The Influence of AI Systems on Military Decision-Making

Sami Qalyoen (s3709868)

Supervisors:
Bas Haring & Zane Kripe

Bachelor Thesis– Informatica

Leiden Institute of Advanced Computer Science (LIACS)
liacs.leidenuniv.nl

July 1, 2026

Abstract

This thesis looks at how artificial intelligence decision support systems (AI-DSS) influence military decision-making and what role bias plays in this process. The main case study is Lavender, an AI system reportedly used by the Israeli military to identify individuals as possible targets in Gaza, together with related systems such as Gospel and Where's Daddy?. Because the internal design and training data of these systems are not public, this thesis does not attempt to determine exactly how they function internally. Instead, it analyses their reported use, based on investigative journalism, academic work and official statements. The analysis follows the decision-making chain from data input, to AI processing and recommendation, to human decision and military action. At each stage it examines where bias may enter or where the conditions for bias are created. It distinguishes between technical biases, such as data and algorithmic bias, and cognitive biases, such as automation and action bias, and shows how these may interact. A major finding of this thesis is that AI-DSS do not need to make the final decision themselves in order to influence military decision-making. They can already shape the process by influencing what information is presented and how it is presented.

Contents

1 Introduction

On 28 February 2026, during the opening phase of Operation Epic Fury, a girls' elementary school in Minab, Iran was reportedly struck by United States forces. The attack on the Shajareh Tayyebeh school killed more than 150 people, most of them young girls [1]. According to reports, artificial intelligence may have played a role in the targeting process. The school was allegedly misidentified as a military base because the system may have relied on outdated satellite imagery from a time when a military facility was located near the same site. Combined with limited human review, this reportedly contributed to the target being approved and struck [2]. The US Department of Defense told CNN [3] that the incident is under investigation and that it would be inappropriate to comment on it yet.

The strike took place during an operation in which thousands of targets were reportedly attacked within the first 24 hours[4]. This extremely fast pace of targeting was made possible in part by AI decision support systems, or AI-DSS: systems that process large amounts of military data and generate outputs such as target lists or recommendations.

In this case such a system, Project Maven, was used by the U.S. military, but these systems are being used by militaries of different countries as well, such as Israel's "Lavender" and "The Gospel" in the Gaza conflict. It looks like the use and relevance of these systems is only going to increase.

The use of AI systems like these raise questions about human oversight, bias and accountability in military targeting. The central research question of this thesis is therefore: To what extent do AI systems such as Lavender and The Gospel influence military decision-making, and what role does bias play in this process? This thesis aims to examine how AI-assisted targeting systems shape human decision-making in warfare and analyse forms of bias that may emerge in the reported use of these systems.

2 Theoretical Framework

This chapter provides the theoretical framework for analysing the role of AI decision support systems in military decision-making and the ways in which bias may influence this process. Before examining Lavender, Gospel and other similar systems as case studies, it might be useful to understand what AI decision support systems are, how they function in military contexts and how the potential biases work.

The chapter first explains the concept of (AI) decision support systems and their relevance in warfare. It then discusses different forms of bias, including technical biases such as algorithmic and data bias, as well as cognitive biases such as automation bias and action bias. This framework will later be used to analyse how bias may emerge in the reported use of Lavender and Gospel.

2.1 What is an AI-Decision Support System?

AI decision support systems, or AI-DSS for short, are systems that use artificial intelligence to support humans making decisions. They are not necessarily designed to make decisions completely independently, but instead to assist a human user by processing information and presenting outputs that may be useful for making a decision. These outputs can have different forms such as alerts, classifications, predictions, risk scores, rankings or recommendations.

The concept of a decision support system existed before modern artificial intelligence [5]. Decision support systems (DSS) are computerized tools designed to provide information to assist decision-making [6]. Blanchard and Bruun describe DSS as tools that can indicate relevant information, predict possible outcomes and recommend possible actions.

AI-DSS are a specific type of decision support system because they use AI techniques, such as machine learning, deep learning or large language models to process and analyse data [7].

The biggest purpose of an AI-DSS is to transform large or complex amounts of data into a form that is easier for a human operator to use and make decisions with. Instead of requiring a person to manually examine all available information, the system can filter, structure and prioritise that information. For example an AI-DSS may detect a pattern in a dataset, flag an unusual case or assign a probability score to a certain outcome.

AI-DSS are used in many different fields such as healthcare, finance, transport, administration and security. In healthcare for example, AI-DSS can assist doctors or radiologists by identifying possible abnormalities, prioritising urgent cases or acting as a second reader. A study on AI-DSS in radiology describes how such systems were integrated into clinical workflow to support radiologists, for example by alerting them to potentially positive findings and prioritising imaging studies for review [8]. This example shows that AI-DSS do not necessarily have to replace the judgement of professional individuals. Instead, they can support experts by helping them manage complex workflows and large amounts of information.

This distinction is important. An AI-DSS is not the same as a fully autonomous decision maker. For AI-DSS, in most cases the human user remains formally responsible for the final decision. The system provides support, but the human decides how to interpret and use that support. However, this does not mean that the system is neutral or passive. Even when an AI-DSS does not make the final decision, it can still influence the decision-making process by shaping which information is presented, how it is ordered and what appears most relevant.

A useful way to understand AI-DSS is as a form of machine generated advice. Van Rossum describes advice as a recommendation that favours a specific option and explains that advice can simplify decision making by reducing cognitive load [9].

AI-DSS can perform different functions depending on their design and purpose. Some systems are mainly descriptive: they identify and present relevant information. Others are predictive: they predict what might happen in the future based on patterns in data. Some systems are prescriptive: they recommend a possible course of action. These functions can also overlap as they are not exclusive[5]. A single AI-DSS may be able to classify data, predict risk and recommend a certain action at the same time.

