Eksamen i Emne OPG3401

Bacheloroppgave

Implementing Lethal Autonomous Weapon Systems Into Mission Command Leadership

av

Kadett Mikkelsen

Godkjent for offentlig publisering

Publiseringsavtale

En avtale om elektronisk publisering av bachelor/prosjektoppgave

Kadetten har opphavsrett til oppgaven, inkludert rettighetene til å publisere den.

Alle oppgaver som oppfyller kravene til publisering vil bli registrert og publisert i Bibsys Brage når kadetten har godkjent publisering.

Oppgaver som er graderte eller begrenset av en inngått avtale vil ikke bli publisert.

Jeg gir herved Luftkrigsskolen rett til å gjøre denne oppgaven tilgjengelig elektronisk, gratis og uten kostnader	[X] Ja	[ ] Nei
Finnes det en avtale om forsinket eller kun intern publisering? (Utfyllende opplysninger må fylles ut)	[ ] Ja	[X] Nei
Hvis ja: kan oppgaven publiseres elektronisk når embargoperioden utløper?	[ ] Ja	[ ] Nei

Plagiaterklæring

Jeg erklærer herved at oppgaven er mitt eget arbeid og med bruk av riktig kildehenvisning.

Jeg har ikke nyttet annen hjelp enn det som er beskrevet i oppgaven.

Jeg er klar over at brudd på dette vil føre til avvisning av oppgaven.

Dato: 01 – 05- 2023

Contents

1 Introduction

1.1 Actualization

Artificial intelligence (AI) is a field of research still in its infancy, even a commonly accepted definition for the term has yet to be decided. Yet experts within the field, such as Paul Scharre, Michael Horowitz and Robert Work, believe a new technological revolution on the scale of the industrial revolutions is underway, starring AI in the lead role (Horowitz, Scharre, & Work, 2018). AI is an enabling technology with a vast number of applications. Kevin Kelly has famously compared AI to electricity. Giving objects autonomy will vastly increase their efficiency, much like electrifying them did over a century ago (Kelly, 2014). Much like the industrial revolutions reshaped society and thereby the battlefield, shaping the history of warfare, the AI revolution may come to shape the future of warfare. Regardless of whether this is true, it is not a matter one can afford to overlook. Time and time again history has shown the consequences of falling behind during revolutions in warfare. German interwar capitalization on new technology such as aircraft, tanks and radio led to astonishing results during the invasion of France during WWII. The battlefield rarely grants second chances (Scharre, 2018, pp. 93-94).

1.2 Topic

This thesis will be centred around the leadership of artificial intelligence (AI) in a military context. Based on mission command leadership the thesis will provide perspectives on what a military leader should consider when implementing AI. The purpose of the thesis is to provide a better understanding of how to implement and interact with AI in a mission command based military hierarchy. As well as discovering the most central ethical and practical implications of doing so.

1.3 Constraints.

This thesis will be centred around lethal autonomous weapon systems (LAWS) as a specific subset of AI. Utilizing AI to deliver weapons is a controversial topic which has been debated ever since the idea became feasible. This makes it the most interesting subject for this thesis, in the authors opinion. Additionally, commanding LAWS is a challenging prospect. Commanding a weapon system capable of taking lives already comes with an array of challenges and responsibilities. Layering AI on top of it adds a whole new set of challenges and responsibilities on top of the previous ones. Creating a better understanding of considerations and knowledge necessary to utilize LAWS safely and effectively in a mission command structure will hopefully contribute to dampening the difficulty of this task.

1.4 Definitions

1.4.1 Artificial Intelligence

Anyone who starts researching artificial intelligence (AI) will quickly discover that the term has no commonly agreed definition. The term artificial intelligence is used to refer to a wide range of computer technologies able to draw some conclusion based on large amounts of data. The term does not distinguish between different ways to accomplish this task (Davis, 2019, p. 116). However, this is not unusual within science, many scientific concepts do not get proper definitions before the research has matured. Given the infancy of AI research it may be unrealistic to expect a common definition as of now (Wang, 2019). For this thesis the definition of AI as *a system's ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation* (Kaplan, 2021), will suffice.

1.4.2 Lethal Autonomous Weapon Systems

Much like AI itself and for the same reasons, there is no international consensus on the term Lethal Autonomous Weapon System (LAWS). Even within the US military the term is also referred to as Lethal Autonomous Weapons or Autonomous Weapon Systems. This thesis will be utilizing the term Lethal Autonomous Weapon Systems as defined by the US Department of Defence in DoDD 3000.09:

A weapon system that, once activated, can select and engage targets without further intervention by a human operator. This includes human-supervised autonomous weapon systems that are designed to allow human operators to override operation of the weapon system, but can select and engage targets without further human input after activation. (US Department of Defense, 2012, pp. 13-14)

In essence LAWS are weapons given the autonomy to act on their own. How restrictive this autonomy is can vary and as such there are a couple different subdivisions within LAWS. This will be further elaborated on under subchapter 3.2 Autonomy.

1.4.3 Mission Command

Mission command is a form of military leadership philosophy based on the combination of centralized intent and decentralized execution. It promotes initiative, freedom and speed of action for subordinates within constraints defined by their superiors.

The core of mission command is to allow subordinate units to act as they see fit, believing that they have a better understanding of a given situation as they experience it first-hand. This enables quicker decisions that can utilize opportunities effectively as they arise. It does however require subordinates to have a firm understanding of the mission's objective along with the commander's intent and reasoning behind it. Ensuring that the subordinate units indeed achieve the goals as required (Norwegian Armed Forces, 2020, p. 13).

While Mission Command is advocated by many militaries around the globe, it has proved difficult to fully achieve in practice. The intricacies of why this is the case, and the current state of mission command is interesting in its own rights. However, these subjects will not be discussed in this thesis. The thesis will be examining the ideal of mission command and how LAWS fit into it.

1.5 Research Question

To what extent are lethal autonomous weapon systems compatible with mission command philosophy and what are the most important ethical and practical considerations when implementing them?

2 Method

This chapter will explain the authors approach to answering the research question. This includes how the data was collected and analysed as well as the authors own influence on the study.

2.1 Analysis of the Research Question

The research question itself forms a practical limitation on how it can be answered. It seeks an answer to the degree of compatibility between the potential of AI and the philosophy behind mission command. The author deems it difficult to answer based on numerical data. The validity of the answer is instead based on the author drawing parallels between recorded perspectives and opinions on mission command and the use of LAWS separately. The author therefore deems qualitative research superior to quantitative research in attempting to answer the research question of this thesis.

2.2 Research Method

The author has deemed qualitative research to be the best approach to answering the research question. Qualitative research utilizes opinions and experiences recorded in non-numerical data (Jacobsen, 2015, p. 65). The aim of the thesis is to describe the theory and concepts within two separate fields of research, the relation of which have not been studied enough. By utilizing a qualitative research approach the author intends to uncover a more nuanced perspective on the two fields of research combined. Revealing more details and context to enable a deeper discussion. This makes the thesis a qualitative literature review with a thorough analytical component in its discussion.

