learning techniques for the neural network.

Lower level rules are least flexible and require high granularity, as these are the parameters of the algorithm and how the DNN classifies what it 'sees'. They can be encoded via fully automated supervised learning. Lower level rules help it 'know' things like where objects are located within a frame, the colour of those objects, their relative location to other objects, the classification of what the object is, etc. These also include a number of assumptions that need to be programmed. These rules are based on assumptions that allow object recognition to become a learnable problem. The assumptions include, amongst others, the size of search region that will be necessary, the resolution that will be necessary, what the degree of occlusion is that the machine will face, what features will be distinguishable, the cascade effects of a number of algorithms acting alongside one another and feeding into one another etc

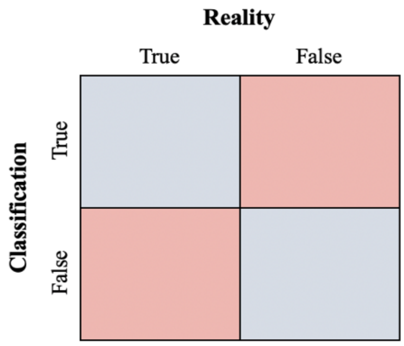

Figure 3: Classification matrix of a DNN classification vs the reality. If the reality matches what the DNN has classified, then it falls within the 'True' and 'True' area.

(Andreopoulos & Tsotsos, 2013). These rules and assumptions at the granular level are essentially the algorithm itself. The breakthrough and promise of DNN is that these lower level rules no longer have to be written by a human, they are written by the DNN itself. The eventual classification of an object by a DNN will fall into one of the four areas in Figure 3. The True-True and False-False area are not problematic as the DNN classification is correct. However, the two areas of True-False classification are the problem areas where middle and higher-level human-made rules for the DNN become important. For example, in the context of DNN providing suppressive fire support, 'True' may mean 'legitimate target' which can be fired at, whilst 'False' may mean 'non-target' which cannot be fired at. These areas of True-False classification therefore mean that either the DNN will fire at something that should not be fired at, or not fire at something that should be fired at. Both of these will have repercussions for the military, either leading to violations of IHL or to soldiers lacking the suppressive fire support they expect and need. The mid- and high-levels rules that the humans wish to uphold in the combat situation will guide how an AI robotic system based on DNN automated image recognition deals with these classification areas. Thus, the AI can either be 'trigger-happy' where humans have opted for broader rules in order to avoid the AI freezing and soldiers being left without support, or it can become 'scared to shoot' where humans have decided on narrower rules risking the AI not providing suppressive fire.

Data

A functional element necessary for AI, especially DNN, is data. This data is not only intrinsically linked to building assumptions and parameters for lower level rules, but also training and testing the DNN. It is the foundation of supervised learning and the promise of DNNs. The recent increase in data generation and collection, alongside the increased computing power of recent technological developments, increases the abilities of DNN for automated image recognition. A well-trained neural network must be trained on as large as possible relevant, labelled data sets. Although the size of data set a DNN can process is still limited by computing power, it is key that the data set be as large as possible. This ensures the highest level of accuracy in image and/or object identification. Besides training the DNN to a high level of accuracy, adequate data is vital to test the network. Standard method is to keep a portion of the data set for testing whether programming and training maintains accuracy on new unseen data. If the DNN classifies a target quickly but incorrectly, the problem may be rooted in either an insufficient amount of data resulting in an inadequate number of training examples of the object/image, or incorrect parameters that are overly broad or overly restrictive (Andreopoulos & Tsotsos, 2013; The MITRE Corporation, 2017; Russell & Norvig, 2010).

With the development of DNN, the interconnection between rules/parameters that make up the networks and the data used to train them becomes more important. These networks are able to discover the parameters on their own through supervised learning. Researchers argue this may be the solution to some challenges that exist around programming the assumptions and rules, arguing that truth lies in the data. A useful example of this is The MITRE Corporation's argument that rather than attempting to "hardwire the laws of aerodynamics into an autopilot application", a neural network is better off learning these laws of aerodynamics from data provided about these laws in action (The MITRE Corporation, 2017). Therefore, it becomes clear that data has an increasingly important role in development of AI technologies such as DNNs and automatic image recognition (The MITRE Corporation, 2017).

In summary, to use DNN in a military setting one needs large labelled data sets to develop the lower level rules and test these rules. Once these lower level rules have been established, a human must define the middle and higher-level rules which dictate how AI will act based upon these lower level classifications.

Chapter 4: Operation Anaconda

To explore the use of automated image recognition based on DNNs for mid- to high-intensity combat, this chapter provides an overview of Operation Anaconda. By discussing the DNN strengths and weaknesses within the context of how they would manifest in the past, it allows clearer understanding whether this technology has the potential to shift relative combat power according to Stephen Biddle's theory of modern system force employment.

4.1. Background

In response to the well-known 2001 "9/11" Al Qaeda attack, the United States (US) and allies began an expeditionary war in Afghanistan. It evolved into drawn-out counter-insurgency, but early phases called Operation Enduring Freedom aimed to destroy and remove the Taliban from power (Kuglar, et al., 2009; Grau & Billingsley, 2011).

Operation Anaconda was one major initial phase battle. In February 2002, CIA and Special Forces began reporting a number of Al Qaeda fighters in the Shahi-Kot Valley (Pakti province). Terrain, caves and Soviet era fortifications provided optimal Al Qaeda defences. Furthermore, its border location provided easy escape routes into Pakistan (Hastert, 2005; Naylor, 2006).