Traditional DSS are based on deterministic computer models, meaning that they will always produce the same output, be predictable and do not involve randomness [10]. This keeps the output simple, explainable and understandable, but it often limits their potential to be used for complex tasks [11]. Now AI-DSS are often based on non deterministic (stochastic) models. These models are trained using datasets that contain examples of inputs and the desired outputs. So unlike deterministic models, these stochastic models are not coded by humans, but essentially code themselves[11]: they learn patterns from the training data and adjust their internal parameters based on those patterns. This allows AI-DSS to handle more complex real world situations. This does not come without a catch, the outputs of these systems will become less straightforward and explainable.

The output of an AI-DSS is often just a probability rather than certainty. This means that the system may not give a simple "yes" or "no" answer, but instead produces a likelihood, score or confidence level[10]. For example, a system may point out that a medical scan is likely to show a certain condition, or that a financial transaction has a high probability of being fraudulent. Such outputs can be useful, but they should not automatically be treated like they are the truth, as they are only probabilities.

Another important characteristic of AI-DSS is that they depend heavily on data quality. If the data used by the system is incomplete, outdated, unrepresentative or incorrectly labelled, the system's output may also be unreliable. AI-DSS may appear objective because they are technical systems, but their outputs are still shaped by human choices. These choices include what data is collected, how categories are defined, what the system is trained to recognise and what kind of outcome the system is designed to prioritise.

For this reason, AI-DSS should be understood as tools that support decision making, not as systems that automatically produce truth. They can improve efficiency by processing information faster than humans and by identifying patterns that may otherwise be missed.

2.2 AI-DSS in military

In military contexts, AI decision support systems are used to support humans in processing and interpreting large amounts of information. Modern warfare produces huge amounts of data, for example through satellite images, drone footage, intercepted communications, sensors, intelligence reports and open source information. This makes military decision-making increasingly complex, because human operators and commanders may not be able to examine all available information manually. Andersin explains that modern militaries have to deal with an overload of information, complexity and time pressure, which has increased interest in AI tools that can process large quantities of data more quickly than ever before[10].

This is not only a matter of future interest. AI-DSS are already being used in different conflicts around the world. For example in the Russia-Ukraine war, both sides on the battlefield have deployed AI[12]. In this war, AI has been used to process large amounts of battlefield data, including satellite imagery, drone footage, geolocation and faster intelligence analysis in order to support target and object recognition. Reuters reports that AI in Ukraine is used as a tool for analysing volumes of information that are too large for humans to process quickly. According to Oleksander Dmitriev Ukraine has collected more than two million hours, that is 228 years, of drone footage since 2022 to train AI models for tasks such as target identification as mentioned before but also for battlefield assessment and strategic decision making. For Russia there is less public evidence about its actual use of AI in the war, although some sources do provide insight into Russian views on military AI. Similar to Ukraine, the Russian Ministry of Defense aims to use AI to support soldiers by improving data analysis and decision making on the battlefield. This approach keeps a human involved in the process(human in the loop), helping military operators understand the situation more quickly and make faster decisions [13]. Some military experts in Russia even believe that in the future decision making in warfare will be handled completely by AI and even the role of a human operator would be removed[14]. Also back in 2018, adviser to the Department for Nonproliferation and Arms Control (DNAC) of the Russian Foreign Ministry, Andrey Malov, was against the ban of lethal autonomous weapon systems (AWS) when such an initiative was discussed in the Group of Governmental Experts of the UN Convention on Inhumane Weapons [15].

Also more recently in the war on Iran have we seen AI-DSS being used. In Operation Epic Fury executed by the United States and Israel. In the first 24 hours of the attack the United States struck over a 1000 targets[2]. America's AI-DSS Maven Smart System developed by Palantir Technologies Inc. played a crucial part in this. It collected information on enemy positions from a range of different sources such as drone and satellite images, radar signals, intelligence data, electric communications and more and combined it into "a common operating picture" of the warzone, says Cameron Stanley, the Pentagon's chief digital and AI officer for the Department of Defense. "Instead of having eight or nine systems" that military commanders use when making decisions in warfare, Maven fuses everything "into a single visualization tool" Stanley elaborated [16]. Myriam Dunn Cavelty, Senior Scientist and Deputy Director Think Tank at the Center for Security Studies (CSS) at ETH Zurich says "AI-enabled decision support systems are now deeply embedded in military infrastructures, particularly on the side of Israel and the United States. Their primary contribution is the capacity to process data, identify patterns, and generate recommendations at speeds and scales beyond human capacity."[17]

Similar developments can be seen in Israel's war in Gaza, where systems such as Lavender, Gospel and more have been reported to support target generation by identifying suspected individuals and buildings or infrastructure as potential targets. According to multiple sources Lavender (which we will dive deeper into later on) was reportedly used as an AI-DSS to identify individuals in targeting situations, while Gospel was reportedly used to mark buildings and structures connected to alleged militant activity [18].

The role of AI-DSS in the military is, like most AI-DSS in other fields, not usually described as replacing the human experts in making decisions entirely. Instead, these systems are meant to support military personnel by organising information, identifying patterns, predicting developments or generating recommendations. Andersin describes AI-DSS as tools that assist militaries in collecting data, predicting patterns and making recommendations based on large amounts of information. [10] In this sense, military AI-DSS work in a similar way to AI-DSS in other fields: they turn complex data into outputs that humans can use more easily.

A useful way to understand AI-DSS in warfare is to look at the decision-making process as a chain. In the context of military targeting, the process does not begin with the final strike decision. Before military action takes place, information must first be collected, analysed, structured and presented to the relevant human decision makers. This is where AI-DSS become important. They can operate in the space between raw intelligence data and human military judgement. Instead of replacing the commander, the system can shape the information environment in which the commander makes a decision.