2.3 Research Design

Discovering how well LAWS implement into mission command is an open-ended problem. Furthermore, given the infancy of the LAWS, relatively little research on the field and even less on its correlation to mission command. Therefore, to answer the research question, this thesis is based on exploratory research. Exploratory research is used when a problem is not clearly defined. It is used to broaden a research field and attempt to highlight new factors within it (Jacobsen, 2015, pp. 79-81). Exploratory research is a necessity for this thesis due to the lack of previous research on the topic of LAWS utilized through mission command leadership. The reflections that derive from this thesis will be new factors within the field formed by drawing parallels between two separate fields of research: Mission command and LAWS.

2.4 Research Data

Autonomous weaponry is a relatively new concept and as such there has been a limited amount of time to study and write about it. Additionally, as new advanced weaponry is involved, such studies, or certain details about them, may be considered national and/or trade secrets, limiting the open access to them.

The source I have relied the most on through this thesis is the book *Army of None* by Paul Scharre (2018). Another book, *Mission Command* by Donald Vandergriff (2017) makes up most of the remaining theory section. The rest of my sources are mostly various journal articles and some research reports sourced from JSTOR. Examples include; *Artificial Intelligence on the Battlefield* by Zachary Davis (2019), *Artificial Intelligence and keeping humans "in the loop"* by Robert Mazzolin (2020) and *Towards Responsible Autonomy* by Amit Arkhipov-Goyal and Esther Chavannes (2019).

When factoring in the infancy of autonomous weaponry, as well as the limited access to study data about them, the possible perspectives this thesis can highlight become somewhat limited. The thesis will still however, be able to highlight some useful insight which could inspire further reflections tied to data this thesis has not covered. The perspectives covered in this thesis will be based on connections the author has drawn between open access, mainly secondary source, literature. Secondary sources come applied with interpretation, analysis and thereby bias. While the bias itself is not desirable, this thesis is still mostly based on secondary sources due to the author having no formal expertise in the field. The interpretation and analysis that secondary sources deliver is therefore crucial to providing a proper answer to the research question. However, this also means the author might have a hard time decerning the bias involved in the collected data.

2.5 The Author's Positioning

The thesis author's lack of formal expertise on autonomous weaponry implies a low amount of personal bias on the topic. At the same time, it may increase susceptibility to bias tied to the expert opinions incorporated into the thesis data. The author does, however, hold a positive view of implementing autonomous weapons to some degree, given that humans retain some form of control. The author has a decent theoretical understanding of mission command leadership philosophy. As well as some experience of its use in practice, though never with artificial intelligence involved. The author is generally very positive to mission command philosophy and its ideal.
Details regarding mission command often differ from nation to nation. It is therefore worth mentioning that the author is mainly educated in, and has only utilized, Norwegian mission command philosophy.

3 Theory

3.1 Mission Command Leadership

The term Mission Command has always proved somewhat elusive. It's often been misunderstood from its origin (Wandergriff & Webber, 2017, p. 51). This chapter provides a general description of Mission Command before taking a look at its historical origins in order to communicate the details of Mission Command, or "Auftragstaktik" as it was originally known as. Researching this Prussian/German concept reveals two compelling aspects. The first aspect answers why mission command was used by highlighting its strengths. It describes how the Germans thought they could mentally outmanoeuvre their adversaries by making better decisions and acting quicker than the enemy due to superior training and doctrine. The second aspect describes the types of officers and soldiers required for Mission Command to function well. This is the key to grasping what type of culture Mission Command truly is (Wandergriff & Webber, 2017, p. 52).

3.1.1 What is Mission Command?

The goal of Mission Command is to enable all personnel to be creative, so that should a situation differ from assumption, one can adapt. In times of emergency or war, when speed is critical, sending instructions up the chain of command in hopes that someone can decide is often too slow. Having a command structure in place that allows swift action is vital in such instances. Mission Command therefore involves issuing orders that allow subordinates to exercise their own judgment in achieving its objectives. It is based on the belief that once the commander's intent is made clear, initiative of action is better off firmly in the hands of the subordinate, who has a first-person perspective on the situation (Wandergriff & Webber, 2017, pp. 19, 50).

However, Mission Command has often been misunderstood purely as a way of formulating and carrying out orders or a form of military doctrine—in truth it's a form of personal mindset and organizational culture (Wandergriff & Webber, 2017, pp. 19, 50-51). This is emphasized in the Norwegian Chief of Defence's core beliefs on leadership as an example (Norwegian Armed Forces, 2020). Mission Command is a command culture that empowers the individual through trust and professional development; providing a personal will to strive for an optimal solution to a given problem. This kind of command culture cannot be comprehensively conveyed in an official instruction. Instead, Mission Command must be integrated into all education and training from the very beginning of basic training (Wandergriff & Webber, 2017, p. 50).

3.1.2 Why use Mission Command?

One of the clearest examples of Mission Command's strength is Nazi Germany's assault on France in 1940. There is a myriad of reasons for why the French defence fell so quickly, but Mission Command, or Auftragstaktik, played an important role in Germany's tactical success. The German victories are commonly attributed to superior technology, concentration of forces and physical speed against a foe tied mostly to the Maginot Line. In reality, they held few advantages in numbers nor weapon technology (Wandergriff & Webber, 2017, p. 53). They attained an advantage in speed; however, this was not, as commonly believed, due to German mechanization. In the interwar period the Germans were excellent at linking lessons learned from WWI, particularly the impact of mobility on the eastern front and the lack of mobility to exploit openings in the 1918 offensive on the western front, with their already established Auftragstaktik. By 1940, partly due to their experience against Poland the previous year, Germany managed to achieve superior mobility to the French – not due to their vehicles, but their command structure, allowing initiative and independence on the frontlines. Once they managed to get inside the French Army's decision-making cycle and disrupt their deployment plan, the French could no longer keep up; they were not adaptable enough (Wandergriff & Webber, 2017, pp. 70, 72, 75). In conclusion, according to Long, the Germans' success in the beginning of the Second World War was not caused by an intimidating new concept of warfare nor superior mechanisation or use of tanks and tactical bombers. It was due to their superior training methods which drew upon their doctrine, borne out of careful and accurate conclusions drawn from WWI. A doctrine closely tied to a proper understanding of Auftragstaktik, Mission Command (Wandergriff & Webber, 2017, p. 84).

