

# **Joint Master in Global Economic Governance and Public Affairs**

***Behind the Target: Palantir, Project Maven  
and Epistemic Dependency in AI-Enabled  
Warfare***

**Supervised by Dr. Frascani**

**Bébel Oijevaar**

**2025/2026**



Co-funded by  
the European Union

**Thesis Pitch:** [Link](#)

### **Statutory Declaration**

I hereby declare that I have composed the present thesis autonomously and without use of any other than the cited sources or means. I have indicated parts that were taken out of published or unpublished work correctly and in a verifiable manner through a quotation. I further assure that I have not presented this thesis to any other institute or university for evaluation and that it has not been published before

21<sup>st</sup> of June, 2026

Bébel Charlotte Oijevaar



### **Acknowledgements**

I would like to express my sincere gratitude to my supervisor, Professor Frascani, for his valuable guidance, feedback and encouragement throughout this research. I am also deeply grateful to my grandmother Marietta, for her support and for the help she provided during my research on Palantir and AI-enabled warfare. Her encouragement and involvement have meant a great deal to me throughout this process.

### **Abstract**

Artificial intelligence is changing military targeting through systems that integrate, classify and operationalise large volumes of information. Many of these systems are developed and maintained by private technology companies, raising a central accountability problem: states may retain formal authority over military decisions while becoming dependent on private actors for the knowledge through which those decisions are made. This thesis examines how Palantir Technologies' role in Project Maven produces epistemic

dependency within U.S. military targeting and considers the consequences for accountability under international humanitarian law.

The thesis conceptualises epistemic dependency as a condition in which state actors remain formally responsible for decisions but cannot independently reproduce the privately controlled processes through which operational knowledge is produced. A qualitative, theory-guided case study analyses procurement and oversight documents, policy documents and investigative journalism. The analysis focuses on platform inspection, the organisation of target-relevant knowledge and the reconstruction of targeting decisions.

This thesis finds that while the Department of Defence controls military data and strike authorisation, it cannot independently review Maven's integrated technical setup. Palantir's systems influence how data are organised and processed. These dependencies matter when fast decisions limit verification and investigations depend on private records. The Minab school strike highlights how outdated intelligence may still be used in fast-paced targeting. Ultimately, effective accountability in AI-enabled warfare requires the state to independently verify and reconstruct the knowledge driving military actions.

**Keywords:** Palantir Technologies; Project Maven; Maven Smart System; epistemic dependency; artificial intelligence; algorithmic opacity; private authority; international humanitarian law; accountability

## Table of Contents

|                                                                                                   |           |
|---------------------------------------------------------------------------------------------------|-----------|
| <b>1. Introduction.....</b>                                                                       | <b>5</b>  |
| 1.1 Big Data and the Rise of Palantir.....                                                        | 7         |
| 1.1.1 The rise of big data .....                                                                  | 7         |
| 1.1.2 Founding and evolution of Palantir.....                                                     | 8         |
| 1.1.3 Project Maven .....                                                                         | 10        |
| <b>2. Literature review .....</b>                                                                 | <b>12</b> |
| 2.1 Introduction.....                                                                             | 12        |
| 2.2 Stream One: Private Authority, Security Governance and the Limits of Existing Frameworks..... | 13        |
| 2.2.1 <i>Foundations: Private Authority in International Relations</i> .....                      | 13        |
| 2.2.2 <i>AI Companies as Private Authorities in International Relations</i> .....                 | 14        |
| 2.2.3 <i>Security Privatisation: From Private Military Companies to Data Infrastructure</i> ..... | 15        |
| 2.2.4 <i>Delegation, Dependency and the Erosion of State Capacity</i> .....                       | 16        |
| 2.3 Stream Two: Algorithmic Governance, Opacity, and the Accountability Gap .....                 | 16        |
| 2.3.1 <i>The Algorithmic Turn in Governance</i> .....                                             | 16        |
| 2.3.2 <i>Opacity and the Structural Prevention of Accountability</i> .....                        | 17        |
| 2.4 The Accountability Gap under International Humanitarian Law .....                             | 18        |
| 2.4.1 <i>The Foundational Assumptions of IHL</i> .....                                            | 18        |
| 2.4.2 <i>Where the Framework Breaks Down</i> .....                                                | 19        |
| 2.5 Theoretical Bridge: Epistemic Dependency and the Accountability Gap .....                     | 20        |
| <b>3. Methodology .....</b>                                                                       | <b>22</b> |
| 3.1 Research questions.....                                                                       | 22        |
| 3.2 Research design .....                                                                         | 22        |
| 3.3 Research Strategy.....                                                                        | 23        |
| 3.4 Operationalisation Epistemic Dependency .....                                                 | 24        |
| 3.5 Data collection .....                                                                         | 25        |
| 3.6 Limitations .....                                                                             | 26        |
| <b>4. Analysis .....</b>                                                                          | <b>26</b> |
| 4.1 Algorithmic Opacity .....                                                                     | 26        |
| 4.2. Epistemic Authority .....                                                                    | 31        |
| 4.3 Accountability gap .....                                                                      | 35        |
| <b>5. Discussion.....</b>                                                                         | <b>39</b> |
| <b>6. Conclusion .....</b>                                                                        | <b>43</b> |
| <b>Bibliography .....</b>                                                                         | <b>45</b> |

## **List of Abbreviations**

AI — Artificial intelligence

BAS-T — Broad Area Surveillance-Targeting

CIA — Central Intelligence Agency

DoD — Department of Defence

DoDI — Department of Defence Instruction

GAO — Government Accountability Office

ICRC — International Committee of the Red Cross

IHL — International humanitarian law

IP — Intellectual property

IRGC — Islamic Revolutionary Guard Corps

ISR — Intelligence, surveillance, and reconnaissance

LLM — Large language model

MCA — Mission command application

MSS — Maven Smart System

NASEM — National Academies of Sciences, Engineering, and Medicine

NATO — North Atlantic Treaty Organization

PMC — Private military company

R&D — Research and development

UN — United Nations

U.S. — United States

## 1. Introduction

On 28 February 2026, Shajareh Tayyebbeh elementary school in Minab, Iran, was reportedly struck during the first day of the U.S. and Israeli military operations in Iran (Pearson & McNeill, 2026). This catastrophe killed at least 170 school children and injured more (Baker, 2026). Internal U.S. investigations suggest the strike resulted from the school being incorrectly identified as a military base during the targeting process due to outdated data (Baker, 2026). The school was located near a long-standing Islamic Revolutionary Guard Corps naval installation and had reportedly once fallen within the perimeter of that compound. Yet by the time of the strike, it had allegedly operated as a school for several years and displayed public indicators of civilian use, including online traces and business listings (Baker, 2026). The strike therefore stands as an extreme example of the civilian harm that can follow when outdated target information is carried into a fast, technology-mediated targeting process. If U.S. responsibility is confirmed, Reuters reports that the incident would rank among the worst cases of civilian casualties in decades of U.S. military operations in the Middle East (Ali & Stewart, 2026).

It is important to understand how a school became a target. This involves reviewing available information, ignored warnings, and verification of its military or civilian status. Modern targeting increasingly depends on automated classifications and fast human approvals. Many of these systems are developed by private technology companies that provide the infrastructure for collecting and analysing military information. These concerns matter because recent conflicts have seen a growing use of artificial intelligence in warfare, from AI-enabled drone systems in Ukraine to reported AI-supported targeting in Gaza. (Hunder, 2024; McKernan & Davies, 2024). As a result. A tension emerges between formal state responsibility for military actions and the growing role of private actors in shaping the information on which those actions depend.

Palantir Technologies is one of those companies that develops data platforms that help military actors to organise and interpret large volumes of information (Palantir, 2025). Its role is examined through Project Maven; a U.S. Department of Defence programme launched in 2017 to integrate artificial intelligence into military intelligence workflows. Through Project Maven, Palantir's software became part of the infrastructure through which battlefield information is processed and turned into operational knowledge. This creates a difficult accountability problem: the state remains formally responsible for the use of force, yet the information behind that force may be structured through systems developed and owned by a private company.

The Minab strike shows the human stakes of this problem. A reported targeting error did not merely produce a technical failure; it produced mass civilian harm, killing children in a place that functioned as a school not a military site. If outdated information about the building remained within a target package, the problem is not confined to the final approval of the strike. It also concerns the earlier process through which the school was classified, reviewed, and made actionable. In such a setting, accountability cannot be assessed only by identifying the person who authorised force. It requires reconstructing how a civilian object became visible as a target, which information was included or excluded, whether the classification was current, and whether human reviewers had the time and capacity to challenge the information presented to them. The ethical concern in AI-enabled warfare is that human judgement may remain formally present while the informational conditions of that judgement are increasingly shaped by systems beyond military actors' full inspection.

This creates the central research puzzle of this thesis. Under international humanitarian law, responsibility is assigned to human and state actors: commanders are required to differentiate between civilian and military targets, evaluate proportionality, and undertake feasible precautions prior to any attack (ICRC, n.d., Rule 15). However, AI-enabled targeting increasingly relies on privately developed systems. As a result, while states retain formal responsibility for the use of force, the foundational knowledge guiding such actions may be influenced by private infrastructure. This thesis explores this tension by examining Palantir Technologies' involvement in Project Maven and investigates whether dependence on Palantir's AI-enabled data infrastructure establishes epistemic dependency for U.S. military personnel. It further considers the implications for accountability within international humanitarian law and assesses whether protections for civilians can be properly evaluated when targeting decisions rest on information systems whose classification methods, verification protocols, and post-strike documentation may be challenging to access or review.

The thesis proceeds as follows. Chapter 1 provides the background to big data, Palantir, and Project Maven. Chapter 2 develops the theoretical framework by linking literature on private authority, algorithmic governance, and accountability under international humanitarian law. Chapter 3 sets out the methodology. Chapter 4 analyses Project Maven through the three dimensions of algorithmic opacity, private epistemic authority, and the accountability gap. Chapter 5 discusses the ethical and future implications of these findings, before Chapter 6 concludes by answering the research question.

## **1.1 Big Data and the Rise of Palantir**

Over the past two decades, there has been a significant transformation in state governance and warfare. The rise of big data and algorithmic technologies has fundamentally altered security and public administration by enabling advanced methods for prediction, control, and decision-making (Kitchin, 2014; Zuboff, 2019). Artificial Intelligence (AI) is rapidly changing the nature of warfare, becoming an integral part of military operations. Today, AI-powered systems are being woven into the "kill chain", the process that spans from identifying targets to carrying out lethal actions. This evolution prompts important questions about the role of human judgment in life-and-death decisions (Roff, 2014). Recent developments in Gaza have highlighted these issues, with reports of AI targeting tools used on an unprecedented scale during combat, reducing the time from analysing data to executing a strike to just seconds (Abraham, 2024). This raises fundamental questions about responsibility: when decisions of war are mediated by privately owned algorithmic systems, it becomes increasingly unclear who bears accountability for their consequences (El-Baroudi, 2023).

At the heart of this transformation is Palantir, a company that defies easy categorisation. Unlike traditional defence contractors or standard software firms, Palantir operates at the nexus of private investment and governmental authority. Palantir supplies data integration and analytical solutions to intelligence agencies, military organizations, and governments worldwide (Abrahamsen & Williams, 2011). Palantir's recent participation in conflicts such as the war in Gaza, where its platforms have facilitated battlefield decision-making, demonstrates the integration of private data enterprises into contemporary warfare (Palantir, 2024). This chapter examines the evolution of big data and the rise of Palantir as a security actor, providing a foundation for subsequent discussions concerning contemporary warfare.

### ***1.1.1 The rise of big data***

The emergence of actors like Palantir is inseparable from a broader transformation in how data itself is produced and mobilised as a resource for power. Data are generally seen as raw elements formed by categorizing aspects of the world, serving as the foundation for information and knowledge (Kitchin, 2014). The true value of data comes from how they are gathered and handled. Before big data became common, creating and examining data was slow, expensive, and narrow in focus, often resulting in insights that were disconnected (Kitchin, 2014). The concept of big data refers to both the vast quantity and the intricate complexity of newly available information, as well as the technical challenges associated

with processing it (Dumbill, 2013). In various contexts, big data has introduced significant challenges to data management, which are commonly classified into three key dimensions: (1) the substantial volume of data; (2) the extensive variety arising from diverse sources; and (3) the rapid velocity at which data must be processed (Niebel et al., 2018). The rise of big data marked a major shift in this dynamic. Advances in digital infrastructure enabled the continuous collection, storage and analysis of vast quantities of data, transforming them into a central resource (Kitchin, 2014). Rather than simply describing the world, data increasingly became a means of predicting and intervening in it. Although not all knowledge comes directly from data, like opinions and beliefs, data are undeniably fundamental to how we interpret the world. Data serve as the essential raw material for activities such as gathering, sorting, categorising, matching, profiling, and modelling, all of which help transform data into information and knowledge to better understand, predict and manage various phenomena (Kitchin 2014). Gathering data over time enables to track changes across time, space, and scale. As data drive insights, access to extensive, high-quality datasets and analytical skills becomes a powerful asset. Those who collect and integrate data hold an advantage in influencing decisions and defining knowledge (Kitchin, 2014). Big data thus signals not just technological progress but a shift toward governance focused on data-driven control and prediction.

In this context, companies as Palantir have emerged. By specialising in the integration and analysis of large, disparate datasets, Palantir establishes itself as an essential intermediary for converting data into intelligence (Iliadis & Acker, 2022). Its activities highlight how private entities can connect the potential of big data to shape decision-making processes, especially within areas traditionally managed by the state, such as security. Understanding how this came to be requires examining Palantir's origins.