Figure 4: Approximate location of the Shahi-Kot Valley, Afghanistan (Unknown, 2002)

Operation Anaconda differentiated from previous operations as the first land battle since the birth of practically applied IT-RMA theory, that did not follow IT-RMA tactical patterns. In earlier operations of Enduring Freedom, the US made use of CIA operatives alongside Special Forces. These two would call in precision-strike bombing, without inserting conventional army battalions, brigades or divisions. However, these IT-RMA tactics failed during an earlier operation in Tora Bora, because the Taliban and Al Qaeda used aspects of pre-IT-RMA modern system force employment to avoid detection and escape aerial bombardments. They obtained relative victory by undermining the US combat power maintained until then, proving that the character and execution of an operation had not changed as significantly as US scholars and strategists believed. Furthermore, they adopted 'older' tactical patterns to combat the 'new' (Grau & Billingsley, 2011; Biddle, 2004).

Al Qaeda believed that despite the US technological superiority, technology would not provide advantage in the Shahi-Kot Valley. Al Qaeda estimated the US would not use precision-strike bombing because of high collateral damage risks, in terms of high civilian numbers and above-mentioned defences. Additionally, precision-strike bombing risked repeating the Tora Bora failure. This meant the US needed to deploy conventional forces, arguably not done since the Vietnam War. Thus, the US chose to resort to pre-IT-RMA tactics during Operation Anaconda (Naylor, 2006; Czarnecki, 2005; Hastert, 2005; Grau & Billingsley, 2011). The research question of this paper explores whether an automated image recognition system may have altered Operational success.

Only the initial D-Day[8] of the operation is discussed below. Important events allow for analysis of how AI may (or not) have been able to assist soldiers with suppressive fire. These events demonstrate the complexity of fighting at the tactical level during mid- to high-intensity engagement. Only specific occurrences are selected for elaboration. For clarity, events are divided into operational evolution from the perspective of soldiers on the ground, and in the air (helicopters or airplanes).

4.2. D-Day

On the ground

When US forces and Afghan allies entered Shahi-Kot Valley on March 2, 2002, they expected Operation Anaconda would be short and quick. Initial information estimated 200-500 Al Qaeda fighters in the valley. The US determined these fighters were light infantry, but not whether these were located in the valley itself or in the surrounding villages (Hastert, 2005; Grau & Billingsley, 2011).

As the US and allied Afghan forces began to move upwards into the valley, their aim was driving Al Qaeda towards other US troops inserted further up the valley. When the troops began to move, they came under heavy fire, killing a US Special Forces advisor and a number of

[8] D-Day covers the first 24 hours of the Operation.

Afghan troops. As a result, this group chose to fall back and reassess their movements. It was later discovered that the heavy fire had come from an AC-130 gunship providing US cover. Those on board later stated an inability to distinguish between the US / Afghan alliance convoy and Al Qaeda fighters. The first US engagement of the operation that lead to failure of the initial battleplan, turned out to be fratricide (Hastert, 2005; Grau & Billingsley, 2011).

Investigations after Operation Anaconda, revealed a number of failures, some more concrete than others. Firstly, the internal navigation computer system of the AC-130 malfunctioned earlier in the night. The crew thought the issue resolved after a computer reset, and continued to rely on it during these first stages of the engagement. However, it had not reset correctly. The crew were not seeing correct grid coordinates to align with ground forces communications. Therefore, when air crew and ground crew communicated coordinates and locations to fire at, they were unaware they would fire at their own forces (Naylor, 2006).

To prevent such mistakes there are checks for both ground and aircrew to identify friendly forces. Firstly, the aircrew use a light to glint ground forces they are looking at. Each friendly force vehicle supposedly had reflective glint tape indicating the AC-130 should not fire at them. The light is outside the spectrum of the human eye but visible through night vision goggles. Therefore, the US soldiers wearing goggles should have been able to see the AC-130 was 'glinting' them and communicate that they had seen it. Each vehicle also supposedly had an orange and purple marker to indicate friendly forces. It is unclear why these markers and signs were seen by neither ground troops nor aircrew (Naylor, 2006). Two other questions remain unanswered: Why the aircraft flight times do not match the narrative, and "why the navigation system seemed to work perfectly in the immediate aftermath of the attack […] steering the [AC-130] home to K2 [its home base in Uzbekistan] through bad weather exactly as it should have" (Naylor, 2006).

That same day, US forces were inserted by Chinook helicopters at designated landing zones throughout the valley. Their aim was to establish the line onto which the US and Afghan troops would drive Al Qaeda. These Chinooks did not come under fire initially, however pilot CW2 Stephen Brissit noticed behaviour that caught his attention. CW2 Brissit observed a man simply sitting on the roof of a house watching the helicopters dropping the soldiers off. Shortly after, this same man opened fire. Despite this man having been spotted by the pilot, neither the pilots nor the troops could fire upon him due to his classification as non-combatant until the moment he showed a weapon and became hostile (Grau & Billingsley, 2011).

As one group of soldiers inserted by Chinook, led by Captain Frank Baltazar, took in their surroundings some were surprised to be 150 meters away from a walled-off compound. Capt. Baltazar was not surprised and had determined through overhead satellite images that the compound was empty in the days leading up to the Operation. However, he still asked the battalion commander to clear it. As the team of soldiers were clearing the compound they identified that Al Qaeda forces who had been there (it was not empty as Baltazar had thought) were in possession of surprisingly high tech equipment and artillery including a "U.S. style tent", "U.S. PVS-7 night-vision goggles" and "binoculars" amongst other things (Naylor, 2006; Grau & Billingsley, 2011).

In the air

One major suppressive fire method of US troops during Operation Anaconda was from the air. Above, it was shown how the AC-130 gunship provided this. However, these gunships were not allowed to fly in daylight therefore, during the day suppressive fire was provided by Apache helicopters. The ways in which ground troops and the Apache helicopter pilots identified targets or target areas provide interesting analytical points on the AI role in similar situations.