3.1.3 Who is required for Mission Command?

The ideal Officer as with Mission Command, Auftragstaktik emphasizes the commander's intent. Intent forms a basis for subordinates to make decisions that align with the overarching goal. As such the core principle of German military education was cultivating independent thinking. This meant teaching officers to think for themselves. The German term for this is Selbständichkeit, meaning independence, and it was emphasized more than the concept of Auftragstaktik. Auftragstaktik was a culture that created the necessary preconditions for Selbständichkeit. The rationale behind Selbständichkeit being to make a good decision quickly, rather than waiting for the perfect answer and possibly missing a crucial chance. As a result, commanders were encouraged not to reprimanded subordinates for exhibiting initiative but chose the wrong action. In this culture the greatest sin was inaction. This philosophy spanned the ranks, from the commanders all the way down to the individual soldier. When communication with a commanding officer was lost, subordinates were trusted to take the appropriate action on their own, rather than stopping until communication could be re-established. This aggressive approach allowed units to capitalize on short-term opportunities (Wandergriff & Webber, 2017, pp. 54-55).
Thus, the key to the success of Auftragstaktik was the stringent selection and training of military leaders. In particular, three traits were highly valued in officers: knowledge, independence, and the joy of taking responsibility. Knowledge was essential for knowing what to do, as well as instil confidence in their subordinates. The ability to act independently was also critical as officers could often be tasked with making decisions on their own. Finally, emphasised as the most important virtue; joy in taking on the burden of responsibility. Independence prepares an officer to handle uncertainty and make good decisions in the absence of direction. However, in order for an officer to reach their full potential, to face the horrors of the battlefield, the German military believed something more was needed: a feeling of responsibility, that no one else can determine the outcome. The German military trained their officers to not only shoulder responsibility, but to shine when it fell onto them (Wandergriff & Webber, 2017, pp. 55-56).

3.2 Autonomy

Autonomous, as defined in the US DoD UAS 2018 Roadmap, is “the ability of an entity to independently develop and select among different courses of action to achieve goals based on the entity’s knowledge of the world, itself and the situation.” (US Department of Defence, 2018, s. 17). This definition requires an autonomous system to have a form of cognition, an ability to make decisions based on knowledge. Autonomy is intrinsically linked to AI. As touched on in the definition of artificial intelligence earlier in this thesis, artificial intelligence is a broad term referring to a collection of technologies. AI today mainly comprises pattern

recognition algorithms, which, when combined with powerful computing power, allow operators to make sense of large data sets. Neural networks enhance the algorithms' ability to identify patterns by training them to associate specific patterns with desired outcomes. The use of neural networks, a set of algorithms attempting to identify underlying relationships within a data set by mimicking the operation of the human brain, makes current machine learning approaches possible. However, AI encompasses more than just neural networks, such as language processing, knowledge representation, and inferential reasoning. These advancements in software, hardware, data collection, and storage enable AI to find valuable information within data sets, provided that the desired outcome is known (Davis, 2019, p. 116).
This chapter aims to provide a better understanding of the term autonomy before delving into artificial intelligence and LAWS themselves.

3.2.1 The Three Dimensions of Autonomy

Paul Scharre splits autonomy into three separate dimensions based on three concepts for autonomous systems: The task performed by the machine; the machine's relation to the human during the task; and the sophistication of the machines decision-making. He argues that increasing the amount of autonomy along any of these three spectrums would increase the machines overall autonomy (Scharre, 2018, pp. 27-28). For the purpose of this thesis these three dimensions have been named; "the machine's purpose"; "the human-machine relationship"; and "the machine's intelligence".

The machine's purpose:
The purpose of the machine, the complexity of the decisions it must make and the consequences of failure, is an important context to consider when regarding a machine's autonomy according to Scharre. While fully autonomous systems are somewhat uncommon today, especially when complexity and risk is involved, many systems divide tasks between autonomous functions and human control. Examples include modern cars which have a heap of autonomous functions both for safety and ease of use. Autopilots on commercial airliners can be switched on and off by the user. In some cases, the autonomous features might override human interface completely when necessary. Such as the ground collision avoidance system on modern U.S. fighter aircraft which takes control of the aircraft and pulls up last minute in case the pilot is disoriented and about to crash. Therefore, Scharre argues, it is meaningless to refer to a system as autonomous without referring to the specific task being automated. There is a difference between autonomous features that aid the user as opposed to the user simply directing the autonomous system. This relation is the topic of the next dimension of autonomy (Scharre, 2018, p. 28).

The human-machine relationship:
Scharre splits this dimension further into three distinct forms of relationship, from least to most autonomous. Firstly, semiautonomous operation or "human in the loop": The machine performs its given task, but at one or more points in its decision cycle is broken by a human and must seek approval before continuing. The system can make a recommended course of action but cannot execute the action without human approval (Scharre, 2018, p. 29).

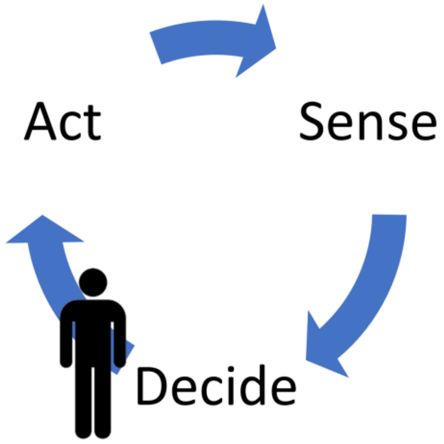

Figure 1: Semiautonomous operation (human in the loop)

Next is supervised autonomous operation or "human on the loop": Once the operation has begun the machine or system may complete the entire decision cycle on its own, without requiring human approval. A human user can observe the cycle however and may abort the operation if desired (Scharre, 2018, p. 29).

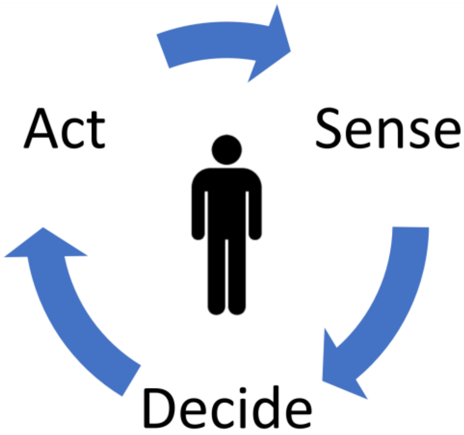

Figure 2: Supervised autonomous operation (human on the loop)

Lastly there is fully autonomous operation or "human out of the loop": The system can sense, decide and act all on its own without human intervention once activated. The system is not required to communicate back to the user while conducting its task (Scharre, 2018, p. 30).

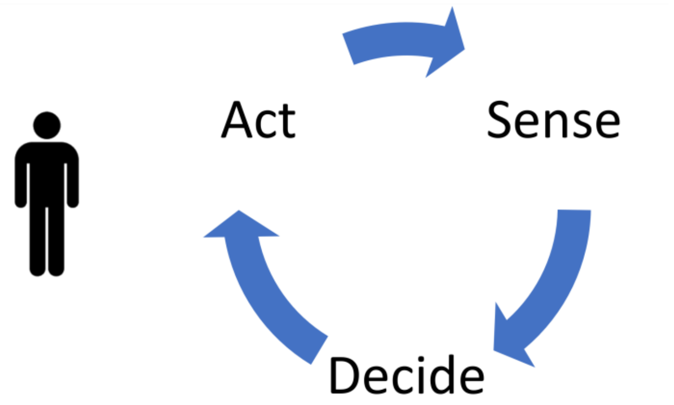

Figure 3: Fully autonomous operation (human out of the loop)

The machine's intelligence:

The last dimension of autonomy is intelligence. More advanced and intelligent machines can be used to accomplish more complex tasks in difficult environments with a plethora of variables to consider. Scharre points to how words like automatic, automated and autonomous are often used to describe increasing levels of intelligence in machines. Automatic machines are unable to make any decisions they simply act when they sense a specific trigger. Automated machines may consider a range of options and as such has some capacity of decision making. The phrase autonomous, according to Scharre, is often used to refer to machines so sophisticated that their cognitive process is not understood by the user. While the user might know the task and whether they succeeded, they often don't understand the process that led to the result. Autonomous systems are created to operate in environments with so many variables that the creators of the system cannot specify how the system should react to each of them. The system must therefore be able to "think" and learn for itself.