### ***1.1.2 Founding and evolution of Palantir***

Founded in 2003, Palantir emerged as a response to the growing demand for data-driven security solutions in the aftermath of 9/11. Palantir was established by Peter Thiel, Alex Karp, Joe Lonsdale and Nathan Gettings (Iliadis & Acker, 2022). The goal was creating software that leverages intelligence to prevent terrorist attacks (Chiang, 2011). Peter Thiel, known for founding PayPal, played a key role in launching the company and was at the same time its biggest investor. Thiel identified a major turning point after 9/11. At the time, his company PayPal used sophisticated pattern recognition algorithms to spot fraudulent transactions (Chiang, 2011). He found that the same technologies could be used to track

money flows linked to terrorism, which inspired the creation of Palantir. Thiel believed that technology had the potential to enhance national security (Chiang, 2011).

At its core, Palantir is a data integration and analytics company. The company supplies information technology solutions for data integration and tracking to government agencies, humanitarian organisations and corporations (Iliadis & Acker, 2022). Palantir stands out technically because it leverages *metadata* and *ontologies* to connect datasets that would otherwise remain isolated (Iliadis & Acker, 2022). Metadata assigns labels to individual entities—such as people, locations, or objects—along with their characteristics like height or position, and their connections, whether those are family ties, organizational links, or physical relationships. This approach enables different types of data across various formats and domains to connect, creating bridges between points that would not be interoperable without such tools (Iliadis & Acker, 2022). The standardisation of metadata necessitates an ontology: a defined set of formal metadata rules, terms, values, and relationships that enables various individuals or systems to model their data. This approach ensures that diverse datasets can be integrated effectively for browsing and utilization (Iliadis & Acker, 2022). Palantir refers to this as an “enterprise-wide semantic layer” that translates raw data into a language users can easily understand (Palantir, 2022). This system links information across multiple databases. For example, matching a name in one with a location in another and a communications trace in a third, to enable pattern detection and predictive analytics far beyond human capability (Iliadis & Acker, 2022). Palantir’s technology not only identifies unknown individuals by connecting them to various records, but also visualises their position and potential actions, anticipating outcomes and making these actionable.

Initially, In-Q-Tel, the investment arm of the CIA, showed interest in the company and invested \$2 million, complemented by Peter Thiel's contribution of \$30 million, enabling Palantir to develop its anti-terrorism software (Gorman, 2009). This early investment positioned Palantir within U.S. intelligence networks, with the CIA serving as its primary client between 2005 and 2008. Gradually, Palantir expanded from intelligence agencies into the defence, law enforcement, and, eventually, commercial sectors. Its core platforms demonstrate a dual focus: *Gotham*, designed for government and military purposes, facilitates the integration and analysis of extensive datasets for intelligence operations. And *Foundry* which is directed at commercial clients, offering similar capabilities for corporate data management and decision-making (Palantir, 2025).

What makes Palantir politically significant is not the scale of its operations alone, but the position it occupies between private enterprise and public authority in decisions that were

once the exclusive domain of the state. The company does not merely supply tools to states but structures the informational environment within which state decisions are made (Krishnaswamy, 2024). Formally, sovereignty over decisions of war and governance remains with the state. In practice, however, those decisions increasingly flow through exclusive systems that the state does not own and cannot fully audit. Abrahamsen and Williams (2011) argue that private security actors now play a much more influential role, becoming vital to security governance and blurring the lines between public authority and private enterprise. The political implications of this are intensified by the characteristics of algorithmic systems themselves. Rather than merely carrying out decisions, these systems actively shape the environment in which choices are made, subtly relocating real power without any official transfer of authority (Yeung, 2018).

This raises a puzzle that existing frameworks struggle to resolve. Classical accounts of state authority rest on the assumption that the legitimate use of force and the tools through which it is exercised remain under public control (Weber, 1919). Palantir's position complicates this assumption in ways that are difficult to map onto conventional categories. It is neither a contractor executing clearly defined state instructions, nor an independent actor pursuing its own agenda. It occupies an ambiguous position that Avant, Finnemore and Sell (2010, p. 2) describe as characteristic of the new landscape of global governance, in which authority is not given but made, and in which private actors increasingly set the terms under which states operate.

### ***1.1.3 Project Maven***

Project Maven was initiated in 2017 as the determination by the U.S. Department of Defence to incorporate artificial intelligence and machine learning into military intelligence processes (Pellerin, 2017). Officially established through the Algorithmic Warfare Cross-Functional Team, the program arose from concerns within the Department regarding the increasing burden on human analysts caused by the vast amounts of imagery and sensor data produced during modern military operations (Pellerin, 2017). The project's initial scope was focused: to implement computer-vision algorithms for identifying objects in both full-motion video and still imagery, enabling analysts to manage substantial volumes of intelligence, surveillance, and reconnaissance data with greater efficiency (Pellerin, 2017). According to Colonel Drew Cukor, the goal was to foster collaboration between "people and computers" in object detection, with algorithms assisting rather than replacing human analysts (Pellerin, 2017).

From the beginning, Maven was subject to political scrutiny due to its integration of commercial artificial intelligence expertise with military targeting operations. Google, which was an early technology partner for the program, announced in 2018 that it would not renew its contract for Maven following significant internal resistance, which included employee resignations and a petition signed by thousands (Griffith, 2018). This event provided a notable example of the ethical challenges associated with Silicon Valley's involvement in military AI. While Google's exit did not terminate Project Maven, it represented a turning point in the program's contractor landscape, enabling other defence technology companies, such as Palantir, to take on a more prominent role in its progression.

Maven gradually moved beyond its original purpose of analysing drone imagery. The programme evolved into a comprehensive system for data fusion and decision support, linking various types of intelligence information to operational activities (Russo, 2026). Today, the Maven Smart System (MSS) is described as a platform that pulls together data from sources like drones, satellites, sensors, radar, geolocation, and other intelligence streams to help identify personnel, equipment, and items relevant to military operations (Russo, 2026). Rather than being just one algorithm, Maven acts as an infrastructure to organise, analyse, and apply target-focused information. Its results can include areas for operators to examine, suggested actions for commanders to consider, or useful data that fits into military decision-making processes (Uppal, 2021).

The Maven Smart System (MSS) serves as the principal operational interface within this comprehensive programme. According to Palantir Technologies (2024), MSS facilitates AI-driven battlespace awareness, global integration, force management, contested logistics, and joint fires and targeting workflows across military branches. Department of Defence reports similarly identify MSS as an AI-based decision-support tool developed through the XVIII Airborne Corps' Scarlet Dragon exercise series. The system aggregates sensor data, applies advanced algorithms, aids personnel in identifying and engaging military targets, supports chain-of-command approvals, and documents battle-damage assessments post-strike (U.S. Department of Defence, 2024c). In this capacity, MSS not only presents operational information but also organises data into a unified picture, reveals potential targets to users, and manages the flow of outputs through approval and evaluation processes.

Palantir's involvement in Project Maven has expanded significantly since 2024. The company secured a \$480 million contract in May 2024 for the Maven Smart System prototype, followed by a \$99.8 million agreement in September for software licenses and support (U.S. Department of Defence, 2024). In May 2025, an additional \$795 million was

allocated for software licenses, increasing the program's scale (U.S. Department of Defence, 2025). Reuters reported that Maven was designated as a core military system, with AI-driven decision-making identified as crucial to future Pentagon operations (Jeans, 2026).

This development positions Maven at the forefront of contemporary discussions regarding AI-enabled warfare. Palantir asserts that its software does not render lethal decisions, emphasizing that human operators retain responsibility for target selection and approval (Jeans, 2026). Nevertheless, the increasing integration of AI within targeting processes has prompted legal, ethical, and security concerns, particularly in situations where operational speed, data integrity, and human oversight converge (Baker, 2026; Jeans, 2026). For the purposes of this thesis, Project Maven is interesting as it represents neither a fully autonomous weapon nor merely an administrative instrument. Rather, it functions as a socio-technical infrastructure, facilitating the intersection of proprietary software, military data, and governmental decision-making in the generation of target-specific intelligence. The following literature review therefore engages two bodies of literature: work on the privatisation of security and the conditions under which private actors acquire *de facto* authority in domains traditionally reserved for the state; and work on algorithmic governance and the political implications of AI-driven decision-making. Together, these provide the conceptual framework necessary to examine the conditions under which Palantir's authority has gone largely uncontested, and what this reveals about the changing nature of power in contemporary governance and warfare.

## **2. Literature review**

### **2.1 Introduction**

The integration of artificial intelligence into armed conflict has generated debate concerning operational risk, compliance with international humanitarian law and accountability (Boutin, 2023; International Committee of the Red Cross [ICRC], 2024). Although states retain legal responsibility for military operations, the assumption that they exercise effective control over the AI systems involved is increasingly difficult to sustain. Private software companies now play a central role in designing and developing military AI and decision-support systems, complicating the relationship between technical control and public responsibility (Boutin, 2023; Vanderborght & Nadibaidze, 2025).

This thesis examines that breakdown through the case of Palantir Technologies and its role in Project Maven. The central research question is:

*“To what extent does Palantir Technologies’ role in Project Maven create a condition of epistemic dependency for U.S. military actors, and how does this condition complicate accountability under international humanitarian law?”*

The review proceeds in three stages. First, it maps the literature on private authority and security privatisation, identifying its limits when applied to algorithmic actors. Second, it maps the literature on algorithmic governance and the opacity of AI decision-making systems, identifying its limits when applied to questions of sovereign authority and international law. Third, it constructs the theoretical bridge between these two streams, introducing the concept of *epistemic dependency* as the analytical thread that connects private authority, algorithmic opacity, and the accountability gap this thesis seeks to explain. The two streams are not treated as parallel bodies of knowledge to be summarised separately: they are developed in dialogue with each other from the outset, because the argument of this thesis depends on their intersection.

## **2.2 Stream One: Private Authority, Security Governance and the Limits of Existing Frameworks**

### ***2.2.1 Foundations: Private Authority in International Relations***

Scholars of international relations have long recognised that private actors such as corporations and industry groups can take on governance roles that were once seen as belonging exclusively to the state. The foundational literature on private authority established a basic but important distinction: authority is not the same as power (Hall & Biersteker, 2002). Power forces compliance; authority means that others follow the lead not because they have to, but because they accept the judgement of the leader as legitimate or simply unavoidable (Hall & Biersteker, 2002). Private actors can acquire this kind of authority through their market position or their knowledge that others depend on (Cutler, Haufler and Porter, 1999).

Constructivist scholarship deepened this insight by demonstrating that authority is not simply a property of states but is socially produced (Wendt, 1999). It is dependent on recognition, legitimacy, and the shared understandings that make certain actors’ judgements count as authoritative (Wendt, 1999). Neoliberal institutionalism showed that international outcomes are shaped by the rules and information environments that structure how states

interact, not only by their material capabilities (Keohane, 1984). Both traditions opened the door to the idea that whoever controls the frameworks within which decisions are made holds a form of power that can be more consequential than the capacity to compel action directly.

These foundations are useful but limited. The private actors described in this literature exercised their authority in visible ways: through contracts and market dominance. They did not anticipate a private actor that exercises authority by designing the systems through which states themselves perceive the world and make decisions. AI and data infrastructure companies do something qualitatively different and understanding that difference requires extending the existing framework considerably.

### ***2.2.2 AI Companies as Private Authorities in International Relations***

Srivastava's 'Algorithmic Governance and the International Politics of Big Tech' (2023) is a key work examining how AI and data companies act as private authorities in international relations. She suggests that major tech firms like Google, Facebook and Amazon gain private authority in a modern way. These systems process and classify large volumes of data in ways that may make public institutions reliant on privately developed technical capabilities. Srivastava (2023, p. 989) conceptualises this as private authority exercised through control over “critical bottlenecks of knowledge, connection, and desire.” In doing so, they shape what governments and militaries can see and decide. What makes this form of authority particularly difficult to challenge is that it is largely invisible. These companies do not openly claim to govern. But the categories built into their systems determine what counts as a threat and what actions become possible. Authority is exercised through design, not declaration (Srivastava, 2023).

Two consequences of this arrangement are directly relevant to the argument of this thesis. First, AI companies create new private authorities precisely because they control knowledge infrastructure that states cannot easily replicate. When a government agency depends on a private company to integrate and analyse its security data, the company that designed that platform holds the informational foundation of state decision-making. Second, this relationship transforms state–corporate relations in ways that conventional contractor models cannot capture. States become structurally dependent on a platform as the internal knowledge required to operate independently is progressively eroded (Srivastava, 2023).

Abrahamsen and Williams (2011) make a related argument in the context of private security. They show that private security actors have become so deeply embedded in how states produce security that the boundary between public authority and private function is no

longer clear. To capture this, they develop the concept of *security assemblages*: arrangements in which public and private actors are so intertwined in practice that security outcomes cannot be attributed to either alone (Abrahamsen and Williams, 2011, p. 17). Rather than a state purchasing a service from an outside provider, the two become structurally fused. This is directly applicable to Palantir's position: when a military's targeting process runs through a private platform, the decision about who to strike is no longer purely a state act. It is produced through an arrangement in which the private actor's system is a constitutive part of the outcome.

### **2.2.3 Security Privatisation: From Private Military Companies to Data Infrastructure**

The privatisation of security is not a new phenomenon. The literature on private military companies (PMCs) established that states have long outsourced functions once considered central to sovereign authority, from logistics to direct combat support (Singer, 2009; Krahmann, 2010).