This can be seen in figure 1, which shows a simplified AI decision support flowchart. The first stage is data input. In military contexts, this data comes from all kinds of sources. For example from surveillance data, satellite imagery, drone footage, intercepted communications, intelligence reports, sensors or open-source information. The purpose of this stage is to collect information about the environment that is operated in. However, the amount of information available in modern warfare is often too large for humans to analyse manually. Militaries are increasingly interested in machine learning based decision support systems because they can help analyse large amounts of information and shorten the time between identifying a target and taking action against it. [11]

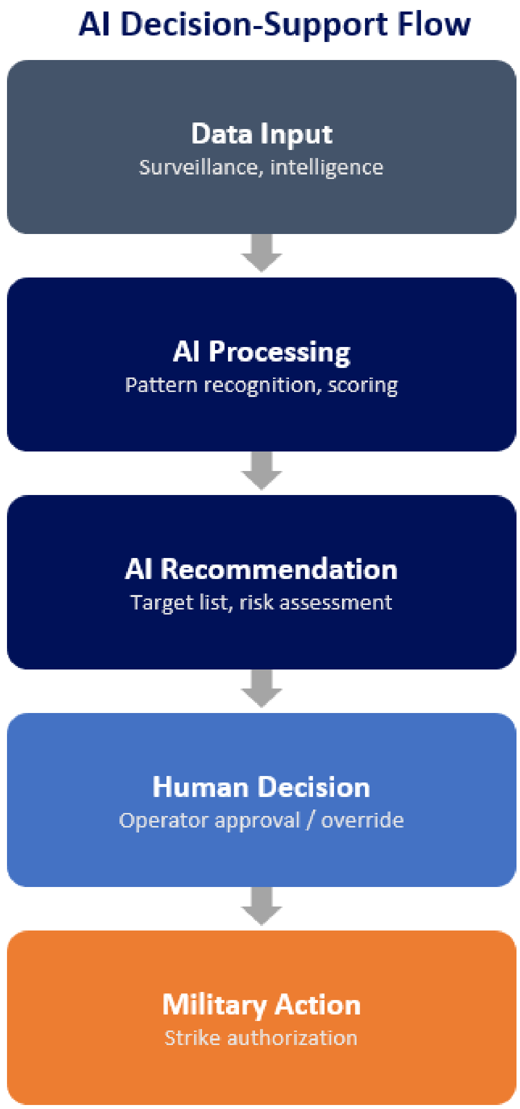

Figure 1: AI Decision-Support flow.

The second stage is AI processing. At this stage, the system analyses the data and looks for patterns. This may involve classification, object recognition, risk scoring, geolocation, prediction or the linking of different data points. Blanchard and Bruun explain that AI-DSS can organise, filter and present information in a way that informs assessments about military action. They also explain that AI-DSS are not limited to one moment in the targeting cycle, but can be used across several stages of targeting, see Figure 2. This means that AI-DSS may influence military decisions before the final moment of attack, for example by shaping which targets are identified, prioritised or presented for review.

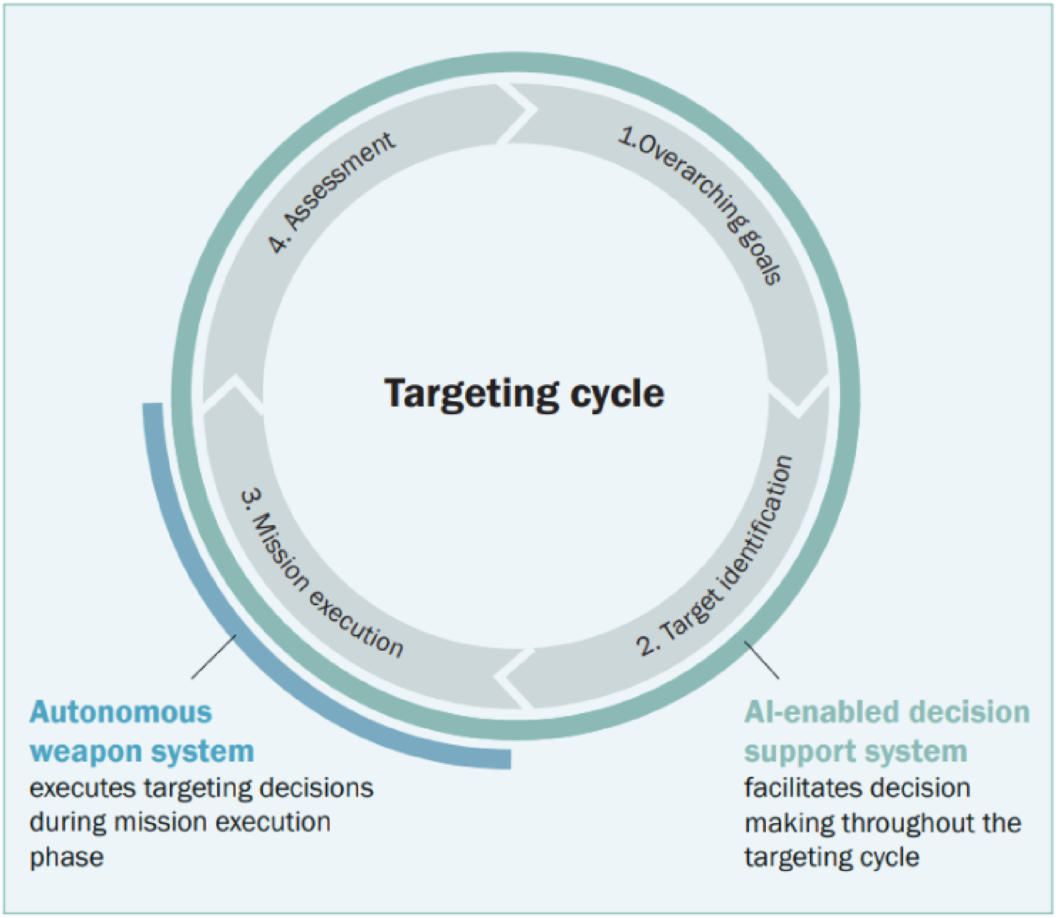

Figure 2: Comparison between autonomous weapon systems and AI-enabled decision support systems in the targeting cycle. Source: SIPRI [5]

The third stage is AI recommendation. After processing the data, the system produces an output that is useful for the human user. This output may be a target list, a risk assessment, a confidence score, a warning or a suggested course of action.