In practice however, the lines between these three levels of intelligence are quite blurry. Scharre points out that future advanced systems that are still conceptual or being built are often referred to as autonomous. However, once they are being used and the users begin to get an understanding of them the same system is instead described as automated. In practice the distinction between an automated and autonomous system then depends on the user's understanding of the system rather than the system itself. On the other hand, the ability to "think" in an entirely different method than humans are the autonomous system's greatest strength. Allowing the system to arrive at different conclusions than a human and often reaching a conclusion at a faster pace. This separate method of "thought" is however also the autonomous system's greatest weakness. When the user does not understand the systems decision-making process, the user might not be able to predict the outcome of a specific order. While the system might have followed out the order to the letter, it might have taken a different path than the user intended. The user was unable to foresee the possible course of action and therefore unable specify that it should not be used. (Scharre, 2018, pp. 31-32).

3.2.2 Narrow, General and Super Intelligent AI

Another way to classify todays AI technology is to divide it into narrow and general AI. Narrow AI refers to specific problem-solving tools used in order to solve a certain task. The AI will be able to learn and improve through trial and error as it continues to attempt the task at hand, but if you were to use the same AI for a separate task it would most likely be unable to use any experience it gained from the previous task. General AI on the other hand aims to replicate human brain functions, in its purest form there would be no difference between general AI cognition and human cognition. While narrow AI has made significant progress, it is still far from replicating human-like reasoning, which is the goal of general AI. Though breakthroughs have been made in replicating human-like reasoning, such as IBM's Watson and Google's DeepMind, they are still far from replicating the full performance of the human brain (Davis, 2019, p. 116). Historically the progress of AI technology has been unpredictable, but projections based on current trends indicate that artificial intelligence with human-level cognition soon (Mazzolin, 2020, p. 50). Once general AI has met the level of human intelligence the next obvious step is to surpass it, creating AI superintelligence: AI systems that are truly self-aware and conscious with the potential to surpass humans in almost all areas, including general knowledge, scientific innovation, and social abilities. This could lead to humans becoming obsolete (Kaplan, 2021, p. 26). Some experts believe that AI capabilities could greatly expand by 2045, providing significant benefits to human society by enhancing general efficiency. This technological advancement could have a significant impact in politics, economics, and military operations, making it a matter of global importance that requires the attention and resources of leading nations (Mazzolin, 2020, p. 50). In his article Davis mentions the concept of "The Singularity," which refers to a point in time when humans merging with the superintelligence, they themselves created will increase their effective intelligence a billionfold. This is perhaps an alternative if humans do not wish to be made redundant by their own creations. Despite the quest for superintelligence, recent advancements in brain enhancement are currently only able to replenish impaired functions and are a long way from providing superhuman abilities to citizens, soldiers, or robots. While general AI and AI superintelligence may inspire science fiction about future cyborgs and robot armies, narrow AI is already a reality today (Davis, 2019, p. 116).

3.2.3 Autonomous weapons

As discussed in the introduction, this thesis will be utilizing the following definition of LAWS from DoD 3000.09:

A weapon system that, once activated, can select and engage targets without further intervention by a human operator. This includes human-supervised autonomous weapon systems that are designed to allow human operators to override operation of the weapon system, but can select and engage targets without further human input after activation. (US Department of Defense, 2012, pp. 13-14)

DoD 3000.09 is the world's first formal policy on autonomy in weapon systems. Making the United States the first country to adopt such a policy. This first attempt managed to cover all the most important details. At this point the policy is a decade old, however. The policy is currently being reviewed and updated as this thesis is being written. (Allen, 2022).

LAWS could conjure the image of “smart weapons” also known as precision guided munitions (PGMs). Most PGMs are not truly autonomous weapons, however. PGM’s make little to no decisions themselves, they simply act out the task that has already been decided for them. The exception to the rule is homing munitions which are in some cases fired while blind, before it then activates its sensors and picks a target. Even this is a very limited form of autonomy, however(Scharre, 2018, pp. 40-42).
The autonomy of a weapon, more often than not, makes its impact before the weapon is fired and indeed in the decision of whether the weapon should fire at all. In fact, the US DoDs definition specifies “A weapon system…” as opposed to a weapon on its own.

A weapon system consists of a sensor to detect targets and a decision-making element to choose to engage the target in addition to the weapon itself that directly affects the target. An autonomous weapon system must be able to carry out all three of these phases on its own. In many modern weapons the sensory and weapon phases are already automated: A machine picks up a target and reports it to a human who then decides to engage, firing a PGM for instance which executes the weapon phase on its own. The decision phase can in some cases be automated, though a human usually maintains the right to accept or deny the machines decision. Meaning that most LAWS that exist today are semi-autonomous. However, there are cases of supervised autonomous and even fully autonomous weapon systems. Many of these systems come in the form of various defensive weapons in use by a wide range of nations today. In an age of PGMs the most effective way to neutralize an incoming attack is often to counter it with PGMs of your own. These weapon systems are designed to target incoming missiles, rockets or aircraft. Due to the speed of these threats, humans may not always be able to react fast enough or have the stamina to remain alert over extended periods of time. This is a scenario where a machine’s tirelessness and reaction time shine. Humans determine the parameters of the weapon; which threats to engage and which to ignore or special rules based on specific situations such as certain timings or angles of fire. After that, the system operates independently. The systems are usually supervised autonomous, allowing a user to abort the process. Though if nobody is monitoring the system it is effectively fully autonomous. These forms of systems are generally not very controversial, explaining their widespread use. They are used for immediate defence; they usually target objects such as missiles, rockets and aircraft rather than directly targeting humans; and they are supervised by humans who are also colocated with the system, able to shut it down if it malfunctions.
Fully autonomous weapon systems are not in wide use today. However, there are a few select examples such as certain forms of loitering munitions. Once activated, these weapons can circle over an area for an extended period, searching for a target which it can select and engage all on its own. Meaning that unlike PGMs, loitering munition is a weapon system all on its own. The user can launch the loitering munition into an area without knowledge of any specific targets beforehand. While a PGM requires a specific target, or simply engages the first target it spots for some homing munitions, a fully autonomous piece of loitering munition can choose an ideal target within the search area itself (Scharre, 2018, pp. 43-49).