Leander (2005) identified the key mechanism through which PMCs acquire influence beyond their formal contractual role: the control of specialised knowledge that is inaccessible or unverifiable by their clients. Since governments rely on this information without the means to evaluate it independently, the power of the private actor extends well beyond what the contract formally authorises (Leander, 2005). This dynamic helps explain how PMCs acquire *de facto* authority in domains that formally remain under state control.

AI and data infrastructure companies replicate this mechanism, but with a crucial difference. The dependency Leander described as a feature of specialised professional expertise is now built into the architecture of the system itself. It is created by a proprietary system whose logic is legally and technically closed to the client (Bloch-Wehba, 2021). The Carnegie Endowment (2025) documents this directly, showing that Palantir's software becomes embedded in how militaries target and conduct operations, making a private company a structurally essential component of the warfare process itself.

This comparison shows what Palantir shares with PMCs and what makes it different. Like PMCs, Palantir sits in an uncertain position between a service provider and an authority. But unlike PMCs, its influence is not visible. It does not send personnel into the field or provide expertise that a government could, in principle, evaluate. Instead, its power lies in the design of systems that states use but cannot look inside. The existing literature on security privatisation was written with actors like PMCs in mind. These are actors whose role could

be seen and whose knowledge could at least be questioned. It was not written for a company whose influence is built into software that governments depend on but do not control.

#### ***2.2.4 Delegation, Dependency and the Erosion of State Capacity***

A key mechanism through which private authority deepens over time is delegation. Delegation theory in international relations has primarily addressed the transfer of authority to international organisations, but its core logic applies here (Hawkins et al., 2006). When a state delegates decision-relevant functions to an external actor, it typically accepts a degree of informational asymmetry: the delegate knows more about the technical operation of the function than the principal (Hawkins et al., 2006). This asymmetry is manageable when the principal retains the capacity to monitor, evaluate, and if necessary, reclaim the delegated function (Hawkins et al., 2006). The problem that arises with AI infrastructure companies is that this capacity erodes over time.

Gjesvik (2023) shows this process clearly in the context of digital infrastructure. Once government agencies organise their work around a private system, they gradually lose the ability to function without it. The staff who once understood how to do the job independently are replaced by people trained to use the platform. The knowledge of how the system works stays with the company, not the government. Over time, switching to a different system becomes too costly and too disruptive to be a realistic option. What starts as a straightforward contract slowly becomes something the state cannot walk away from. This is the core problem: the state remains legally responsible for the decisions made, but it no longer fully controls the process through which those decisions are reached. That gap between legal responsibility and actual control is what this thesis calls an accountability gap.

The privatisation stream establishes three foundational claims. First, AI companies gain private authority by controlling the knowledge infrastructure states depend on. Second, this authority works through system design, which makes it invisible and hard to challenge. Third, the longer a state relies on such a system, the deeper the dependency grows. What this literature does not provide is a way to assess these dynamics in the context of lethal force. That is where the second stream becomes necessary.

### **2.3 Stream Two: Algorithmic Governance, Opacity, and the Accountability Gap**

#### ***2.3.1 The Algorithmic Turn in Governance***

Algorithms have always played a role in governance. Categories and scoring systems have long been used to manage populations and allocate resources at scale (Bowker and Star 1999). What has changed is the nature of those systems. Earlier decision-support tools were,

in principle, legible: a human analyst could trace a recommendation back through a defined set of rules and explain the output. Machine learning algorithms do not work this way. They identify patterns across vast datasets through processes that cannot be reduced to explicit rules, producing outputs whose internal reasoning cannot be reconstructed even by the engineers who designed them (Pasquale, 2015; Amore, 2020). This shift from legible to opaque decision-making is the central governance problem this stream of literature addresses (Yeung, 2018, p. 505).

Yeung (2018) identifies what she terms *hypernudging* as the specific governance problem created by algorithmic systems: the capacity to shape the choice environment of decision-makers in ways that are pervasive and largely invisible to those being influenced. This is directly applicable to military targeting systems. When a commander receives an AI-generated targeting recommendation, the system has already structured the informational environment within which that decision is made. This determines what data is visible, how threats are categorised, and which targets are flagged. The human formally decides, but the decision space has been defined by the algorithm. Algorithmic regulation, as Yeung (2018) argues, thus represents a qualitatively new form of governance that existing accountability frameworks were not designed to handle.

### **2.3.2 Opacity and the Structural Prevention of Accountability**

The foundational argument on algorithmic opacity is Pasquale's (2015). He contends that opacity is not accidental. Algorithmic systems are protected by secrecy and legal mechanisms that prevent independent scrutiny. Opacity is a deliberate strategy that concentrates authority in the hands of system designers while insulating them from accountability for outcomes (Pasquale, 2015). Knowledge is power, and the capacity to scrutinise others while avoiding scrutiny oneself is among the most consequential forms of power in contemporary governance (Pasquale, 2015).

Amore (2020) takes this further, arguing that opacity is not a flaw to be corrected through greater transparency but the fundamental condition of machine learning systems. Algorithms give incomplete accounts of themselves, learning through interactions with data in ways that exceed their source code. It is therefore insufficient to respond to the injustices of an algorithm by demanding access to the source code, because the reasoning behind any given decision cannot be fully reconstructed after the fact. Responsibility cannot be traced not because of a lack of access, but because the system is designed such that a unified, traceable account of any decision does not exist (Amore, 2020).

Bloch-Wehba (2021) extends this analysis to government-deployed AI specifically. She identifies what she terms intentional opacity: the deliberate structuring of AI systems deployed in public functions in ways that prevent scrutiny even by the agencies using them (Bloch-Wehba, 2021). When a private vendor supplies an AI system under a proprietary contract, the state typically has limited insight into how the system works or how its outputs should be interpreted. Accountability is not merely difficult; it is structurally prevented. Crucially, this applies not only to external observers but to the state itself: the contracting agency may be formally responsible for the system's outputs while being practically unable to assess whether those outputs are accurate or compliant with legal requirements (Bloch-Wehba, 2021).

Taken together, these three accounts establish what this thesis will treat as the opacity problem: the structural condition in which the reasoning behind AI-assisted decisions is unavailable to those who are formally accountable for their consequences.

## **2.4 The Accountability Gap under International Humanitarian Law**

### ***2.4.1 The Foundational Assumptions of IHL***

International Humanitarian Law (IHL) was constructed around a specific model of decision-making in armed conflict. The core principles of distinction, proportionality, and precaution each presuppose a human decision-maker whose reasoning process can be assessed, challenged, and held to account (International Committee of the Red Cross, n.d., Rules 1 and 7). The principle of distinction requires a commander to judge whether an object or person constitutes a legitimate military target. The principle of proportionality requires a commander to weigh anticipated civilian harm against expected military advantage (ICRC, Rules 1, 7 and 14). The principle of precaution requires a commander to take all feasible steps to verify the nature of a target before attack (ICRC, Rule 15). Each of these requirements is a cognitive and deliberative standard: they assume a mind that judges and bears responsibility for the judgement made (Schmitt and Thurnher, 2013).

The doctrine of command responsibility reinforces this structure. Under customary international law and its codification in the Rome Statute (Article 28, 1998), commanders bear legal responsibility for the actions of forces under their effective control, including when they fail to prevent or punish violations of IHL. This doctrine assumes that the commander is in a position of genuine informational and decisional authority: that they know, or should know, what actions are being taken in their name and on the basis of what information (Sassoli, 2019). The entire accountability architecture of IHL, in other words, presupposes a

traceable chain from decision to decision-maker, and a decision-maker who holds the epistemic basis for the decision.

These assumptions were not designed with algorithmic targeting systems in mind. They were designed for human commanders making judgements with information available to them and their staff.

#### ***2.4.2 Where the Framework Breaks Down***

Roff and Moyes (2016) identify the core problem. They argue that meaningful human control, the standard IHL effectively requires, cannot be satisfied by the formal presence of a human in the decision loop if that human does not hold the epistemic basis for the decision. A commander who authorises a strike based on a targeting recommendation generated by an opaque algorithm has not exercised the kind of judgement that proportionality and precaution require (Roff and Moyes, 2016). They have authorised an output whose reasoning they cannot access. The formal act of authorisation is present; the substantive judgement that IHL demands is not (Roff and Moyes, 2016).

Schmitt and Thurnher (2013) extend this analysis in the context of increasingly autonomous weapons systems. They argue that the degree of human control over targeting decisions cannot be assessed in the abstract. It must be evaluated in relation to the specific characteristics of the system in use: its reliability and the degree to which its outputs can be verified by a human operator before action is taken. When a system produces targeting recommendations at the speed and volume documented in recent conflicts, the time available for meaningful human review may be measured in seconds. The Brennan Center (2026) documents this gap in empirical terms. Even basic information about military AI systems such as their performance metrics and compliance with IHL requirements is routinely unavailable to legislative bodies and the public (Toh and Ayoub, 2026). The accountability chain runs through a private actor that bears no formal legal responsibility under IHL. This creates a potential mismatch between formal legal responsibility and technical influence. States and military personnel retain responsibility for complying with international humanitarian law, while private suppliers make design choices that shape the recommendations presented to operational decision-makers (Bruun & Goussac, 2026).

El-Baroudi (2023) identifies a further dimension of this problem: the question of corporate accountability in international law. Current frameworks do not impose direct legal obligations on private companies for the role their systems play in targeting decisions. Corporate liability under international law remains limited and contested, and the contractual

relationship between a private AI company and a state military does not in itself create legal responsibility for the outcomes of that system's recommendations. The result is a distribution of responsibility in which the state is formally liable for decisions it does not fully control, and the company that shapes those decisions bears no formal liability at all (El-Baroudi, 2023).

## **2.5 Theoretical Bridge: Epistemic Dependency and the Accountability Gap**

The two literature streams above highlight related aspects of the governance issue discussed in this thesis. The first shows that private actors can gain influence over security through expertise, even when states hold formal authority (Abrahamsen & Williams, 2011; Leander, 2005). The second addresses how proprietary and AI-based systems complicate transparency and accountability under international humanitarian law, as human operators remain legally responsible for decisions these systems inform (Amoore, 2020; Bloch-Wehba, 2021; Dorsey, 2026; ICRC, 2024; Pasquale, 2015; Roff & Moyes, 2016; Sassòli, 2019; Yeung, 2018). Neither fully explains cases where private tech firms supply AI-enabled military infrastructure to states for target-relevant decision-making.


To connect these two streams, this thesis develops the concept of *epistemic dependency*. The concept of epistemic dependency originally comes from social epistemology, where it refers to situations in which individuals must depend on the knowledge or judgment of others when they cannot assess all relevant evidence themselves (Hardwig, 1985). This thesis reinterprets the idea in a new institutional and political context: the way state military actors depend on privately provided AI infrastructure to create and use knowledge that is important for targeting decisions. In this context, the thesis links epistemic dependency to topics such as private authority, the opacity of algorithms, and issues of accountability within International Humanitarian Law. For the purposes of this thesis, epistemic dependency refers to a possible relationship in which a state's capacity to assess and act upon target-relevant information becomes reliant on technical infrastructure supplied by a private actor, in this case, Palantir Technologies.

The approach does not assume that Palantir independently choose targets, possess military data, or take over state decision-making power. Instead, it focuses on how privately provided infrastructure contributes to integrating, classifying, prioritising, verifying, and transforming military information into actionable proposals.

As Figure 1 illustrates, epistemic dependency is constituted through three interconnected dimensions: algorithmic opacity, private epistemic authority and the

accountability gap. Algorithmic opacity concerns the degree to which proprietary and contractual features of Maven Smart System limit inspection. Private epistemic authority concerns the extent to which Palantir’s infrastructure shapes how military information is organised. The accountability gap concerns whether reliance on such infrastructure complicates the reconstruction and attribution of responsibility under international humanitarian law. The arrows in the figure indicate that these dimensions are mutually reinforcing. Opacity can strengthen private epistemic authority by making system-mediated knowledge harder to challenge; private epistemic authority can influence accountability problems by shaping the informational basis of military decisions; and the accountability gap persists where the relevant knowledge chain cannot be reconstructed after civilian harm.

**Figure 1. Epistemic Dependency**



**Figure 1.** Nodes represent structural features of AI-mediated knowledge environments; arrows indicate directional and bidirectional relationships between concepts.

This framework does not assume in advance that Project Maven creates epistemic dependency. Instead, it provides the analytical criteria through which that possibility can be assessed. A strong form of epistemic dependency would exist if state actors relied on Palantir’s infrastructure for target-relevant knowledge without sufficient capacity to inspect, verify, challenge, or reconstruct its outputs. A weaker form would exist if Palantir helps

organise military information while state actors retain meaningful control over data, verification, and decision-making. If Palantir's role is unrelated to the production of target-relevant knowledge, then epistemic dependency has limited explanatory value.

### **3.0 Methodology**

This chapter sets out the research questions, research philosophy, approach, strategy, and data collection methods that guide this thesis. It also addresses the operationalisation of the central concept and the limitations of the study.

#### **3.1 Research questions**

**Main RQ:** *“To what extent does Palantir Technologies’ role in Project Maven create a condition of epistemic dependency for U.S. military actors, and how does this condition complicate accountability under international humanitarian law?”*

**SRQ1:** To what extent does the procurement architecture of Project Maven complicate state actors’ capacity to inspect and verify the technical systems through which target-relevant knowledge is processed?

**SRQ2:** How does Palantir’s Maven Smart System structure the production, organisation and operationalisation of target-relevant knowledge within U.S. military targeting workflows?