As the ground troops inserted by Chinook oriented themselves after the initial surprise of intense fire by Al Qaeda, they began to spot Al Qaeda locations. Once locating the approximate target areas, they would communicate the location to Apache pilots in the valley.

One method to spot Al Qaeda fighters was by identifying the puff of smoke when a mortar is fired. Based on location information from the ground troops, the pilots would search for a mortar pit, which was difficult as Al Qaeda fighters wore "traditional Afghan brown woollen pakhul hats, with brown scarves" that camouflaged well in the natural landscape. Another sign of enemy mortar positions, was the kick-back of sand behind heavier weapons when fired. However, Al Qaeda was professional enough to wet this sand, making detecting their location more difficult (Naylor, 2006; Grau & Billingsley, 2011). Another method used by the helicopter pilots to identify the general area Al Qaeda was located, was by identifying the arches of tracer rounds fired at other helicopters and remembering the approximate location it came from (Naylor, 2006).

Not only were Al Qaeda challenging to locate, but also adept at avoiding the close air support – they were learning how to adjust rapidly to US methods. Whenever an aircraft appeared they would cease firing and retreat into the Soviet-era bunkers and caves. Once these aircraft left they would re-emerge, firing on the US troops. In order to overcome this, US troops had to adjust their own methodology to adapt to the unexpected way in which Al Qaeda was fighting and avoiding their technological superiority (Grau & Billingsley, 2011).

The troops inserted into the valley experienced heavy fire from the moment they stepped off the helicopters. They observed a number of enemy fighters moving along ridgelines surrounding the landing zones, although unknown to some of the inserted troops, there were also Special Forces located along the ridgelines. During the day these Special Forces fired down into the valley, attempting to provide supported to their own. However, they forgot to display VS-17 panel markers identifying them as 'friendly' for the inserted troops, who then began to fire on them. Once these soldiers displayed markers, it became clear they were friendly (Grau & Billingsley, 2011).

During the first day it became clear that Al Qaeda were fighting from prepared positions placed higher on the ridgelines, that provided them a level of superiority. These locations were placed in such a way to avoid standard US military procedures for providing suppressive fire. They were purposely placed outside the areas in which the US Air Force normally conducts pre-operations bombing. Furthermore, enemy fire was highly accurate due deployment of forward observers, as well as effective communication between groups of enemy fighters located in different areas along the ridgelines. US troops were unable to identify the precise enemy location. Even when the enemy was located, US troops could not fire upon them as most had been separated from their equipment. As a result, Al Qaeda effectively conducted suppressive fire that pinned down the US troops, rendering them unable to move and set up their own suppressive fire. Thus, Al Qaeda demonstrated that effective suppressive fire is a major factor in establishing relative combat power and for the main part of D-Day maintained the upper hand this way (Grau & Billingsley, 2011).

Chapter 5: Automated Image Recognition, Suppressive Fire and Operation Anaconda

5.1. Suppressive Fire

Throughout this first day of the operation there are instances where the conduct of suppressive fire is described in a manner that allows for the exploration of AI taking over the role. First the concept of suppressive fire needs exploration.

Suppressive fire can be defined as "rifle fire, which is precisely aimed at a definite point or area target" in order to "control enemy and area that they occupy" (McCullough, 2010). By using effective suppressive fire, soldiers move more easily around the battlefield and achieve their objective. There are three main elements to suppressive fire. Firstly, the fire must be aimed at the general area in which the enemy is known to be located. Secondly, it is important that the fire is well aimed within the area, to ensure the enemy remains suppressed. Thirdly, the fire must be sustained, meaning it must be consistent. Determining the most effective rate of fire to achieve consistency, depends on the tactical situation on the ground. Achieving all three elements becomes more difficult as the stress of combat increases (McCullough, 2010; Biddle, 2004). The advances in automated image recognition can play a role in the first and second elements of this type of fire. Therefore, the discussion in this chapter focusses on the role of automated image recognition in identifying a target or target area.

5.2. Discussion

5.2.1. Identifying a target

As highlighted in the definition of suppressive fire, the first element of target selection is identifying a target. During Operation Anaconda, target selection was carried out by soldiers on the ground and in the air.

On the ground

As the first US troops were inserted into the valley by Chinooks, CW2 Brissit noticed a man sitting on a roof watching as the helicopter offloaded the soldiers. This provides the first point for discussion, how an AI based on DNNs would manage such a situation. The main issue is the ability to distinguish between a combatant and a non-combatant. The higher-level rules of

IHL provide guidance regarding what the lower level of rules or parameters should distinguish, about who or what can be targeted when carrying out military action including suppressive fire. In practice this is a challenge – that human soldiers also face in these circumstances. In this case CW2 Brissit classified the man on the roof as a civilian that could not be fired upon but the many dimensions that go into this classification must be explored.

There are two main elements to the principle of distinction as established by IHL. Firstly, distinguishing between civilians and combatants, and secondly, distinguishing the status of combatants as legitimate targets[9] or *hors de combat*[10]. The first element is clearly demonstrated by the interaction with the man on the roof, whilst the second becomes apparent when exploring the challenges around the first element.

Distinction: Combatant vs. Non-Combatant

CW2 Brissit could not act against the man sitting on the roof as he was classified civilian, until the moment he drew a weapon - the man was not wearing a uniform to indicate he was active member of an 'armed force'. Al Qaeda in Afghanistan did not have distinctive uniforms but wore traditional Afghan clothing as worn by civilians. Thus, distinguishing between civilians and combatants based on the appearance of the individuals in this operation was a major challenge and would be a similar challenge for any DNN with automated image recognition. Establishing the parameters the DNN would use to distinguish between a civilian and an Al Qaeda fighter would rely on the ability to train its recognition of a human holding a weapon, and a large enough data set with these images would be needed for training and testing prior to deployment of the system for suppressive fire. The lack of accurate intelligence[11] in the lead-up to Operation Anaconda suggests that the necessary training data was not and would not be available.