4 Discussion

4.1 Autonomous weapon's role in Mission Command Leadership

In the context of mission command leadership, the role of LAWS is a complex and controversial issue. On one hand, LAWS could enhance the capabilities of military commanders by providing advanced autonomous weapons that can quickly and accurately engage targets in complex and rapidly changing battlefield environments. On the other hand, the use of LAWS could pose significant risks and ethical dilemmas for military leaders, who must balance these benefits of increased effectivity with the risks of unintended consequences.

As military leaders contend with the role of LAWS in mission command leadership, this technology is becoming a rapidly evolving area of concern that will require careful consideration and discussion to ensure its responsible use. To that end this chapter aims to discuss what role LAWS can and/or should fill within a mission command leadership hierarchy.

4.1.1 Humans leading AI

Decentralized execution based on the commander's intent is the foundation of mission command. In a similar manner Scharre recommends that humans should have a leading role in the human-machine relationship. Humans should provide goals and let the AI figure out how to reach those goals (Scharre, 2018, p. 17). This is the least controversial opinion on the utilisation of LAWS, one that has been implemented to a certain degree already. It is also a method that pairs very well with Narrow AI as it is designed to solve a specific problem and is more niche in application. Thus, it requires assistance from someone with a broader perspective in order to make full use of the results it produces. Humans could be a good fit for that role, however they might encounter some complications in doing so.

With narrow AI's poor ability to maintain a wider situational awareness, as detailed in chapter 3.2.2, this task may instead be given to human elements. However, with humans involved in the loop of every AI system, the loop would be getting crowded. On a battlefield ever increasing in scale and pace, maintaining situational awareness in order to know how and where to utilize LAWS may be a daunting task. Especially when there are many AI-driven battlefield systems operating quickly, each with its own chain of command. Coordinating the humans who are in the loop for fast-moving battlefield operations will be a challenge. Especially when it involves multiple adversaries, domains, agencies, clearance levels, contractors, allies, and organizational cultures. Particularly if the goal is to maintain a competitive edge through fast decision-making, which is one of the greatest strengths of LAWS. With many people in many roles who affect AI contribution to individual and shared goals, achieving strategic objectives will be a challenge. Extraordinary cooperation and communication between the human elements will be necessary (Davis, 2019). This requires sharing real-time information and intelligence across multiple agencies and organizations, as well as developing common operating procedures and protocols. With many AI-driven battlefield systems operating quickly and independently, it is important to establish clear lines of communication between human operators, so that they can coordinate their efforts and avoid unintended consequences. Above all else, effective cooperation and communication between human operators requires a high level of trust, both between each other, but also towards the LAWS.

Military leaders should consider how the operators will interact with the autonomous weapon system, particularly in terms of providing oversight and control. A key element that can affect the system's effectiveness and safety is human-machine interaction. The system's human operator should be able to comprehend the systems choices and take appropriate action if necessary. The System's manufacturer therefore needs to ensure the system can provide the operator with intelligible insight on its own decision-making process. A leader in charge of the system should use this information to evaluate whether the system's recommendations match the commander's goal as well as how the system might influence human decision-making. One possible consequence is that the use of LAWS could lead to a form of automation bias. This refers to the common human behaviour of disregarding or failing to actively search for information that contradicts the results produced by an automated system, based on the assumption that machines are more accurate and reliable than humans (Arkhipov-Goyal & Chavannes, 2019, s. 30). Another is that the employment of LAWS can give people the impression that they are removed or detached from the real conflict, which might enable them to make choices that they would not otherwise make if they were physically present on the battlefield. Their ethical judgment is compromised because of their desensitization to the ethical ramifications of their behaviour.

4.1.2 AI leading AI

In the previous sub-chapter, the demand for extraordinary cooperation and communication between human elements operating LAWS was addressed as a concern. A possible solution to this issue is to replace the human element attached to each single LAWS. Instead, LAWS are connected by AI that allow for faster and more efficient coordination and communication between systems. In essence, this approach would involve fewer, larger weapon systems comprised of multiple smaller weapon systems. While such a system might be unthinkable today, the future evolution of general AI might make it more feasible.

The goal of General AI is to mimic the human thought process. Therefore, it should be able to provide the wider situational awareness necessary to effectively provide goals and utilize results of narrow AI in accordance with an overall objective. This general AI would then receive its directives from a human commander. General AI as a system made up of multiple LAWS could perhaps do better coordination than LAWS working separately with humans having to coordinate between them. This could allow LAWS to better utilize their advantage of faster decision making as they are free to interact with each other, creating a more comprehensive dataset to base their decisions on.

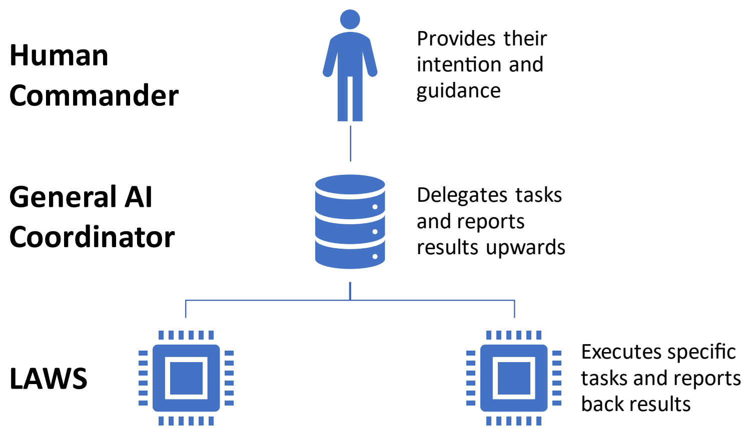

Figure 4: An example of hierarchy where LAWS are directly led by AI and indirectly led by humans

However, this approach also puts a large responsibility on the human element overseeing the general AI. They would need to keep track of decisions made by multiple LAWS, but this should be made easier as the general AI compiles and presents the necessary data. Nevertheless, the integration of multiple AI systems on the battlefield can lead to unexpected outcomes, with some even having strategic consequences. The issue of unexpected convergences arises as separate AI-infused platforms operate in a shared battlespace. It is unclear how these systems will interact with each other, and the outcomes resulting from friendly interactions could be compounded by interactions with foreign AI systems. In chapter 3.2, the difficulty of predicting the actions of even a single advanced autonomous system was addressed. Interactions between multiple such systems would be even more prone to producing unanticipated and unexplainable results due to the uncertainty surrounding the mechanisms that produce the outcomes.

Furthermore, an unintended result in the form of a malfunction in the general AI coordinating the separate LAWS could have large repercussions. Any malfunction whether at accidental or targeted would affect multiple LAWS thereby increasing the impact. This is in and of itself not different from targeting any human command and control element. However, the increased risk of the system executing an unsanctioned action due to a malfunction is still present. This is a risk that is removed by utilizing semiautonomous and mitigated by supervised autonomous LAWS. A semiautonomous system might hamper much of the speed the AI contributes with, however.

Lastly, the notion that LAWS may create distance and detachment from conflict, leading to compromised ethical judgment from the human elements, as discussed in the last chapter, would be perpetuated and perhaps worsened.