**SRQ3:** To what extent does reliance on Palantir’s Maven Smart System complicate the reconstruction, assessment and attribution of accountability for targeting decisions under international humanitarian law?

#### **3.2 Research design**

This thesis adopts an interpretivist research philosophy. This means that reality is understood as shaped by legal, institutional, and political practices rather than as something fixed and objective (Bryman, 2016). Consequently, knowledge is produced through the interpretation of documents, statements, and frameworks rather than through quantitative measurement or statistical inference. This approach suits the thesis, as the relationship between Palantir and the DoD involves dynamics of power, knowledge, and authority that cannot be measured. They can only be interpreted. The three analytical dimensions used to examine this relationship are algorithmic opacity, private epistemic authority, and the accountability gap. These are interpretive constructs that require analytical judgment rather

than direct observation. Epistemic dependency itself cannot be directly observed either. It must be inferred from the traces it leaves in procurement structures, legal frameworks, and institutional arrangements. One reflexivity consideration needs to be addressed. Because the analysis interprets rather than quantifies a relationship, there is a risk that epistemic dependency is confirmed by the material rather than genuinely assessed against it. To address this, the framework is treated as a candidate analytical lens rather than a predetermined conclusion. The analysis remains open to finding the relationship absent or partial.

The research approach is abductive and theory guided. The thesis moves between existing theoretical debates and the case of Project Maven; it does not start with a fixed hypothesis to be tested. The concept of epistemic dependency is developed through engagement with the literature on private authority, algorithmic governance, and international humanitarian law, and is then used as an analytical lens through which the case is interpreted. This approach is abductive because the thesis does not seek to prove causality in a positivist sense, but to assess whether the relationship between Palantir Technologies and U.S. military actors can be best understood through the concept of epistemic dependency. The framework is therefore applied to evaluate whether three co-present dimensions are observable in the case: algorithmic opacity, private epistemic authority, and the accountability gap. The language of testing is deliberately avoided. Instead, the analysis assesses whether epistemic dependency is analytically useful for explaining how target-relevant knowledge is produced, mediated, and made accountable within Project Maven.

The overall focus of the thesis is to determine whether, and in what form, the role of Palantir Technologies in contemporary AI warfare can be characterised as one of epistemic dependency, using Project Maven as the empirical basis for the analysis.

### **3.3 Research Strategy**

The research strategy is a qualitative, theory-guided, single-case study. Project Maven is selected as the primary case because it is among the most extensively documented and operationally significant instances of a U.S. military AI targeting programme developed in partnership with a private technology company. It is a bounded case in that it concerns a specific programme, a specific institutional context, and a specific legal setting. This boundedness makes it suitable for the kind of close interpretive analysis this thesis employs. The case study is supported by reference to comparable instances, including the use of AI-assisted targeting in Gaza, Ukraine, and Iran. These do not constitute additional cases. They serve as contextualising examples that suggest the dynamics observed in Project Maven may

not be unique to this programme alone. They are treated as illustrative rather than comparative, in keeping with the single case study design, and no claims about generalisability are made on their basis. The time horizon of this study is contemporary. The thesis examines the role of Palantir Technologies as it has developed from 2017 to the present, drawing on the most recent available evidence. It does not trace change over time in a longitudinal sense but characterises that role as it currently stands.

### 3.4 Operationalisation Epistemic Dependency

**Table 1. Operationalisation of Epistemic Dependency**

| <b>Dimension</b>                          | <b>Core meaning</b>                                                                                                                   | <b>Key indicators</b>                                                                                                                                                                                                                                                         | <b>Guiding question</b>                                                                                             |
|-------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| <b><i>Algorithmic opacity</i></b>         | The extent to which Maven's technical and procurement structure limits inspection, verification, or reconstruction of system outputs. | <ul style="list-style-type: none"> <li>– Proprietary architecture</li> <li>– Unclear audit rights</li> <li>– IP &amp; data rights restrictions</li> <li>– Model/platform access</li> <li>– Model testing vs. platform inspectability</li> </ul>                               | Can state actors independently inspect and verify the systems through which target-relevant knowledge is processed? |
| <b><i>Private epistemic authority</i></b> | The extent to which Palantir's infrastructure shapes how military information becomes operationally meaningful.                       | <ul style="list-style-type: none"> <li>– Data integration &amp; classification</li> <li>– Prioritisation &amp; alerts</li> <li>– Intelligence attachments</li> <li>– No-Strike List integration</li> <li>– Target lifecycle management</li> <li>– Workflow routing</li> </ul> | How is target-relevant knowledge produced and operationalised between Palantir's MSS and U.S. military actors?      |
| <b><i>Accountability gap</i></b>          | The extent to which reliance on MSS complicates the reconstruction and attribution of responsibility under IHL.                       | <ul style="list-style-type: none"> <li>– Human-in-the-loop procedures</li> <li>– Decision compression</li> <li>– Verification capacity</li> <li>– Audit trails &amp; target packages</li> <li>– Approval records</li> <li>– Post-strike investigation access</li> </ul>       | Can the chain from detection, verification, approval, and execution be reconstructed after civilian harm?           |

*Note.* Epistemic dependency is assessed by examining whether these three dimensions appear together.

This table operationalises epistemic dependency as a relationship that can be assessed through three connected dimensions. Algorithmic opacity examines whether state actors can inspect and verify the systems through which target-relevant knowledge is processed. Private epistemic authority examines whether Palantir's infrastructure shapes how military information is organised and made actionable within targeting workflows. The accountability gap then considers whether this reliance complicates the reconstruction and attribution of

responsibility under international humanitarian law. Taken together, these dimensions allow the analysis to assess epistemic dependency without assuming in advance that Palantir replaces state authority or directly controls targeting decisions.

### **3.5 Data collection**

Data collection is based on desk research of publicly available sources. The thesis does not involve primary fieldwork, interviews, or access to classified material. The empirical material is drawn from four categories of publicly available sources, which are assigned different evidential weight. First, peer-reviewed academic literature and scholarly books provide the theoretical foundations concerning private authority, epistemic dependency, algorithmic opacity and international humanitarian law. Second, official and institutional documents provide the main empirical basis for the Project Maven analysis. These include U.S. Department of Defence contract announcements and policy documents, reports by the Government Accountability Office and Department of Defence Inspector General, National Academies assessments, Marine Corps guidance, congressional correspondence and ICRC legal materials. Third, Palantir publications, technical documents and investor announcements are used to examine how the company describes MSS, its Ontology and its targeting-related applications. Because these materials present the company's own account, they are assessed alongside independent sources. Fourth, investigative journalism and policy research from organisations such as Reuters, the Associated Press, the Guardian, the Washington Post, the Brennan Center, the Carnegie Endowment and CSIS are used where official information is incomplete, particularly regarding operational developments, civilian-harm allegations and the Minab strike. Claims from journalistic and corporate sources are corroborated where possible and presented cautiously where independent verification remains unavailable.

Data analysis proceeds through qualitative document analysis, specifically an interpretive content analysis guided by the three operationalised dimensions of epistemic dependency. For each dimension, the relevant sources are read for the presence of the observable indicators specified in the operationalisation framework. The goal is not to count indicators but to assess whether the empirical material consistently and credibly supports the characterisation of Palantir's role in Project Maven as exhibiting algorithmic opacity, private epistemic authority, and an accountability gap simultaneously. The analysis is relational rather than causal. It does not claim that one dimension produces another, but that all three co-characterise the role Palantir Technologies plays in contemporary AI warfare.

### 3.6 Limitations

This study faces several interconnected limitations. The most significant is that Project Maven is a classified programme, meaning the internal decision logic of the Maven Smart System is proprietary to Palantir and inaccessible to the researcher. As argued throughout this thesis, however, this inaccessibility is itself a manifestation of the phenomenon under examination rather than merely an external constraint. The study relies entirely on publicly available material, with no interviews or primary fieldwork conducted, meaning the role of Palantir Technologies is necessarily analysed from the outside and claims about how the system is experienced in practice remain inferential. As a single case study, the findings are not generalisable. The references to Gaza, Ukraine, and Iran are illustrative rather than comparative, and the thesis cannot rule out that what it identifies as epistemic dependency is a feature of all complex contractor relationships. A distinction of this kind could only be established empirically through a comparative design. The empirical material reflects the state of the case in mid-2026, a moment in which the academic literature has not yet caught up with operational and legal developments, meaning some claims rest on policy and journalistic sources rather than peer-reviewed scholarship. Finally, the analysis is shaped by the researcher's conceptual choices, including the decision to use epistemic dependency as the analytical frame and to select Project Maven as the case. The thesis does not claim to offer a neutral account but an analytically grounded interpretation through an explicitly stated framework. Transparency about these choices, rather than a claim to objectivity, is the appropriate standard for an interpretivist study of this kind.

## 4. Analysis

### 4.1 Algorithmic Opacity

This chapter addresses the first sub-question:

*To what extent does the procurement architecture of Project Maven complicate state actors' capacity to inspect and verify the technical systems through which target-relevant knowledge is processed?*

As discussed in the literature review and following Pasquale (2015) opacity is understood as a structural condition produced through legal restrictions and unequal access to the internal workings of algorithmic systems. In the context of Project Maven, the primary concern extends beyond the technical explainability of its models. It also encompasses whether the procurement framework establishes epistemic dependency by making the state reliant on

privately developed systems that could potentially hinder governmental oversight. In this section, opacity is therefore treated as a technical and institutional condition: it concerns whether algorithms are explainable and able by the state to inspect and reconstruct the broader infrastructure through which target-relevant knowledge is produced.

As discussed in the background chapter, Project Maven was created to integrate artificial intelligence and machine learning into military intelligence workflows. An important feature of the procurement structure was the recognition of military software as an evolving capability and not a static technical product. According to Palantir's CTO Sankar (2026), Project Maven departed from conventional acquisition approaches in which software purchases resembled those of hardware. Hardware usually undergoes sequential stages: research, development, testing, evaluation, production, and maintenance. Software, on the other hand, needs ongoing updates after deployment. As a result, Colonel Drew Cukor procured Maven's software using Broad Agency Announcements and classified the programme under research and development thereby enabling modifications to the software during development and operational phases (Sankar & Hart, 2026). This procurement architecture points out an important element of Maven's acquisition system: state oversight was required to address systems that could be modified continuously. This shift is relevant to algorithmic opacity because it means that the system subject to oversight was not a fixed artefact, but an evolving software environment. Effective oversight requires recurring evaluation of model updates and changes to operational workflows. Continuous modification does not itself prove epistemic dependency, but it creates a potential dependency mechanism if the knowledge and access required to understand and evaluate those modifications remain concentrated within the private company. The Department of Defence Inspector General's 2022 evaluation provides an example of this problem. The report found that Maven's contracts and cooperative agreements were actively monitored through reporting requirements. However, these practices had not been formally documented, creating risk to programme continuity as Maven expanded and personnel changed (DoD IG, 2022). The Inspector General's finding is relevant to epistemic dependency because it shows that Maven's oversight capacity was partly embedded in undocumented practices (DoD IG, 2022). As the programme expanded and personnel changed, this created a risk that the knowledge required to evaluate contractor performance and maintain consistent scrutiny would be lost or fragmented. This shows that the government's ability to supervise an evolving technical system depended partly on informal knowledge held by personnel. In a continuously changing software environment, weak documentation limits the transfer of

oversight knowledge across officials and over time. This institutional vulnerability can contribute to epistemic dependency because the state may retain formal monitoring authority while lacking a sufficiently durable internal knowledge base to evaluate how the system changes and whether contractor-provided modifications remain operationally acceptable.

Another important issue concerns intellectual property. Sankar and Hart (2026) explain that the Department's general position was that, if the U.S. government pays for research and development, it should own the intellectual property produced through that work. Project Maven complicated this assumption because contractors such as Palantir, Microsoft, and Amazon did not arrive as firms conducting new research from scratch; these companies brought pre-existing commercial products into the programme, developed through substantial prior private investment (Sankar & Hart, 2026). In this sense, the government was not simply funding new R&D but purchasing access to and adaptation of already-developed commercial software infrastructures. Sankar and Hart (2026) emphasise that companies did not own government data and remained constrained by legal restrictions such as the International Traffic in Arms. Yet they also state that firms retained intellectual property over proprietary model weights and core platforms, while the government received rights only to Maven-specific logic configured on top of those systems (Sankar & Hart, 2026). In Palantir's case, this meant that the company retained ownership of its core platform, while the government obtained rights to the mission-specific configurations built on top of it (Sankar & Hart, 2026).

This division of ownership is relevant to algorithmic opacity because opacity can arise through legal and institutional restrictions on access. Where core model weight remains proprietary, government users may be able to operate the system and assess its outputs without possessing equivalent access to the processes through which those outputs are produced. The relationship with epistemic dependency emerges when such opacity prevents the state from developing an independent understanding of the system. The publicly available evidence does not establish that Maven's licence rights gave the government equivalent control over Palantir's proprietary model architecture. The documented allocation of rights therefore demonstrates an asymmetry in which the government-controlled Maven's operational application while Palantir retained the underlying infrastructure and associated technical knowledge. This asymmetry constitutes a link from algorithmic opacity to epistemic dependency. The opacity arose from the legal and institutional separation between the mission-specific logic available to the government and the proprietary platform retained by Palantir. Epistemic dependency followed from the state's reliance on that vendor-

controlled infrastructure to explain the system through which target-relevant data was transformed into operational knowledge.