The fourth stage is human decision. In theory, this is where the human operator or commander reviews the AI output and decides whether to approve, reject or question it. This is an important distinction between AI-DSS and autonomous weapon systems. Autonomous weapon systems may select and apply force without further human intervention, while AI-DSS provide information or recommendations to humans.

The final stage is military action, such as strike authorization. At this stage, the AI-DSS has not necessarily "pulled the trigger" itself, but it may have influenced the process that led to military action. This is why AI-DSS are important for this thesis. Their role is not direct in the same way as a weapon, but they can still shape the pathway from information to action. A system that organises data, generates possible targets and presents recommendations can influence what humans see as relevant, urgent or legitimate.

The main issue is therefore not only whether a human remains formally involved. The more important question is whether the human decision is still meaningful.

2.3 Bias

Bias is an important concept when discussing AI decision support systems, because these systems do not simply present neutral information. They collect, process, filter, classify and present data in a certain way. In this thesis, bias can be understood as a systematic distortion that affects how information is produced, interpreted or acted upon[19]. Bias can appear in the technical system itself, but also in the human response to the system's output [11]. The National Institute of Standards and Technology (NIST), a United States federal agency, describes bias in AI as a broad problem that can occur across technology processes and can create harmful impacts even when this is not intended.

For this thesis, it is useful to distinguish between technical biases and cognitive biases. Technical biases are connected to the data, model, design or classification process of an AI system. Cognitive biases are connected to the human user and the way they interpret or rely on AI-generated outputs. This distinction is important because AI-DSS are not fully independent decision makers. They produce outputs that are later interpreted by humans. Therefore, bias can exist both in the system and in the human-machine interaction.

2.3.1 Technical biases

Algorithmic/data bias Algorithmic bias refers to systematic errors in the output of an AI system. These errors may arise because of the data used to train the system, the assumptions built into the mode, or the way the system classifies information. Algorithmic bias is essentially repeated errors in outcomes that can occur when a system is trained on unrepresentative data, making it more accurate for one group than for another[10].

This is closely connected to data bias. AI systems depend heavily on the data they are trained on and the data they later process. This matters because AI outputs can appear objective, even though they are shaped by human choices about what data is collected, how categories are defined and what the system is trained to recognize.

Sampling bias Sampling bias is a specific type of data bias. It occurs when the data used by a system does not properly represent the wider population or situation to which the system is applied. We speak of sampling bias when there arises a situation where some groups in a population are more likely to be selected than others [10]. This can create a distorted picture of reality. For example, if a dataset contains more examples from one group or environment than another, the system may learn patterns that are not generally reliable.

An example of this problem is a study by Buolamwini and Gebru. Their research found large accuracy differences in commercial facial analysis systems, with darker skinned women being misclassified far more often than lighter skinned men. The study showed that some benchmark datasets were overwhelmingly composed of lighter skinned subjects, which contributed to unequal performance across groups[20]. This example is not about warfare, but it clearly shows how biased or unbalanced data can lead to technical bias in the system's outputs.

2.3.2 Cognitive biases

Group attribution bias Group attribution bias occurs when characteristics associated with a group are applied to individuals. Instead of judging someone based only on individual evidence, the person is interpreted through assumptions about the group they are believed to belong to. Andersin describes group attribution bias as involving stereotyping of an "out-group" or favouring an "in-group."

Action bias Action bias refers to the tendency to favour action over inaction. It means that people may feel pressure to "do something", even when waiting or gathering more information would be more careful. And especially in the military during warfare, this bias may be amplified by speed, because very fast AI-generated recommendations in the urgency of war can create pressure to act immediately rather than delay action for further verification.

The idea of action bias has been studied outside military contexts as well. A study on elite football goalkeepers during penalty kicks found that goalkeepers often jumped left or right even when staying in the centre would statistically have been better[21]. It explains this as a bias toward action when action is seen as the normal or expected response.

Automation bias Automation bias is one of the most important cognitive biases for AI-DSS. It refers to the tendency of humans to over rely on automated systems and to accept their outputs without enough critical evaluation. Humans would not have the cognitive abilities to override the authority of the machine's decision[2].

Humans may overly trust computer based decision support systems and may ignore or fail to search for information that contradicts the machine's output. This is important because AI-DSS often appear objective. A human user may assume that the system has processed more information than they could themselves and therefore may place too much trust in the output.

The ICRC (International Committee of the Red Cross) warns that complex decision support systems can create a "rubber stamp" problem, where the human role is reduced to approving or rejecting a machine-generated output. In such cases, the human may remain formally involved, but the actual decision-making process becomes heavily shaped by the system.[11]

Automation bias does not mean that humans always trust AI blindly. Sometimes the opposite can happen, known as algorithm aversion, where people distrust AI after seeing it make mistakes[9]. Van Rossum's research on AI advice shows that time pressure, advice quality and advice source can all affect how much people rely on advice. Her study found that advice can shift decision making and that time pressure can increase reliance on advice. This shows that human reliance on AI is not fixed, but instead depends on the context in which the decision is made.

Confirmation bias Another cognitive bias that might be relevant is confirmation bias. Confirmation bias is the tendency to interpret information in a way that confirms existing beliefs or expectations. In the context of AI-DSS, this could happen when a user already expects a certain result and then treats the AI output as confirmation.[22]

3 Methodology

This thesis researches how AI decision support systems may influence military decision-making, with a specific focus on the role of bias. The main question is how these systems may shape the decisions made by humans and where technical and cognitive biases may enter that process.

Lavender will be used as the main case study, while Gospel, Where's Daddy? and potentially other systems may be discussed to show that Lavender operates within a wider AI-assisted targeting environment. These specific Israeli systems were chosen because they are recent examples of AI-DSS that have reportedly been used during the Gaza conflict over the past few years, making them highly relevant to current debates about AI in warfare. Their reported use has also been discussed more extensively than many other military AI systems, which makes them suitable for a case study based on available reporting, academic literature and official statements.

Since the training data and exact internal design of these systems are not publicly available, this thesis will not attempt to prove how they function internally. Instead, the analysis will focus on reported practices and decision-making structures. The chapter will follow the decision-making chain from data input, to AI processing, to AI-generated recommendation, to human decision and finally to military action. At each stage, the analysis will examine where bias may enter, how serious this risk is and how one type of bias may influence another. In practical terms, each stage of the chain will be analysed by linking reported features of the systems to the bias concepts introduced in the theoretical framework.

The chapter will also focus on the interaction between technical biases (such as sampling or algorithmic bias) and cognitive biases such as automation or action bias.

Finally, the thesis can include a discussion of the broader implications of the findings. This section will consider why the identified risks matter, but also whether AI-DSS may still offer benefits when used under appropriate conditions. The aim is to move beyond simply identifying bias and to explore what the findings mean for human oversight and the future use of AI-DSS in military decision-making.

4 AI-DSS used by Israel

Israel has been reported to use various AI decision support systems (AI-DSS) for warfare in recent years. Each one working differently and serving a different purpose. This chapter will examine AI systems Lavender and The Gospel. It is important to note that information on these systems is limited. They are systems used for the Israeli military so their exact design and training data are not publicly available. In consequence this thesis analyses their reported use rather than claiming to prove exactly how the systems function internally. This chapter is based on investigative journalism, academic analysis and official statements from the IDF (Israeli Defense Forces). Lavender and Gospel both qualify as AI decision support systems meaning they do not officially make the final military decisions by themselves. In the next two sub chapters we will go over Lavender, which focuses on people. And Gospel, which focuses on buildings and infrastructure.

4.1 Lavender

Lavender is reported to be an AI-DSS used to identify individuals suspected of being connected to Hamas or PIJ (Palestinian Islamic Jihad). After October 7th when Israel was invaded by Hamas and in response launched an offensive on Gaza, Lavender was relied upon heavily by the IDF according to +972 Magazine, an independent magazine run by a group of Palestinian and Israeli journalists. [18]. At large scale Lavender was used to generate potential human targets. According to these reports, Lavender works by analysing large amounts of surveillance and intelligence data. It compares individuals in Gaza with known characteristics of Hamas and PIJ operatives. It then assigns people a rating based on the likelihood that they are connected to these groups. This rating is a score from 1 to 100 and according to sources of +972 Magazine these ratings are given to almost every person in Gaza. These recommendations are then given to military commanders.

But what data does Lavender use and how does it obtain it? According to the IDF [23]"Intelligence used to identify objects or persons as military targets may come from a range of sources – including geospatial intelligence (such as satellite and aerial imagery), signals intelligence (such as communications interceptions), human sources, multi-sensor platforms (such as aerial assets and ground-based sensors), enemy military documents and open-source information". Reports from The New York Times [24] in 2024 stated that the Israeli Military established a mass surveillance program on all of Gaza based on facial recognition. It uses facial recognition to recognize an individual without their consent through cameras, military checkpoints and from drone footage [10]. This thesis will not examine facial recognition in detail, but it is important to mention because surveillance data from cameras, checkpoints and drone footage may be used for the Israeli military's AI systems.

According to +972 Magazine [18], Lavender generated at least 37,000 target recommendations during the first six weeks of the conflict. These included both high and low-ranking operatives. The system was also reported to have around 90 percent accuracy. According to the same report, the IDF manually checked the accuracy of a random sample of several hundred targets selected by Lavender to obtain this figure. Even if this 90 percent accuracy figure is accepted, in a military context a 10 percent error rate can still be very serious. Especially when the system is used on such a large scale. If thousands of people are recommended as targets, even a small percentage of mistakes can lead to many wrongly identified individuals.

The final decision to approve an attack still officially relies on the military commander. However, this does not mean that Lavender has no influence on the decision-making process. Because the system generates recommendations, it already shapes which individuals are seen as possible targets and therefore influences what human operators are asked to review. Later in this thesis we will go deeper in this.

Lavender is often used in conjunction with other AI systems such as the Gospel and

Where's Daddy?. Where's Daddy? is reported to be a tracking system that follows suspected militants and sends a signal to the Israeli military when they enter their homes [25]. According to reports [18], this meant that a person identified by Lavender could later be tracked through Where's Daddy?, after which their house could become a potential strike location. This is important because it shows that Lavender was not only connected to identifying possible targets, but also to the timing and location of attacks.

Lavender may identify a person as a possible target, while Where's Daddy? may help locate that person. In this way, the decision-making process is not based on one system alone, but on a larger technological structure in which different systems support different stages of targeting. This is important for later in this thesis when analysing bias, because errors or assumptions in one part of the system can influence later decisions.

4.2 The Gospel

The Gospel, also known as Habsora, is another AI decision support system reportedly used by the IDF. Unlike Lavender, The Gospel does not focus mainly on individuals, but on buildings and infrastructure. It is reported to mark buildings and structures as potential military targets, especially when alleged militants are believed to operate from them.

According to another report by +972 Magazine [26], The Gospel is used to generate targets at a much faster speed than traditional human intelligence processes. It analyses large amounts of intelligence data and surveillance data which were obtained like how was mentioned earlier. It helps identify possible locations for strikes. These locations can include buildings, command centres, tunnels, weapons storage sites or other infrastructure that the military considers relevant. Just like Lavender, The Gospel does not officially make the final decision itself, so they are not autonomous systems and humans are still in the loop. However, it might still be possible that the Gospel gives recommendations that can influence what human operators and commanders review.

5 Analysis

This analysis follows the decision-making chain developed in the theoretical framework: data input, AI processing, AI-generated recommendation, human decision and military action. For each stage, the chapter examines reported examples from investigative journalism, academic literature and policy reports to identify where bias may appear or where conditions for bias may be created. The aim is not to prove the internal functioning of Lavender, Gospel or related systems, because their data and algorithms are not publicly available. Instead, the analysis focuses on the observable decision-making structure around these systems. This makes it possible to assess how technical bias and cognitive bias may interact across the wider AI-assisted targeting process.

5.1 Data input

Before an AI-DSS produces any recommendation, it has already done something less visible which is it has taken in data. This is the first stage of the decision-making chain set out earlier in this thesis and it is easy to overlook because nothing seems to happen yet. No target is generated and no human has acted. This is exactly why the input stage matters however. Whatever the system later processes, ranks or recommends is built on the data it was given. If that data is selective, incomplete or already shaped by earlier assumptions, then the conditions for bias are created here. Long before an operator sees an output.

A first problem is that we often do not know what goes into these systems. In the case of Lavender, the exact data input is not public. According to the IDF, recommendations may come from a wide variety of sources, namely geospatial intelligence, signal intelligence, human sources and open-source information [23]. That list sounds comprehensive, but it is also a description of a very mixed data environment. Each of these sources carries its own limitations.

Signal intelligence, such as communications interceptions, may show what actors are saying or doing but not necessarily what they intend.

Human sources can be partial or politically shaped. For example, in the Israel-Gaza context, human intelligence has sometimes relied on Palestinian informants whose information likely has been shaped by fear and intimidation. Reuters reported back in 2008 that Physicians for Human Rights-Israel documented cases where Gazan patients said they were denied permission to leave Gaza for medical treatment if they refused to provide information to Israeli security agents[27]. This shows that while human sources can provide valuable context, the intelligence they provide may also be influenced by personal vulnerability or political pressure.

Open-source information may also create bias at the data input stage, because publicly available data does not represent every group, language or behaviour equally. If some communities are more visible online than others, an AI-DSS may learn from an uneven picture of the environment. This is an example of sampling bias. As researchers at SIPRI Bruun and Bo argue[28], biased military AI datasets can lead to civilians or civilian objects being wrongly seen as threats or not detected at all.

When these different sources are brought together, they may create the impression of a complete picture of the battlefield, while in reality the data may still be incomplete and shaped by earlier selection choices.

Facial recognition technology is a useful example of how unreliable data can enter at this stage. Andersin discusses reports that Israel used facial recognition in Gaza through cameras, military checkpoints and drone footage and that this kind of technology can be used as input for AI-based systems to identify human targets [10]. The difficulty is that facial recognition is far less reliable in uncontrolled environments. Poor lighting, awkward angles, distance, movement, dust and crowded or damaged surroundings all reduce its accuracy [10]. A warzone is close to a worst-case environment for this technology. This is not only a theoretical concern, one Israeli military official reportedly said the technology sometimes wrongly flagged civilians as wanted Hamas militants [24].

Labelling is also a place where bias may quietly enter. Modern military AI depends heavily on labelled data and the labels are produced by people. The Pentagon's algorithmic-warfare team, Project Maven, used drone video feeds from Iraq and Syria for object identification and data labelling [29]. The word labelling sounds purely technical, but in a military context it involves judgement. Deciding whether an object, a vehicle, a movement or a person is suspicious, hostile or target-relevant is not the same as labelling a cat in a photograph. These categories carry assumptions about what normal and abnormal behaviour look like in a particular place. If those assumptions are skewed, the skew is built directly into the training data and the system might later treat it as ground truth. This connects to the ICRC's warning that machine learning-based DSS involve uncertainties, assumptions and biases in the system, in the human user and in the interaction between the two [11]. The labelling stage is exactly where human assumptions and the technical system meet.

Taken together, these examples suggest that bias may enter the decision-making chain before the AI-DSS produces an output, because the system can only ever process the information it receives. If that information is incomplete, unevenly collected, wrongly labelled, or shaped by earlier human assumptions, then the later recommendation may already be affected before the human operator sees it.

5.2 AI processing and AI recommendation

After the data input stage, the AI-DSS begins to process the information it has received. This is where the system turns data into a classification, score or recommendation. In the case of Lavender, the system reportedly gives individuals a score from 1 to 100 based on the likelihood that they are connected to Hamas or Palestinian Islamic Jihad [18]. This is important because a score can make suspicion look more precise than it really is. A number gives the impression that the system has measured someone's status objectively, while the score is still based on patterns and assumptions.

The risk becomes clearer when looking at how patterns are interpreted. AI systems can identify correlations in data, but correlation is not the same as causation. A person may use a certain phone, communicate with certain people, move through a certain area or share behavioural patterns with known militants, but this does not automatically prove that the person is a militant. AI-DSS can identify correlations, but cannot establish causation in the way a human judgement may require [10].

This is also where group attribution bias may enter. Lavender reportedly works by comparing individuals with known characteristics of Hamas or PIJ operatives. The problem is that the person may then be judged through their similarity to a group rather than through individual evidence. This does not prove that Lavender definitely produces group attribution bias, because the internal criteria are not public. But the reported logic of scoring people based on similarity to known operatives creates a risk that group-based assumptions become part of the classification process.

Bruun and Bo's article "Constant care must be taken to address bias in military AI" gives a useful example of this problem. They explain that an AI system might treat the use of prepaid phone cards as an indicator of enemy status, even though prepaid cards may also be common among groups such as migrant workers [28]. This shows how a system can correctly detect a fact, but still interpret it wrongly. The data point itself may be accurate, but the meaning attached to it may be biased. The relevant bias here is mainly group attribution bias, because a behaviour associated with one group is applied to individuals more broadly. If this assumption is built into the system's classification process and repeatedly affects outputs, it may also become algorithmic bias.

Facial recognition creates a similar issue during AI processing. At this stage, the system compares a face with a database and produces a possible match or confidence score. Andersin explains that facial recognition does not produce absolute certainty, but works with probabilities and confidence levels [10].

The recommendation stage is where these classifications become more operationally important. A score, match or classification can become a target recommendation or part of a target list. Lavender reportedly generated at least 37,000 target recommendations during the first weeks of the war [18]. This scale matters because even a small error rate can affect many people when the system is used on thousands of cases. At this point, the bias is no longer only technical. The AI output begins to shape what human operators see as suspicious, relevant or worth reviewing.

This subchapter therefore shows that AI processing and AI recommendation are not neutral middle steps. They transform uncertain information into scores, matches and targetable outputs. The main biases at this stage are algorithmic bias, data bias and group attribution bias. Algorithmic bias may appear in the way the system weighs and interprets indicators. Data bias may continue from the input stage into the output. Group attribution bias may appear when individuals are judged through similarity to a wider profile. The next section will examine what happens when these outputs are passed to human operators and commanders for review.

5.3 Human decision and military action

After the AI-DSS has produced a recommendation, the next stage is human review. This is the stage where the IDF argues that the system does not make the final decision by itself, because a human still has to approve the target [23]. However, the question is not only whether a human is formally involved. The more important question is whether that human is able to make an independent judgement.

This is where automation bias becomes relevant. If a person has already been selected and scored by Lavender, the human reviewer is not starting from a neutral position. The person is already presented as a possible target. This may make the recommendation feel more reliable, especially because it comes from a technical system that appears to have processed much more information than a human could. In that situation, the human may rely too much on the system's output instead of actively looking for reasons to question it.

The risk becomes stronger when the review process is fast. An IDF official explained that target review, which used to take much longer, could now take only minutes. [30]. In some cases human review of Lavender-generated targets was extremely limited [18], sometimes as little as 20 seconds according to former operators within the IDF. This does not prove that every decision was automatically accepted, but it does suggest a decision-making environment where automation and action bias could become more likely. If the system produces targets quickly and the human only has a short time to review them, the human role may become more about confirming the output than independently assessing it.

Confirmation bias may also play a role at the human review stage. If an operator already expects certain people, neighbourhoods, communication patterns or behaviours to be connected to militant activity, an AI-generated recommendation may appear to confirm that expectation.

The final stage is military action. At this point, action bias is important. In war, there is already pressure to act quickly because a suspected target may move, escape or become harder to locate. AI-DSS can intensify this pressure by producing recommendations at high speed and by making targets appear immediately actionable. Where's Daddy? is relevant here, because it reportedly alerted the IDF when suspected militants entered their homes [18]. This kind of alert can create a sense of urgency: the target is available now, so waiting may feel like losing the opportunity.

This does not mean that every strike following an AI recommendation is the result of bias. But it does show how automation bias and action bias may interact. The system first presents someone as suspicious, which may encourage reliance on the recommendation. Then the military context creates pressure to act quickly, which may reduce time for further verification. In this way, the human decision may still exist formally, but it can be strongly shaped by the structure and urgency created by the AI-assisted targeting process.

6 Conclusions and Further Research

This thesis has analysed how bias may appear in the reported use of AI decision support systems in military decision-making, with Lavender as the main case study and Gospel, Where's Daddy? and other related systems as part of the wider targeting environment. Because the internal design and training data of these systems are not publicly available, the thesis did not try to prove exactly how they function. Instead, it focused on the reported decision-making chain from data input, to AI processing and recommendation, to human decision and military action.

The analysis showed that bias can enter at several points in this chain. At the data input stage, incomplete or uneven intelligence may already shape the system's later output. During AI processing and recommendation, uncertainty can be turned into scores, matches or target lists, making suspicion appear more precise and objective than it really is. At the human decision and military action stages, automation bias and action bias may become relevant, especially when operators review AI-generated outputs under time pressure and in the urgency of war.

The main conclusion is that AI-DSS can influence military decision-making even when they do not make the final decision themselves. Their power lies in shaping what humans see as relevant, suspicious or urgent. This means that simply keeping a human "in the loop" is not enough. The more important question is whether that human has enough time, information and the right state of mind to make a meaningful and independent judgement.

Does this mean that AI-DSS should automatically be rejected? I think this is difficult to conclude. They may help militaries process large amounts of information and improve situational awareness. However, their use in targeting requires strong safety measures, especially where civilian harm is possible. These measures should include better verification, clearer accountability, more transparency about system limitations and enough time for human review. Without these conditions, there is a risk that uncertain or biased recommendations become too easily turned into military action. These systems do assist in killing thousands of humans, many of which were innocent civilians that were not meant to be killed. If you look at the other side however, if militaries were not to use these systems, would there be more or less bias and innocents dying? That could be something to look into. Further research is also needed when more information becomes available about how these systems are developed, trained and used in practice. Much of the information used in this thesis only came to light in recent years. As AI-DSS are likely to become more common in future conflicts, understanding their risks and limits will only become more important.

References

[1] BBC. At least 153 dead after reported strike on school, Iran says, 2026. https://www.bbc.com/news/articles/c1l7rvqq51eo.

[2] The Times. The first AI war: US and Israel use Iran to test autonomous tech, 2026. https://www.thetimes.com/world/middle-east/article/ai-war-us-israel-iran-school-drone-sttq7jqt0.

[3] CNN. US Tomahawk struck Iranian base next to school destroyed in deadly attack, video appears to confirm, 2026. https://edition.cnn.com/2026/03/09/middleeast/iran-minab-school-airstrike-new-footage-intl.

[4] The Telegraph. 1,000 targets a day in Iran: How AI is accelerating war, 2026. https://www.telegraph.co.uk/news/2026/05/01/1000-targets-day-how-ai-accelerating-america-iran-war/.

[5] Alexander Blanchard and Laura Bruun. Autonomous Weapon Systems and AI-enabled Decision Support Systems in Military Targeting: A Comparison and Recommended Policy Responses, 2025. https://www.sipri.org/publications/2025/other-publications/autonomous-weapon-systems-and-ai-enabled-decision-support-systems\protect\penalty\z@-military-targeting-comparison-and.

[6] Marko Bohanec. Data Mining and Decision Support. pages 23–35, 1990. https://link.springer.com/chapter/10.1007/978-1-4615-0286-9_3#auth-Marko-Bohanec.

[7] Anthony King. Artificial Intelligence and Military Decision-Making, 2024. https://academic.oup.com/jogss/article/9/2/ogae009/7667104.

[8] Clinical Imaging. Artificial intelligence decision-support systems in radiology workflow. 2022. https://www.sciencedirect.com/science/article/abs/pii/S0899707122002352.

[9] Emma van Rossum. Under Pressure: Examining the Influence of AI and Human Expert Advice on Decision-Making Under Time Constraints, 2025. https://studenttheses.uu.nl/items/ba6ae284-5423-47f3-9b63-119630ae05f3.

[10] Emelie Andersin. The Use of the 'Lavender' in Gaza and the Law of Targeting: AI-Decision Support Systems and Facial Recognition Technology. *Journal of International Humanitarian Legal Studies*, pages 336–370, 2025. https://scholarlypublications.universiteitleiden.nl/access/item%3A4287028/view.

[11] International Committee of the Red Cross. Decisions, Decisions, Decisions: Computation and Artificial Intelligence in Military Decision-Making, 2024. https://library.icrc.org/library/docs/DOC/icrc-4757-002.pdf.

[12] Reuters. Ukraine collects vast war data trove to train AI models, 2024. https://www.reuters.com/technology/ukraine-collects-vast-war-data-trove-train-ai-models-2024-12-20/.

[13] Russia Matters. The Roles and Implications of AI in the Russian-Ukrainian Conflict, 2024. https://www.russiamatters.org/analysis/roles-and-implications-ai-russian-ukrainian-conflict.

[14] CNA. Artificial Intelligence and Autonomy in Russia, 2021. https://www.cna.org/reports/2021/05/Artificial-Intelligence-and-Autonomy-in-Russia.pdf.

[15] Kommersant. Reporting on Russia and lethal autonomous weapon systems at the UN Convention on Certain Conventional Weapons, 2018. https://www.kommersant.ru/doc/3714110?from=doc_vre.

[16] Michael Klare. AI Plays Major Role in the War on Iran. *Arms Control Today*, 2026. https://www.armscontrol.org/act/2026-05/news/ai-plays-major-role-war-iran.

[17] Myriam Dunn Cavelty. What We Know and Don't Know About AI in Iran, 2026. https://css.ethz.ch/en/center/CSS-news/2026/04/what-we-know-and-dont-know-about-ai-in-iran.html.

[18] Yuval Abraham. 'Lavender': The AI Machine Directing Israel's Bombing Spree in Gaza, 2024. https://www.972mag.com/lavender-ai-israeli-army-gaza/.

[19] National Institute of Standards and Technology. Towards a Standard for Identifying and Managing Bias in Artificial Intelligence. Technical Report Special Publication 1270, NIST, 2022. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.1270.pdf.

[20] Joy Buolamwini and Timnit Gebru. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of Machine Learning Research*, volume 81, pages 1–15, 2018. https://proceedings.mlr.press/v81/buolamwini18a.html.

[21] Michael Bar-Eli, Ofer H. Azar, Ilana Ritov, Yael Keidar-Levin, and Galit Schein. Action Bias Among Elite Soccer Goalkeepers: The Case of Penalty Kicks. *Journal of Economic Psychology*, 28(5):606–621, 2007. https://www.sciencedirect.com/science/article/abs/pii/S0167487006001048.

[22] Wikipedia. Confirmation bias. https://en.wikipedia.org/wiki/Confirmation_bias.

[23] Israel Defense Forces. The IDF's Use of Data Technologies in Intelligence Processing, 2024. https://www.idf.il/210062.

[24] Sheera Frenkel. Israel Deploys Expansive Facial Recognition Program in Gaza, 2024. The New York Times. https://www.nytimes.com/2024/03/27/technology/israel-facial-recognition-gaza.html.

[25] Roos Creyghton. The Lavender Precedent: Automated Kill Lists and the Limits of International Humanitarian Law, 2025. https://aoav.org.uk/2025/the-lavender-precedent-automated-kill-lists-and-the-limits-of-international-humanitari

[26] Yuval Abraham. 'A Mass Assassination Factory': Inside Israel's Calculated Bombing of Gaza, 2023. https://www.972mag.com/mass-assassination-factory-israel-calculated-bombing-gaza/.

[27] Reuters. Israel demands Gaza patients be informants-report, 2008. https://www.reuters.com/article/economy/israel-demands-gaza-patients-be-informants-report-idUSL3135222/.

[28] Laura Bruun and Marta Bo. 'Constant care' must be taken to address bias in military AI, 2025. ICRC Humanitarian Law and Policy Blog. https://blogs.icrc.org/law-and-policy/2025/08/28/constant-care-must-be-taken-to-address-bias-in-military-ai/.

[29] Dustin A. Lewis, Naz K. Modirzadeh, and Gabriella Blum. The pentagon's new algorithmic-warfare team, 2017. Lawfare. https://www.lawfaremedia.org/article/pentagons-new-algorithmic-warfare-team.

[30] Marissa Newman. Israel Quietly Embeds AI Systems in Deadly Military Operations, 2023. Bloomberg. https://www.bloomberg.com/news/articles/2023-07-16/israel-using-ai-systems-to-plan-deadly-militaryoperations.