4.1.3 AI Leading Humans

The prospect of artificial intelligence (AI) commanding humans is becoming increasingly plausible as AI technology develops. In this chapter we will examine the potential advantages and disadvantages of AI commanding military operations, as well as the ethical issues raised by such a change in power dynamics.

One potential outcome is that AI could create orders for humans to execute or assist them by providing possible solutions. This would be a controversial solution today, but once again the rapid advance in AI technology could soon make the statement seem less preposterous. In some cases, a general AI system that mimics human functions may be a more effective commander than a human. The notion of an AI Superintelligence as addressed in subchapter 3.2.2 may only be about 20 years ahead of us and could introduce an AI commander superior to any human. It is unlikely to think that any nation would allow an AI system such as this to go fully unsupervised. Therefore, it would still be subordinate to a human commander, however, certain human elements might answer directly to this artificial superintelligence responsible for coordinating assets in a similar manner to the one seen in the previous sub chapter.

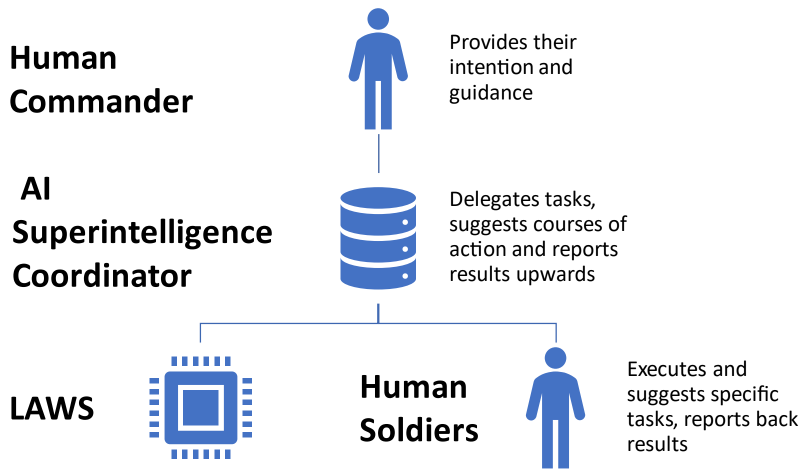

Figure 5: An example of a hierarchy where LAWS and human soldiersboth recieve orders directly from AI superintelligence and indirectly from a human commander

Utilizing the strength of AI, processing speed, one could create a common situational awareness between the multitudes of assets operating in the battlespace. This would be powered by AI superintelligence compiling and distributing necessary information to assets in real time. At the same time humans remain directly attached to the conflict and are thus able to share their perspectives, including ethical judgement AI may lack, into the situational picture.

However, the risks of malfunction as explored in the previous subchapter cannot be ignored. In the case of the AI superintelligence coordinator the consequences of malfunction might be even greater than with the general AI depending on the level of responsibility and dependence pinned on the AI superintelligence. Though, in the case of a malfunction an added advantage with having human subordinates to the AI occurs. If the soldiers are educated in accordance with mission command philosophy, they should have a good basis for independent thought. This ability combined with a sound ethical judgement and a good grasp of the overall mission objective, should in many cases allow the soldier to notice a fault in their orders. This is especially true if they also educated in the AI’s greater weakness to malfunction. That is unless they fall into the trap of automation bias, addressed in subchapter 4.1.1, and become overly reliant on the system's recommendations.

Another perspective regarding the concept of Mission Command, which emphasizes the joy of accepting responsibility, raises questions about whether LAWS can embody this same spirit. The joy of responsibility was exemplified in subchapter 3.1.3 as the foremost qualification for a leader in mission command's origin, Auftragstaktik. AI are after all famously incapable of inhabiting feelings, joy included. On the other hand, joy of responsibility is exemplified as a virtue of mission command as it allows a commander to keep their morale despite facing the horrors of the battlefield. In the case of an AI, they do not, for better or for worse, conjure any feelings when faced with the reality of war and their morale therefore does not falter. One could therefore argue that the necessity for this joy of responsibility is lessened if not entirely removed.

However, a mission command requirement whose necessity is not lessened, but perhaps increased is trust. While mission command in a traditional sense requires commanders to trust soldiers' ability to make the right call based on their perspective, the integration of AI requires them to trust machines to do the same. In the scenario suggested in this subchapter soldiers may also need to trust an AI to give them orders and instructions, which could, if wrong, cost the soldiers their life. On the other hand, soldiers need to be careful not to trust the AI blindly to the point where they do not notice a faulty order and must be able to act independently. Trusting AI is a balancing act that deserves a more detailed discussion.

4.2 Trusting the Machine

Trust is a pivotal requirement in the execution of mission command leadership. If AI systems are to be implemented into a mission command structure, being able to put one's trust in them is essential. Trust was brought forward as an important prerequisite in two separate occasions in the last chapter. Overall, this chapter aims to provide a comprehensive overview of the concept of trust when it comes to AI systems, including its importance, the factors that contribute to building trust, the challenges that arise, and the methods that can be used to address these challenges.

4.2.1 AI Intent

Intent is a central component to trust. In order to trust someone, or in this case something, you must be able to believe that their intent is to support, or at least not negate the goal you wish to accomplish. One question that arises is whether AI have an intent at all or if you just don't fully understand it. Furthermore, if the AI has no intent or we don't understand their intent, will a military leader still be able to impart their intent behind a goal to the AI. After all the very definition of mission command underlines the importance of communicating intent.

When considering the deployment of LAWS in military operations, it's crucial to assess the nature of AI Intent. In general, AI lacks intent in the sense that it is not aware of the results of its actions or their ramifications. The rules or algorithms that AI systems are trained to follow only allow them to make proper judgments based on data and patterns they have learned. However, this absence of intent in a conventional sense does not by itself imply that the behaviour of AI cannot be trusted. Trust can instead be established by having a clear understanding of the intention behind the AI programming and how it will behave in different situations. In a sense you would be trusting the LAWS' creator rather than the intelligence within the system itself, approaching it in a more similar manner to any other traditional weapon. Naturally, like with other weapons, faults may still occur.

As discussed earlier in this thesis the nature of complex AI and machine learning means that AI systems are not always entirely predictable, even for the manufacturer. Since imparting intent is a key requirement in order to achieve the benefits of mission command. The commander must non the less be able to communicate the mission objectives and intent to the AI system in a way that the system can understand and use to make decisions. This necessitates an in-depth understanding of how the AI system functions and how decisions are programmed into it. The system must also be built to give the human operator feedback so they may verify the system's judgments and intervene if required.

In conclusion, developing trust in LAWS necessitates a clear understanding of the intent behind the AI programming and how it will act in various scenarios. While AI systems lack intent in the traditional sense, their decision-making processes can still be complex and difficult to understand. Thus, imparting human intent on AI systems is a critical requirement for the effective use of autonomous weapon systems in military operations. This requirement highlights the importance of communication and the need for the commander to convey the mission objectives and intent to the AI system. This communication aspect becomes even more critical when considering the level of autonomy and freedom allowed for the autonomous weapons.