In addition, oversight reports indicate that Maven's procurement process encountered initial contractual deficiencies, followed by subsequent actions aimed at remedying these issues. According to the U.S. Government Accountability office (GAO, 2026) Maven obtained funding prior to establishing well-defined requirements. The prioritisation of rapid capability development complicated the ability to hold vendors accountable for software deliverables and to define AI related data rights and performance criteria. In response, Maven officials later introduced more explicit contractual provisions addressing data ownership and model performance standards (GAO, 2026). The GAO further notes that Maven officials continued to experience challenges in determining appropriate levels of intellectual property and data rights before contract award, necessitating both internal legal consultation and programme-level coordination (GAO, 2026). These findings imply that procurement opacity was attributable to difficulties in translating AI oversight requirements into enforceable contractual arrangements.

Simultaneously, the evidence does not indicate complete opacity. According to the National Academies of Sciences, Engineering and Medicine (NASEM, 2023), Maven established performance benchmarks for commercial algorithms and evaluated vendor systems using Department of Defence data, including withheld datasets inaccessible to vendors. These assessments determined whether models correctly identified relevant objects and assessed the incidence of false or missed detections (NASEM, 2023). A verification episode during 2019–2020 illustrated this capability: following a decline in performance by an updated commercial computer-vision model, Maven postponed deployment, conducted a thorough investigation, and resumed implementation only after corrections and additional testing (NASEM, 2023). This demonstrates that the Department maintained substantial model-level verification capabilities.

Initial oversight of Project Maven mainly concentrated on individual algorithms and early contract requirements. According to evidence from NASEM and the Department of Defence Inspector General, government officials were able to assess certain commercial algorithms' performance and keep track of early Maven contracts (DoD IG, 2022; NASEM, 2023). As Maven expanded, its operational framework grew beyond computer-vision models. Palantir's Maven Smart System became an integrated platform for data management and command-and-control (U.S. Department of Defence, 2019). It brings together intelligence from sources such as satellite imagery, drone footage and geolocation data. This information is presented

through a common interface that link applications like Foundry, Gaia, Target Workbench, Maverick and LogX, and moves detections into targeting workflows (U.S. Marine Corps, 2025).

This progression highlights the need to differentiate between verifying algorithms at the model level and ensuring inspection at the broader platform level. While testing a single algorithm shows how it performs under set conditions, it does not guarantee that authorities can independently follow data, models, interfaces, and workflow configurations throughout the entire platform. Public records show some capability for testing individual models but offer limited insight into whether the state can examine the full process by which MSS generates actionable operational knowledge (NASEM, 2023; GAO, 2026).

The conclusion drawn here is limited since parts of Maven's technical architecture are classified or withheld from public disclosure. This thesis used what can be established from publicly available procurement, oversight, and test-and-evaluation records. These records do not indicate a complete absence of state oversight. They show that Maven developed meaningful monitoring and test-and-evaluation practices, including DoD-led testing of commercial algorithms on withheld data and later efforts to define model performance benchmarks (NASEM, 2023, pp. 33–34). At the same time, the same records show that these practices evolved in response to earlier institutional gaps. GAO reports that Maven was initially funded without well-defined or documented requirements and that early contracts lacked AI-related requirements needed to hold vendors accountable and more specific terms on data ownership and data rights in later contract actions (GAO, 2026, pp. 24–29). Similarly, the DoD Inspector General found that Maven contracts were actively monitored, but that the reporting procedures used for this monitoring had not been formally documented, creating risks as the programme expanded and personnel changed (DoD IG, 2022).

The available evidence therefore supports that project Maven did not operate outside state monitoring, but its procurement architecture made verification dependent on evolving contractual safeguards, operational testing practices, and negotiated vendor disclosures. This matters because MSS is now described as a broader data-management and battle-management platform (U.S. Marine Corps, 2025). As a result, the relevant verification problem also concerns whether state actors can inspect how different datasets and targeting workflows interact to produce operationally actionable knowledge. Platform-level inspection is therefore more demanding than model-level testing: it requires the capacity to reconstruct how information moved through the system before it appeared to military users as a possible target.

Based on the public record, Project Maven complicates state verification because the Department remains formally responsible for targeting decisions, while its publicly demonstrated access to the full technical architecture of MSS remains limited. The available evidence shows that the Department could test selected models and monitor contract deliverables, but it does not show that state actors could fully inspect or reproduce how the integrated platform organised data, generated outputs, and made information actionable for targeting. The next section examines how MSS structures the production of target-relevant knowledge once information enters the platform itself. This limited inspection creates algorithmic opacity and may produce epistemic dependency where state actors must rely on private expertise to understand or verify system-generated knowledge.

#### **4.2. Epistemic authority**

This section addresses the second sub-question:

*How does Palantir's Maven Smart System structure the production, organisation and operationalisation of target-relevant knowledge within U.S. military targeting workflows?*

Expanding upon the procurement conditions discussed previously, this section examines the distribution of epistemic authority between Palantir and U.S. military personnel within Maven Smart System. In this context, epistemic authority refers to the ability to shape how information is prioritized and made meaningful during the targeting process. While military personnel retain the formal authority to authorize attacks, these decisions are increasingly made through MSS, which structures the gathering and presentation of relevant information. Consequently, Palantir's influence stems from its critical role in organizing the knowledge that supports such operational decisions.

To examine this distribution of epistemic authority more closely, the analysis investigates the targeting workflow of Maven Smart System (MSS). As explained earlier MSS is an AI-driven warfighting platform. It combines intelligence from over 179 diverse sources and uses large language models for decision support. It reduces sensor-to-shooter timelines from hours to minutes (Telasang, 2026). Palantir reframed MSS from a specialized image-analysis tool into a mission command application (MCA) (Telasang, 2026). The foundation of MSS is its ontology—a semantic integration framework that aligns digital assets such as datasets, machine learning models, and virtual tables with their corresponding real-world entities (Telasang, 2026). This means that separate personal records, access log

entries and geospatial ISR coordinates can be automatically associated with the same entity. This ontology-based grouping allows analysts to perform analysis across both classified and unclassified data via a unified interface (Telasang, 2026). This process helps generate knowledge relevant to specific goals by having MSS convert raw data into organized ontology objects and relationships, making them accessible for analytical workflows. In a NATO MSS exercise described by Palantir, satellite images and AI-generated detections were transformed into supporting datasets and categorized as objects like “Satellite Image” and “Detection,” which then became essential for analyst applications and AI reasoning (Palantir, 2026). Palantir’s technical materials also note that their ontology organizes existing data sources into objects, properties, and links that represent real-world items and define the connections between them (Palantir Technologies). Consequently, MSS does more than just pull individual records; it structures information so that detections, images, locations, and other data points are presented as interconnected objects for analysis (Palantir Technologies).

From an analytical standpoint, this architecture affects how analysts interpret information. Yeung (2017) notes that data-driven decision-support systems generate knowledge by connecting data points and emphasizing certain correlations within the user's context. For MSS, its ontology determines how records are grouped, which connections are made clear, and what information can be searched about a potential target. This process is an example of epistemic mediation, because the meaning given to a detection relies in part on how MSS presents relationships. Public documents provide a general understanding of the platform’s objects, links, and processes, but do not explain entity-matching and association rules in detail. As a result, evidence shows that MSS influences the information used for military assessments, however the technical specifics behind some associations remain unknown to the public.

Palantir’s role is assessed through the way MSS and Target Workbench manage target-related data, present detections, organise priorities, and channel information that could lead to military action. The available evidence indicates that Palantir operates primarily at the platform and workflow level (Chabrier-Montijo, 2024). MSS serves as an interface through which detections become operationally usable through Broad Area Surveillance–Targeting workflows. In these workflows, detections from imagery are published in real time to MSS, where military intelligence analysts can access them (Chabrier-Montijo, 2024). Public evidence shows that Palantir does not set military targeting standards. Military personnel use commander-approved tools, such as a High Payoff Target List and Attack Guidance Matrix, to determine target importance and procedures (Chabrier-Montijo, 2024). Authority remains

with military decision-makers. Palantir's UK and Europe lead, Louis Mosley, confirms that state actors decide how AI targeting tools are used, while Palantir only supplies the technology (Cieslak & Murphy, 2026).

The distinction between formal authority and epistemic mediation is an important consideration in this context. While military personnel are responsible for defining the criteria used to validate targets, MSS's infrastructure influences how detections are compiled, evaluated, correlated with supporting intelligence, and presented for subsequent action. Chabrier-Montijo (2024) notes that detections are transmitted to Target Workbench within MSS, where they undergo prioritisation and are routed for possible prosecution via systems such as the Advanced Field Artillery Tactical Data System. According to Palantir's Target Workbench documentation, the platform serves as a comprehensive software environment for managing targets throughout their lifecycle, encompassing identification, prioritisation, approval, execution, and assessment (Palantir, 2023, pp. 1–3). In addition, Palantir indicates that users can customise workflow stages, specify approval requirements, generate alerts, attach intelligence data, and oversee target management and prioritisation using the High Payoff Target List infrastructure (Palantir, 2023, pp. 1–3). Additionally, the platform can display warnings regarding the currency of intelligence information and incorporating No-Strike List data (Palantir, 2023, p. 3).

Official Marine Corps sources reinforce the platform-level interpretation of MSS. The Marine Corps defines MSS as a data-centric command-and-control system that consolidates information into a real-time, synchronised battlespace view, facilitating AI-enabled target-management workflows (U.S. Marine Corps, 2025a). An additional implementation document identifies Target Workbench as an integral component of the MSS platform and details MSS's support for sensing, engagement, and battle-management capabilities across various command levels (U.S. Marine Corps, 2025b). Collectively, these documents indicate that Palantir does not determine lawful military objectives or approve attacks. However, Palantir's platform is instrumental in structuring the technical environment, thereby enabling target-relevant knowledge to become operationally actionable. Through MSS and Target Workbench, detections can be linked with intelligence and warnings and directed towards potential engagement. At this stage, the platform extends beyond presenting intelligence; it seamlessly integrates data analysis and target evaluation with subsequent operational systems (Telasang, 2026). This comprehensive workflow enables a significant reduction in the response time between detection and potential military action—from hours or days to just minutes (Telasang, 2026).

The Minab school strike illustrates why this distinction matters in practice. On the first day of the U.S. and Israel's military operations against Iran, the 28th of February, Shajareh Tayyebbeh elementary school in Minab was reportedly struck twice. The catastrophe killed between 175 and 180 people, most of them girls under the age of twelve (Livingstone, 2026). The school was located less than 100 yards from a long-standing Islamic Revolutionary Guard Corps naval installation and had previously been within the perimeter of that compound. However, it had operated as a school for several years before the strike and had public indicators of civilian use, including social media activity and its own website (Livingstone, 2026). Baker (2026) adds that the target package presented the site as a military facility, even though the school appeared in Iranian business listings and was visible on Google Maps. Heise (2026) similarly reports that the building was allegedly still listed as a military object in a U.S. database and that Maven was said to have adopted this incorrect classification into an automated decision-making process without sufficient verification. The problem, according to Baker, was that an outdated classification remained embedded in the target package and appears to have been carried forward through an AI-enabled targeting process, while contradictory civilian indicators outside the formal package were not incorporated (Baker, 2026). The Minab case serves as a stress test for knowledge production. Reporting does not confirm if MSS was used in the strike, but it highlights a key failure: rapid processing cannot fix outdated classifications when conflicting information is excluded from the target package. Even efficient platforms can perpetuate errors if they are embedded in the base data. It is critical to differentiate between processing and validating information, especially when target packages are quickly moving toward operational use.

The analysis gives a partial answer to the second sub-question. MSS uses Palantir's ontology to turn scattered data into connected knowledge about targets. Target Workbench organises this knowledge by sorting targets into priority categories, review stages, warnings, and approval steps. MSS links these target packages with engagement systems. Military personnel set targeting standards, validate targets, and approve actions. Palantir indirectly influences how information is moved through the workflow. When military personnel depend on MSS to carry out important tasks but cannot verify or replicate how the integrated platform works, this creates epistemic dependency. In such cases, the state remains officially responsible for targeting decisions but relies heavily on private systems to process much of the information needed to make those decisions. The next section looks at how this dependency makes it harder to determine and assign accountability under international humanitarian law.

### 4.3 Accountability gap

This section examines the third sub-question:

*How does reliance on Palantir's Maven Smart System (MSS) affect the assessment and attribution of accountability for targeting decisions under international humanitarian law?*

As discussed in the literature review, the accountability gap considers whether dependence on privately sourced AI infrastructure complicates the evaluation and assignment of responsibility under international humanitarian law, particularly when incidents involving civilian harm necessitate reconstruction of the informational processes that preceded an attack.

The discussion of the first sub-question established that Project Maven is not entirely opaque, since the Department of Defence has developed some ability to test vendor algorithms and oversee initial contract deliverables. However, it also highlighted ongoing uncertainty regarding the overall inspection of MSS as an integrated platform. Addressing the second sub-question revealed that Palantir does not formally define military objectives or authorize attacks, but that both MSS and Target Workbench shape the technical environment where detections are organized and processed within targeting workflows. The question of accountability stems directly from these prior findings. When target-relevant knowledge is generated through a privately developed and only partly inspectable system, responsibility for civilian harm cannot be determined solely by identifying the human commander who authorized the strike. Instead, it requires reconstructing how actionable targets were produced within the software-mediated targeting process. This section explains that MSS complicates accountability in two main ways. First, it speeds up the process and changes the conditions needed for thorough human verification. If operators do not have enough time to review target information, simply having formal human approval might not fulfil the deeper requirements of Article 57 from Additional Protocol I (Geneva Conventions, 1977). Second, post-strike accountability becomes dependent on access to the privately structured targeting workflow through which the target was produced.