[9] The Geneva Conventions offer protection to both medical units and religious actors attached to the armed forces (Geneva Convention I Article 19 and Article 24, 1949) but is unclear on the identification of religious actors.
[10] *Hors de combat* includes combatants who are "in the power of an adverse Party", "clearly express an intention to surrender", or have been "rendered unconscious or [are] otherwise incapacitated by wounds or sickness, and therefore is incapable of defending himself" (Additional Protocol I to the Geneva Conventions of 12 August 1949 | Article 41, 1977)
[11] The lack of accurate intelligence was simply because a lack of knowledge of how many Al Qaeda fighters were in the valley but also due to lack of data and intelligence sharing between different intelligence agencies (Grau & Billingsley, 2011).

Although the training data to positively identify an Al Qaeda fighter may not exist, the data does exist for identifying US soldiers. Therefore, the DNN could use a process of elimination to identify and classify an object or person. If the object/person is not a US soldier and not a civilian (i.e. is holding a weapon) then it could classify by default as legitimate target. This ability to classify friendly forces will be explored later in this section.

Even with access to adequate data sets for training the DNN and establishing the correct parameters, additional challenges include weaknesses around automated image recognition systems ability to identify targets on a 'noisy' background as well as human behavioural challenges. If a background is relatively static, it can be subtracted from the image being processed. In doing so a DNN is able to identify which pixels are the foreground and can over time track the movements of these foreground pixels. For highly structured behaviour such as ballet or tai chi taking place on relatively static backgrounds, this means DNN classification will be more accurate. However, high-intensity combat is not highly structured behaviour and the background relatively 'noisy', thus being able to detect the movement of the man on the roof will be a challenge to a DNN in this setting - especially once active combat begins (Russell & Norvig, 2010).

Related to this is the "lack [of] simple vocabulary of human behaviour" (Russell & Norvig, 2010). We do not have the ability to determine what feature of human behaviour indicates human intentions, nor do we have the ability to define what a human is currently doing. To counteract this, the strength of DNN and supervised learning is based on large labelled data sets that the algorithm uses to identify the necessary features for classification, without a human having to determine and program them. Thus, given the ability to overcome the challenge of 'noisy' backgrounds and a large enough data set, a DNN could identify human behavioural features thus far out of the realm of possibility.

Distinction: Legitimate and non-legitimate combatant targets

The second type of distinction a DNN needs, is ability to distinguish between categories of combatants. A DNN needs the ability to identify combatants who have surrendered or are indicating their intention of surrendering, combatants who are incapacitated and incapable of defending themselves, combatants who are medical personnel and combatants who are religious chaplains. Although in the case of Operation Anaconda there is little mention of these distinctions occurring or being necessary, an automated image recognition DNN needs to have these parameters established regardless. For these parameters, the training and testing data set (albeit the same as the one above or different), would need all the variations how these actions or features manifest themselves and/or are communicated.

In the case of Operation Anaconda elements described in the previous chapter, one can hypothetically discuss having to cease fire if Al Qaeda fighters surrender or are incapacitated and unable to defend themselves. What are the lower level rules that a DNN would use to determine that the target was no longer ‘a target’, and how do these need to be trained? There are some commonly understood signals for surrendering, such as waving a white ‘truce’ flag, visibly throwing away one’s weapon, or putting up one’s hands (International Committee of the Red Cross, Unknown). The International Committee of the Red Cross (ICRC) has established the norm of wearing a red cross, red crescent or red crystal as identifiers of medical personnel on the battlefield (International Committee of the Red Cross, Unknown). These parameters are likely to fall within the mid-level rules that need to be established and decided by a human. Because they are not universally accepted, it is likely that they will depend on the context - for instance the likelihood of what features will indicate a religious chaplain is likely to change based on the region or religious make-up of the opposing forces. However, programming these may provide an easy way for an opposing force to prevent themselves being targeted, due to the rigidity of the algorithms (discussed later), and increase the amount of human bias in the algorithm.

Distinction: Friendly vs. unfriendly combatants

In determining the distinction between combatant and non-combatant, it can be concluded that in the context of Operation Anaconda in Afghanistan, the main and possibly only way is to identify targets with weapons. However, in terms of automated image recognition software, another element needed is the ability to distinguish whether the target is friendly or unfriendly. As the US military would be able to build or readily have a large data set, one method would be training the DNN to identify US troops and have parameters to ensure it did not fire upon friendly troops. This may seem like a viable solution to this distinction challenge as mentioned above. However, the instance in which troops found US military equipment being used by unfriendly forces points to the problems of this approach. Should the DNN have identified the presence of the US issue tent, or an individual target wearing the other military equipment (which may be easily retrievable off deceased soldiers) such as the night vision goggles, two problems would have occurred based on the narrowness or broadness of the parameters. If parameters are set such that even one US issued piece of equipment leads to a friendly classification, the system may not provide suppressive fire when necessary. If the parameters are set such that there need to be a high number of indications of US issued equipment (such as uniform, guns, helmet etc.) then it may lead to friendly targets being classified as unfriendly when these features are obscured by dirt, simply not being worn, or came off in battle [12].

Exploring the intricacies of (legitimate) target selection and nuances of who can and cannot be targeted, it becomes clear any automated image recognition system deployed in this combat context needs training by data sets that contain all features playing a role in target identification. The aforementioned distinction elements that would need to take place, are only three of many and do not account for all manifestations of opposing force behaviour, dress, or identification. Furthermore, these features are those understood and used by human ways of identification - but the number of dimensions used by a DNN in order to 'see' is far greater than a human is even capable of imagining. The DNN will have to be trained on large data sets, and even more importantly tested on a number of large data sets to ensure its parameters for identifying a target not only hold up IHL but also enhance the capabilities of the mission rather than hindering it.

The question remains whether the algorithm will be viable in circumstances where a rare phenomenon not present in the training and testing data set, may be unidentifiable by the algorithm and render it unable to carry out what it was programmed to do (Russell & Norvig, 2010). This becomes an even larger challenge because the very nature of the human behaviour in warfare, is to surprise the other by doing something novel and unpredictable. Sun Tzu said "[s]trategies that have brought victory should not be repeated" (Tzu, 2009) (Clausewitz, 1968), but if something is not repeated there cannot be data to train and test how a DNN will 'see' it. In this case the ability to program parameters for features enabling identification are near impossible, and finding a large enough data set for the neural network to learn on its own may be equally impossible.

[12] For example, US Navy Seals in the valley were misidentified as enemy because their battle dress utility trousers could not be identified and they only displayed the VS-17 panels (for friendly force identification) after they were already fired upon (Grau & Billingsley, 2011).

From the air

Aboard the US AC-130 gunship a number of methods were used in unison, to identify a target for suppressive fire. Firstly, troops on the ground would communicate their grid coordinates to the personnel on board, providing guidance where the gunship should look to ensure the troops could continue their moves into the valley. Secondly, use of glinting and colour markers indicated what is not a target and should not be fired upon. In addition to these methods, the explanations provided after Operation Anaconda indicate that on board the gunship, US soldiers analysed geographical features to determine their location in relation to the ground troops and the communications received from the ground.

Key elements for exploration of automated image recognition abilities, are the glinting and colour marking, and the use of geographical markers. A data set for training a DNN on these may be already available or put together easily. Geographical markers are static in comparison to the other identification features discussed thus far. Furthermore, the necessary data can be obtained from satellite imagery and does not require data collection on the ground or in the area. Maps of the area also exist that could be used for training and testing the DNN. As highlighted in the case, US intelligence were aware of the terrain and fortifications as the Soviet War was fought here and Capt. Baltazar used satellite imagery in his planning.

The method of glinting or colour markers also provides a set of clear features to help automated image recognition identify friendly combatants, and preliminarily identify targets for further verification as legitimate (discussed above). Friendly forces can be informed of the DNN parameters and ensure they place markers to be easily identifiable by AI. The use of markers by US troops and allies is negative elimination, and could reduce some misidentification that occurred even with human oversight in Operation Anaconda. It does not necessarily alleviate a problem, but shifts it onto another medium. The issue in this case was that markers and signals were not seen at all – a problem that is also possible when using a DNN.

Trust and 'Understandability'

The experiences of the AC-130 personnel their computer malfunction point towards a larger reliance issue with digital technologies, especially AI based tech. They believed that they had solved the issue by restarting the system, and therefore trusted their displays without fully understanding what was wrong. Even post operation investigations did not pinpoint the flaw. This issue raises major concerns about trust and 'understandability' of autonomous systems in combat.

Human-machine teaming authors argue that there are two types of trust, 'over-trust' and 'under-trust'. Human over-trust readily accepts machine decisions, believing machines must inherently be right. Under-trusting causes humans to continuously check and take over tasks. However, this trust is closely linked to ability of understanding what the machine does and how. In this case, the personnel had no means of knowing that their computer reset did not fix the earlier errors. The system was either too opaque, a common criticism of 'black-box' AI, and/or the personnel had insufficient knowledge. This demonstrated the risks involved when those working with the machine do not know the workings behind the algorithm output, especially with a number of different AI technologies functioning together (Roff & Danks, 2018) As highlighted in chapter 4, a DNN is highly complex and requires highly sophisticated knowledge in order to train and understand. Even those who build and develop these DNN do not fully understand the way in which the weighting and cascading effects between nodes and layers truly works. The following challenge of adversarial images clearly demonstrates this (Roff & Danks, 2018).

Tricking the AI: Adversarial Images

A major challenge and weakness, directly related to inability of fully understanding how DNN works, are 'adversarial images'. These appear normal to the eye but contain elements leading to misclassification by DNN. For example, a DNN trained to differentiate between gibbons and pandas[13], could identify pandas in images with 57,7% accuracy. After an alternate vector was added over the panda, this same DNN misidentified the image as a 99.3% chance of being gibbon. This image still appeared as a panda to the human eye (with 100% accuracy). This demonstrates a fundamental difference between 'computer vision' and 'human vision'. The number of features a computer uses to classify an image is far greater than humans do consciously, or can even imagine. This highlights how easily this system could be 'tricked' without human awareness and understanding. This weakness shows up in most neural networks regardless of parameters or training methods.

[13] Two similarly coloured monkeys.

DNN "tricks" would lead to unintentional mistakes such as the fratricide during Operation Anaconda, and allows for purposeful manipulation of the system by an opposing force. This could lead to detrimental effects such as attacking civilian targets, friendly forces, or not attacking at all (The MITRE Corporation, 2017; Scharre, 2018).

5.2.2. Identifying a target area

Distinction: Friendly vs. unfriendly combatants

During the Operation soldiers used various methods to identify areas from which mortars were being fired. Many points discussed above apply to area identification by humans, but have corollaries in DNN application. The first example is the ground soldiers identifying plumes of smoke and sand that follow a mortar being fired (which they fired at and/or communicated approximate locations to helicopter pilots) and secondly, the Apache pilots following the arches of the tracer rounds. This solved the distinction between friendly and non-friendly target areas. Active firing of a weapon automatically makes the person a legitimate combatant.

Programming lower-level rules is easier as one can build a data set for training by firing these weapons in various contexts. As the sand blow-back, smoke and even tracer rounds, are generally applicable features, distinguishing whether it is friendly or unfriendly fire is a major challenge but not unique to DNN. A distinction method already discussed is use of markers such as glinting tape, colour panels and VS-17 panels. Previous points apply equally to target area target area identification.

Following of tracer rounds fired at a US troop presents a dangerous automated loop conundrum, should there be multiple DNNs. If algorithm parameters are set that tracer rounds fired towards US troops indicate an unfriendly combatant target area, it takes only one mistake (such as firing at the Special Forces who had not displayed their VS-17 panels) to start a chain reaction of DNNs firing at what they classify as unfriendly targets. Although (as during the Operation) the loop could be broken by a soldier moving or more clearly displaying friendly markers, the speed of high-intensity combat would make fratricide unavoidable. This type of DNN functioning would need rigorous testing, and the parameters at each of the levels of rules accurately enabling soldiers on the ground to prevent themselves being fired on, and how to stop the DNN once it starts firing.

Tricking the AI: Working Outside Its Parameters

The identification technique using the kick-back of sand behind a mortar provides an important point regarding parameters based on lower, middle and higher-level rules, and their rigidity. Al Qaeda was aware of the sand as means to identify their location, and simply wet the sand. Another example is locating their firing positions outside the zones normally bombed prior to the military operation. Finally, the enemy also adapted to US close air-support, successfully avoiding it. They clearly demonstrated the ability to learn US parameters and advantageously functioned outside these. This could be even more applicable to AI parameters.

Once a human soldier detects that enemy functioning nullifies their methods, he/she is able to adapt. The rigidity of AI parameters as a result of it having to be pre-programmed and tested, and its inability to adapt on the go without readjusting the neural net training make it less adaptable in combat than a human soldier. This brings forth another basic principle of warfighting, related to surprise as discussed earlier, as established by Sun Tzu: "The Shape of water is never the same and military force should be infinitely adaptable" (Tzu, 2009). Therefore, once an enemy fighter learns the parameters, the AI providing suppressive fire would cease to function as it should, thus reducing adaptability of the force and reduce rather than enhance relative combat power.

Conclusion

"What would a machine have done in Petrov's place? The answer is clear: the machine would have done whatever it was programmed to do, without ever understanding the consequences of its actions" – Paul Scharre, Army of None

Although Deep Neural Networks have shown the ability to perform to the same level or higher than humans in certain tasks, the use of automated image recognition systems based on these networks as an element of a weapon system to supply suppressive power will (to a large extent) not shift the balance of relative combat power between two opposing forces engaged in mid- to high-intensity ground combat. Although DNNs have shown promise in self-driving vehicles and other technological developments, there are a number of barriers to their ability for improving the military effectiveness of a force deploying an AI robotic system to provide suppressive fire. Firstly, the feasibility of the system autonomously distinguishing between types of targets on the battlefield is limited, including the ability to differentiate between a combatant, non-combatant, legitimate combatant target, and friendly and unfriendly combatants. This challenge is not novel nor limited to DNNs. It is an existing challenge experienced by the soldiers in Operation Anaconda. The novelty of the challenge for a DNN is the ability to train and test its distinction ability and reliability. A DNN would need a large labelled data set to begin classifying a target (a data set that is not currently available), and would need to be tested on a sufficient data set before it can be deployed. This testing also poses a novel challenge as the aim in warfare is to surprise the enemy, but how does one test, or even train, a DNN on a data set that does not exist because of the surprise factor?

One of the routes that can be taken to assist with distinction, is the one already in use by soldiers, namely the placing of markers that enable 'friendly' identification and establishing the parameters such that this allows for negative elimination. When looking closely at the reason the AC-130 gunship fired on its own troops, these are not problems unique to humans that would be alleviated by an automated image recognition system taking over. Furthermore, the malfunction of the computer system in this operation highlights that these systems may introduce new problems that are unique to DNNs and automated systems.

Unique problems also include trust, understandability, deception and the rigidity of the parameters within which the DNN functions. Human interaction with these systems brings to light the tendency to ‘over-trust’, or ‘under-trust’ an automated system because it must either ‘always be right’ or ‘always need double-checking’. This is closely related to the inherent lack of ‘understandability’ of the system due to its complexity and opacity. This issue becomes increasingly challenging when discussing the ease with which an adversarial image can trick an automated image recognition system without a human being able to detect that the deception is taking place. However, the largest challenge of a DNN functioning in the context of Operation Anaconda would be its rigidity and inability to adapt to the rapidity and complexity of mid- to high-intensity ground combat.

Finally, although the DNN ability to build its own models through supervised learning have removed the challenge that humans have to build the models for classification of objects and images, this does not remove the difficulties of humans having to develop the parameters within which the DNN will work. This supervised learning will only work for the lower level rules, whereas for the mid- to higher level rules a human will have to decide whether the system will be ‘trigger happy’ or ‘scared to shoot’. Either of these will provide challenges and problems for how the system will act in context such as Operation Anaconda, which may have deadly detrimental effects for the soldiers on the ground.

This paper has its limitations. The analysis only supplies insights into the potential of AI within mid- to high-intensity combat involving the use of ground forces engaging in modern system force employment. The paper begins to explore the potential, or lack thereof, for AI to bring about the IT-RMA that many argued had already come about as a result of the 1991 Gulf War. In order for an RMA to begin, major military innovations need to be translated into increased effectiveness that can be seen by fundamental changes in tactical warfare patterns. However, from the perspective taken in this paper, an AI system based on DNN at the current level of development lacks adaptability, which can be argued as the most important capability at the tactical level (it has remained fundamental since its delineation in Sun Tzu’s writings). Without adaptability, a system cannot handle the element of surprise – another key foundation of all warfighting. Therefore, without the ability to realise these fundamental aspects of warfighting, it is questionable how simply making use of AI would bring about an RMA.

It does not claim that AI has no potential to bring about the next RMA. The way in which the WW1 RMA was brought about was adaptation to increased lethality in weapons. Should AI technology be employed in a manner that does lead to increased lethality in mid- to high-intensity ground combat, that requires another form of force employment than the modern system, and leads to universal adoption of this new form of force employment, then it will have led to an RMA. However, this initial analysis indicates that it cannot, at the current level of development, increase the lethality of fighting and thus this military innovation is not being translated into sufficiently increased military effectiveness for an RMA.

The implications of the findings of this paper are that the way in which AI is approached may need to be altered in order to fully understand its potential. This paper aimed to explore what it is that humans do to achieve a higher relative combat power than their opponents - and then explore the ability of an AI technology to replicate this. What this methodology begins to explore is what are the key elements that an AI must be able to achieve, and what should be left to humans. War is a human endeavour and in turn, the technology used to fight it must be an extension of or support what a human is capable of. An example of a similar technology is that of gun powder. Gun powder alone was not enough to revolutionise fighting, but adapting gunpowder to meet the needs of humans in different types of combat, allowed humans and the technology (in human hands) to reach its full potential. In order to build a full picture of this, further methodical research is required to understand what is it that makes a technology revolutionary, how precisely to turn certain technologies into military effectiveness, and understanding the human-technology relationship balance that is needed.

Further Research

1. In order to fully complete the understanding of DNNs in Operation Anaconda and modern system force employment each of the characteristics need to be researched and analysed as this paper does. Furthermore, to understand the full extent of AI capabilities the robotics hardware, as well as other AI technologies, should be added to the analysis. No single AI technology will act on its own and the dynamics of the different technologies need to be researched.

2. This paper only looked at mid- to high-intensity ground combat using modern systems force employment. To gain a full understanding of the capabilities and focus of AI, the methodology used in this paper can be applied to other types of combat in order to gain a full understanding of the necessary military capabilities as a whole.

3. The key to ensuring DNN inserted at this level of combat is useful and safe, is adequately large data sets including for rare phenomena, for training and testing – these need to start being built. Coupled with research into best ways to deploy AI, and which types of AI technology would lead to increased military effectiveness, further research into data collection methods and needs can be done. Thus, when the way in which militaries will use AI is determined, the necessary data sets already exist thus, speeding up the process for training and testing.

Bibliography

Adamsky, D. & Bjerga, K. I., 2010. Introduction to the Information-Technology Revolution in Military Affairs. *Journal of Strategic Studies,* 33(4), pp. 463 - 468.

Additional Protocol I to the Geneva Conventions of 12 August 1949 | Article 41 (1977).

Andreopoulos, A. & Tsotsos, J. K., 2013. 50 Years of object recognition: Directions forward. *Computer Vision and Image Understanding,* Volume 117, pp. 827 - 891.

Anthes, G., 2017. Artificial Intelligence Poised to Ride a New Wave. *Communications of the ACM,* 60(7), pp. 19-21.

Arquilla, J. & Ronfeldt, D., 1993. Cyberwar is Coming!. *Comparative Strategy,* 12(2), pp. 141 - 165.

Benioff, M. R., 2017. Forward. In: K. Schwab, ed. *The Fourth Industrial Revolution.* London: Penguin Random House, p. v.

Biddle, S., 2004. *Military Power: Explaining Victory and Defeat in Modern Battle.* Princeton: Princeton University Press.

Biddle, S., 2005. Military Power: A Reply. *Journal of Strategic Studies,* 28(3), pp. 453 - 469.

Brose, C., 2019. *The New Revolution in Military Affairs: Wars Sci-Fi Future.* [Online] Available at: https://www.foreignaffairs.com/articles/2019-04-16/new-revolution-military-affairs [Accessed 27 August 2019].

Brun, B. G. I., 2010. 'While You're Making Other Plans' - The 'Other RMA'. *Journal of Strategic Studies,* 33(4), pp. 535 - 565.

Clausewitz, C. v., 1968. *On War.* London: Penguin Books.

Cohen, E. A., 2004. Change and Transformation in Military Affairs. *Journal of Strategic Studies,* 27(3), pp. 395 - 407.

Czarnecki, J. E., 2005. Not a Good Day to Die: The Untold Story of Operation Anaconda (review). *Project Muse.*

Davenport, T. H. & Ronanki, R., 2018. *Artificial Intelligence for the Real World,* Unknown: Harvard Business Review.

De Spiegeleire, S., Maas, M. & Sweijs, T., 2017. *Artificial Intelligence and the Future of Defense,* The Hague: The Hague Centre for Strategic Studies.

Department of the Army & United States Marine Corps, 1996. *FM 3-100 MCWP.* Unknown: US Army Chemical School.

Ellis, J., 1976. *Eye Deep in Hell: Trench Warfare in World War 1.* Baltimore: Johns Hopkins University Press.

FitzGereld, B., Levin, P. L. & Parziale, J., 2016. *Open Source Software & The Department of Defense,* Washington D.C.: Center for a New American Security.

Franck, R. E., 2004. Innovation and the Technology of Conflict During the Napoleonic Revolution in Military Affairs. *Conflict Management and Peace Science,* Volume 21, pp. 69 - 84.

Geneva Convention I Article 19 (1949).

Goure, D., 2017. *The Next Revolution in Military Affairs: How America's Military Will Dominate.* [Online]
Available at: https://nationalinterest.org/blog/the-buzz/the-next-revolution-military-affairs-how-americas-military-23833 [Accessed 27 August 2019].

Grau, L. W. & Billingsley, D., 2011. *Operation Anaconda: America's first major bettle in Afghanistan.* Lawrence: University Press of Kansas.

Hastert, P. L., 2005. Operation Anaconda: Perception Meets Reality in the Hills of Afghanistan. *Taylor & Francis,* pp. 11-20.

Headquarters of the Army, 1993. *FM 100-5.* Washington DC: United States Army.

Howard, M., 1986. Men against Fire: The Doctrine of the Offensive in 1914. In: P. Paret, ed. *Makers of Modern Strategy: From Machiavelli to the Nuclear Age.* Princeton: Princeton University Press, pp. 510 - 526.

International Committee of the Red Cross, Unknown. *Medical Personnel.* [Online]
Available at: https://casebook.icrc.org/glossary/medical-personnel [Accessed 29 August 2019].

International Committee of the Red Cross, Unknown. *Surrender.* [Online]
Available at: https://casebook.icrc.org/glossary/surrender [Accessed 29 August 2019].

Kott, A., 2008. Introduction. In: A. Kott, ed. *Battle of Cognition.* Westport: Greenwood Publishing Group, pp. 1 - 10.

Kuglar, R. L., Baranick, M. & Binnendijk, H., 2009. *Operation Anaconda: Lessons for Joint Operations,* Washington DC: National Defense University.

Launchbury, J., 2017. *A DARPA Perspective on Artificial Intelligence.* Unknown: DARPAtv.

Luttwak, E. N., 2003. *Strategy: The Logic of War and Peace.* Second ed. Cambridge: Harvard University Press.

McCulloch, W. S. & Pitts, W., 1990. A Logical Calculus of the Ideas Immanent in Nervous Activity. *Bulletin of Mathematical Biology,* 52(1/2), pp. 99 - 115.

McCullough, J., 2010. Chapter 4 Combat Fire Techniques. In: J. McCullough, ed. *The Ultimate Guide to U.S. Army Combat.* New York: Skyhorse Publishing, pp. 227 - 245.

Microsoft News Center, 2016. *Democratising AI.* [Online]
Available at: https://news.microsoft.com/features/democratizing-ai/ [Accessed 29 August 2019].

National Science and Technology Council , 2016. *Preparing for the Future of Artificial Intelligence,* Washington D.C.: U.S. Government: Executive Office of the President.

Naylor, S., 2006. *Not a Good Day to Die.* London: Penguin Books.

Negnevitsky, M., 2005. *Artificial Intelligence: A Guide to Intelligent Systems.* Second Edition ed. Harlow: Pearson Education Limited.

Nye, J. S. J. & Owens, W. A., 1996. America's Information Edge. *Foreign Affairs,* 20(2), pp. 21 - 36.

Roff, H. M. & Danks, D., 2018. "Trust but Verify": The Difficulty of Trusting Autonomous Weapons Systems. *Journal of Military Ethics,* 17(1), pp. 2-20.

Rose, P., 2012. Divinising Technology and Violence: Technopoly, the Warfare State, and the Revolution in Military Affairs. *Journal of Contemporary Religion,* 27(3), pp. 365 - 381.

Russell, S. & Norvig, P., 2010. *Artificial Intelligence: A Modern Approach.* Third Edition ed. Upper Saddle River: Pearson Education Inc..

Scharre, P., 2018. *Army of None: Autonomous Weapons and the Future of War.* First Edition ed. New York: W.W. Norton & Company.

Schousboe, L., 2019. *The Pitfalls of Writing About Revolutionary Defense Technology.* [Online]
Available at: https://warontherocks.com/2019/07/the-pitfalls-of-writing-about-revolutionary-defense-technology/ [Accessed 28 August 2019].

Schwab, K., 2017. *The Fourth Industrial Revolution.* London: Penguin Random House.

Shapiro, J., 1999. Information and War: Is it a Revolution?. In: J. P. W. Zalmay M. Khalilzad, ed. *The Changing Role of Information in Warfare.* Santa Monica: RAND.

Stanley-Mitchell, E. A., 2001. Technology's Double-Edged Sword: The Case of US Army Battlefield Digitisation. *Defense Analysis,* 17(3), pp. 267 - 288.

Strachan, H., 2012. Strategy and War. In: J. Lindley-French & Y. Boyer, eds. *The Oxford Handbook of War.* New York: Oxford University Press Inc., pp. 30 - 43.

The MITRE Corporation, 2017. *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD,* McLean: JASON The MITRE Corporation.

The National Interest Staff, 2019. *A Frightening Fact the F-22 and F-35 Share (And Why Chinese Stealth Fighters Are in the Skies).* [Online]
Available at: https://nationalinterest.org/blog/buzz/frightening-fact-f-22-and-f-35-share-and-why-chinese-stealth-fighters-are-skies-54392 [Accessed 29 August 2019].

Tzu, S., 2009. The Art of War. In: A. Sadler, ed. *The Chinese Martial Code.* North Clarendon: Tuttle Publishing, pp. 79 - 120.

United States Marine Corps, 1998. *Media Inlfuence and its Effects on Military Operations,* Quantico: United States Maring Corps.

Unknown, 2002. *Battles yet to come.* [Online]
Available at: https://www.economist.com/unknown/2002/03/04/battles-yet-to-come
[Accessed 29 August 2019].