4.2.2 Autonomy and Freedom

The complex nature of the tasks LAWS is involved in, and the risks associated with failure often require a high level of autonomy for the weapons to function. However, this also raises concern over how much freedom humans should allow autonomous weapons to have. In order to establish trust in these weapons, while still utilizing them effectively, it is necessary to strike a balance between the autonomy and freedom given to LAWS. This requires a clear understanding of the different dimensions of autonomy and how autonomy differs from freedom.

Scharre argues there is a common misconception that autonomy is directly linked to intelligence (Scharre, 2018, p. 50). While autonomy is commonly defined as the ability to make decisions, intelligence is not the only factor impacting AI's decision-making. As discussed in subchapter 3.2.1 autonomy is split into three separate dimensions all of which impact a machines ability to make decisions. However, as weapons used in life-or-death situations autonomous weapons must inherently be able to solve complex problems where the consequences of failure may be severe. As described in subchapter 3.2.1 this already places LAWS near the top of one dimension, the machine's purpose. Due to the complexity of the tasks involved and risks associated with failure it would also be a requirement for most LAWS to place in the upper scale of another dimension of autonomy, the machine's intelligence. In other words, two of the three dimensions of autonomy described in subchapter 3.2 generally need to be scaled to the maximum for LAWS to even meet acceptable standards.

This puts many LAWS in a situation where the only feasible option for altering their autonomy is the final dimension, the human-machine relationship. This is often the essence of the debate regarding autonomous weapons. This is the dimension that physically restricts a machine. The debate surrounding it therefore boils down to a relatively simple question, with a difficult answer. How much freedom should humans allow LAWS to have? Should they be semiautonomous, supervised autonomous or fully autonomous? With too little freedom one risks limiting LAWS to the point of infectivity. However, if one trusts LAWS too much, there is a risk of them running amok.

4.2.3 Weapons on the Loose

According to Scharre, humans tend to have two intuitions about autonomous systems. The first is that they are reliable and precise, while the second is the fear that they might malfunction and cause harm. Today's autonomous systems can perform better than humans in many situations. However, they lack common sense to adapt once they encounter an issue beyond their programming. This makes them both reliable and untrustworthy. One solution to this problem is to maintain human control over the system so that if it fails, humans can take over quickly. However, as discussed in the last chapter humans may still be prone to automation bias and might not intervene even when the system is wrong. The alternative, however, are fully autonomous systems where humans have no chance to intervene immediately (Scharre, 2018, pp. 145-146).

When activating an autonomous system, there is a level of trust involved that the system will function as expected. However, this trust should not become blind faith. To maintain a balanced level of trust, users must understand the system's capabilities and limitations, which requires rigorous testing and evaluation. But with increasing intelligence, and thus complexity, in autonomous systems, it becomes increasingly impossible to test every possible scenario. The largest disadvantage of LAWS as discussed in the previous chapter was the risk of malfunction in the form of misunderstanding intent. To reduce this risk and teach AI intent, AI needs to test and train with humans. If AI is ever going to learn how to coordinate with human soldiers, AI might have to be put in charge of human soldiers. Testing, evaluation, and experience is vital for building confidence. However, it cannot guarantee the absence of failure, there is always a high risk of failure when an untested scenario occurs (Scharre, 2018, p. 149). While testing cannot make humans immune to failure either, trust is made easier if the participants are familiar with each other and believe their intentions align. The difficulty behind establishing the right amount of trust therefore loops back to the difficulty of understanding an AI's intentions.

Over the years, humans have developed extensive experience in designing, testing, and operating complex, high-risk systems such as nuclear plants, aircraft, and spacecraft. This experience has formed a strong research field focused on enhancing the safety and durability of complex systems. However, this experience also suggests that avoiding failure completely is impossible. Major incidents have still occurred in all these systems despite their robust design and testing. The complexity simply makes anticipating all possible outcomes impossible and an unanticipated incident is bound to happen as a result. Engineers refer to these incidents as "normal incidents" as they are exactly that, normal (Scharre, 2018, p. 151). As a result, despite their failures, the example systems are all still in use today due to their benefits outweighing the risks. The crucial question regarding LAWS is therefore whether their strengths outweigh their inevitable failures. To put it bluntly, will LAWS be a better way to kill?

4.3 Making AI Kill

It is heavily disputed whether it is morally acceptable for humans to allow AI to kill. Some argue that doing so raises severe ethical concerns, while others suggest that they have the potential to reduce harm to both soldiers and non-combatants. This chapter will delve into these perspectives and examine the difficult problems associated with the usage of AI as lethal weapons.

4.3.1 Who is Accountable for Killing?

The loss of human control over the use of force is one ethical concern with the deployment of LAWS. The concern is targeted at, the machine, not a human operator, deciding when and how to use lethal force. This raises questions regarding accountability and responsibility for the results of these decisions. To ensure the moral application of LAWS, it may consequently be necessary to establish a clear framework for accountability. The framework needs to have processes for looking into incidents and holding individuals accountable for any legal or ethical transgressions.

However, such a framework may be difficult to implement in practice. The use of fully autonomous LAWS in warfare could create what Major Zhiyuan calls an "accountability gap". A situation where neither the programmers, commanders of the LAWS, nor the LAWS themselves can be held responsible for any unforeseen atrocities committed by the LAWS. Defence contractors are typically shielded from civil liability, and it would be difficult to prove any malicious intent as the LAWS did not act as intended. In the case of an unforeseen action the commanders could not have reasonably predicted it. Thus, commanders cannot be held responsible for the LAWS actions. For the LAWS itself, as discussed in the previous chapter, it is difficult to attach intent to AI itself, holding it responsible for its own actions would therefore be futile (MAJ Zhiyuan, 2021, p. 2).

The accountability gap highlights a significant distinction between fully autonomous LAWS and human soldiers. According to Uwe Steinhoff, there is no moral difference between the conduct of human soldiers and those of a fully automated LAWS, and commanders cannot be held accountable for unforeseeable actions committed by either one (Steinhoff, 2013, p. 185). A soldier is, however, not subject to the same accountability gap because they may be held responsible for their actions. Thus, an issue arises when replacing human soldiers with fully autonomous LAWS. However, as Major Zhiyuan himself points out, the accountability gap does not exist in the case of a semi or even supervised autonomous weapons where a commander is expected to halt an unintended action before it is executed. Thus, the commander may be held responsible should they fail to do so (MAJ Zhiyuan, 2021, pp. 2-3).

4.3.2 A Moral Obligation to Deploy Autonomous Weapons?

In accordance with international humanitarian law, the military is required to defend non-combatants from injury, but they also have a moral commitment to defend their own troops. LAWS could help reduce fratricide and collateral damage by providing faster and more accurate targeting. They could also access dangerous areas without setting human life at risk. Therefore, if LAWS can perform the mission as effectively and safer than using human soldiers, it could be argued that the military has a moral obligation to use LAWS instead of soldiers.

The effectiveness of LAWS therefore becomes an important topic regarding their use as lethal weapons. Throughout the thesis LAWS' susceptibility to unintentional acts has been well established and is perhaps the largest concern towards their use. This was further complicated in the previous subchapter where it was acknowledged that accountability gap could occur regarding such actions. In both cases the risk is mitigated by utilizing semi or supervised autonomous LAWS, allowing humans to intervene as opposed to fully autonomous LAWS. However, as we established in subchapter 4.2.3 the risk can never be removed entirely. Still, this highlights the importance of testing and evaluating the system's capabilities and limitations, no different from any other complex invention, in order to achieve an effective balance between AI and Humans elements of the LAWS.

While the current state of AI technology may not be adequate, LAWS will keep improving with robust design and testing. Militaries will likely start deploying semi, and in some cases supervised, autonomous LAWS on a larger scale as they could replace soldiers with an acceptable amount of risk. With a longer perspective, the rise of general and even super intelligent AI technology could eventually justify the use of fully autonomous LAWS, though this would require more study and research and is difficult to assess today.

5 Conclusion

Based on the discussion throughout this thesis it can be argued that LAWS are compatible with mission command philosophy, to an extent. The discussion of this thesis has analysed various components of the research question. Thus, it has reached the following conclusions on the most important ethical and practical considerations when implementing LAWS:

The first chapter of the discussion examines the possible role of LAWS in mission command leadership. Focusing on the practical implementation and possible challenges that arise, acknowledging that it is a complex and controversial issue requiring careful consideration. Three main approaches are discussed: humans leading AI and AI leading AI/humans. While the former has already been implemented to some degree, the latter two approaches are still largely theoretical, but are likely becoming more feasible as AI technology progresses. It is emphasized that the human element must remain present in all cases to provide goals and maintain ethical judgment. However, AI bring quite a few advantages especially when speed and accuracy is involved. Still, the integration of LAWS brings with it a new set of risks and ethical dilemmas, such as the potential for automation bias and unexpected outcomes, that must be addressed.

The second chapter discusses how the use of LAWS in military operations requires the development of trust in the systems by understanding their intent and the extent of their autonomy and freedom. This chapter highlights the practical challenges of establishing trust in automated systems, a key component of mission command philosophy. It is crucial to strike a balance between the autonomy and freedom given to LAWS to establish trust and utilize them effectively. While trust can be established by understanding the AI programming, it is important to consider that AI systems are not entirely predictable, and faults may still occur. Testing and evaluation of LAWS in coordination with humans will be vital for building confidence in the systems, but it cannot guarantee the absence of failure. Ultimately, the success of LAWS will depend on whether their strengths outweigh their inevitable failures.

In the third chapter the discussion examines the use of LAWS from an ethical and legal perspective. It highlights the serious ethical and legal questions raised by the deployment of LAWS, including the accountability of the use of force and the moral obligation to deploy them. We acknowledge that while the current state of AI technology may not be adequate for the use of fully autonomous LAWS, semi and supervised autonomous LAWS could soon be deployed on a larger scale with an acceptable amount of risk. However, robust design and testing is required to ensure their effectiveness and safety. The decision to use LAWS must consider the potential risks and benefits, and any deployment must be accompanied by a clear system of accountability and responsibility.

Ultimately, regarding the extent of LAWS' compatibility, the very freedom of action central to mission command is the biggest challenge to implementing todays AI. As todays AI is based on rigid structure and specific commands. However, as detailed many times through this thesis AI is an ever-evolving field of research. In theory LAWS have the potential to vastly increase the efficiency of mission command. Especially when exploring the possibility of LAWS as an integrated system of multiple weapon systems and/or humans. The entire scale of benefits LAWS may provide is difficult to comprehend. However, so is the scale of potential unintended consequences. This thesis has focused on how LAWS should be implemented rather than if. However, the precarious nature of LAWS and their potential for failure has been a central topic for discussion. In conclusion, implementing LAWS into a mission command hierarchy could allow them the freedom to truly unleash their potential. However, regardless of the amount of testing and experience, there is always a risk. As dictated by normal incidents, the question is not if, but when and how bad. Even then LAWS is not a field of research one can afford to overlook. The risk of ignoring LAWS may be even greater than implementing them. As AI technology advances its implications for the battlefield and thus strategic stability itself will increase. Even those who don't use it will need to know how to counter it. Because the battlefield rarely grants second chances.

Referanser

Allen, G. C. (2022, June 6). *DOD Is Updating Its Decade-Old Autonomous Weapons Policy, but Confusion Remains Widespread*. Retrieved from Center For Strategic and International Studies: https://www.csis.org/analysis/dod-updating-its-decade-old-autonomous-weapons-policy-confusion-remains-widespread

Arkhipov-Goyal, A., & Chavannes, E. (2019). Human Agency. I A. Arkhipov-Goyal, & E. Chavannes, *Towards Responsible Autonomy* (ss. 12-35). Hague Centre for Strategic Studies.

Congressional Research Service. (2022). *Defense Primer: U.S. Policy on Lethal Autonomous .*

Davis, Z. (2019). Artificial Intelligence on the Battlefield: Implications for Deterrence and Surprise. *PRISM*, pp. 114-131.

Horowitz, M. C., Scharre, P., & Work, R. O. (2018). *ARTIFICIAL INTELLIGENCE: What Every Policymaker Needs to Know.* Center for a New American Security.

Jacobsen, D. I. (2015). *Hvordan Gjennomføre Undersøkelser? [How to Conduct Research?]* (3 ed.). Oslo: Cappelen Damm.

Kaplan, A. (2021). Artificial Intelligence (AI): When Humans and Machines Might Have to Coexist. In P. Verdegem, *AI for everyone?* (pp. 21-32). University of Westminster Press.

Kelly, K. (2014, October 27). *The Three Breakthroughs That Have Finally Unleashed AI on the World*. Retrieved from WIRED: https://www.wired.com/2014/10/future-of-artificial-intelligence/

MAJ Zhiyuan, B. O. (2021). Ethical Problems Posed by the Use of. *Pointer, Journal of the SAF*.

Mazzolin, R. (2020). *Artificial Intelligence and keeping humans "in the loop".* Centre for International Governance Innovation.

Norwegian Armed Forces. (2020). *Forsvarets grunnsyn på ledelse [The Norwegian Armed Forces' View on Leadership].*

Scharre, P. (2018). *Army of None.* New York: W.W. Norton.

Steinhoff, U. (2013). *Killing Them Safely: Extreme Asymmetry and Its Discontents, In Bradley Jay, Strawser, Killing by remote control: the ethics of an unmanned military.* Oxford University Press.

US Department of Defence. (2018). *Unmanned Systems Integrated Roadmap, 2017-2042.*

US Department of Defense. (2012). *Department of Defense Directive: Autonomy in Weapon Systems (DoDD 3000.09).*

Wandergriff, D., & Webber, S. (2017). *Mission Command: The who, what, where when and why.* Amazon.

Wang, P. (2019, March 29). On Defining Artificial Intelligence. *Journal of Artificial General Intelligence*, ss. 1-37.