The first accountability problem concerns the operational tempo associated with AI-enabled targeting. This development is part of the “warification” of international humanitarian law, referring to a process through which technologies and activities not traditionally associated with armed conflict become normalised as components of warfare

(Dorsey & Moffett, 2026). Systems such as Palantir's Maven Smart System are framed through the language of precision and efficiency. The concept of "warification" challenges the belief that increasingly precise technologies inherently lead to more discriminate forms of warfare. In addition, the rapid progression of the kill-chain creates accountability problems (Dorsey & Moffett, 2026). The term for this is decision compression, which describes the streamlining of multi-stage military planning and targeting into AI-supported outputs (Tsofniasvili, 2026). Decision compression moves targeting from a slower, deliberative process into a high-speed workflow. This situation results in a direct conflict with Article 57 of Additional Protocol I (Geneva Conventions, 1977), which obliges those responsible for planning or authorising an attack to undertake all feasible measures to confirm that the target qualifies as a military objective. Proper verification necessitates adequate time and information to assess available evidence. If the targeting process outpaces the ability of operators to fulfil this, the involvement of a human reviewer may maintain procedural approval yet diminish the depth of legal scrutiny. Consequently, decision compression poses an international humanitarian law issue by eliminating the practical opportunity to identify and rectify mistakes prior to engagement (ICRC, 2024, pp. 64–67; Tsofniasvili, 2026).

Palantir positions MSS as a tool to mitigate operational risks. According to the company, AI-enabled systems have the potential to enhance the accuracy and precision of military decision-making, while also minimising targeting errors and supporting compliance with international humanitarian law (Palantir Technologies, 2023). Additionally, Palantir highlights that its software can integrate civilian harm data, No-Strike Lists, collateral-concern assessments, user annotations, and audit trails into existing operational workflows (Palantir, 2024). This capability provides an important perspective: increased targeting speed may enhance protection if it supplies decision-makers with information that is more comprehensive and timelier than fragmented human-led processes. Nevertheless, the validity of Palantir's safety claims relies on the use of accurate data and effective human oversight for assessment and challenge of outputs generated by the system. Decision compression has the potential to undermine these essential conditions.

These problems become concrete in the case of Maven's use in Iran. The Center for Strategic and International Studies (CSIS) reports that, during the first 24 hours of the war in Iran, the United States used MSS to help strike more than 1,000 targets. This is a tenfold increase over previous targeting capacity (Mande & Allen, 2026). Maven rapidly ingested drone and satellite data and flagged potential targets and scored them by priority, even assessing likelihood of collateral harm (Copp et al., 2026). During a twelve-hour period

beginning on March 1, U.S. and Israeli forces launched roughly 900 airstrikes across Iran. These strikes hit regime targets including military bases, missile sites and even Iran's Supreme Leader. U.S. sources revealed that hundreds of targets had been pre-processed by AI (Copp et al., 2026). Palantir's system supposedly proposed hundreds of targets, prioritized them and provided coordinates essentially automating work that would previously have required weeks of staff activity. After that, human commanders executed many of the AI's suggested strikes in rapid succession. However, the speed and scale of target production increased their epistemic dependence on information already prioritised by MSS, while reducing the time available to independently verify the underlying evidence. The Minab school strike illustrates the consequences that can follow when an accelerated targeting process fails to identify and correct inaccurate target information. The attack on Shajareh Tayyebbeh elementary school in Minab shows how decision compression can hinder effective verification. The evidence does not confirm that MSS chose the school autonomously or that there was no human oversight. Rather, preliminary findings indicate that U.S. Central Command relied on target coordinates derived from outdated Defence Intelligence Agency data (Madhani et al., 2026). Preliminary findings show that the targeting decision was based on outdated coordinates and intelligence, which still regarded the site as a military facility despite changes to its physical appearance and public identity (Madhani et al., 2026; Reuters, 2026). Satellite images available to the public revealed that by around 2017, a wall had separated the school from the adjacent military compound and an old watchtower had been dismantled (Livingstone, 2026). Overhead views showed colourful murals and outdoor play areas and online maps listed the building as a school, and its website shared information about its students, teachers, and activities (Madhani et al., 2026; Amnesty International, 2026). Nevertheless, the site continued to be treated as a military object, implying that conflicting evidence did not alter how targeting staff classified it. In a fast-paced targeting workflow, this rigidity leads to greater epistemic dependence, as operators depend on the established representation of the location. If the process moved too quickly for personnel to spot or resolve these contradictions, formal human approval would provide only weak assurance that the verification required by Article 57 of Additional Protocol I was carried out. The Minab case therefore demonstrates how compressed decision-making can turn an outdated classification into a live target when the underlying information is insufficiently examined.

As a comparison, the same accountability problem appears in reporting on Israel's use of Lavender in Gaza. Investigations report that AI system "Lavender" processed large

volumes of surveillance data and identified up to 37,000 Palestinians as potential targets (McKernan & Davies, 2024). One Israeli intelligence officer told in an interview that they spent around twenty seconds on each target and had “zero added-value as a human” apart from acting as a stamp of approval (McKernan & Davies, 2024). This example supports the broader point that decision compression can turn human oversight into procedural confirmation, which complicates later assessment if verification under IHL was meaningful.

The second accountability issue concerns traceability following incidents of civilian harm. Palantir describes its platforms as enabling lawful and responsible targeting by integrating safeguards against civilian harm, collateral-concern checks, No-Strike List data and logging into military decision-making processes (Palantir, 2024). This assertion is significant, as Palantir’s justification for MSS extends beyond mere speed or efficiency; the company contends that software-mediated targeting can enhance accountability by increasing operators’ access to relevant information and providing a comprehensive record of decision-making processes (Palantir, 2024). This approach is consistent with DoDI 3000.17, which mandates the integration of civilian harm mitigation and response measures into planning, targeting, assessment, and investigative procedures (DoD, 2023, pp. 5, 25). In theory, if MSS’s safeguards and audit functions enable investigators to reconstruct the identification, review, approval, and execution of targets, the accountability gap should be reduced. The Minab strike serves as a case study to evaluate the practical effectiveness of this model.

The aftermath of Minab exposes the limits of this model. Following the strike, Congress formally requested clarification from Secretary of Defence Pete Hegseth regarding the role of artificial intelligence in Operation Epic Fury. A congressional letter led by Representative Jason Crow and endorsed by 120 Democratic members identified the Minab school strike as a catalyst for concerns about AI-enabled targeting, civilian harm mitigation and accountability protocols (Crow et al., 2026). The letter specifically asked whether Maven Smart System had been used to identify the school as a target and whether human verification had been conducted regarding the target’s accuracy (Crow et al., 2026). It also linked the strike to weakened civilian-protection institutions, noting that budget cuts and staff changes had affected the Civilian Protection Centre of Excellence and reduced the number of personnel reviewing targeting decisions (Crow et al., 2026). Since DoDI 3000.17 depends on institutional capacity, civilian-harm safeguards only matter if there are personnel, procedures and review structures capable of applying them in practice. Reuters reports that the U.S. military heightened its investigation after media coverage suggested possible U.S. involvement, noting that confirmation would make the strike one of the most severe civilian-

casualty incidents attributed to U.S. military activity in the Middle East in recent decades (Reuters, 2026). Yet by June 2026, no substantial public response appears to have answered the congressional questions concerning MSS's role in the strike. Hegseth's response remained procedural: he stated that a general officer from outside CENTCOM would investigate, but did not clarify whether MSS was involved, how the school became a target, whether human operators confirmed its civilian status, or whether the platform's audit mechanisms could reconstruct the targeting chain (Reuters, 2026). This absence of detail is important because Palantir's own accountability claim depends on traceability and reviewability. Investigative reporting has partially filled the public record by documenting the persistence of outdated target information despite publicly available evidence of the school's civilian function (Livingstone, 2026). Baker adds that the target package presented the site as a military facility, even though the school appeared in Iranian business listings and was visible on Google Maps (Baker, 2026).

The Minab case highlights the vulnerability of an accountability framework that relies solely on technical traceability without accompanying public transparency. While Palantir's platforms are presented as tools for enhancing the lawfulness, reviewability, and accountability of targeting operations, the lack of a clear public explanation following a mass-casualty civilian event is important. This indicates that accountability enabled by software remains contingent upon the willingness or obligation of state and corporate entities to disclose the specifics of how the targeting process was executed. Decision compression during the war in Iran also intensified epistemic dependency. With MSS supporting strikes against more than 1,000 targets during the first 24 hours. Military personnel became increasingly reliant on target information already selected and prioritised by the platform. This scale raises doubt about whether each target could be independently examined and challenged before authorisation. Where such compressed review prevents meaningful verification, formal human approval provides limited assurance of compliance with Article 57 of Additional Protocol I (Geneva Conventions, 1977).

## **5. Discussion**

Taken together, the findings demonstrate a condition of epistemic dependency. The first part of this thesis examined algorithmic opacity, highlighting a distinction between what the Department of Defence (DoD) can test directly and what it can independently inspect. Maven officials conducted tests on select commercial models using withheld military data

and paused deployment if performance dropped. Palantir, however, kept ownership of its core platform and proprietary model weights, with the DoD only gaining rights to Maven-specific configurations. Initial contract requirements related to AI performance were incomplete, and as the program grew, some monitoring procedures went undocumented. As a result, public records show that model-level verification took place but provide limited evidence that the government can fully replicate how data functions across the MSS integrated platform.

Private epistemic authority is evident in how the platform operates. MSS collects data from over 179 sources and Palantir's ontology connects imagery, geolocation information, detection events, and various records into unified objects and relationships. Target Workbench advances these targeting packages through steps such as identification, prioritisation, approval, execution, and assessment, adding alerts, intelligence files, and No-Strike List details along the way. Although military personnel maintain control over targeting guidelines and authorisation of attacks, they do so within a system whose organisation and workflow are privately engineered.

The accountability gap results from this dependency. During operations in Iran, the MSS reportedly supported over 1,000 targets within a 24-hour period, thereby increasing reliance on information that had already been selected and prioritized by the platform. The Minab case illustrates the ramifications when outdated target data remains actionable despite clear civilian indicators. In the aftermath of harm, identifying the source of such errors depends on the availability of target packages, approval documentation, and system logs. Consequently, epistemic dependency emerges through the interplay of limited platform transparency, Palantir's influence in structuring targeting knowledge, and the use of the same infrastructure for both verifying and reconstructing targeting decisions.

The findings of this thesis point to a broader ethical problem in the future of AI-enabled warfare. Military judgement increasingly depends on technical systems that present target-relevant knowledge before a human decision is made. This shifts the ethical debate away from the familiar image of fully autonomous weapons acting without human involvement. The more immediate concern is that humans may remain formally responsible while the information behind their decisions is shaped by systems they cannot fully inspect. This matters for civilian protection. International humanitarian law depends on the ability of commanders to distinguish between civilian and military objects, assess proportionality, and take feasible precautions before an attack (Sassòli, 2019; Schmitt & Thurnher, 2013). These duties require reliable information, sufficient time for verification, and the practical ability to question the basis on which a target has been presented. The analysis of Project Maven

suggests that AI-enabled systems may support military decision-making by processing large amounts of data, but they may also create new forms of dependence. If a platform organises the information through which a target becomes visible, then civilian protection depends on the quality, currency, and contestability of that information.

The Minab case illustrates the ethical stakes of this problem. It shows why accountability after civilian harm requires reconstruction of the informational chain that preceded the attack. If a school remained embedded in a target package, the ethical failure would not be limited to the final authorisation of force. It would also concern the earlier process through which the school was made actionable. This is where epistemic dependency becomes ethically important. The risk is that harm can emerge from the organisation of knowledge before it appears as a legal decision. The future of human control in warfare must therefore be assessed more carefully. The phrase “human-in-the-loop” can create a false sense of reassurance. A human approval process has limited ethical value if the reviewer lacks the time, information, or authority to challenge the output placed before them. AI-enabled targeting systems are often justified through speed and scale, but those same qualities can compress deliberation. As targeting cycles accelerate, human judgement may become a procedural checkpoint rather than a meaningful act of verification. This problem is likely to become more urgent because recent conflicts have accelerated the military use of AI. In Ukraine, AI-enabled drone systems have been developed to support target identification and autonomous navigation. In Gaza, investigations have reported the use of AI-supported systems to identify large numbers of potential targets (McKernan & Davies, 2024; Reuters, 2024a, 2024b). These examples are different from Project Maven, but they indicate a broader trajectory. AI is becoming part of the ordinary infrastructure of warfare. The important question in this debate is what forms of accountability, restraint, and civilian protection can survive when warfare becomes faster, more data-driven, and more dependent on private technical systems.

The arms race logic intensifies this problem. States increasingly frame military AI as a strategic necessity: if adversaries develop faster targeting systems, democratic states claim they must do the same. This creates pressure to prioritise speed, scale, and operational advantage over transparency and restraint. In that environment, ethical limits can start to appear as military weaknesses. Project Maven should be understood within this wider political context. It is not only a technical programme, but part of a broader transformation in which military power depends on private technological capacity. The danger is that the race to remain ahead technologically may reduce the space for democratic debate about what

kinds of warfare should be permitted. This concern becomes sharper in a shifting world order. AI-enabled warfare develops at a moment of geopolitical rivalry, weakening confidence in multilateral restraint, and increasingly personalised forms of political leadership. The risks are especially serious when leaders use extreme rhetoric about mass violence. Trump's statement that "a whole civilization will die tonight" demonstrates the danger of combining aggressive political language with military systems designed to accelerate the use of force (Reuters, 2026). This raises a democratic accountability problem. In democratic systems, the use of force is supposed to be constrained by law. Yet, as analysed in this thesis, AI-enabled targeting systems often operate through classified programmes, proprietary platforms, and public-private partnerships that are difficult for citizens, courts, and parliaments to inspect. The state remains the actor formally responsible under international law, but the knowledge infrastructure behind its decisions may be built and maintained by private companies. This creates a gap between public responsibility and private technical power. If democratic publics cannot understand how targets are produced, they cannot meaningfully evaluate whether force is being used lawfully.

Private companies therefore cannot be treated as neutral suppliers. Palantir's role in Project Maven does not make it a battlefield commander, but it does make the company part of the infrastructure through which military knowledge is produced. That position carries ethical responsibility. Companies that design systems for targeting-related workflows should have obligations that extend beyond contract performance. They should be expected to conduct human rights due diligence, support meaningful human review and cooperate with accountability processes when civilian harm occurs. Without such obligations, private firms can shape the conditions of military judgement while leaving formal responsibility entirely with the state.

The contrast with other technology companies shows why voluntary corporate ethics are insufficient. Google's withdrawal from Project Maven after employee protests demonstrated that internal resistance could place limits on corporate participation in military AI (Griffith, 2018). However, later developments also show that corporate commitments can shift under political, commercial, and national-security pressure. Anthropic presents itself as a safety-oriented AI company, yet it has partnered with Palantir and Amazon Web Services to bring Claude models to U.S. intelligence and defence operations and later entered into a Department of Defence agreement framed around responsible AI capabilities (Anthropic, 2025; Palantir Technologies, 2024). This shows that voluntary principles alone cannot

provide a stable basis for accountability in a sector shaped by strategic competition and state demand.

The future of accountability in AI-enabled warfare therefore requires stronger institutional standards. First, states should be able to demonstrate how AI-supported targets are verified before attack, including how outdated or contradictory information is flagged. Second, human reviewers should have enough time, information, and authority to challenge system outputs. Third, platforms used in targeting-related workflows should preserve audit trails that allow investigators to reconstruct how a target was identified, prioritised, reviewed, and approved. Fourth, private companies should face clearer duties to cooperate with civilian harm investigations where their systems helped structure the relevant targeting process.

## **6. Conclusion**

This thesis asked to what extent Palantir Technologies' role in Project Maven creates epistemic dependency for U.S. military actors and how this condition complicates accountability under international humanitarian law. It concludes that Project Maven creates a qualified form of epistemic dependency. U.S. military actors retain control over targeting standards, validation, and strike authorisation, but rely increasingly on Palantir's infrastructure to organise, prioritise, and reconstruct target-relevant knowledge. This finding bridges the two bodies of literature examined in the thesis. The literature on private authority explains how corporations acquire influence through specialised knowledge and infrastructure, while scholarship on algorithmic governance shows how proprietary systems can restrict scrutiny and shape decision-making environments. Project Maven demonstrates how these dynamics intersect in military targeting. Private authority can expand through the organisation of knowledge even where the state retains formal legal authority.

The strongest counterargument is that MSS may improve targeting by integrating fragmented intelligence, reducing human error, incorporating civilian-protection safeguards, and preserving a clearer record of decisions. The analysis accepts that these benefits are possible and that the Department of Defence retains meaningful model-testing capacity. However, these safeguards do not eliminate epistemic dependency where military actors cannot independently inspect the integrated platform or where decision compression limits their ability to challenge the information presented. This condition complicates accountability under IHL because the state remains responsible for targeting decisions while the knowledge

required to verify and later explain those decisions is partly organised through private infrastructure. Where review becomes too compressed, formal human approval provides limited evidence that Article 57 verification was meaningfully performed. Where post-strike reconstruction depends on MSS records, accountability also depends on access to the platform through which the target was produced.

The broader implication of this thesis is that accountability in AI-enabled warfare begins before the final decision to strike. It begins in the systems that collect, classify, prioritise, and present military knowledge. Project Maven shows how public authority and private infrastructure can become intertwined: the state remains legally responsible for force, while private systems help shape the information through which force becomes possible. As military AI becomes more central to warfare, civilian protection will depend on the ability to inspect how targets are produced, how errors are identified, and how harm is explained after the fact.

## Bibliography

- Abraham, Y. (2024, April 3). 'Lavender': The AI machine directing Israel's bombing spree in Gaza. *+972 Magazine*. <https://www.972mag.com/lavender-ai-israeli-army-gaza/>
- Abrahamsen, R., & Williams, M. C. (2011). *Security beyond the state: Private security in international politics*. Cambridge University Press.
- Ali, I., & Stewart, P. (2026, March 11). US may have struck Iranian girls' school after using outdated targeting data, sources say. *Reuters*. <https://www.reuters.com/world/middle-east/us-may-have-struck-iranian-girls-school-after-using-outdated-targeting-data-2026-03-11/>
- Amoore, L. (2020). *Cloud ethics: Algorithms and the attributes of ourselves and others*. Duke University Press.
- Amnesty International. (2026, March 16). *USA/Iran: Those responsible for deadly and unlawful U.S. strike on school that killed over 100 children must be held accountable*. <https://www.amnesty.org/en/latest/news/2026/03/usa-iran-those-responsible-for-deadly-and-unlawful-us-strike-on-school-that-killed-over-100-children-must-be-held-accountable/>
- Anthropic. (2025, July 14). *Anthropic and the Department of Defense to advance responsible AI in defense operations*. <https://www.anthropic.com/news/anthropic-and-the-department-of-defense-to-advance-responsible-ai-in-defense-operations>
- Avant, D. D., Finnemore, M., & Sell, S. K. (Eds.). (2010). *Who governs the globe?* Cambridge University Press.
- Baker, K. T. (2026a, March 21). *Kill chain*. Artificial Bureaucracy. <https://artificialbureaucracy.substack.com/p/kill-chain>
- Baker, K. T. (2026b, March 26). AI got the blame for the Iran school bombing. The truth is far more worrying. *The Guardian*.

<https://www.theguardian.com/news/2026/mar/26/ai-got-the-blame-for-the-iran-school-bombing-the-truth-is-far-more-worrying>

Bienvenue, E., Kelton, M., Rogers, Z., Sullivan, M., & Ford, M. (2025, December 1). *Private tech companies, the state, and the new character of war*. Carnegie Endowment for International Peace. <https://carnegieendowment.org/research/2025/12/ukraine-war-tech-companies>

Bloch-Wehba, H. (2021, June 17). *Transparency's AI problem*. Knight First Amendment Institute at Columbia University. <https://knightcolumbia.org/content/transparencys-ai-problem>

Boutin, B. (2023). State responsibility in relation to military applications of artificial intelligence. *Leiden Journal of International Law*, 36(1), 133–150.  
<https://doi.org/10.1017/S0922156522000607>

Bowker, G. C., & Star, S. L. (1999). *Sorting things out: Classification and its consequences*. MIT Press.

Bruun, L., & Goussac, N. (2026, May 26). *Collaboration without over-reliance: The role of industry in making military AI 'lawful by design'*. ICRC Humanitarian Law & Policy Blog. <https://blogs.icrc.org/law-and-policy/2026/05/26/collaboration-without-over-reliance-the-role-of-industry-in-making-military-ai-lawful-by-design/>

Bryman, A. (2016). *Social research methods* (5th ed.). Oxford University Press.

Chabrier-Montijo, C. A. (2024, March 13). *XVIII Airborne Corps BAS-T employment*. Line of Departure. <https://www.lineofdeparture.army.mil/Journals/Field-Artillery/FA-2024-Issue-1/Airborne-Employment/>

Chiang, O. (2011, February 23). Super crunchers. *Forbes*.  
<https://www.forbes.com/forbes/2011/0314/technology-facebook-palantir-thinkprogress-super-crunchers.html>

- Cieslak, M., & Murphy, M. (2026, April 1). Palantir UK boss says it's up to militaries to decide how AI targeting is used in war. *BBC News*.  
<https://www.bbc.com/news/articles/cdrm52g4pl2o>
- Copp, T., Dwoskin, E., & Duncan, I. (2026, March 4). *Anthropic's AI tool Claude central to U.S. campaign in Iran, amid a bitter feud*. The Washington Post. <https://www.washingtonpost.com/technology/2026/03/04/anthropic-ai-iran-campaign/>
- Crow, J., & Van Hollen, C. (2026, March 11). *Letter to Secretary of Defense Pete Hegseth on the Minab bombing and civilian casualties in Iran*.  
<https://www.vanhollen.senate.gov/download/letter-to-hegseth-on-minab-bombing-civcas-iran>
- Cutler, A. C., Haufler, V., & Porter, T. (Eds.). (1999). *Private authority and international affairs*. State University of New York Press.
- Department of Defense Office of Inspector General. (2022). *Evaluation of contract monitoring and management for Project Maven* (Report No. DODIG-2022-049). U.S. Department of Defense.
- Dorsey, J., & Bo, M. (2026). AI-enabled decision-support systems in the joint targeting cycle: Legal challenges, risks, and the human(e) dimension. *International Law Studies*, 107(1), 184–227. <https://digital-commons.usnwc.edu/ils/vol107/iss1/7/>
- Dorsey, J., & Moffett, L. (2026). The warification of international humanitarian law and the artifice of artificial intelligence in decision-support systems: Restoring balance through the legitimacy of military operations. In H. Krieger, P. Kalmanovitz, E. Lieblich, & J. P. Sexton (Eds.), *Yearbook of international humanitarian law: Volume 27, 2024: International humanitarian law under pressure* (pp. 125–156). Springer.

- Dumbill, E. (2013). Making sense of big data. *Big Data*, 1(1), 1–2.  
<https://doi.org/10.1089/big.2012.1503>
- Frankel, J., & Biesecker, M. (2026, March 6). *Evidence suggests the deadly blast at an Iranian school was likely a U.S. airstrike*. Associated Press. <https://apnews.com/article/c3095dc9729881b567277a1c5c47efb2>
- Gjesvik, L. (2023). Private infrastructure in weaponized interdependence. *Review of International Political Economy*, 30(2), 722–746.  
<https://doi.org/10.1080/09692290.2022.2069145>
- Gorman, S. (2009, September 4). How team of geeks cracked spy trade. *The Wall Street Journal*. <https://www.wsj.com/articles/SB125200842406984303>
- Griffith, E. (2018, June 1). Google won't renew controversial Pentagon AI project. *Wired*.  
<https://www.wired.com/story/google-wont-renew-controversial-pentagon-ai-project/>
- Hall, R. B., & Biersteker, T. J. (Eds.). (2002). *The emergence of private authority in global governance*. Cambridge University Press.
- Hardwig, J. (1985). Epistemic dependence. *The Journal of Philosophy*, 82(7), 335–349.  
<https://doi.org/10.2307/2026523>
- Harper, J. (2025, September 12). Marine Corps releases new guidance on Maven Smart System integration. *DefenseScoop*. <https://defensescoop.com/2025/09/12/marine-corps-maradmin-maven-smart-system-mss-palantir-rollout/>
- Hawkins, D. G., Lake, D. A., Nielson, D. L., & Tierney, M. J. (Eds.). (2006). *Delegation and agency in international organizations*. Cambridge University Press.
- Heise Online. (2026, March 30). *Alleged US attack on school in Iran: Palantir system in focus*. <https://www.heise.de/en/news/Alleged-US-attack-on-school-in-Iran-Palantir-system-in-focus-11229159.html>

Hunder, M. (2024, July 18). Ukraine rushes to create AI-enabled war drones. *Reuters*.

<https://www.reuters.com/technology/artificial-intelligence/ukraine-rushes-create-ai-enabled-war-drones-2024-07-18/>

Iliadis, A., & Acker, A. (2022). The politics of metadata: Using ontologies to make data

interoperable. *Big Data & Society*, 9(1). <https://doi.org/10.1177/20539517221097861>

International Committee of the Red Cross. (n.d.-a). *Practice relating to Rule 15: Principle of*

*precautions in attack*. Customary IHL Database. Retrieved June 21, 2026, from

<https://ihl-databases.icrc.org/en/customary-ihl/v2/rule15>

International Committee of the Red Cross. (n.d.-b). *The principle of distinction between*

*civilian objects and military objectives (Rule 7)*. Customary IHL Database. Retrieved

June 21, 2026, from <https://ihl-databases.icrc.org/en/customary-ihl/v1/rule7>

International Committee of the Red Cross. (n.d.-c). *The principle of distinction between*

*civilians and combatants (Rule 1)*. Customary IHL Database. Retrieved June 21,

2026, from <https://ihl-databases.icrc.org/en/customary-ihl/v1/rule1>

International Committee of the Red Cross. (n.d.-d). *Proportionality in attack (Rule 14)*.

Customary IHL Database. Retrieved June 21, 2026, from [https://ihl-](https://ihl-databases.icrc.org/en/customary-ihl/v1/rule14)

[databases.icrc.org/en/customary-ihl/v1/rule14](https://ihl-databases.icrc.org/en/customary-ihl/v1/rule14)

International Committee of the Red Cross. (n.d.-e). *Target verification (Rule 16)*. Customary

IHL Database. Retrieved June 21, 2026, from [https://ihl-](https://ihl-databases.icrc.org/en/customary-ihl/v1/rule16)

[databases.icrc.org/en/customary-ihl/v1/rule16](https://ihl-databases.icrc.org/en/customary-ihl/v1/rule16)

International Committee of the Red Cross. (2024a). *Artificial intelligence and related*

*technologies in military decision-making on the use of force in armed conflicts:*

*Current developments and potential implications*.

<https://www.icrc.org/en/publication/expert-consultation-report-artificial-intelligence-and-related-technologies-military>

International Committee of the Red Cross. (2024b). *International humanitarian law and the challenges of contemporary armed conflicts: Building a culture of compliance for IHL to protect humanity in today's and future conflicts*.

<https://www.icrc.org/en/publication/international-humanitarian-law-and-challenges-contemporary-armed-conflicts-building>

International Committee of the Red Cross. (2024c). *Private businesses and armed conflict: An introduction to relevant rules of international humanitarian law*.

<https://www.icrc.org/en/publication/private-businesses-and-armed-conflict-introduction-relevant-rules-international>

Jeans, D. (2026, March 20). Pentagon to adopt Palantir AI as core US military system, memo says. *Reuters*. <https://www.reuters.com/technology/pentagon-adopt-palantir-ai-as-core-us-military-system-memo-says-2026-03-20/>

Keohane, R. O. (1984). *After hegemony: Cooperation and discord in the world political economy*. Princeton University Press.

Kitchin, R. (2014). *The data revolution: Big data, open data, data infrastructures and their consequences*. SAGE Publications. <https://doi.org/10.4135/9781473909472>

Krahmann, E. (2010). *States, citizens and the privatisation of security*. Cambridge University Press.

Krishnaswamy, A. (2024, January 4). *Connecting AI to decisions with the Palantir Ontology*. Palantir Blog. <https://blog.palantir.com/connecting-ai-to-decisions-with-the-palantir-ontology-c73f7b0a1a72>

Leander, A. (2005). The power to construct international security: On the significance of private military companies. *Millennium*, 33(3), 803–826.  
<https://doi.org/10.1177/03058298050330030601>

- Livingstone, K. (2026, March 24). Deadly Iran school strike casts shadow over Pentagon's AI targeting push. *Military Times*. <https://www.militarytimes.com/news/your-military/2026/03/24/deadly-iran-school-strike-casts-shadow-over-pentagons-ai-targeting-push/>
- Madhani, A., Frankel, J., Biesecker, M., & Tucker, E. (2026, March 11). *Outdated intel likely led U.S. to carry out deadly strike on Iranian elementary school, AP sources say*. Associated Press. <https://apnews.com/article/iran-us-school-hegseth-trump-2ffff06808f7a584b0a03831897ab0b8>
- Mande, M., & Allen, G. C. (2026, June 2). *What is Maven Smart System, and what does it do?* Center for Strategic and International Studies. <https://www.csis.org/analysis/what-maven-smart-system-and-what-does-it-do>
- McKernan, B., & Davies, H. (2024, April 3). 'The machine did it coldly': Israel used AI to identify 37,000 Hamas targets. *The Guardian*. <https://www.theguardian.com/world/2024/apr/03/israel-gaza-ai-database-hamas-airstrikes>
- National Academies of Sciences, Engineering, and Medicine. (2023). *Test and evaluation challenges in artificial intelligence-enabled systems for the Department of the Air Force*. The National Academies Press. <https://doi.org/10.17226/27092>
- Nexus Insight. (n.d.). *The "speed of thought" era: AI and decision compression in modern warfare*. <https://nexusinsight.org/blogDetails99.php>
- Niebel, T., Rasel, F., & Viete, S. (2019). BIG data—BIG gains? Understanding the link between big data analytics and innovation. *Economics of Innovation and New Technology*, 28(3), 296–316. <https://doi.org/10.1080/10438599.2018.1493075>

- Palantir. (2025, February 25). *Empowering the warfighter: Palantir's partnership with Microsoft*. Palantir Blog. <https://blog.palantir.com/empowering-the-warfighter-palantirs-partnership-with-microsoft-820255fa04a5>
- Palantir. (2026, March 5). *Maven Smart System: Innovating for the Alliance*. Palantir Blog. <https://blog.palantir.com/maven-smart-system-innovating-for-the-alliance-5ebc31709eea>
- Palantir Technologies. (n.d.-a). *Ontology overview*. Palantir documentation. Retrieved June 21, 2026, from <https://www.palantir.com/docs/foundry/ontology/overview/>
- Palantir Technologies. (n.d.-b). *Object and link types: Link types*. Palantir documentation. Retrieved June 21, 2026, from <https://www.palantir.com/docs/foundry/object-link-types/link-types-overview/>
- Palantir Technologies. (2022). *Enabling interoperability and preventing lock-in with Foundry*. [https://www.palantir.com/assets/xrfr7uokpv1b/7BxLPkTqJU9QhLTQCjJMo6/eed1457949dc2d1cd6b6e71936c0aa9c/Enabling\\_Interoperability\\_and\\_Embracing\\_Openness\\_with\\_Foundry.pdf](https://www.palantir.com/assets/xrfr7uokpv1b/7BxLPkTqJU9QhLTQCjJMo6/eed1457949dc2d1cd6b6e71936c0aa9c/Enabling_Interoperability_and_Embracing_Openness_with_Foundry.pdf)
- Palantir Technologies. (2023). *Target Workbench*. [https://www.palantir.com/assets/xrfr7uokpv1b/1IqzwzpemtBSm98TNCczao/49bbc30cbe4d2d4d189ab27bd07376c/Palantir\\_Target\\_Workbench\\_\\_1\\_.pdf](https://www.palantir.com/assets/xrfr7uokpv1b/1IqzwzpemtBSm98TNCczao/49bbc30cbe4d2d4d189ab27bd07376c/Palantir_Target_Workbench__1_.pdf)
- Palantir Technologies. (2024a). *Maven Smart System*. <https://www.palantir.com/offerings/defense/air-space/>
- Palantir Technologies. (2024b, November 13). *Safeguarding freedom: How defense efforts align with human rights*. Palantir Blog. <https://blog.palantir.com/safeguarding-freedom-87080c30a712>

Palantir Technologies Inc. (2025). *2024 annual report (Form 10-K)*.

<https://investors.palantir.com/files/2024%20FY%20PLTR%2010-K.pdf>

Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.

Pearson, J., & McNeill, R. (2026, March 12). Bombed Iranian girls school had vivid website and yearslong online presence. *Reuters*.

<https://www.reuters.com/investigations/bombed-iranian-girls-school-had-vivid-website-yearslong-online-presence-2026-03-12/>

Pellerin, C. (2017, July 21). *Project Maven to deploy computer algorithms to war zone by year's end*. U.S. Department of Defense. <https://www.defense.gov/News/News-Stories/Article/Article/1254719/project-maven-to-deploy-computer-algorithms-to-war-zone-by-years-end/>

Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the protection of victims of international armed conflicts (Protocol I), June 8, 1977, 1125 U.N.T.S. 3. <https://treaties.un.org/doc/publication/unts/volume%201125/volume-1125-i-17512-english.pdf>

Reuters. (2026, March 13). Hegseth says US military has designated officer to complete probe on Iran school strike. *Reuters*. <https://www.reuters.com/world/hegseth-says-us-military-has-designated-officer-complete-probe-iran-school-2026-03-13/>

Reuters. (2026, April 7). *Trump says "a whole civilization will die tonight" if Iran does not make a deal*. <https://www.reuters.com/world/middle-east/trump-says-a-whole-civilization-will-die-tonight-if-iran-does-not-make-deal-2026-04-07/>

Roff, H. M. (2014). The strategic robot problem: Lethal autonomous weapons in war. *Journal of Military Ethics*, 13(3), 211–227.

<https://doi.org/10.1080/15027570.2014.975010>

- Roff, H. M., & Moyes, R. (2016). *Meaningful human control, artificial intelligence and autonomous weapons*. Article 36.
- Russo, A. (2026). Does AI affect the democratic conduct of war? Analyzing US and Israeli military AI deployment. *Global Policy*. Advance online publication.  
<https://doi.org/10.1111/1758-5899.70117>
- Sankar, S., & Hart, J. (2026). *Mobilize: How to reboot the American industrial base and stop World War III*. Post Hill Press.
- Sassòli, M. (2019). *International humanitarian law: Rules, controversies, and solutions to problems arising in warfare*. Edward Elgar Publishing.
- Schmitt, M. N., & Thurnher, J. S. (2013). ‘Out of the loop’: Autonomous weapon systems and the law of armed conflict. *Harvard National Security Journal*, 4, 231–281.
- Singer, P. W. (2009). *Wired for war: The robotics revolution and conflict in the 21st century*. Penguin Press.
- Srivastava, S. (2023). Algorithmic governance and the international politics of Big Tech. *Perspectives on Politics*, 21(3), 989–1000.  
<https://doi.org/10.1017/S1537592721003145>
- Stewart, P., & Ali, I. (2026, March 6). US investigation points to likely US responsibility in Iran school strike, sources say. *Reuters*. <https://www.reuters.com/world/middle-east/us-investigation-points-likely-us-responsibility-iran-school-strike-sources-say-2026-03-06/>
- Telasang, N. (2026, May 9). *AI-driven intelligence fusion in military cyber-forensic systems: A case study of Palantir’s Maven Smart System (MSS)* [Preprint]. SSRN. <https://ssrn.com/abstract=6738482>

- Toh, A., & Ayoub, E. (2026, March 11). *The business of military AI*. Brennan Center for Justice. [https://www.brennancenter.org/media/15340/download/bcj-167\\_business\\_of\\_military\\_ai\\_final.pdf?inline=1](https://www.brennancenter.org/media/15340/download/bcj-167_business_of_military_ai_final.pdf?inline=1)
- Tsotniashvili, Z. (2026). *Algorithmic warfare in the Iran conflict: AI-driven decision compression, the erosion of human oversight, and accountability gaps in contemporary military operations*. ResearchGate. [https://www.researchgate.net/publication/401535600\\_Algorithmic\\_Warfare\\_in\\_the\\_Iran\\_Conflict\\_AI-](https://www.researchgate.net/publication/401535600_Algorithmic_Warfare_in_the_Iran_Conflict_AI-)
- Uppal, R. (2021, February 25). *Project Maven accelerated weaponizing artificial intelligence for the US military: Was it used to kill Iran's nuclear scientist?* International Defense Security & Technology. <https://idstch.com/project-maven-accelerated-weaponizing-artificial-intelligence-for-the-us-military-was-it-used-to-kill-irans-nuclear-scientist/defence-military-systems-cat/air-force-aviation/>
- U.S. Department of Defense. (2019, August 9). *Department of Defense enterprise cloud and its importance to the warfighter media roundtable*. <https://www.defense.gov/News/Transcripts/Transcript/Article/1931163/departement-of-defense-enterprise-cloud-and-its-importance-to-the-warfighter-med/>
- U.S. Department of Defense. (2023). *Civilian harm mitigation and response* (DoD Instruction 3000.17). <https://www.esd.whs.mil/Portals/54/Documents/DD/issuances/dodi/300017p.pdf>
- U.S. Department of Defense. (2024a, May 29). *Contracts for May 29, 2024*. <https://www.defense.gov/News/Contracts/Contract/Article/3790490/>
- U.S. Department of Defense. (2024b, September 3). *Defense leaders combine top-down guidance with frontline expertise to develop capabilities at speed*.

<https://www.defense.gov/News/News-Stories/Article/Article/3892427/defense-leaders-combine-top-down-guidance-with-frontline-expertise-to-develop-c/>

U.S. Department of Defense. (2024c, September 18). *Contracts for September 18, 2024*.

<https://www.defense.gov/News/Contracts/Contract/Article/3910169/>

U.S. Department of Defense. (2025a). *Intellectual property guidebook for DoD acquisition*.

<https://www.acq.osd.mil/asda/dpc/api/docs/intellectual%20property%20guidebook%20for%20dod%20acquisition%20signed.pdf>

U.S. Department of Defense. (2025b, May 21). *Contracts for May 21, 2025*.

<https://www.defense.gov/News/Contracts/Contract/Article/4194643/>

U.S. Government Accountability Office. (2026). *Artificial intelligence acquisitions: Agencies should collect and apply lessons learned to improve future procurements* (GAO-26-107859). <https://www.gao.gov/assets/gao-26-107859.pdf>

U.S. Marine Corps. (2025, September 11). *Announcement of Maven Smart System licensing for Marine Corps* (MARADMIN 424/25).

<https://www.marines.mil/News/Messages/Messages-Display/Article/4299351/announcement-of-maven-smart-system-licensing-for-marine-corps/>

Vanderborght, R., & Nadibaidze, A. (2025). Military demonstrations as digital spectacles: How virtual presentations of AI decision-support systems shape perceptions of war and security. *European Journal of International Security*, 1–20. Advance online publication. <https://doi.org/10.1017/eis.2025.10015>

Weber, M. (1946). Politics as a vocation. In H. H. Gerth & C. W. Mills (Eds. & Trans.), *From Max Weber: Essays in sociology* (pp. 77–128). Oxford University Press. (Original work published 1919)

Wendt, A. (1999). *Social theory of international politics*. Cambridge University Press.

Yeung, K. (2017). 'Hypermudge': Big Data as a mode of regulation by design. *Information, Communication & Society*, 20(1), 118–136.

<https://doi.org/10.1080/1369118X.2016.1186713>

Yeung, K. (2018). Algorithmic regulation: A critical interrogation. *Regulation & Governance*, 12(4), 505–523. <https://doi.org/10.1111/rego.12158>

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs.