



University
of Exeter

ECIL
EXETER CENTRE FOR
INTERNATIONAL LAW

Target Selection by Autonomous Weapons Systems: What Does the Duty of Target Verification Require?

Aurel Sari

Exeter Centre for International Law

Working Paper Series

2026/2



University
of Exeter

ECIL

EXETER CENTRE FOR
INTERNATIONAL LAW

The **Exeter Centre for International Law** builds on a long and distinguished tradition of international legal scholarship at Exeter Law School. The Centre's mission is to provide an intellectual environment for the study and development of international law and to stimulate discussion and collaboration in response to the most pressing challenges facing the international community. As part of this mission, the Centre publishes the present Working Paper Series.

Centre Director:	Dr Chris O'Meara
General Editor:	Professor Aurel Sari
Editor in Chief:	Professor Kubo Mačák

Exeter Centre for International Law
Exeter Law School, Amory Building
Rennes Drive, Exeter, EX4 4RJ, United Kingdom

 <http://www.exeter.ac.uk/ecil>
 [@ExeterCIL](https://twitter.com/ExeterCIL)

© All rights reserved. No part of this paper may be reproduced in any form without the permission of the author.

Cite as Aurel Sari, "Target Selection by Autonomous Weapons Systems: What Does the Duty of Target Verification Require?", ECIL Working Paper 2026/2.

Target Selection by Autonomous Weapons Systems: What Does the Duty of Target Verification Require?

Aurel Sari*

This article examines what the precautionary duty of target verification under Article 57(2)(a)(i) of Additional Protocol I requires when target selection is delegated to autonomous weapons systems (AWS). It argues that target selection, as understood in military doctrine, is not a single act of choice, but a structured, multistage process in which a pool of candidate entities is progressively narrowed through the application of operational, legal and other criteria. This means that even in a genuinely autonomous system, neither target identification nor target selection is a fully automated function. The article traces the debates within the UN Group of Governmental Experts on lethal AWS over inclusive and narrow readings of ‘identification’ and ‘selection’, arguing that the definitional weight should shift from the connectors joining these critical functions to the division of labour between operator and system. It then disaggregates the duty of target verification into four constituent elements—information gathering, reliability assessment, legal characterization and individuation of the target—to show that the duty attaches to the actual person or object selected for engagement and thus remains engaged until that entity has been identified. Consequently, the duty of target verification cannot be discharged in full at the point an AWS is activated. Drawing on two contrasting configurations modelled on real systems, the article maps how the different constituent elements of target verification are distributed between human operators and machines, drawing out the implications for other precautionary obligations, the GGE’s definitional debates and the proposition that legal assessments may never be delegated to a machine.

1. INTRODUCTION

Automated weapons systems are not a recent invention. Radar-guided gun platforms designed to detect and engage airborne threats have been in service since the early 1980s. The Phalanx system,

* Professor of Public International Law, University of Exeter.

developed by the United States Navy and subsequently adopted by more than twenty other States, is perhaps the best-known example.¹ So far, these systems have been deployed primarily against targets that are military objectives by nature, such as enemy missiles and military aircraft, in environments where civilians are mostly absent. However, developments in artificial intelligence (AI), sensors and information technology are beginning to change this picture. Machine learning systems enable the real-time processing of multiple data streams to recognize targets in more dynamic circumstances. The likely result is a proliferation of autonomous targeting against a wider set of targets in a wider range of environments.

Most definitions of autonomous weapons systems (AWS) identify the capacity to select and engage targets without human intervention as their key distinguishing feature.² The United States Department of Defense (DoD) Directive 3000.09 defines an autonomous weapon system as one ‘that, once activated, can select and engage targets without further intervention by an operator’.³ Similarly, the International Committee of the Red Cross (ICRC) describes AWS as

-
- ¹ The Phalanx Close-In Weapon System (CIWS) entered production in 1978 and operational service in 1980. The system autonomously conducts search, detection, threat evaluation, tracking, engagement and kill assessment without operator input during the engagement cycle. For further detail, see US Department of the Navy, ‘MK 15 Phalanx Close-In Weapon System (CIWS)’ (*Navy.mil*, 20 September 2021) <<https://www.navy.mil/resources/fact-files/display-factfiles/article/2167831/mk-15-phalanx-close-in-weapon-system-ciws/>> accessed 19 August 2026. Similar systems include the Goalkeeper (Thales, Netherlands), Harpy and HAROP loitering munitions (Israel Aerospace Industries) and the MANTIS Counter-Rocket, Artillery and Mortar system (Rheinmetall, Germany). Generally, see Vincent Boulanin and Maaïke Verbruggen, *Mapping the Development of Autonomy in Weapons Systems* (Stockholm International Peace Research Institute 2017).
- ² Ingvild Bode and Hendrik Huelss, *Autonomous Weapons Systems and International Norms* (McGill-Queen’s University Press 2022) 26; Diego Mauri, *Autonomous Weapons Systems and the Protection of the Human Person: An International Law Analysis* (Edward Elgar 2022) 28; Dan Saxon, *Fighting Machines: Autonomous Weapons and Human Dignity* (University of Pennsylvania Press 2022) 8; Afonso Seixas-Nunes, *The Legality and Accountability of Autonomous Weapon Systems: A Humanitarian Law Perspective* (CUP 2022) 82; Tim McFarland, *Autonomous Weapons Systems and the Law of Armed Conflict: Compatibility with International Humanitarian Law* (CUP 2020) 8; Nathan Gabriel Wood, ‘Autonomous Weapons Systems: A Clarification’ (2023) 22 *Journal of Military Ethics* 18, 20; Daniele Amoroso and Guglielmo Tamburrini, ‘The Ethical and Legal Case Against Autonomy in Weapons Systems’ (2018) 18 *Global Jurist* 1, 1; Christopher M Ford, ‘Autonomous Weapons and International Law’ (2017) 69 *South Carolina Law Review* 413, 417; Rebecca Crootof, ‘The Killer Robots Are Here: Legal and Policy Implications’ (2015) 36 *Cardozo Law Review* 1837, 1857.
- ³ United States Department of Defense, *Directive 3000.09, Autonomy in Weapons Systems* (25 January 2023) (hereinafter DoD Directive 3000.09) 21. ‘Operator’ is defined at 22 as ‘[a] person who operates a platform or weapon system’, so the assessment reserved to the operator is in any event a human one.

‘weapons systems that select and engage [...] one or more targets without human intervention’.⁴ More recently, the Chair of the Group of Governmental Experts (GGE) on emerging technologies in the area of lethal autonomous weapons systems, constituted within the framework of the Convention on Certain Conventional Weapons (CCW),⁵ characterized lethal AWS as ‘a combination of one or more weapons and functionally integrated technological components that can identify, select, and engage a target, without intervention by a human operating the system’.⁶

Despite the centrality of target selection to most definitions of AWS, there is no authoritative or universally accepted account of what the selection of a target actually entails. The phrase is generally understood to mean the process whereby a person or object is chosen for attack.⁷ However, there is no agreed catalogue of specific steps or actions that this involves. The widespread working assumption is that AWS do not define their own target sets.⁸ Instead, they select targets by matching information about the operating environment, obtained through their sensors and other inputs, against a target profile predefined by their operators. For example, an operator might instruct an AWS to find enemy vehicles matching certain parameters, whereupon the system searches its zone of operations for members of that category. The boundary of target selection nonetheless remains unclear. Does selecting a target refer both to the operator’s selection of a target category and to the system’s subsequent action in finding and identifying an actual instance of that category? If so, target selection is not a fully automated function, since it combines human and machine inputs.⁹ Alternatively, does it refer only to the system’s action in finding and identifying a specific instance of a human-selected category? On this reading, target selection is a

⁴ International Committee of the Red Cross, *Autonomous Weapons Systems and International Humanitarian Law: Selected Issues* (ICRC 2026) (hereinafter ICRC Selected Issues) 6.

⁵ Convention on Prohibitions or Restrictions on the Use of Certain Conventional Weapons which may be Deemed to be Excessively Injurious or to Have Indiscriminate Effects (adopted 10 October 1980, entered into force 2 December 1983) 1342 UNTS 137 (hereinafter CCW).

⁶ Chairperson of the GGE on LAWS, ‘Rolling Text’ (5 June 2026) (hereinafter Rolling Text, 5 June 2026) para 1. On the successive formulations put to the Group, see Chairperson of the GGE on LAWS, ‘Non-exhaustive Compilation of Definitions and Characterizations’ (10 March 2023) UN Doc CCW/GGE.1/2023/CRP.1.

⁷ eg Eliav Lieblich and Eyal Benvenisti, ‘The Obligation to Exercise Discretion in Warfare: Why Autonomous Weapons Systems are Unlawful’ in Nehal Bhuta and others (eds), *Autonomous Weapons Systems: Law, Ethics, Policy* (CUP 2016) 245, 249; Ford (n 2) 419; Christof Heyns, Report of the Special Rapporteur on Extrajudicial, Summary or Arbitrary Executions (9 April 2013) UN Doc A/HRC/23/47, para 38.

⁸ eg International Committee of the Red Cross, ‘Ethics and Autonomous Weapon Systems: An Ethical Basis for Human Control?’ (29 March 2018) (hereinafter ICRC Ethics Paper) UN Doc CCW/GGE.1/2018/WP.5, para 35; Crootof (n 2) 1854.

⁹ See also Mauri (n 2) 26–27; McFarland (n 2) 34–35.

fully automated activity, but only because the definition ignores the critical contribution made by humans.

The lack of clarity about the meaning and scope of target selection has significant implications under the law of armed conflict. Those who plan or decide upon an attack are subject to a precautionary duty to do everything feasible to verify that any target to be attacked is a military objective liable to attack. The obligation is set out in Article 57(2)(a)(i) of Additional Protocol I (AP I)¹⁰ and is recognized as a rule of customary international law.¹¹ It is not immediately clear at what stage in the targeting sequence this duty must be discharged when an AWS is used for target selection.¹² Where a human operator deploys a weapon system to attack a specific target they have preselected,¹³ the law requires the operator to verify that this target is a military objective before they deploy the system to launch the attack.¹⁴ The position is different when target selection is delegated to an autonomous system. Here, the operator instructs the system to find members of a category of targets and the system proceeds to identify specific instances of that category against the operator's criteria. In such cases, must the operator discharge the duty of target verification in full before launching the AWS, as the ICRC suggests?¹⁵ Or does the fact that the task of finding and identifying specific members of a category of targets has been delegated to the AWS mean that the duty of target verification remains engaged after system activation?

The answer to these questions determines when the duty of target verification must be discharged and by whom: the human operator or by the autonomous system itself. The answer also has implications for the implementation of other precautionary obligations that depend on target verification, such as the duty not to launch attacks expected to cause disproportionate

¹⁰ Protocol Additional to the Geneva Conventions of 12 August 1949, and Relating to the Protection of Victims of International Armed Conflicts (8 June 1977) 1125 UNTS 3 (hereinafter Additional Protocol I).

¹¹ Jean-Marie Henckaerts and Louise Doswald-Beck, *Customary International Humanitarian Law* (CUP 2005) vol 1, Rule 16 (hereinafter ICRC Customary Study).

¹² See Vincent Boulanin, Laura Bruun and Netta Goussac, *Autonomous Weapons Systems and International Humanitarian Law: Identifying Limits and the Required Type and Degree of Human-Machine Interaction* (Stockholm International Peace Research Institute 2021) 24 (treating as open questions how far in advance the assessments mandated by the law of armed conflict may be made, how long an *ex ante* assessment remains valid after activation of an autonomous system and whether an assessment must be kept current after the system has been activated).

¹³ See DoD Directive 3000.09 (n 3) 23 (defining 'semi-autonomous weapons systems' as a 'weapon system that, once activated, is intended to only engage individual targets or specific target groups that have been selected by an operator').

¹⁴ Mauri (n 2) 143.

¹⁵ ICRC Selected Issues (n 4) 11.

incidental civilian harm.¹⁶ Since the proportionality rule is only engaged in attacks that are directed against military objectives, the process of target verification must precede any proportionality assessment.¹⁷ This means that a proportionality assessment cannot be carried out by human operators at the time of activating an AWS for attack unless they have fully discharged their duty to verify the target by that point.

The purpose of this article is to resolve these questions by examining how the duty of target verification under Article 57(2)(a)(i) AP I maps onto the use of autonomous systems for target selection. Section 2 examines target selection as military doctrine understands it, showing that it does not involve a single act of choice, but constitutes a structured process in which a pool of candidate entities is progressively narrowed. While the use of automated systems to accelerate the targeting cycle is rapidly expanding, so far this automation remains confined to discrete tasks, rather than extending to the entire process. The section also distinguishes AWS from AI decision-support systems and identifies the type of systems of greatest legal interest, namely systems tasked with finding and engaging individual members of a target category chosen in advance by a human operator. Section 3 turns to the efforts of the GGE to define identification and selection as critical functions of AWS. It traces the oscillation between inclusive and narrow readings of selection and argues that neither identification nor selection is, even in a genuinely autonomous system, a fully automated function. The definitional weight therefore shifts from the connectors that join the enumerated functions to the division of labour between operator and system. Section 4 analyses the content and structure of the duty of target verification, establishing that it combines factual and legal determinations, attaches to the specific entity selected for attack and remains engaged until that entity has been identified. It then relates the constituent elements of the duty to the stages of the targeting cycle, examines how those elements fall to be distributed between human operators and autonomous systems across two contrasting configurations and draws out the implications, including for the implementation of other precautionary obligations, the definitional work of the GGE and the ICRC's position that legal assessments may never be delegated to a machine. Section 5 concludes.

2. TARGET SELECTION AND AUTOMATION

¹⁶ Additional Protocol I, art 57.

¹⁷ Geoffrey S Corn, 'War, Law, and the Oft Overlooked Value of Process as a Precautionary Measure' (2015) 42 *Pepperdine Law Review* 419, 436–437.

In everyday usage, to select means to choose between alternatives.¹⁸ Selection is therefore relational and comparative: it presupposes the existence of alternative candidates or courses of action, a set of criteria by which to discriminate between them and a decision by which some are preferred and others are set aside.¹⁹ The ordinary meaning of the word selection suggests that ‘target selection’ refers to the process of choosing something as a target for attack among possible alternatives, based on certain parameters. This is the meaning that Article 57(3) AP I ascribes to the term when it demands that targets meeting certain criteria must be selected in cases where alternative choices are possible. The same meaning is also confirmed by military practice, as illustrated by the joint targeting cycle that serves as a framework for target selection in operations conducted by the North Atlantic Treaty Organization (NATO). Target selection in the military doctrinal sense is a structured process during which a pool of candidate entities is progressively narrowed through the repeated application of various criteria, until some of these candidate targets are eventually approved for engagement.²⁰

The targeting process is becoming increasingly automated, driven largely by developments in AI technologies.²¹ Much of this automation is the work of decision-support systems, which generate assessments, predictions and recommendations to inform human decision-making

¹⁸ Oxford English Dictionary, ‘select (v.), sense 1.a. To choose or pick out in preference to another or others’, June 2026, <<https://doi.org/10.1093/OED/5485793048>>.

¹⁹ On selection and choice as the identification of a preferred action over a set of alternatives, see Daniel A Levinthal, *Evolutionary Processes and Organizational Adaptation: A Mendelian Perspective on Strategic Management* (OUP 2021) 27–42. On the structure of choice as elimination followed by selection, Keshav Sureka and Debabrata Pal, ‘Choice by Elimination then Selection’ (2024) 78(3) *Research in Economics* 100966. For the classical decision-theoretic account of choice as the selection of an alternative that maximizes expected utility, see Ward Edwards, ‘The Theory of Decision Making’ (1954) 51 *Psychological Bulletin* 380.

²⁰ For a general account and different national versions of the targeting process, see Phillip R Pratzner, ‘The Current Targeting Process’ in Paul A L Ducheine, Michael N Schmitt and Frans P B Osinga (eds), *Targeting: The Challenges of Modern Warfare* (TMC Asser Press 2016) 77; Michael Schmitt and others, ‘Joint and Combined Targeting: Structure and Process’ in Jens David Ohlin, Larry May and Claire Finkelstein (eds), *Weighing Lives in War* (OUP 2017) 298; Mitch Ferry, ‘F3EA — A Targeting Paradigm for Contemporary Warfare’ (2013) 10 *Australian Army Journal* 49.

²¹ On the automation of the targeting process, see Anthony King, ‘Digital Targeting: Artificial Intelligence, Data, and Military Intelligence’ (2024) 9 *Journal of Global Security Studies* oga009, 4–6; Christopher J Lowrance, Elliott R Harris and C Anthony Pfaff, ‘Expediting the Targeting Process Responsibly Using Artificial Intelligence’ in *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications IV* (Proc SPIE vol 12113, 2022) art 1211307; and, for a treatment of AI across the phases of the targeting cycle, Arthur Holland Michel, *Decisions, Decisions, Decisions: Computation and Artificial Intelligence in Military Decision-Making* (International Committee of the Red Cross 2024).

without themselves applying force.²² Such systems must be distinguished from AWS. An AWS does not merely inform a decision to use violence, but directly engages targets with kinetic force. This has two consequences. Normatively, it means that AWS are weapons within the meaning of the law of armed conflict to which the rules governing weapons and means of warfare apply.²³ Operationally, it means that AWS are not an intelligence enterprise, but deployed platforms designed to find, track and physically engage their targets. Although AWS must be equipped with sensors, processing power and other technical features to acquire and process the data necessary to perform their functions,²⁴ their organic capabilities are far more limited than the processing capacity that a deployed headquarters can bring to bear. For the foreseeable future, AWS should therefore be expected to perform only some of the target selection tasks arising during the targeting process, such as matching a candidate entity against a predefined target profile, which fall towards the end of the targeting process and build on selection activities already performed by humans or other systems. How far an AWS can be entrusted with target selection, and which elements of that process it performs, is where the meaning of target selection turns from a doctrinal question into a legal one.

2.1. The joint targeting cycle

In contemporary military operations, target selection forms part of a process that governs how military objectives are identified, legally assessed, approved and engaged. The process is operationalized through the joint targeting cycle, which NATO doctrine defines as ‘the process of selecting and prioritising targets [...] and matching the appropriate response to them, taking

²² On the use of AI decision-support systems in the military domain, see Alexander Blanchard and Laura Bruun, *Autonomous Weapons Systems and AI-Enabled Decision Support Systems in Military Targeting: A Comparison and Recommended Policy Responses* (Stockholm International Peace Research Institute 2025); Emelia Probasco and others, *AI for Military Decision-Making: Harnessing the Advantages and Avoiding the Risks* (Center for Security and Emerging Technology 2025); and Anna Nadibaidze, Ingild Bode and Qiaochu Zhang, *AI in Military Decision Support Systems: A Review of Developments and Debates* (Center for War Studies 2024).

²³ For a general overview of the rules governing weapons, see William H Boothby, *Weapons and the Law of Armed Conflict* (2nd edn, OUP 2016) (hereinafter Boothby, *Weapons and the Law*); Yoram Dinstein, *The Conduct of Hostilities under the Law of International Armed Conflict* (3rd edn, CUP 2016) 1–38; and Stuart Casey-Maslen and Sharon Weill, ‘The use of weapons in armed conflict’ in Stuart Casey-Maslen (ed), *Weapons under International Human Rights Law* (CUP 2014) 240.

²⁴ See United States, ‘The Autonomy Aspect of Lethal Autonomous Weapons Systems’ (2 September 2025) UN Doc CCW/GGE.1/2025/WP.6, para 18 (analysing autonomous weapons systems by reference to their character as weapons systems that apply force).

account of operational requirements and capabilities, with a view to creating desired effects in accordance with the commander's objectives'.²⁵ The cycle operates through two methodologies. Deliberate or preplanned targeting goes through six phases to produce a plan for engaging targets known to exist in the area of operations.²⁶ Candidate targets are advanced through successive stages of refinement until some are formally approved for engagement. Dynamic targeting, by contrast, is employed against targets that were not detected, located or selected for action in time to enter the deliberate cycle, but which require an immediate response.²⁷ Their development takes place in a more compressed timeframe, using abbreviated procedures, and may not reach the same level of detail.²⁸

Target selection in the deliberate targeting cycle begins with the commander's guidance, including the formulation of guidelines for target selection and discovery, which involves locating and identifying entities of interest in the operating environment.²⁹ Candidate entities proceed through target development, encompassing basic target development, which consists of the initial identification, location, description and functional characterization of an entity of interest, and intermediate target development, which completes this characterization and assesses the entity's significance.³⁰ An entity that satisfies these standards is placed on the Candidate Target List, pending validation.³¹ Quality control assesses the accuracy and currency of the supporting intelligence.³² Targets are then subject to formal validation. The validation review considers whether the target constitutes a lawful military objective and whether engaging it would comply with the applicable rules of engagement. Validated targets are placed on the Joint Target List, the Restricted Target List or the No-Strike List.³³ Validated targets are then nominated and prioritized for engagement. Once approved, they are placed on the Joint Prioritized Target List, which is

²⁵ North Atlantic Treaty Organization, *Allied Joint Doctrine for Joint Targeting* (AJP-3.9, Edition B, Version 1, 2021) (hereinafter AJP-3.9) 1-1.

²⁶ AJP-3.9 (n 25) 1-14.

²⁷ AJP-3.9 (n 25) 1-12.

²⁸ See Joint Chiefs of Staff, *Joint Targeting*, Joint Publication 3-60 (2018) II-3 (recording that '[d]ue to dynamic targeting's compressed timeline, development must be accomplished quickly' and that such targets 'might not be characterized to the same level of detail before execution').

²⁹ AJP-3.9 (n 25) 1-9 and 1-15.

³⁰ AJP-3.9 (n 25) 1-5 and 1-7.

³¹ AJP-3.9 (n 25) 1-5.

³² AJP-3.9 (n 25) 2-4 and 2-3 (quality control assessing the accuracy of the supporting target intelligence, and the credibility of sources).

³³ AJP-3.9 (n 25) 2-4–2-5 and 1-16.

handed down to component commands for execution. Before engaging a target on that list, tactical operators must acquire positive identification (PID), using visual recognition, electronic support systems or other available means to determine that the person or object about to be attacked is in fact the validated target.³⁴

The selection of targets involves the application of several sets of criteria throughout the targeting cycle, each functioning as a filter that progressively reduces the population of candidate entities. The outermost filter is political and strategic. The North Atlantic Council defines which target sets may enter the target development process at all.³⁵ Sensitive targets that require strategic-level or national-level review are identified at the outset as requiring additional approval before planning proceeds.³⁶ Operational criteria govern which entities become relevant candidates for target development. Target system and intelligence analysis prioritize entities by reference to their contribution to the adversary's ability to conduct operations, distinguishing between high-value, high-payoff and time-sensitive targets.³⁷ Later stages in the cycle assess the practical feasibility of target engagement, including by considering weaponeering solutions.³⁸ Further operational criteria are applied when validated targets are ranked against one another by reference to the commander's intent, objectives and guidance, with the aim of maximizing the effective use of available capabilities.³⁹ Intelligence criteria are applied at different points to determine whether enough is known about a target to support its further development and ultimately its approval. Legal criteria run throughout, including the application of the principle of distinction before candidate targets enter the Joint Target List.⁴⁰ During advanced target development, proportionality assessments are conducted to decide whether targets are placed on the Joint Prioritized Target List.⁴¹ In addition, information and perception criteria are applied to ensure that target engagement is consistent with approved messaging and information operations objectives.⁴² Gender analysis constitutes another discrete input.⁴³

³⁴ AJP-3.9 (n 25) 1-8 and 1-25.

³⁵ AJP-3.9 (n 25) 1-15.

³⁶ AJP-3.9 (n 25) 1-9.

³⁷ AJP-3.9 (n 25) 2-1.

³⁸ AJP-3.9 (n 25) 1-17–1-18.

³⁹ AJP-3.9 (n 25) 1-18.

⁴⁰ AJP-3.9 (n 25) 1-16.

⁴¹ AJP-3.9 (n 25) 1-17–1-18 and 1-25.

⁴² AJP-3.9 (n 25) 1-18. See also North Atlantic Treaty Organization, *Allied Joint Doctrine for Information Operations* (AJP-10.1, Edition A, Version 1, 2023).

⁴³ AJP-3.9 (n 25) 1-16.

Target selection in the deliberate targeting cycle thus constitutes a formal, multistage process that involves identifying a pool of potential targets and choosing some of those targets for engagement by repeatedly applying a set of political, strategic, operational, epistemic, legal, informational and gender-based selection criteria. Dynamic targeting differs from the deliberate process primarily in how these selection criteria are applied. Dynamic targeting prosecutes targets of opportunity that require an immediate response because they present a threat to the force or mission and their engagement supports the commander's objectives.⁴⁴ Such targets may be known, but were not detected or located until that point; they may have been previously identified, but were not selected for action during deliberate planning; or they may be wholly unanticipated targets that emerge and are recognized only during the execution of current operations.⁴⁵ In all of these cases, there is insufficient time to complete the multiphase selection process that governs deliberate targeting. Instead, the selection criteria are applied in a less formal and more compressed manner, relying on the find, fix, track, target, engage, exploit and assess (F2T2E2A) process.

In both the deliberate and the dynamic process, the selection criteria are applied repeatedly and across successive stages. Target selection in the joint targeting cycle is therefore not a single decision point, but a multistage process comprising successive activities, including the sensory detection of entities, the identification and nomination of candidate targets, target development, validation and approval and, finally, positive identification before engagement. Each of these activities is, in principle, a task that can be delegated to a machine.

2.2. Automating target selection

Many target selection tasks are already being delegated, in whole or in part, to computer-based systems. Target discovery is a prime use-case. Because discovering and identifying a target depends on correlating large volumes of all-source intelligence, it is particularly well suited to machine learning techniques capable of processing vast quantities of data at speeds that exceed human analytical capacity. One of the best-known examples is the United States Department of Defense's Project Maven programme, which applies computer vision and deep learning to full-motion video, satellite imagery and other intelligence sources to detect, classify and track objects.⁴⁶ Other use-

⁴⁴ United States Air Force, *Air Force Doctrine Publication 3-60: Targeting* (2026) 6 (defining targets of opportunity as those 'identified too late or not selected for action in time to be included in deliberate targeting').

⁴⁵ AJP-3.9 (n 25) 1-12.

⁴⁶ On Project Maven, see Richard H Shultz and Richard D Clarke, 'Big Data at War: Special Operations Forces, Project Maven, and Twenty-First-Century Warfare' (*Modern War Institute*, 25 August 2020) <<https://mwi.westpoint.edu/big-data-at-war-special-operations-forces-project-maven-and-twenty-first->

cases include profile-matching, where the technical or physical signatures of a candidate target are matched to the features associated with a category of military objectives.⁴⁷ AI-enabled tools are reportedly also used for target prioritization.⁴⁸ By contrast, quality control is less readily automated, since judging the credibility of a given piece of intelligence typically depends on qualitative assessments that resist straightforward formalization. Automated tools nonetheless increasingly perform functions that support target validation, for instance by corroborating a nomination across multiple independent data sources. Israel's Habsora system offers an example. The system is reported to cross-reference telecommunications, imagery and locational data to confirm a correlation before generating an alert, for example by verifying that a flagged telephone number is co-located with a specific address.⁴⁹

As these examples demonstrate, automation is not concentrated at any single point in the targeting cycle. Current practice applies AI-enabled tools and other computer-based technologies to target selection at every stage, from discovery through to execution.⁵⁰ However, none of the existing systems automate the entire selection process. For now, the application of these technologies is confined to specific activities or stages. The degree to which this partial automation displaces human involvement in the selection of targets varies. In some cases, automated processes predominantly inform human decisions that remain substantively open. In other cases, they may

century-warfare/> accessed 19 August 2026; and Lucy Suchman, 'Algorithmic Warfare and the Reinvention of Accuracy' (2020) 8 *Critical Studies on Security* 175. For the operational description, see also Merel Ekelhof, 'AI Is Changing the Battlefield, But Perhaps Not How You Think: An Analysis of the Operationalization of Targeting Law and the Increasing Use of AI in Military Operations' in Robin Geiß and Henning Lahmann (eds), *Research Handbook on Warfare and Artificial Intelligence* (Edward Elgar 2024) 161, 167.

⁴⁷ Israel's 'Habsora' (the Gospel) system is reported to compile, fuse and cross-reference layers of information from different datasets to generate target suggestions for intelligence analysts, for instance by corroborating that a flagged telephone number is co-located with a specified address: see Harry Davies, Bethan McKernan and Dan Sabbagh, 'The Gospel: How Israel Uses AI to Select Bombing Targets in Gaza' (*The Guardian*, 1 December 2023) <<https://www.theguardian.com/world/2023/dec/01/the-gospel-how-israel-uses-ai-to-select-bombing-targets>> accessed 19 August 2026. For legal analysis, see Michael N Schmitt, 'The Gospel, Lavender, and the Law of Armed Conflict' (*Articles of War*, Lieber Institute West Point, 1 May 2024) <<https://lieber.westpoint.edu/gospel-lavender-law-armed-conflict/>> accessed 19 August 2026.

⁴⁸ Jessica Dorsey and Marta Bo, 'AI-Enabled Decision-Support Systems in the Joint Targeting Cycle: Legal Challenges, Risks, and the Human(e) Dimension' (2026) 107 *International Law Studies* 184, 200.

⁴⁹ Dorsey and Bo (n 48) 196–197.

⁵⁰ On the resulting risk that human involvement becomes nominal, see Arthur Holland Michel, *The Black Box, Unlocked: Predictability and Understandability in Military AI* (United Nations Institute for Disarmament Research 2020) 12–16; and Nadibaidze, Bode and Zhang (n 22).

narrow operator decision-making to one of confirmation.

Most of the systems employed to automate target selection are AI decision-support systems (AI DSS).⁵¹ While no authoritative definition of AI DSS exists, recent scholarship and policy practice converge around a functional understanding of the term. The ICRC and the Geneva Academy, for example, have defined AI DSS as ‘computerized tools that are designed to assist humans at different levels in the chain of command to complete decision-making tasks’.⁵² Other organizations and most of the academic literature describe AI DSS along similar lines.⁵³ What unites these formulations is their focus on function. AI DSS generate cognitive outputs, such as assessments, predictions and recommendations, that support and inform human decision-making, but they produce no physical effects themselves. Project Maven, for example, processes various data streams to detect and classify objects, generating annotations for human analysts. Habsora cross-references intelligence sources to flag a correlation for human verification. Anduril’s Lattice system fuses data from multiple sensors to detect, track and classify objects of interest, producing a common operating picture.⁵⁴ None of these systems directly generate any physical effects. The outputs of AI DSS are, by definition, separated from kinetic action by at least one human decision-making layer.

⁵¹ International Committee of the Red Cross and Geneva Academy, *Artificial Intelligence and Related Technologies in Military Decision-Making on the Use of Force in Armed Conflicts: Current Developments and Potential Implications* (ICRC 2024) (hereinafter ICRC and Geneva Academy) 8; Marta Bo and Jessica Dorsey, ‘The “Need” for Speed: The Cost of Unregulated AI Decision-Support Systems to Civilians’ (Opinio Juris, 4 April 2024); and Holland Michel (n 21) 13.

⁵² ICRC and Geneva Academy (n 51) 8.

⁵³ eg Alexander Blanchard and Laura Bruun, *Autonomous Weapons Systems and AI-enabled Decision Support Systems in Military Targeting: A Comparison and Recommended Policy Responses* (Stockholm International Peace Research Institute 2025) 3–4; Bill Boothby, ‘AI Warfare and the Law’ (2025) 104 *International Law Studies* 1 (hereinafter Boothby, ‘AI Warfare and the Law’) 45; Renato Wolf, ‘Applying Precautions in Target Verification with AI Decision Support Systems’ (2025) 58 *Israel Law Review* 282; Holland Michel (n 21) 13; Klaudia Klonowska, ‘Article 36: Review of AI Decision-Support Systems and Other Emerging Technologies of Warfare’ in Terry D Gill and others (eds), *Yearbook of International Humanitarian Law, Volume 23 (2020)* (TMC Asser Press 2022) 123, 135.

⁵⁴ Anduril Industries, ‘Command & Control’ (Anduril) <<https://www.anduril.com/lattice/command-and-control>> accessed 19 August 2026, describing Lattice for Command & Control as ‘an AI-powered battle management platform that integrates thousands of sensors and effectors to accelerate complex kill chains at machine speed’. The company’s product description, as reported, states that the platform uses artificial intelligence, computer vision, edge computing and sensor fusion to detect, track and classify objects of interest: Jon Harper, ‘Anduril Integrates AI Tech into Army IVAS Headsets’ (DefenseScoop, 19 September 2024) <<https://defensescoop.com/2024/09/19/ivas-anduril-microsoft-lattice-integration-army/>> accessed 19 August 2026. The system is described in the same terms in the scholarly literature: Ekelhof (n 46) 173.

Although AI DSS may perform a variety of target selection functions in both deliberate and dynamic targeting, ranging from detection and profile-matching to corroboration and prioritization, they do not, by definition, engage targets directly. This distinguishes AI DSS from AWS.⁵⁵ As noted earlier, AWS are widely defined as systems that, once activated, are capable of both selecting and engaging a target without further intervention by a human operator. What sets AWS apart from AI DSS is that they are weapons systems designed to deliver physical effects directly.⁵⁶

2.3. AWS as weapons systems

Since the start of its deliberations, the GGE has consistently described the systems under its consideration as weapons systems.⁵⁷ However, since the GGE has never defined a weapon system as a discrete term, we must look elsewhere to uncover its meaning. None of the principal law of armed conflict instruments define the term ‘weapon’ with general effect. The treaty definitions that do exist apply only to specific categories of weapons, such as chemical weapons or cluster munitions.⁵⁸ The meaning of the concept is nevertheless settled. A weapon is generally understood to mean an object or substance designed or intended to inflict death or injury on persons or physical damage to objects, that is to say an instrument of violence. State practice amply confirms this. Australia defines a weapon as ‘any arms, munition, device, whether tangible or intangible,

⁵⁵ For a different view, see Mustafa Can Sati, ‘Military AI Technologies under International Humanitarian Law: Re-Evaluation of Fundamental Principles’ (PhD thesis, Maastricht University 2026) 141–142.

⁵⁶ Chairperson of the GGE on LAWS, ‘Revised Chair’s Paper’ (20 September 2021) Characterizations, para (d) (weapons systems that incorporate autonomy only into functions other than selecting and engaging to apply force against targets ‘cannot be characterized as autonomous weapons systems’).

⁵⁷ The systems under consideration have been described as weapons systems throughout the GGE process: see, eg, Group of Governmental Experts on LAWS, ‘Report of the 2018 Session’ (2018) UN Doc CCW/GGE.1/2018/3, para 21; Group of Governmental Experts on LAWS, ‘Report of the 2019 Session’ (2019) UN Doc CCW/GGE.1/2019/3 (hereinafter Report of the 2019 Session) Annex IV, Guiding Principle (a); and Australia, Canada, Japan, Poland, the Republic of Korea, the United Kingdom and the United States, ‘Draft Articles on Autonomous Weapons Systems’ (15 May 2023) UN Doc CCW/GGE.1/2023/WP.4/Rev.2 (hereinafter Draft Articles on Autonomous Weapons Systems) which throughout characterize LAWS as ‘weapons systems’. On the meaning of ‘weapon’, see Boothby, *Weapons and the Law* (n 23) 4–5 (a weapon is a ‘device, system, munition, implement, substance, object, or piece of equipment’ used, intended or designed for use to apply an offensive capability against a military object or enemy combatant).

⁵⁸ Convention on the Prohibition of the Development, Production, Stockpiling and Use of Chemical Weapons and on Their Destruction (13 January 1993) 1974 UNTS 45, art II(1); Convention on Cluster Munitions (30 May 2008) 2688 UNTS 39, art 2(2).

designed or intended to be used to cause... injury to, or death of, persons; or... damage to, destruction of or neutralisation of objects'.⁵⁹ United States Army regulations adopt a materially identical formulation,⁶⁰ as does the Commander's Handbook of the US Navy, Marine Corps and Coast Guard, which extends the definition to 'weapons and weapons systems' intended to have an effect of 'injuring, damaging, destroying, or disabling personnel or property'.⁶¹ Comparable definitions appear in the guidance or manuals of Belgium, Denmark, Estonia, Germany, New Zealand, Norway and Switzerland, each anchoring the concept of a weapon in its capacity to kill, injure, damage or destroy.⁶² Although neither the International Court of Justice nor the

⁵⁹ Australian Department of Defence, *Australia's Guide to the Legal Review of New Weapons, Means or Methods of Warfare* (10 April 2024) (hereinafter Australian Guide) sec 5(a).

⁶⁰ Department of the Army, *Legal Review of Weapons and Weapon Systems*, Army Regulation 27–53 (23 September 2019) (hereinafter AR 27–53) 8. See also Secretary of the Air Force, *Legal Reviews of Weapons and Cyber Capabilities*, AFI 51-402 (27 July 2011) (hereinafter AFI 51-402) 6.

⁶¹ US Navy, US Marine Corps and US Coast Guard, *The Commander's Handbook on the Law of Naval Operations*, NWP 1-14M/MCTP 11-10B/COMDTPUB P5800.7A (2022) ('Commander's Handbook') sec 9.1, defining weapons and weapons systems as all arms, munitions, materiel, instruments, mechanisms, devices and components required for their operation 'that are intended to have an effect of injuring, damaging, destroying, or disabling personnel or property'. The Handbook excludes launch and delivery platforms from the definition, as does AFI 51-402 (n 60) 6.

⁶² Belgian Ministry of Defence, Commissie voor de Juridische Beoordeling van Nieuwe Wapens en Nieuwe Middelen of Methodes van Oorlogvoering, DGJUR-SPS-CJBCEJ-CXX-001 (21 November 2018) para 2.a; Danish Ministry of Defence, Cirkulære om folkeretlig vurdering i forbindelse med udforskning, udvikling, erhvervelse eller valget af et nyt våben og af metoder eller midler i krigsførelse, CIR1H nr 9494 (29 May 2018) (hereinafter Danish Ministry of Defence) sec 2(1); Estonian Ministry of Defence, Guide to the Implementation of the Legal Analysis of Weapons within the Area of Government of the Ministry of Defence (Annex 2) para 2.1; German Ministry of Defence, Zentrale Dienstvorschrift A-2146/1, Prüfung neuer Waffen, Mittel und Methoden der Kriegführung nach Artikel 36 des I. Zusatzprotokolls von 1977 zu den Genfer Abkommen von 1949 (2016) para 301, reproduced as an annex to Deutscher Bundestag, Drucksache 18/9191 (15 July 2016) 37; New Zealand Defence Force, *Manual of Armed Forces Law: Law of Armed Conflict*, DM 69 (2nd edn, 2022) (hereinafter New Zealand DM 69) vol 4, sec 7.2.1; Norwegian Ministry of Defence, *Manual of the Law of Armed Conflict* (Norwegian Chief of Defence 2013) (hereinafter Norwegian Manual) sec 8.5; Swiss Federal Department of Defence, Civil Protection and Sport, Verordnung des VBS über die Beschaffung, die Nutzung und die Ausserdienststellung von Material (Materialverordnung VBS, MatV), SR 514.20 (26 March 2018) art 11(2)(a)–(b). See, to similar effect, Colombian Ministry of National Defence, *Manual de Derecho Operacional*, FF.MM 3-41 Público (2009); Netherlands Ministry of Defence, *Handboek Humanitair Oorlogsrecht* (2025); French Ministry for Armed Forces, *Manual of the Law of Military Operations* (2022) (hereinafter French Manual); Spanish Ministry of Defence, *Derecho internacional humanitario (DIH) en las FAS*, PDC-02.01 (2022); UK Ministry of Defence, *The Joint Service Manual of the Law of Armed Conflict* (JSP 383, OUP 2004) (hereinafter UK Manual). See also J Kraska and

international criminal tribunals have defined the term expressly, their case-law consistently proceeds on the same understanding.⁶³ Scholarship is to the same effect.⁶⁴

The concept of a ‘weapon system’ is likewise undefined in the law of armed conflict, but it is fairly well settled in military doctrine and scholarship as referring collectively to a weapon and any components, devices and equipment directly necessary for its operation or self-sufficiency.⁶⁵ French doctrine thus describes a weapons system as ‘a combination of interconnected equipment, linked by functional dependencies or the integration of different technologies’,⁶⁶ while Australian and Danish guidance define it, in near-identical terms, as the equipment, materials and delivery mechanisms necessary for a weapon to function as intended.⁶⁷ A weapons system on this understanding comprises two elements: the weapon itself, which causes the relevant violent effect directly, and one or more non-weapon components, such as sensors, launchers, computing and control equipment, communications links, that do not themselves cause that effect, but directly enable the weapon to perform its intended functions.

Nothing in the GGE debates suggests that the Group understands ‘weapon’ and ‘weapon

others (eds), *Newport Manual on the Law of Naval Warfare* (2nd edn, Stockton Center for International Law 2025) (hereinafter Newport Manual) 106; Program on Humanitarian Policy and Conflict Research, *HPCR Manual on International Law Applicable to Air and Missile Warfare* (CUP 2013) (hereinafter Air and Missile Warfare Manual) 49.

⁶³ *Legality of the Threat or Use of Nuclear Weapons* (Advisory Opinion) [1996] ICJ Rep 226, paras 35 and 78; *Prosecutor v Martić* (Trial Judgement) IT-95-11-T, Trial Chamber I (12 June 2007) paras 461–463.

⁶⁴ Boothby, *Weapons and the Law* (n 23) 4; Dinstein (n 23) 2; Marco Longobardo, ‘Means of Combat’ in Dieter Fleck (ed), *The Handbook of International Humanitarian Law* (4th edn, OUP 2021) 129, 129; and Steven Haines, ‘The Developing Law of Weapons: Humanity, Distinction, and Precautions in Attack’ in Andrew Clapham and Paola Gaeta (eds), *The Oxford Handbook of International Law in Armed Conflict* (OUP 2014) 273, 276.

⁶⁵ AR 27–53 (n 60) 9, defining a weapon system as ‘the weapon itself, combined with those components required for its operation that permit the weapon to kill, injure, disable or incapacitate enemy personnel or to destroy, damage, or incapacitate enemy facilities, equipment, property or materiel’; North Atlantic Treaty Organization Standardization Office, NATOTerm, ‘weapons system’; Commander’s Handbook (n 61) sec 9.1; Air and Missile Warfare Manual (n 62) 50; Michael N Schmitt and Jeffrey S Thurnher, ‘“Out of the Loop”: Autonomous Weapons Systems and the Law of Armed Conflict’ (2013) 4 *Harvard National Security Journal* 231, 234; and McFarland (n 2) 34–35.

⁶⁶ Ministère des Armées, *Maintien en condition opérationnelle*, Publication interarmées PIA-4.17_MCO(2020), N° 47/ARM/CICDE/NP (17 March 2020) 10, fn 1 (defining a ‘système d’armes’ as a system ‘ayant vocation à obtenir un effet militaire plus ou moins direct’).

⁶⁷ Australian Guide (n 59) sec 5(b); Danish Ministry of Defence (n 62) sec 2(2). See also Swiss Federal Department of Defence, Civil Protection and Sport (n 62) art 11(1) and (2)(d) (treating weapons, means and methods of warfare collectively as ‘Waffensysteme’ and extending the review to platforms capable of deploying them).

system’ to mean anything other than the mainstream definitions set out above. The GGE has consistently listed target engagement as one of the critical functions of AWS. To engage a target is to execute an attack against it, that is to direct violence at a person or an object. The ICRC’s foundational definition made the point expressly, describing the engagement limb of its formula as the capacity to ‘attack (i.e. intercept, use force against, neutralise, damage or destroy) targets without human intervention’.⁶⁸ The United States has described the function in similar terms: ‘[i]n armed conflict, weapons are used to engage the adversary with force in order to kill, injure, damage, and destroy enemy forces’, the essence of an AWS being that the system can perform this function autonomously, so that ‘[a]fter activation, the system can, without any operator intervention, initiate and execute all actions necessary to engage the target’.⁶⁹ The Chair’s paper of September 2021 likewise tied the critical functions to the application of force, describing them as the functions of ‘selecting and engaging to apply force against targets’.⁷⁰ No participant has proposed a rival understanding of the term and no formulation put to the Group has ever made the engagement function optional. The closest the Group has come to disagreement is in recurring exchanges over whether the systems under discussion must be capable of causing death, but these relate to the ‘lethal’ qualifier attached to the system, not to the meaning of the engagement function.⁷¹

Beginning with the rolling text of 8 November 2024, the GGE has described a lethal autonomous weapon system as ‘an integrated combination of one or more weapons and technological components’ that can ‘identify and/or select, and engage’ a target without human intervention.⁷² The description of AWS as a combination of weapons and technological components has persisted, with variations, through the rolling texts of 12 May 2025 and 18 December 2025,⁷³ and remains current in the rolling text of 5 June 2026, which characterizes a

⁶⁸ International Committee of the Red Cross, ‘Statement to the CCW Meeting of Experts on Lethal Autonomous Weapons Systems’ (Geneva, 13 April 2015) (hereinafter ICRC Statement of 13 April 2015) 2.

⁶⁹ United States (n 24) paras 14–15.

⁷⁰ Revised Chair’s Paper (n 56) para (a).

⁷¹ Chairperson of the GGE on LAWS, ‘Chair’s Summary – Second 2025 Session of the GGE on LAWS’ (20 October 2025) UN Doc CCW/GGE.1/2025/WP.9, para 12. A definition of ‘lethal’ as referring to ‘the capacity to cause death to a person’ appeared in the rolling text of 12 May 2025, sec I, para 2. It was deleted in the revised version of Box I circulated on 3 September 2025, which also substituted ‘loss of life’ for ‘death’, the Chairperson explaining that it was considered unnecessary to define lethality in terms of causing death (para 16 of the same summary).

⁷² Chairperson of the GGE on LAWS, ‘Revised Rolling Text’ (8 November 2024) (hereinafter Revised Rolling Text, 8 November 2024) sec I, para 1.

⁷³ Chairperson of the GGE on LAWS, ‘Revised Rolling Text’ (12 May 2025) (hereinafter Revised Rolling Text, 12 May 2025) sec I, para 1; Chairperson of the GGE on LAWS, ‘Rolling Text’ (18 December 2025) (hereinafter

lethal autonomous weapon system as ‘a combination of one or more weapons and functionally integrated technological components that can identify, select, and engage a target, without intervention by a human operating the system’.⁷⁴ The reference to ‘technological components’ tracks the mainstream distinction between a weapon and the non-weapon components of a weapons system described above. It renders explicit what the undifferentiated references to a ‘weapon system’ during the 2014–2023 sessions had left implicit, namely that the autonomous functions the GGE is concerned with may be distributed across several components of a system. In other words, the critical autonomous functions may be performed either by the weapon or by the non-weapon component of an AWS, or indeed by both. How exactly those functions are distributed will vary from case to case, depending on the type of weapons system concerned.

2.4. Target selection by AWS

In principle, autonomous target selection by an AWS could take one of three forms.⁷⁵ At one extreme, an AWS could be tasked with completing the entire target selection process autonomously, running the full sequence from target discovery through target development and quality control to formal validation and, finally, positive identification immediately before engagement. However, this scenario is not realistic in a deployed system at present. The complete automation of target selection is not taking place even at headquarters level, where computing power, data and time are most abundant. As the preceding section illustrates, automation to date has been directed at discrete tasks within the targeting cycle rather than at combining all functions end to end.⁷⁶ A deployed weapon system tasked with completing the entire target selection process would need on-board computing capacity approaching that of an intelligence fusion cell, real-time access to the wider networked intelligence picture and the capacity to perform qualitative judgments that resist automation because they depend on an assessment of source reliability and context that current machine learning techniques cannot reliably replicate.⁷⁷

The most frequently cited candidate for this scenario is the reported use of Turkish-made

Rolling Text, 18 December 2025) sec I, para 1.

⁷⁴ Rolling Text, 5 June 2026 (n 6) para 1.

⁷⁵ Comparable taxonomies distinguishing degrees of autonomy in target selection appear in United States, ‘Implementing International Humanitarian Law in the Use of Autonomy in Weapon Systems’ (28 March 2019) UN Doc CCW/GGE.1/2019/WP.5, para 5 (three general scenarios for the use of autonomous functions); Boulanin and Verbruggen (n 1) 26–27; and McFarland (n 2) 34–35.

⁷⁶ See section 2.2 above.

⁷⁷ Blanchard and Bruun (n 22) 8 (identifying reliability as a core risk factor for AI-enabled targeting systems); Dorsey and Bo (n 48) 219–220 (automation bias and the limits of algorithmic corroboration).

Kargu-2 loitering munitions against retreating forces in Libya in March 2020, which a United Nations Panel of Experts described as having a true ‘fire, forget and find’ capability.⁷⁸ Even taken at face value, the Panel’s account does not establish that target development, quality control and formal validation were in fact performed on board.⁷⁹ For now, full-cycle autonomy in a deployed weapon system is not operationally realistic due to the limits of on-board processing power and networked connectivity.

At the other extreme, an automated system could be tasked merely with engaging a unique target that has been preselected at an earlier point in the targeting cycle. Functionally similar weapons systems are commonplace. These include precision-guided missiles with terminal seekers and other ‘lock-on-after-launch’ homing munitions directed against a specific, predetermined target.⁸⁰ Some of these weapons, such as missiles which navigate to coordinates fixed in advance and deliver their effect there, have no target selection capability at all.⁸¹ More sophisticated weapons, such as anti-radiation missiles fired at a specific air defence radar,⁸² may be able to detect certain features of the preselected target and rely on these to increase their own accuracy. Although such systems are able to refine their point of impact, they do not select the target of the attack or prioritize between different candidates.⁸³ The same is true for weapons that merely react passively to their targets. For example, a mine is ‘designed to be exploded by the presence, proximity or

⁷⁸ United Nations Security Council, Letter dated 8 March 2021 from the Panel of Experts on Libya established pursuant to Resolution 1973 (2011) addressed to the President of the Security Council, S/2021/229 (8 March 2021) para 63 (‘fire, forget and find’).

⁷⁹ See the account of the manufacturer’s and independent commentators’ responses in Arthur Holland Michel, ‘Are Autonomous Weapons Systems the Scariest Threat in the 2021 UN Coup Attempt Report?’ *Bulletin of the Atomic Scientists* (10 June 2021).

⁸⁰ Examples include coordinate-seeking munitions such as the Joint Direct Attack Munition (JDAM) and cruise missiles navigating to pre-set coordinates, which have no target selection capability, and lock-on-after-launch homing munitions such as anti-radiation missiles: see Boulanin and Verbruggen (n 1) 47–49; and Boothby, *Weapons and the Law* (n 23) 247.

⁸¹ The classic example is the AGM-88 High-Speed Anti-Radiation Missile (HARM), which homes passively on the emissions of a pre-identified radar. Such weapons refine the point of impact against a target already selected by the operator, but do not themselves select the target or choose between candidates: see Boothby, *Weapons and the Law* (n 23) 247; and Boulanin and Verbruggen (n 1) 48–49.

⁸² United States (n 24) para 20(a) and (c).

⁸³ In a lock-on-after-launch engagement, ‘the operator generally fires the weapon in circumstances in which the operator anticipates that weapon will find and “lock on” to the target that the operator intends to strike’: CCW/GGE.1/2025/WP.6 (n 24) para 20(b). Where the release parameters are widened so that the seeker may acquire any member of a target class within a search area, the system shades into the intermediate configuration considered below.

contact of a person or vehicle’, while a booby-trap is likewise actuated by its victim.⁸⁴ Typically, such devices apply no discriminating criteria and make no comparative choice, but merely react to environmental stimuli.⁸⁵

Systems of the type just described are widely regarded within the GGE as falling outside the definition of AWS.⁸⁶ The GGE has treated mines and booby-traps as adequately governed by Amended Protocol II of the CCW.⁸⁷ In more general terms, the rolling text of June 2026 declares in subparagraph 1A that a system is not included in the characterization of AWS ‘when the personnel planning, deciding upon, or carrying out the attack determines the specific target or group of targets to be engaged by the system’.⁸⁸ The reason for this exclusion is straightforward: where a system is tasked with attacking a unique, preselected target at a known location, no target selection is involved at all, since the target would have had to be identified, developed, validated and approved before the system’s activation, leaving the system merely to deliver the kinetic effect. Even if the system were to use sensors to detect the target in order to increase its accuracy, it would not be selecting the target in the sense of exercising choice between possible alternatives.

Whereas the first scenario is not feasible at present, the second involves no autonomous target selection at all. Between these extremes lies a wide range of intermediate configurations, in

⁸⁴ Protocol on Prohibitions or Restrictions on the Use of Mines, Booby-Traps and Other Devices as amended on 3 May 1996 (3 May 1996) 2048 UNTS 93 (hereinafter Amended Protocol II) art 2(1) and (4). A booby-trap is ‘designed, constructed or adapted to kill or injure’ and ‘functions unexpectedly when a person disturbs or approaches an apparently harmless object or performs an apparently safe act’: Amended Protocol II, art 2(4).

⁸⁵ A mine is ‘a munition placed under, on or near the ground or other surface area and designed to be exploded by the presence, proximity or contact of a person or vehicle’, and a booby-trap is ‘any device or material which is designed, constructed or adapted to kill or injure, and which functions unexpectedly when a person disturbs or approaches an apparently harmless object or performs an apparently safe act’: Amended Protocol II, art 2(1) and (4). See also Boothby, *Weapons and the Law* (n 23) 163–164.

⁸⁶ Chairperson of the GGE on LAWS, ‘Chair’s Summary – First 2026 Session of the GGE on LAWS’ (1 April 2026) UN Doc CCW/GGE.1/2026/WP.2 (hereinafter CCW/GGE.1/2026/WP.2) paras 13 and 23; CCW/GGE.1/2025/WP.6 (n 24) para 20(d) (‘it makes sense not to include them within the characterization of LAWS’).

⁸⁷ United States (n 24) para 17, observing that discussions in the GGE have not focused on landmines or booby-traps, ‘for which the CCW Amended Protocol II already provides a sound legal framework’.

⁸⁸ Rolling Text, 5 June 2026 (n 6) sub-para 1A, second sentence. Sub-paragraph 1A first appeared in the rolling text of 18 December 2025, which accompanied the reintroduction of ‘identify’ as a distinct critical function: Rolling Text, 18 December 2025 (n 73) Box I, para 1 and sub-para 1A; CCW/GGE.1/2026/WP.2 (n 86) para 12. The sentence quoted in the text was added only on 5 March 2026, in the form ‘chooses the specific target(s) to be engaged by the system’ and reached its present wording in the rolling text of 5 June 2026: para 20 of the same summary.

which some target selection tasks are performed by human operators before activation and others are delegated to the autonomous system. The paradigmatic example is a system tasked with finding and engaging specific instances of a category of targets preselected by a human operator, such as a certain type of enemy vehicle. The category and its membership criteria having been fixed in advance, the system's task is to search a designated area for entities matching those criteria, determine whether a detected entity is in fact a member of the category, treat a positive match as validation and approval for engagement and then, in a separate step, engage the entity. The IAI Harpy loitering munition illustrates the process.⁸⁹ Once launched into a predefined operational area, the weapon searches autonomously for radiating targets, compares detected emissions against a library of known hostile signatures and commits to engage any instance that matches, without further human confirmation.⁹⁰ Such a system exercises genuine target selection functions, since it performs a scan-and-match action to find individual candidate entities and confirm their membership of the category of military objectives it is designed to engage. Depending on the system and the operational context, it may also have to perform additional tasks, such as prioritizing between multiple qualifying candidates detected within the same area.

In principle, AWS could carry out a wide range of target selection functions autonomously. In practice, they are more likely to perform a narrower range towards the tail end of the targeting cycle, since they lack the technical capabilities needed for the full sequence of functions that characterizes deliberate targeting. Moreover, since their primary purpose is to engage members of preselected categories of targets, they need perform only those selection functions necessary to find, fix and engage individual members of the category, a task corresponding far more closely to the dynamic cycle than to the deliberate one. The purpose of the next section is to consider how the GGE has grappled with this division of labour.

3. TARGET SELECTION AT THE GGE

Target selection has been treated as one of the 'critical functions' of AWS from the earliest discussions of autonomous systems within the GGE framework. The concept of critical functions was introduced by the ICRC to serve a definitional purpose.⁹¹ Modern weapons systems

⁸⁹ On the IAI Harpy, see McFarland (n 2) 36–37 and 124; 'Israel Aerospace Industries Harpy Loitering Munition' (Automated Decision Research) <<https://automatedresearch.org/weapon/israel-aerospace-systems-harpy-loitering-munition/>> accessed 19 August 2026; Boulanin and Verbruggen (n 1) 53–54.

⁹⁰ McFarland (n 2) 36–37; Paul Scharre, *Army of None: Autonomous Weapons and the Future of War* (WW Norton 2018) 27–30.

⁹¹ International Committee of the Red Cross, 'Report of the ICRC Expert Meeting on Autonomous Weapon

incorporate automation across a wide range of functions, including navigation, flight control, sensor management, communications, self-diagnostics and battle damage assessment.⁹² Not all of these raise ethical or legal concerns. The critical functions concept is meant to identify those automated capabilities that do.⁹³ The GGE has endorsed this definitional logic in its work. For example, the Chair's paper of September 2021 characterized a weapon system as autonomous 'if it can, through the use of sensors, computers and algorithms, perform the critical functions of selecting and engaging to apply force against targets without intervention by a human operator or without permanent human involvement or control'.⁹⁴

The consensus that target selection is a critical function has not, however, translated into an agreement about the meaning of the term. The debate within the GGE has revolved around several axes. The first concerns the scope of the concept. While it has proven straightforward to distinguish target selection from target engagement, with the latter denoting the actual execution of an engagement decision through the application of violence against a person or object,⁹⁵ it has been considerably more difficult to agree on what upstream activities target selection encompasses. The debate turns largely on whether target selection is an inclusive concept that includes the identification of targets, broadly tracking the meaning of the term in military doctrine,⁹⁶ or whether it is a narrower concept that excludes target identification. The second axis concerns whether AWS

Systems: Technical, Military, Legal and Humanitarian Aspects, 26–28 March 2014' (Geneva, 9 May 2014) (hereinafter ICRC Expert Meeting Report) 1; International Committee of the Red Cross, 'Statement to the CCW Meeting of Experts on Lethal Autonomous Weapons Systems' (13 May 2014) (hereinafter ICRC Statement of 13 May 2014) 1. See also International Committee of the Red Cross, 'Statement to the CCW Meeting of Experts on Lethal Autonomous Weapons Systems' (14 May 2014) (hereinafter ICRC Statement of 14 May 2014) 1 ('the process of target acquisition, tracking, selection, and attack').

⁹² See Boulain and Verbruggen (n 1) 20–21 (identifying five capability areas in which autonomy is already used in existing military systems: mobility, targeting, intelligence, interoperability and health management).

⁹³ Switzerland, 'Statement to the CCW Group of Governmental Experts on Lethal Autonomous Weapons Systems' (27 August 2018) 3 ('What matters is not "autonomy" per se. What matters is the nature of the tasks that are performed autonomously.'). 'Report of the 2015 Informal Meeting of Experts on Lethal Autonomous Weapons Systems' (2 June 2015) UN Doc CCW/MSP/2015/3, para 37(c)(i) (recording the expert view that '[i]n the search for defining elements of LAWS, it may be useful to focus on the "critical functions" of selecting and attacking targets'); and A Hoepflinger, 'Autonomy and the Critical Functions of Weapon Systems' (Presentation to the CCW Informal Meeting of Experts on Lethal Autonomous Weapons Systems, April 2016) (distinguishing perception and detection functions from the critical select/decide and attack functions).

⁹⁴ Revised Chair's Paper (n 56) para (a).

⁹⁵ See section 3.1 above; United States (n 24) paras 14–15; ICRC Statement of 13 April 2015 (n 68) 2.

⁹⁶ See section 2.1 above.

must exhibit autonomy in their critical functions cumulatively or disjunctively. Fewer systems will qualify as AWS if they must be capable of performing all critical functions autonomously pursuant to a cumulative approach, compared to only having to perform some of those functions under the disjunctive approach. A third axis revolves around the appropriateness of automating these critical functions at all and, if so, under what conditions. This last axis connects the definitional debate to the parallel discussions about meaningful human control and context-appropriate human involvement. For present purposes, we will focus on the first two axes.

3.1. The scope of target selection

The first axis of the debate has dominated the definitional discussion within the GGE. In the initial phase, running from 2014 to 2018, the prevailing characterizations of AWS adopted a wide understanding of selection. At first, the ICRC proposed a four-step sequence of critical functions consisting of ‘acquiring, tracking, selecting and attacking targets’.⁹⁷ This treated the informational activities preceding the actual selection of a target as distinct critical functions.⁹⁸ However, the ICRC’s consolidated definition of 2015 folded target identification into the concept of target selection: to ‘select’ a target was defined as to ‘search for or detect, identify, track’ it.⁹⁹ The ICRC carried this formula through in its submissions of 2017 and 2018,¹⁰⁰ though not always consistently.¹⁰¹ Some States, such as Switzerland, adopted the ICRC position verbatim,¹⁰² while others were more hesitant. This included the United Kingdom, which in 2018 drew a line between target detection and target selection, suggesting that systems such as the Phalanx close-in weapon system are able to detect a target based on parameters specified by the controller, but cannot both

⁹⁷ ICRC Expert Meeting Report (n 91) 1, 3; ICRC Statement of 13 May 2014 (n 91) 1. See also ICRC Statement of 14 May 2014 (n 91) 1 (‘the process of target acquisition, tracking, selection, and attack’).

⁹⁸ ICRC Expert Meeting Report (n 91) 1, 3.

⁹⁹ ICRC Statement of 13 April 2015 (n 68) 2.

¹⁰⁰ See, eg, ICRC Ethics Paper (n 8) para 2 and fn 3, restating the working definition of an autonomous weapon system as one that can ‘select (i.e. search for or detect, identify, track, select) and attack (i.e. use force against, neutralize, damage or destroy) targets without human intervention’.

¹⁰¹ International Committee of the Red Cross, ‘Closing Statement to the CCW Meeting of Experts on Lethal Autonomous Weapons Systems’ (17 April 2015) 1 (‘the critical functions of identifying, selecting, and attacking targets’).

¹⁰² Switzerland, ‘Towards a “Compliance-Based” Approach to LAWS’ (Informal Working Paper submitted to the CCW Informal Meeting of Experts on Lethal Autonomous Weapons Systems, Geneva, 11–15 April 2016, 30 March 2016) para 3.

detect and select a target based on their own reasoning.¹⁰³ In the UK's view, detection is technically automatable and normatively unproblematic in a way that target selection may not be.¹⁰⁴

A second phase opened with the Group's first formally agreed conclusion, adopted in 2019, which declared autonomy 'in the identification, selection or engagement of a target' to be among the essential characteristics of AWS.¹⁰⁵ Two features of this formulation stand out. First, the conclusion named identification as an element distinct from selection, notwithstanding that the ICRC's parallel working definition continued, at this point, to subsume the former within the latter. Second, the three functions were joined disjunctively by 'or', rather than cumulatively by 'and', foreshadowing the second axis of the debate addressed below. The tripartite enumeration of identification, selection and engagement as distinct functions set the template for the years that followed. Submissions by individual participants and groups of States largely tracked this tripartite structure between 2019 and 2023, though with important variations.¹⁰⁶

¹⁰³ United Kingdom, 'Statement for Discussion on Characterisation of the Systems under Consideration' (Meeting of the Group of Governmental Experts on Lethal Autonomous Weapons Systems, 10 April 2018) paras 2 and 4.

¹⁰⁴ United Kingdom, 'Statement for Discussion on the Human Element in the Use of Lethal Force and Aspects of Human-Machine Interaction' (Meeting of the Group of Governmental Experts on Lethal Autonomous Weapons Systems, 11 April 2018) para 3.

¹⁰⁵ Report of the 2019 Session (n 57) para 19(a). The disjunctive connector is noteworthy: at the first 2025 session, the cumulative camp invoked 'the language in para 19 (a) of the 2019 GGE report' as authority for joining the functions with 'and' rather than 'and/or', yet para 19(a) as adopted in fact reads 'or': see Chairperson of the GGE on LAWS, 'Chair's Summary – First 2025 Session of the GGE on LAWS' (7 April 2025) UN Doc CCW/GGE.1/2025/WP.1, para 14.

¹⁰⁶ Panama, 'Commentary on the 11 Guiding Principles' in Chairperson of the GGE on LAWS, 'Chairperson's Summary' (19 April 2021) UN Doc CCW/GGE.1/2020/WP.7, Annex III, 64, 65 ('the critical functions to identify, select and engage targets'); Austria, Belgium, Brazil, Chile, Ireland, Germany, Luxembourg, Mexico and New Zealand, 'Joint Commentary', UN Doc CCW/GGE.1/2020/WP.7, Annex III, 23, 25, paras 14(b) and (d) (twice referring to 'the identification, selection and engagement of targets'), reiterated in the Austria-led submission of 2021; Brazil, Chile and Mexico, 'Elements for a Future Normative Framework Conducive to a Legally Binding Instrument to Address the Ethical Humanitarian and Legal Concerns Posed by Emerging Technologies in the Area of (Lethal) Autonomous Weapons (LAWS)', written contribution submitted to the GGE on LAWS (June 2021) sub-para 2.2.3; Australia, Canada, Japan, the Republic of Korea, the United Kingdom and the United States, 'Principles and Good Practices on Emerging Technologies in the Area of Lethal Autonomous Weapons Systems' (8 August 2022) UN Doc CCW/GGE.1/2022/WP.2 (hereinafter *Principles and Good Practices*) para 9; Draft Articles on Autonomous Weapons Systems (n 57) preamble, first recital, although draft art 2(2) of the same paper joins the functions disjunctively; reissued as UN Doc CCW/GGE.1/2024/WP.10 (26 August 2024). Japan endorsed the triad verbatim in UN Doc CCW/GGE.1/2024/WP.8 (24 July 2024) para 5.

The consolidated rolling texts produced since 2024 have oscillated between the inclusive and the narrow positions. The Group's first provisional consensus characterization, recorded in the rolling text of 28 August 2024, retained the cumulative three verbs to describe an AWS as a system that can 'identify, select and engage a target, or support the accomplishment of these tasks'.¹⁰⁷ The revised text of 8 November 2024 dropped the support clause and introduced a disjunctive qualification, characterizing an AWS as an integrated combination of weapons and technological components that enable the system 'to identify and/or select, and engage a target'.¹⁰⁸ The Chair's summary records note that discussions during the first session of 2025 centred on the exact interpretation of the words 'identify' and 'select'.¹⁰⁹ The rolling text of 12 May 2025 then abandoned identification as a distinct critical function altogether, reducing the formula to a system 'that can select and engage a target'.¹¹⁰ The Chair justified this consolidation on the basis that identification is typically a precondition for the function of selection and that 'select', as used in the characterization, encompasses identification.¹¹¹

This consolidation did not hold. At the second session of 2025, many delegations supported the streamlined formula, arguing that including target identification within the notion of target selection avoided an overly restrictive approach and minimized potential loopholes.¹¹² However, other States maintained that 'identification and selection represent two distinct phases in the targeting process' and expressed concern that the streamlined wording could blur the boundary between AWS and other weapons systems not under the Group's consideration.¹¹³ A proposal to reintroduce 'identify' as a named function was put to the Group,¹¹⁴ coupled with a definition of the identification function.¹¹⁵ The text discussed on the final day of the September

¹⁰⁷ Chairperson of the GGE on LAWS, 'Rolling Text' (28 August 2024) (hereinafter Rolling Text, 28 August 2024) sec 'A working characterization of LAWS'.

¹⁰⁸ Revised Rolling Text, 8 November 2024 (n 72) sec I, para 1.

¹⁰⁹ CCW/GGE.1/2025/WP.1 (n 105) para 14.

¹¹⁰ Revised Rolling Text, 12 May 2025 (n 73) sec I, para 1.

¹¹¹ CCW/GGE.1/2025/WP.9 (n 71) para 10; Chairperson of the GGE on LAWS, 'Additional Thoughts from the Chair on the Critical Functions of LAWS' (Background Paper, 2025) (hereinafter Additional Thoughts from the Chair) 1 ("select" encompasses identification).

¹¹² CCW/GGE.1/2025/WP.9 (n 71) para 11.

¹¹³ CCW/GGE.1/2025/WP.9 (n 71) para 11. See also Sri Lanka, Statement to the first 2026 session of the GGE on LAWS, Boxes 1 and 3 (2–6 March 2026) 1 ("identify" and "select" are two distinct steps in the targeting process).

¹¹⁴ The proposal originated with the ICRC: Additional Thoughts from the Chair (n 111) 1–2. It marks a reversal of the subsumption position that the ICRC had itself established between 2015 and 2018, discussed above.

¹¹⁵ Additional Thoughts from the Chair (n 111) 1–2 (attributing the reintroduction proposal, together with its

2025 session accordingly characterized a lethal autonomous weapon system as one ‘that once activated, can identify, select and engage a target’, and defined identification as ‘the process by which a system scans an environment for potential targets and matches a person or object against a pre-determined target profile’.¹¹⁶

The rolling text of 18 December 2025 restored ‘identify’ as a distinct function on this basis, but dropped the definition of identification itself.¹¹⁷ The Chair’s summary of the first 2026 session confirms that ‘identify’ was reintroduced ‘to reflect delegations’ views that identification constitutes a distinct phase in the targeting cycle’.¹¹⁸ The December 2025 text also introduced subparagraph 1A, clarifying that the involvement of human operators in defining the parameters of the targets to be engaged does not prevent a system from being an AWS. The text thus acknowledges, implicitly, that certain identification functions are carried out by humans and that identification is therefore not a fully automated function.¹¹⁹ At the first session of 2026, delegations welcoming the reintroduction of ‘identify’ argued that it reflects the targeting cycle more accurately and strengthens alignment with the law of armed conflict, particularly the principle of distinction. Others cautioned that the clause as formulated risked an overly expansive characterization that extended to existing weapons systems, such as fire-and-forget and homing missiles, which can be used in compliance with the law of armed conflict.¹²⁰

The latest version of the rolling text of 5 June 2026 maintains the three-verb formula, characterizing a lethal autonomous weapon system as a combination of one or more weapons and

definition of identification and the ‘does not exclude’ clause, to the ICRC); CCW/GGE.1/2025/WP.9 (n 71) para 14. For the ICRC’s earlier subsumption position, see nn 48 and 86.

¹¹⁶ CCW/GGE.1/2025/WP.9 (n 71) para 17, reproducing the text discussed on 5 September 2025: para 1 (‘... that once activated, can identify, select and engage a target, without intervention by a human operator in the execution of these tasks’) and para 1.bis (‘For these purposes, the identification function can be defined as the process by which a system scans an environment for potential targets and matches a person or object against a pre-determined target profile. The fact that a human operator sets the target(s) that the system is programmed to select and engage, or the parameters that comprise a target profile, does not exclude the system from the characterization.’).

¹¹⁷ Rolling Text, 18 December 2025 (n 73) Box I, para 1 and sub-para 1A. The scan-and-match definition of identification proposed in September 2025 was not carried into the December 2025 text, which retained only the anti-exclusion clause.

¹¹⁸ CCW/GGE.1/2026/WP.2 (n 86) para 12.

¹¹⁹ CCW/GGE.1/2026/WP.2 (n 86) para 12 (recording that sub-paragraph 1A was intended to address concerns that the inclusion of ‘identify’ would lead to overly narrowing the scope of the characterization or creating potential loopholes).

¹²⁰ CCW/GGE.1/2026/WP.2 (n 86) para 13.

functionally integrated technological components ‘that can identify, select, and engage a target, without intervention by a human operating the system’.¹²¹ Subparagraph 1A clarifies that the fact that a human designs, programmes or otherwise determines a set of potential targets or their characteristics does not exclude a system from that characterization, but adds that a system is excluded ‘when the personnel planning, deciding upon, or carrying out the attack determines the specific target or group of targets to be engaged by the system’.¹²²

The present position, then, is that the GGE’s working characterization of AWS treats identification as a critical function in its own right, distinct from selection and engagement. This raises an immediate question: understood as distinct critical functions, what exactly do target identification and target selection entail?

3.2. Target identification as a distinct critical function

The GGE has not adopted an agreed definition of identification. Rather, as the Chair has acknowledged, there are ‘multiple versions of identification’ at play in the GGE debates.¹²³ It seems that the term has been used in at least five different senses.

On the weakest reading, to identify means little more than to detect an entity through sensors. For example, the United States noted that systems tasked with identifying and selecting targets must necessarily be equipped with ‘sensors, such as cameras or radar, that enable the system to detect possible targets in the environment’.¹²⁴ Detection in this sense involves the analysis of sensory inputs to identify the presence of discrete entities against environmental background noise, but without establishing the unique identity of those objects.

Identification has also been understood in a specific sense to mean ‘classifying the target as the intended target’, that is confirming that an entity encountered in the battlespace is the unique, predetermined target that the system was tasked to strike.¹²⁵ So understood, identification simply means confirmation of the unique identity of a preselected target and thus collapses into the terminal step of approving the target for engagement. This is why in the rolling text of 3 September 2025, identification in this sense was understood to be subsumed by the function of

¹²¹ Rolling Text, 5 June 2026 (n 6) para 1.

¹²² Rolling Text, 5 June 2026 (n 6) sub-para 1A. For the March 2026 redrafting of the sub-paragraph, see CCW/GGE.1/2026/WP.2 (n 86) paras 19–20.

¹²³ Additional Thoughts from the Chair (n 111) 1 (recording ‘uncertainty about what “identification” in this sense would be, as there are understood to be multiple versions of identification’).

¹²⁴ CCW/GGE.1/2025/WP.6 (n 24) para 18.

¹²⁵ Additional Thoughts from the Chair (n 111) 1.

target selection.¹²⁶

The now dominant sense equates identification with the classification of sensed data against a stored target profile. This is the sense captured in the definition proposed to the Group in September 2025, which described the identification function as ‘the process by which a system scans an environment for potential targets and matches a person or object against a pre-determined target profile’.¹²⁷ Subparagraph 1A of the latest rolling text also presupposes this sense when it provides that human determination of ‘a set of potential targets or characteristics thereof’ does not exclude a system from the characterization.¹²⁸ Identification in this sense is a wider function than the bare detection of an entity against environmental noise, because it encompasses a scan-and-match operation performed over an open population of target candidates to establish that the entity concerned is a member of a certain category of targets. Identification in this sense also differs from the simple identification of a unique, predetermined target, since it establishes membership of a wider category of targets.

Another usage of the term focuses on the identification tasks undertaken before an autonomous system is activated and deployed, including constructing the target profile and membership criteria for the category of targets it is meant to engage. Because these tasks are carried out before deployment, they are necessarily performed by humans. Subparagraph 1A of the rolling text uses identification in this temporal sense when it stipulates that the advance determination of the set of potential targets, or of the characteristics defining that set, does not prevent a system from being characterized as an AWS.¹²⁹

Finally, in a legal sense, identification refers to the legal classification of a target as a military objective liable to attack.¹³⁰ It is this sense that delegations at the first session of 2026 invoked when they welcomed the reintroduction of ‘identify’ on the ground that it ‘strengthens alignment with international humanitarian law, particularly the principle of distinction’.¹³¹ In the deliberate targeting cycle, the legal classification of a target takes place at the point of target validation, which

¹²⁶ Additional Thoughts from the Chair (n 111) 1; Chairperson of the GGE on LAWS, ‘Revised Version of Box I of the Rolling Text’ (3 September 2025) para 1 (‘select and engage a target’).

¹²⁷ CCW/GGE.1/2025/WP.9 (n 71) para 17, Box I, para 1.bis. The Chairperson’s background paper renders the same proposal with the variant wording ‘pre-programmed target profile’: Additional Thoughts from the Chair (n 111) 1–2. The definition was not carried into the consolidated texts of December 2025 and June 2026.

¹²⁸ Rolling Text, 5 June 2026 (n 6) para 1, sub-para 1A.

¹²⁹ Rolling Text, 5 June 2026 (n 6) sub-para 1A, first sentence. The division of labour this codifies is developed further below.

¹³⁰ Additional Thoughts from the Chair (n 111) 1.

¹³¹ CCW/GGE.1/2026/WP.2 (n 86) para 13.

ensures, among other things, compliance with the law of armed conflict.¹³² Legal classification finds another expression in the form of positive identification, which NATO doctrine defines as ‘[a] recognition derived from observation and analysis of target characteristics’ that provides ‘reasonable certainty that a functionally and geospatially defined entity is a valid military objective’, just before the entity is engaged.¹³³

Surveying this usage as a whole, no settled meaning of ‘identify’ emerges from the GGE record, but only a cluster of different meanings that include detection, confirmation of the intended target, profile-matching, pre-deployment construction of target categories and legal classification. The Chairperson’s 2025 background paper on the critical functions goes some way towards remedying this terminological diversity. Rather than imposing a single meaning, it distinguishes three types of identification:¹³⁴

‘1. identification that takes place in the development phase (so pre-activation); 2. identification that takes place after activation and encompasses it in a broad sense, so it would include everything from finding “new targets” (thereby possessing the ability to determine the target to be a legitimate military target or not, so the IHL-based identification), to matching preprogrammed general targets to a specific target that the system encounters, to simply recognizing the intended specific target; 3. identification that only does the latter, so simply classifying the target as the intended target.’

The taxonomy’s premise is that identification is an epistemic activity ascribing a status to an entity and that it does not occur at a single point in the targeting process. Pre-activation identification corresponds to the early phases of the deliberate cycle, which define target sets, construct target profiles and set the criteria by which membership of a category of targets is determined.¹³⁵ The second type includes finding new targets, which corresponds to target discovery, that is ‘the process of locating and identifying entities and objects of interest in the operating environment’, and to basic target development, which begins ‘the process of uniquely identifying, locating, describing, functionally characterizing’ an entity and lists identification as the first of its five standards.¹³⁶ The ‘IHL-based identification’ corresponds to target validation.¹³⁷ ‘Matching

¹³² AJP-3.9 (n 25) para 1.3.32.

¹³³ AJP-3.9 (n 25) para 1.3.18.

¹³⁴ Additional Thoughts from the Chair (n 111) 1.

¹³⁵ See section 2.1 above; AJP-3.9 (n 25) para 1.5.1(a) (Phase 1 identifies appropriate NAC-approved target sets) and ch 2 on intelligence support to target development.

¹³⁶ AJP-3.9 (n 25) paras 1.3.25 and 1.3.2; see also para 2.3.1 (identification as the first of the five basic target development standards).

¹³⁷ AJP-3.9 (n 25) para 1.3.32 (target validation ensures, inter alia, ‘compliance with relevant international law’); see

preprogrammed general targets to a specific target that the system encounters' corresponds to the application of target selection standards at the tactical level, where 'a target must be positively identified against the target selection standards' before engagement.¹³⁸ The third type, 'simply recognizing the intended specific target', has no direct doctrinal equivalent, being concerned with establishing the unique identity of a preselected target. The taxonomy is therefore broadly aligned with military doctrine in treating identification not as one event, but as a family of activities dispersed across the targeting cycle, from the intelligence work of the earliest phases to the final status confirmation that precedes approval for engagement.

Understood in this way, target identification is a multistage process involving activities performed both before the activation of an AWS and during its deployment. Some pre-activation identification tasks are, by definition, carried out by humans, including selecting the criteria that define membership of a category of targets and constructing the target profiles against which the system will match sensed entities. Other identification tasks are delegable to the system, such as the scan-and-match operation performed against a stored profile and the terminal recognition of the intended target. It follows that identification is a critical function of an AWS, but it is not, even in an AWS, a fully automated one.

3.3. Target selection as a distinct critical function

If target identification and target selection are two distinct critical functions, they must describe different types of activities, otherwise they would not be distinct. The concept of target selection must therefore exclude the epistemic activity that generates information about an entity and ascribes a status to it. Instead, selection in a narrow sense must refer solely to the decisional activity that commits to one course of action in relation to a target in preference to other possible courses of action. The law of targeting itself uses the verb in this narrow sense. The one provision of AP I that speaks of selecting a target, Article 57(3), presupposes that a set of entities have already been identified as military objectives and requires the exercise of a choice between them. International jurisprudence reflects the same approach. The International Criminal Court has defined the concept of 'directing an attack' to mean 'selecting the intended target and deciding on the attack', thus treating selection as a decisional act adjacent to, and of the same kind as, the decision to attack.¹³⁹

Targeting doctrine suggests that the process of selection in a narrow sense begins before

also para 1.5.1(a).

¹³⁸ AJP-3.9 (n 25) para 5.4.2 (within the D3A methodology). For target selection standards, see para 1.3.28.

¹³⁹ *Prosecutor v Ntaganda* (Trial Judgment) ICC-01/04-02/06-2359, Trial Chamber VI (8 July 2019) para 744.

any individual target exists. In the opening phases of an operation, the North Atlantic Council approves the target sets that may enter development,¹⁴⁰ while the joint force commander's guidance designates entities and categories of entities for inclusion on the No-Strike List.¹⁴¹ Once discovery yields entities of interest, selection operates as a series of entry and progression gates within target development. Submitting an entity to the Candidate Target List is a choice, made by nations, component commands or other appropriate authorities, to advance it into development in preference to setting it aside.¹⁴² The phases of target development function as successive decision points ensuring that only entities meeting the relevant operational, legal and other criteria move forward. Target validation then places targets approved for prosecution on the Joint Target List, awaiting nomination.¹⁴³ Other targets may be placed on the Restricted Target List or the No-Strike List.¹⁴⁴ From here, selection continues through nomination, where a choice is made as to which validated targets are put forward for the forthcoming operational period,¹⁴⁵ and prioritization, where the order of engagement is fixed. The sequence culminates in the joint force commander's approval of the Joint Prioritized Target List. In dynamic targeting, where targets of opportunity are prosecuted through the compressed sequence of F2T2E2A, the decisional function reappears at the target step, where the commander holding target engagement authority approves the engagement of the specific detected entity against pre-established criteria,¹⁴⁶ and in combat engagement, where the immediacy of the situation leaves the entire choice, whether or not to engage this entity now, to those executing the engagement.¹⁴⁷

Target selection in the narrow sense is therefore no more a single decision point than identification is. It comprises a chain of choices occurring at successive points across the targeting cycle, each applying information generated by identification and committing to a course of action

¹⁴⁰ AJP-3.9 (n 25) para 1.3.29 ('Target sets are approved by the North Atlantic Council') and sec 3.2.2.

¹⁴¹ AJP-3.9 (n 25) paras 1.3.16 (no-strike list) and 1.5.1(a); see also para 1.5.1(b), fn 23 (noting that the NSL may evolve during the operation and that engagement of entities on it may violate the law of armed conflict).

¹⁴² AJP-3.9 (n 25) para 1.3.3 (candidate target list, '[a] list of entities submitted by nations, component command or appropriate organizations, that are in target development and have not yet been validated').

¹⁴³ AJP-3.9 (n 25) para 1.3.15 (joint target list, '[a] list of validated targets not yet approved for inclusion in the joint prioritized target list').

¹⁴⁴ AJP-3.9 (n 25) paras 1.3.21 (restricted target) and 1.3.22 (restricted target list, noting that the restrictions 'do not change the fact that targets on the RTL are valid military targets').

¹⁴⁵ AJP-3.9 (n 25) para 2.3.8 (target nomination) and para 1.3.3, note 1 (validated CATL entities 'can ... be considered for inclusion on the Target Nomination List').

¹⁴⁶ AJP-3.9 (n 25) sec 5.3 and para 1.3.27.

¹⁴⁷ AJP-3.9 (n 25) para 1.3.5.

regarding the target in preference to alternatives. Some of these acts precede activation and involve exclusively human input. Symmetrically with identification, target selection in the narrow sense is therefore not, even in an AWS, a fully automated function.

3.4. Disjunctive or cumulative

The second axis of the debate concerns whether the enumerated critical functions are joined cumulatively, so that a system must be capable of performing all of them autonomously in order to qualify as an AWS, or disjunctively, so that autonomy in some of them suffices. The choice matters because a cumulative approach would cover fewer systems than the potentially more inclusive disjunctive approach.

The GGE debate has moved back and forth between the two approaches. The Group's first agreed conclusion in 2019 joined 'identification, selection or engagement of a target' disjunctively.¹⁴⁸ A paper submitted by a United States-led group of States in 2022 retained the disjunctive formulation,¹⁴⁹ while a set of draft treaty articles submitted in 2023 by a group that included some of the same States adopted the cumulative approach.¹⁵⁰ The rolling text of 28 August 2024 employed the cumulative triad,¹⁵¹ while the revised text of 8 November 2024 adopted a hybrid connector that characterized AWS as systems able 'to identify and/or select, and engage a target'.¹⁵² The hybrid connector confined the disjunction to the first two functions, rather than to all three. In other words, it suggested that systems qualify as AWS either if they autonomously *identify and engage* a target or if they autonomously *select and engage* a target. As the Chair's summary explains, this formulation did not attract universal support during the first GGE session of 2025:

Some delegations were in favour of a cumulative approach, retaining only the word 'and' between the functions (as per the language in para 19 (a) of the 2019 GGE report), whilst others deemed it essential to keep 'and/or' in order to have a broad scope and avoid potential loopholes. This divide reflects two different visions when it comes to the elements of an instrument that the Group is working on: there are the delegations who think we need a broad

¹⁴⁸ Report of the 2019 Session (n 57) para 19(a).

¹⁴⁹ Principles and Good Practices (n 106) para 9 ('the identification, selection, or engagement of a target'); the same disjunctive formulation appears in draft art 2(2) of the Draft Articles on Autonomous Weapons Systems (n 57) as reissued in UN Doc CCW/GGE.1/2024/WP.10 (26 August 2024).

¹⁵⁰ Draft Articles on Autonomous Weapons Systems (n 57) preamble, first recital ('once activated, can identify, select, and engage targets with lethal force without further intervention by an operator'). Draft art 2(2) of the same paper nonetheless joins the functions disjunctively.

¹⁵¹ Rolling Text, 28 August 2024 (n 107) para 1.

¹⁵² Revised Rolling Text, 8 November 2024 (n 72) sec I, para 1.

characterization to ensure that all LAWS will be covered under the working characterization, whilst other delegations prefer a cumulative approach in order to avoid the risk of prohibiting/regulating weapons systems that are already used in accordance with IHL, and which are, in their view, not LAWS per se.¹⁵³

The May 2025 rolling text dissolved the immediate drafting problem by deleting ‘identify’ altogether and reducing the enumeration to ‘select and engage’, on the express understanding that selection encompasses identification.¹⁵⁴ When ‘identify’ returned in the December 2025 text, it returned within a cumulative triad, and the rolling text of 5 June 2026 maintains that structure, describing AWS as a system ‘that can identify, select, and engage a target, without intervention by a human operating the system’.¹⁵⁵

On its face, then, the cumulative camp has prevailed. The principal virtue of the cumulative approach is that it distinguishes AWS from automated systems that do not deliver kinetic effects directly, such as AI DSS, and from weapons systems that merely execute an engagement decision taken by a human, such as ‘fire-and-forget’ missiles. In short, the cumulative approach avoids an overly inclusive definition of AWS. Against this, the disjunctive camp objects that on the cumulative approach, an automated system needs to defeat only one of the three critical functions to avoid qualifying as an AWS. For example, if target selection means final approval of a target for engagement, then the nominal involvement of a human who simply rubber-stamps machine-proposed targets would prevent the system from qualifying as an AWS.¹⁵⁶

This is a valid concern, but the solution does not lie in reverting to a disjunctive approach. The ‘to identify and/or select, and engage a target’ formula applies, first, to systems that autonomously *identify and engage* a target while leaving all selection choices to human operators. Under the latest rolling text, that configuration is expressly excluded by the second sentence of subparagraph 1A, since it involves humans preselecting targets for machine engagement.¹⁵⁷ The formula applies, secondly, to systems that autonomously *select and engage* a target while leaving all identification activities to humans. It is difficult to envisage how such an arrangement would work in practice. Applied to the deliberate cycle, it would involve a constant back and forth between humans performing all epistemic tasks and the system making all the decisions on the basis of their

¹⁵³ CCW/GGE.1/2025/WP.1 (n 105) para 14.

¹⁵⁴ Revised Rolling Text, 12 May 2025 (n 73) sec I, para 1; CCW/GGE.1/2025/WP.9 (n 71) para 10 (recording that the choice ‘was made with the understanding that identification is typically a precondition for the function of selection’); Additional Thoughts from the Chair (n 111) 1 (“‘select’ encompasses identification”).

¹⁵⁵ Rolling Text, 18 December 2025 (n 73) Box I, para 1; Rolling Text, 5 June 2026 (n 6) para 1 and sub-para 1A.

¹⁵⁶ CCW/GGE.1/2026/WP.2 (n 86) para 23.

¹⁵⁷ Rolling Text, 5 June 2026 (n 6) sub-para 1A, second sentence.

input, which is neither an efficient nor a prudent division of labour. The arrangement makes more sense in dynamic targeting, where the system's choice narrows to the formal approval of the target for engagement. However, if that approval is based entirely on the application of pre-programmed criteria to human-made epistemic assessments, including the positive identification of the unique entity chosen for engagement, we would be looking at a system that is automated to engage targets identified by human operators, rather than one which is exercising choice autonomously. For all intents and purposes, such a system would effectively be engaging preselected targets and therefore should be excluded from the characterization pursuant to the second sentence of subparagraph 1A.

The disjunctive approach runs into these difficulties because it assumes that identification and selection are entirely distinct functions, either of which can be performed in full by a human operator, while the other is performed by an autonomous system. Yet whichever function is retained by a human, the operator is in essence preselecting targets for automated engagement, which excludes the arrangement from the definition of AWS. The difficulty is avoided by abandoning the tripartite enumeration and returning to a characterization of AWS as systems that autonomously select and engage targets, with selection understood inclusively to encompass identification and other targeting functions. This addresses the concern over nominal human involvement and also accounts for functions that otherwise fall through the gaps in the tripartite enumeration. Target development does not fit the enumeration well, as it involves both informational activities and the exercise of choice. Quality control is not a comfortable fit either, since it neither generates new target information, nor chooses a course of action relating to the target.¹⁵⁸ It is therefore best seen as a second-order epistemic activity, assessing the reliability of identification's outputs rather than adding to them. The setting of target selection standards, the construction of target profiles and the framing of engagement criteria are also second-order acts, involving choices about how identification and selection will subsequently operate. Weaponizing, capabilities analysis and collateral damage estimation are different again. They are genuinely decisional, but the object of the choice is not the target, but the means, methods and expected effects of the attack.

Accordingly, there are good reasons to follow the Chairperson's suggestion and adopt an inclusive reading of target selection. Understood in this way, target selection encompasses the informational activities of identification together with the successive choices through which candidate targets are advanced, restricted, excluded or approved for engagement. The definitional

¹⁵⁸ AJP-3.9 (n 25) para 1.3.20.

weight then shifts from the connectors that join the enumerated functions to the substantive question of which target selection tasks are performed by humans before an AWS is activated and which are delegated to the system itself.¹⁵⁹ As the next section will explain, it is this division of labour that determines how the duty of target verification maps onto the use of AWS.

4. THE DUTY OF TARGET VERIFICATION

Having traced the evolution of the GGE efforts to characterize AWS, we can now examine how the duty of target verification maps onto the process of target selection. An autonomous weapon system, on the understanding that has emerged from the GGE debates, selects and engages a target without further human intervention. A system that merely engages a target a human has already preselected falls outside this definition, while full autonomy across the entire targeting cycle remains unrealistic. Systems of practical interest therefore occupy the middle ground between these two extremes, where some target selection tasks are performed by the operator before the system is activated and others are delegated to the system itself. This division of labour has an immediate effect on how the duty of target verification is implemented in such cases.

Article 57(2)(a)(i) AP I requires those who plan or decide upon an attack to ‘do everything feasible to verify that the objectives to be attacked are neither civilians nor civilian objects and are not subject to special protection but are military objectives within the meaning of paragraph 2 of Article 52 and that it is not prohibited by the provisions of this Protocol to attack them’.¹⁶⁰ The rule gives operational effect to the principle of distinction by requiring those who plan or decide upon an attack to apply that principle in their decision-making before force is directed against a target.¹⁶¹

Article 57(2)(a)(i) AP I speaks of verification as if it were a single undertaking designed to confirm that an objective selected for attack meets certain conditions. In reality, target verification

¹⁵⁹ The definitional enterprise has its critics. Yasmin Afina, ‘Combat Code Compliance: International Humanitarian Law and the Development Stages of AI for Targeting’ (PhD thesis, University of Essex 2024) 40–41 objects that the concentration on lethal autonomous weapons systems is unduly narrow, there being ‘not even a universally accepted definition of what LAWS constitute’, while other applications raise as many legal considerations for targeting, and urges that scholarship expand beyond that scope so as to ‘move past definitional concerns’.

¹⁶⁰ Additional Protocol I, art 57(2)(a)(i).

¹⁶¹ Jean-François Quéguiner, ‘Precautions under the Law Governing the Conduct of Hostilities’ (2006) 88 *International Review of the Red Cross* 793, 797; William H Boothby, *The Law of Targeting* (OUP 2012) (hereinafter Boothby, *Law of Targeting*) sec 4.2.11; Corn (n 17) 436–437.

emerges only through a sequence of activities spanning information gathering, reliability assessment, legal characterization and the individuation of the target. Just as target selection is, at present, a multistage process not delegated to autonomous systems in its entirety, so the implementation of the duty of target verification also involves a division of labour between human operators and AWS that stretches across multiple phases of the targeting cycle. The purpose of the present section is to clarify this process. It begins by establishing the specific steps the duty of target verification requires, before relating these steps to the different stages of the targeting cycle. It then assesses how the resulting components of the duty may be distributed between human operators and autonomous weapons systems once the latter are tasked with selecting targets and closes by drawing out some of the implications that follow.

4.1. The nature and content of the obligation

The duty of target verification is a complex obligation. Several features of the rule merit attention in the present context. First, on its face, the language of Article 57(2)(a)(i) implies that the attacker must first verify that the objective to be attacked is not a civilian, a civilian object or an entity subject to special protection, before verifying that it qualifies as a military objective and that it is not prohibited by the Protocol to attack it. This sequence is not logical. Since civilians and civilian objects are defined negatively under the law of armed conflict,¹⁶² a better starting point is to determine whether an entity qualifies as a military objective. The wording also passes over the fact that civilians who take a direct part in hostilities are liable to attack.¹⁶³ In practice, the duty is therefore best applied in two steps. Those who plan or decide upon an attack must do everything feasible to verify, first, that the person or object to be attacked qualifies as a military objective and, second, that no provision of the Protocol prohibits directing an attack against that person or object.

Second, the duty of target verification attaches to the specific persons or objects against which force will be directed, not to a category or class of targets in the abstract. Determining that mobile rocket launchers as a class of objects are military objectives by nature says nothing about whether the particular entity detected by a sensor is in fact a rocket launcher, rather than a lorry.

¹⁶² Additional Protocol I, arts 50(1) and 52(1), each defining the protected category by exclusion from the category of military objectives or of the persons referred to in art 4A of the Third Geneva Convention and art 43 of the Protocol. On the significance of the negative definition for the design of automatic target recognition, see Afina (n 159) 172.

¹⁶³ Additional Protocol I, art 51(3).

It is the latter determination at which Article 57(2)(a)(i) is aimed.¹⁶⁴ This is reflected in the language of the rule, which employs the transitive verb ‘verify’ with a specific direct object, ‘the objectives to be attacked’. This individuated formulation runs unbroken through the drafting history, from the 1958 requirement to make sure that ‘the objective, or objectives, to be attacked’ are military objectives and are ‘duly identified’ to the 1973 draft commentary’s statement that identification is required ‘in a given case’.¹⁶⁵ It is also reflected in the language of other provisions, such as the prohibition of attacks not directed against a *specific* military objective and the rules on doubt, which apply at the level of individual and hence specific persons and objects.¹⁶⁶

National practice, case law and expert manuals confirm this reading. The Danish Military Law Manual insists that verification extends to anything capable of shedding more light on the specific nature of the individual or object.¹⁶⁷ The US Department of Defense Manual requires attackers to confirm that the person or object to be attacked is in fact the intended target.¹⁶⁸ The Norwegian Manual demands that a person who is to be attacked first be identified as a lawful

¹⁶⁴ Agnieszka Jachec-Neale, *The Concept of Military Objectives in International Law and Targeting Practice* (Routledge 2015) 143, quotes Kalshoven for the proposition that the requirement that the definition be satisfied in the circumstances ruling at the time ‘effectively precludes military commanders from relying exclusively on abstract categorisations in determination of whether specific objects constitute military objectives’. The point is common ground among the standard commentaries: eg Michael Bothe, Karl Josef Partsch and Waldemar A Solf, *New Rules for Victims of Armed Conflicts: Commentary on the Two 1977 Protocols Additional to the Geneva Conventions of 1949* (2nd edn, Martinus Nijhoff 2013) art 57, para 2.4.1 (describing the precautionary rules as giving specific guidance on feasible measures for ‘verification of the identity of the object of attack as a military objective’).

¹⁶⁵ International Committee of the Red Cross, *Draft Rules for the Limitation of the Dangers Incurred by the Civilian Population in Time of War* (2nd edn, Geneva, 1958) (hereinafter ICRC Draft Rules) art 8(a) and commentary (suggesting that those who plan an attack must ensure that the target to be attacked ‘belongs to a category of objective of generally recognised military importance’); International Committee of the Red Cross, *Draft Additional Protocols to the Geneva Conventions of August 12, 1949: Commentary* (Geneva, 1973) (hereinafter ICRC Draft Protocols Commentary) 65. See also International Committee of the Red Cross, *Conference of Government Experts on the Reaffirmation and Development of International Humanitarian Law Applicable in Armed Conflicts: Report on the Work of the Conference* (Geneva, 1971) (hereinafter ICRC Conference of Government Experts) 82, para 464 (an attack should be launched only if it is quite clear that the objectives concerned are military in character).

¹⁶⁶ Additional Protocol I, art 51(4)(a) and arts 50(1) and 52(3).

¹⁶⁷ Danish Ministry of Defence, Defence Command Denmark, *Military Manual on International Law Relevant to Danish Armed Forces in International Operations* (2016, updated 2020) (hereinafter Danish Military Manual) ch 8, sec 4, Rule 8.6 and Example 8.37a.

¹⁶⁸ United States Department of Defense, *Law of War Manual* (revised edn, July 2023) (hereinafter DoD Law of War Manual) sec 5.5.3.

target, which in a planned operation entails locating ‘the right person’.¹⁶⁹ In *Strugar*, the shelling of the Old Town of Dubrovnik could not be defended by a bare assertion that enemy forces might have been present within its walls. The Trial Chamber examined whether any credible intelligence supported the classification of the specific target and found that none did.¹⁷⁰ The Israeli Supreme Court proceeded on the same understanding in the *Targeted Killings* case, requiring ‘information which has been most thoroughly verified’ as to the identity and activity of the civilian said to be taking a direct part in hostilities, a demand that can be satisfied only in relation to a particular individual.¹⁷¹ The case law of the international criminal tribunals proceeds on the same footing. In *Dragomir Milošević*, the Appeals Chamber rejected an attempt to establish the lawfulness of attacks by reference to the character of the zones in which they were carried out, holding that the distinction between civilians and combatants and between civilian objects and military objectives must be made on a case-by-case basis and that what fell to be established was the status of the population targeted in specific incidents.¹⁷² In *Karadžić*, the Appeals Chamber stated the test for the attacker in equally individuated terms, asking whether the attacker could reasonably have believed, in the circumstances of the person contemplating the attack and on the information available, that the object was being used to make an effective contribution to military action.¹⁷³

The individuated framing is common to the expert manuals.¹⁷⁴ The same premise underlies the discussions at the GGE. Sub-paragraph 1A of the rolling text withholds the characterization of an autonomous weapon system where the personnel planning, deciding upon or carrying out the attack determine the ‘specific target or group of targets’ to be engaged, so that the line between human and machine selection is drawn by reference to the determination of the concrete object of attack, rather than of the category to which it belongs.¹⁷⁵

Third, the duty is addressed to ‘those who plan or decide upon an attack’. The formulation

¹⁶⁹ Norwegian Manual (n 62) sec 2.8.

¹⁷⁰ *Prosecutor v Strugar* (Trial Judgement) IT-01-42-T, Trial Chamber II (31 January 2005) paras 182–214 and 284–289.

¹⁷¹ HCJ 769/02 *The Public Committee against Torture in Israel v Government of Israel* (13 December 2006) para 40.

¹⁷² *Prosecutor v Dragomir Milošević* (Appeal Judgement) IT-98-29/1-A, Appeals Chamber (12 November 2009) paras 53–55 and 139.

¹⁷³ *Prosecutor v Karadžić* (Appeal Judgement) MICT-13-55-A, Appeals Chamber (20 March 2019) paras 488–489.

¹⁷⁴ Michael N Schmitt (ed), *Tallinn Manual 2.0 on the International Law Applicable to Cyber Operations* (CUP 2017) (hereinafter Tallinn Manual 2.0) commentary to Rule 115, para 2 (verifying the status of ‘the person or object to be attacked’); Newport Manual (n 62) sec 8.9.1.

¹⁷⁵ Rolling Text, 5 June 2026 (n 6) sub-para 1A. The March 2026 draft of the sentence read ‘chooses the specific target(s) to be engaged by the system’: CCW/GGE.1/2026/WP.2 (n 86) para 20.

clarifies that the relevant assessments must be made before the attack is executed, rather than by individual combatants at the moment of engagement.¹⁷⁶ Who counts among the planners and decision-makers was contested at the Diplomatic Conference, where Switzerland and Austria took the view that the obligations of Article 57(2) could realistically concern only the higher echelons of command.¹⁷⁷ The better view, and the one now dominant in State practice, rejects any categorical restriction of this kind. The UK Manual provides that responsibility for the precautions ‘will range from commanders-in-chief and their planning staff to single soldiers opening fire on their own initiative’, the determinative criterion being whether the individual has any discretion in the way the attack is carried out.¹⁷⁸

Fourth, the provision does not specify what steps must be taken by those who plan or decide upon an attack in order to verify a target.¹⁷⁹ It simply demands that they do ‘everything’ feasible. The silence is deliberate. At the Diplomatic Conference, the ICRC explained that it had proposed a flexible wording for what later became Article 57 AP I, because it was for the parties concerned to make the rule more precise in the light of the organization of their armed forces and the kind of troops engaged.¹⁸⁰ The operational content of the duty of target verification was thus left to be worked out in practice. Even so, the rule implies at least four distinct steps.

To verify a proposition is to establish it by reference to evidence. Accordingly, the attacker

¹⁷⁶ Marco Sassòli, *International Humanitarian Law: Rules, Controversies, and Solutions to Problems Arising in Warfare* (2nd edn, Edward Elgar 2024) para 8.343.

¹⁷⁷ Official Records, vol VI, CDDH/SR.42, 212, paras 43 (Switzerland: the obligations ‘could concern the high commands only – the higher grades of the military hierarchy’) and 46 (Austria: the precautions ‘could only be taken at a higher level of military command’). Switzerland’s reservation to art 57(2), made on ratification on 17 February 1982, was withdrawn in 2005.

¹⁷⁸ UK Manual (n 62) para 5.32.9 (responsibility ranging ‘from commanders-in-chief and their planning staff to single soldiers opening fire on their own initiative’, the criterion being whether the person ‘has any discretion in the way the attack is carried out’). To the same effect, Air and Missile Warfare Manual (n 62) Rule 34 and commentary, paras 2–4; Tallinn Manual 2.0 (n 174) commentary to Rule 115, para 2; and Ian Henderson, *The Contemporary Law of Targeting: Military Objectives, Proportionality and Precautions in Attack under Additional Protocol I* (Martinus Nijhoff 2009) 159–161.

¹⁷⁹ Official Records, vol VI, 219 (Afghanistan, written explanation of vote, regretting the absence from art 57(2)(a) of more specific guidance for commanders in the field). The silence is noted in the literature. Loren Voss, ‘The Overlooked Importance of Intelligence Analysis in IHL’ (2025) 107 *International Review of the Red Cross* 287, 302 (the provision requires those who plan or decide upon an attack to do everything feasible to verify the objective, ‘but it does not specify what those feasible actions may be’).

¹⁸⁰ Official Records, vol XIV, CDDH/III/SR.21, 182, para 6. The same rationale is given at 183, para 7, an impracticable rule being said to promote neither the credibility nor the effectiveness of humanitarian law.

must first gather information to establish the status of the prospective target. This includes information both about the target itself, for instance whether it is a person or an object, and about the wider operational situation where the relevant rules, such as the definition of military objectives,¹⁸¹ require these to be considered. The 1973 commentary enumerates land and aerial reconnaissance conducted by human or by technical means as the methods through which identification is achieved and treats the absence of any information about the existence or whereabouts of a military objective as a bar to attack.¹⁸² The United Kingdom proceeded on the same footing on ratification, recording its understanding that those responsible for planning, deciding upon or executing attacks necessarily reach their decisions on the basis of an assessment of the information from all sources reasonably available to them at the relevant time.¹⁸³ The attacker must then assess the reliability of what has been gathered. Verification in its ordinary sense is a test of accuracy, rather than just an act of collection. Similarly, the rules on doubt in Articles 50(1) and 52(3) AP I presuppose that the attacker forms a view of the confidence the available material warrants. The drafters of AP I again proceeded on the same assumption. The obligations of those who launch an attack were separated from those of the persons who plan or decide upon it precisely because earlier observations may prove to have been partly erroneous and the circumstances prevailing when the attack was planned may since have changed.¹⁸⁴ Practice treats the evaluative step as distinct from collection. The first of the feasible precautions enumerated in the US Department of Defense Manual is a review of the accuracy and reliability of the information supporting the assessment that a potential target is a military objective.¹⁸⁵ The third step is the legal characterization of the entity. This is the only step the provision states expressly. The fourth step is individuation, the confirmation that the entity to which the preceding steps relate is the entity against which force will in fact be directed. The direct object of the verb requires as much, the rule addressing itself to ‘the objectives to be attacked’ and closing by requiring that it not be prohibited

¹⁸¹ Additional Protocol I, art 52(2).

¹⁸² ICRC Draft Protocols Commentary (n 165) 65 .

¹⁸³ United Kingdom, declaration (c) made on ratification of Additional Protocol I, 28 January 1998, certified by the Geneva Conventions Act (First Protocol) Order 1998, SI 1998/1754, Schedule.

¹⁸⁴ ICRC Draft Protocols Commentary (n 165) 65 (explaining the separate treatment of those who launch an attack on the grounds that earlier observations made at the time of planning and deciding may be partly erroneous and that the circumstances then prevailing may have undergone a profound change).

¹⁸⁵ DoD Law of War Manual (n 168) sec 5.5.3. See also Danish Military Manual (n 167) ch 8, sec 4.2.1; Air and Missile Warfare Manual (n 62) commentary to Rule 32(a), para 2; Yves Sandoz, Christophe Swinarski and Bruno Zimmermann (eds), *Commentary on the Additional Protocols of 8 June 1977 to the Geneva Conventions of 12 August 1949* (Martinus Nijhoff 1987) para 1952.

to attack ‘them’. So does the drafting history, the earliest formulation of the duty having held the legal and the individuating limbs apart in requiring that objectives be military objectives and be duly identified.¹⁸⁶

Verification accordingly combines factual and legal determinations and incorporates both an identification function and an exercise of selection in the narrow sense. Australian doctrine disaggregates the obligation along these lines into the sequential steps of locating and observing the potential target and assessing it as a valid military objective not otherwise protected by the law of armed conflict.¹⁸⁷ The committee established to review the NATO bombing campaign against the Federal Republic of Yugoslavia likewise concluded that a military commander must both establish an effective intelligence-gathering system and direct his forces to use available technical means to identify targets properly during operations.¹⁸⁸ Henderson captures the same structure as a two-stage process of information gathering followed by information assessment, during which the applicable legal tests are applied to the material gathered.¹⁸⁹

Fifth, in discharging these functions, the attacker is bound by an obligation of conduct, not of result. The Diplomatic Conference rejected as unrealistic the absolute duty proposed in the ICRC’s draft texts to ensure that objectives are duly identified as military objectives and settled instead on the compromise language of feasibility.¹⁹⁰ Both the steps to be taken and the level of effort required to verify the target are governed by that standard. Feasible precautions are those which are practicable or practically possible, taking into account all circumstances ruling at the

¹⁸⁶ ICRC Draft Rules (n 165) art 8(a); ICRC Draft Protocols Commentary (n 165) 65 (requiring identification ‘in a given case’).

¹⁸⁷ Australian Defence Force, *Australian Defence Doctrine Publication 06.4: Law of Armed Conflict* (2006) paras 5.61–5.62.

¹⁸⁸ Final Report to the Prosecutor by the Committee Established to Review the NATO Bombing Campaign Against the Federal Republic of Yugoslavia (13 June 2000) 39 ILM 1257 (hereinafter Final Report) para 29.

¹⁸⁹ Henderson (n 178) 162–165. See also Newport Manual (n 62) sec 8.9.1.

¹⁹⁰ ICRC Draft Protocols Commentary (n 165) 64–65; setting draft art 50(1)(a) out in two columns, Proposal I requiring that those who plan or decide upon an attack ‘shall ensure that the objectives to be attacked are duly identified as military objectives’ and Proposal II that they ‘shall take all reasonable steps to ensure’ the same, the commentary noting at 65 that Proposal II ‘allows the military command greater scope’; Official Records, vol XIV, CDDH/III/SR.21, 186, para 30 (United Kingdom, adopting the second alternative because ‘shall ensure’ implied an unrealistic absolute obligation) and 190, paras 60–61 (Belgium, drawing the civil-law distinction between an obligation as to means and one as to result and concluding that ‘it was the obligation as to means that was applicable in the case of precautionary measures’). The identification requirement had been criticized at the 1971 Conference of Government Experts as placing too heavy a burden on commanders lacking the intelligence to be certain of an objective’s character: ICRC Conference of Government Experts (n 165) 82, para 464.

time, including humanitarian and military considerations.¹⁹¹ Military considerations legitimately factored into the assessment include the risk to friendly forces involved in gathering further information, competing operational demands on scarce intelligence, surveillance and reconnaissance assets and the availability of weapons systems.¹⁹² The standard is therefore relative, requiring the most effective and reasonably available means in the circumstances prevailing at the time to be used systematically to obtain the most reliable information possible before an attack.¹⁹³ Feasibility also bears on the level of certainty the verification effort must produce before an attack may proceed. At one end of the spectrum, the ICRC Commentary asserts that in case of doubt, ‘even if there is only slight doubt’, those who plan or decide upon an attack must call for additional information,¹⁹⁴ a standard widely rejected as unrealistic on the ground that intelligence in armed conflict will rarely be complete.¹⁹⁵ National practice occupies various intermediate positions, ranging from the United Kingdom’s insistence on ‘substantial doubt’ through the more likely than not standard adopted by Norway and the ‘reasonable certainty’ demanded by France and Spain to the honest and genuine good-faith belief required by the United States, which excludes attacks

¹⁹¹ See the declarations of, among others, the United Kingdom (declaration (b)), Canada, Germany, Italy and the Netherlands made upon ratification of Additional Protocol I; Amended Protocol II, art 3(10); Protocol on Prohibitions or Restrictions on the Use of Incendiary Weapons (adopted 10 October 1980, entered into force 2 December 1983) 1342 UNTS 137 (hereinafter Protocol III) art 1(5); Official Records, vol XV, CDDH/215/Rev.1, 285, para 98.

¹⁹² New Zealand DM 69 (n 62) vol 4, para 8.7.9 (directing a commander deciding whether everything feasible has been done to verify the nature of an objective to take into account the ‘availability of reconnaissance and intelligence resources’, the ‘completeness and reliability of information already held’ and the ‘risk to the force involved in gathering more information’); Norwegian Manual (n 62) sec 2.6 (account being taken of ‘tactical and operational considerations, available means, and the need to protect one’s own forces’); UK Manual (n 62) para 5.32.2 and fn 202.

¹⁹³ Final Report (n 188) para 29; Quéguiner (n 161) 797–798 (adding that an attack launched on the basis of mere suspicion as to the military nature of the target amounts ipso facto to a violation of the principle of distinction). See also Kimberley Trapp, ‘Great Resources Mean Great Responsibility: A Framework of Analysis for Assessing Compliance with API Obligations in the Information Age’ in Dan Saxon (ed), *International Humanitarian Law and the Changing Technology of War* (Martinus Nijhoff 2013) 153, 158–159 (deriving a systemic obligation on States to develop the necessary information-gathering capabilities, calibrated to their resources).

¹⁹⁴ Sandoz, Swinarski and Zimmermann (n 185) para 2195.

¹⁹⁵ Henderson (n 178) 163–165; Michael N Schmitt and Eric W Widmar, “‘On Target’: Precision and Balance in the Contemporary Law of Targeting” (2014) 7 *Journal of National Security Law and Policy* 379, 401–402 (observing that slight doubt will almost always exist on a counterinsurgency battlefield where the adversary conceals its activities among the civilian population).

based on ‘mere speculation’ as to the target’s status.¹⁹⁶ Judicial formulations are pitched at a comparable level of generality, the Supreme Court of Israel requiring well based information before a civilian may be categorized as taking a direct part in hostilities and describing the burden borne by the attacking army as a ‘heavy’ one.¹⁹⁷ The expert manuals decline to fix a threshold at all.¹⁹⁸ Indeed, the better view is that the search for a fixed threshold is misconceived, as the level of certainty required to verify a target is dynamic and depends on the interplay between military advantage and the risk posed by misidentification.¹⁹⁹

Sixth, although verification is a planning-stage duty, it is not exhausted by a single determination made at the outset of the targeting process, but imposes a continuing obligation to collect, collate and evaluate timely target intelligence up to the execution of the attack.²⁰⁰ National

¹⁹⁶ UK Manual (n 62) para 5.3.4 (benefit of the doubt only in cases of ‘substantial doubt’); Norwegian Manual (n 62) sec 2.5 (‘an attack may only be carried out if it is more likely than not that a person or object is a lawful target’); French Manual (n 62) sec 8.1.2.1 (positive identification providing ‘the reasonable certainty that a functionally and geospatially defined entity (individual or object) is a valid military objective’, carried out at the target analysis stage and again before any engagement); Spain, Estado Mayor de la Defensa, *PDC-02.01: Derecho Internacional Humanitario (DIH) en las Fuerzas Armadas Españolas* (2022) paras 0138 and 0588 (military objectives must be identified as such before an attack is initiated and their characterization as such presupposes moving from doubt, which calls for further information, to ‘certeza razonable’ on the basis of the information available at the time); DoD Law of War Manual (n 168) sec 5.4.3.2 (the fact that a person is a military-aged male, without more, is an insufficient basis for attack). Operational standards of positive identification are best understood as policy instruments operationalizing the principle of distinction rather than as evidence of a legally required threshold: Michael N Schmitt and Michael Schauss, ‘Uncertainty in the Law of Targeting: Towards a Cognitive Framework’ (2019) 10 Harvard National Security Journal 148, 160–161; John J Merriam, ‘Affirmative Target Identification’ (2016) 56 Virginia Journal of International Law 83, 132–135.

¹⁹⁷ *Public Committee against Torture in Israel v Government of Israel* (n 171) para 40.

¹⁹⁸ Air and Missile Warfare Manual (n 62) commentary to Rule 12(a) (asking whether the degree of doubt would cause ‘a reasonable attacker in the same or similar circumstances’ to abstain); Tallinn Manual 2.0 (n 174) 424 (the requisite level of doubt is ‘unsettled’).

¹⁹⁹ Schmitt and Schauss (n 196) 151–152 and 163–166. The same intuition runs through the earlier literature and the ICRC’s own guidance: Henderson (n 178) 163–165; Geoffrey S Corn, ‘Targeting, Command Judgment, and a Proposed Quantum of Information Component’ (2012) 77 Brooklyn Law Review 437, 476–494; Adil Ahmad Haque, ‘Killing in the Fog of War’ (2012) 86 Southern California Law Review 63, 91 and 102–103, developed in Adil Ahmad Haque, *Law and Morality at War* (OUP 2017) 119–120; Nils Melzer, *Interpretive Guidance on the Notion of Direct Participation in Hostilities under International Humanitarian Law* (ICRC 2009) 76 (the standard of doubt ‘must reflect the level of certainty that can reasonably be achieved in the circumstances’).

²⁰⁰ Bothe, Partsch and Solf (n 164) art 57, para 2.5 (a ‘continuing obligation to assign a high priority to the collection, collation, evaluation and dissemination of timely target intelligence’); Dinstein (n 23) para 445 (the information relied upon must be up to date, ‘collated and interpreted just prior to the execution of the attack’).

practice confirms the point. The UK Manual requires target lists to be constantly reviewed in the light of fresh information and changing circumstances,²⁰¹ and the Danish and Norwegian manuals are explicit that the placement of a person or object on an approved target list does not discharge the duty, the designation having to be supported by sufficiently current intelligence and continually reassessed until the objective is attacked.²⁰² Target verification accordingly is an ongoing duty that runs from the start of the planning of or decision upon an attack to the actual commencement of the engagement.

4.2. *Mapping the duty of verification onto the targeting process*

Evidently, the duty of target verification must be discharged in full at some point in the targeting cycle before the target is engaged. It is not immediately clear, however, where that point lies. Mapping the legal obligation onto the doctrinal process answers this question, as it reveals where, and by whom, the four constituent elements of the target verification duty, namely information gathering, reliability assessment, legal characterization and individuation, are performed in the targeting cycle. It also establishes the baseline against which the delegation of target selection functions to autonomous systems can be assessed.²⁰³

Target verification begins with target discovery.²⁰⁴ Discovery generates candidate entities and produces the initial detection, location and identification data on which everything that follows depends, but it does not itself determine that any entity is a military objective. Discovery is conducted against criteria fixed in advance. Target selection standards prescribe the accuracy and timeliness required of detection systems, including target location error, minimum size and the distinction between static and moving entities.²⁰⁵ Doctrine thereby calibrates, before any individual target is in view, the informational quality of the data that will feed into the verification process.

²⁰¹ UK Manual (n 62) para 5.32.3 (giving the example of an enemy headquarters that ceases to be a military objective once the formation deploys elsewhere). See also New Zealand DM 69 (n 62) vol 4, paras 8.7.17–8.7.19.

²⁰² Danish Military Manual (n 167) ch 8, sec 4.2.2, including Examples 8.43 and 8.44; Norwegian Manual (n 62) sec 2.8 (requiring, for attacks against persons, that the individual be re-identified if lost from sight or if doubt otherwise arises as to their location).

²⁰³ The targeting cycle is used here to locate where in practice the components of the duty are performed, not to derive the content of the duty from doctrine.

²⁰⁴ AJP-3.9 (n 25) 1-9. The requirement is difficult to satisfy where the match is made against proxy features rather than against the legally relevant ones. See Afina (n 159) 152.

²⁰⁵ AJP-3.9 (n 25) 1-10 (target selection standards).

Discovered entities pass into target development.²⁰⁶ Target development is where the informational and legal assessment limbs of the verification duty converge. On the information-gathering side, it assembles the intelligence collected about an entity into target folders.²⁰⁷ This is the point in the process where substantiation requirements found in national practice, such as the Danish requirement that the intelligence basis comprise multiple and preferably independent types of information, are satisfied.²⁰⁸ On the assessment side, target development expressly examines whether engaging an entity in question would be lawful and for this purpose foresees ‘consultation between targeteers and legal advisors (LEGADs) in order to ensure the selection of lawful targets’.²⁰⁹ Target development therefore contributes to the verification duty in both of its dimensions, but it does not conclude either, as its output is a candidate for validation, not a verified target.

Before a developed target is validated, its intelligence foundation is tested in the form of quality control.²¹⁰ This stage merits particular attention because Article 57(2)(a)(i) AP I requires not merely that information be gathered, but that it be assessed. The first of the feasible precautions enumerated in the DoD manual is ‘reviewing the accuracy and reliability of information supporting the assessment that a potential target is a military objective’.²¹¹ The Air and Missile Warfare Manual similarly requires verification to be based on sources that are ‘sufficiently up-to-date and reliable, according to sound military practice’, taking into account the possibility of enemy disinformation.²¹² Quality control institutionalizes this requirement.²¹³

²⁰⁶ AJP-3.9 (n 25) 1-15.

²⁰⁷ AJP-3.9 (n 25) 1-16 and 1-17.

²⁰⁸ Danish Military Manual (n 167) ch 8, sec 4.2.1. See also French Manual (n 62) sec 8.1.2 (the target intelligence folder having to provide ‘available, sufficiently accurate, recent and time-consistent information’, consistent with the folder content and the target delimitation); and Henderson (n 178) 162–165.

²⁰⁹ AJP-3.9 (n 25) 1-15–1-16.

²¹⁰ AJP-3.9 (n 25) 1-8.

²¹¹ DoD Law of War Manual (n 168) sec 5.5.3 (listing as the first of the feasible precautions ‘[r]eviewing the accuracy and reliability of the information supporting the assessment that a potential target is a military objective’). Positive identification is not to be equated with compliance with the principle of distinction as a whole. As Ekelhof (n 46) 172 observes, a formal positive identification may in some circumstances be established through a single intelligence source which does not convey the object’s status as a lawful military objective, so that ‘the law and the procedure are not the same’.

²¹² Air and Missile Warfare Manual (n 62) commentary to Rule 32(a), para 2.

²¹³ Voss (n 179) 301 (the target vetting process leverages the expertise of the wider national intelligence community ‘in order to verify the fidelity of the intelligence and analysis used to develop the target’, whereas the target validation process determines, among other things, that the proposed target complies with international

Validation is the stage at which the legal characterization of the target is formally made. During target validation, the facts assembled through discovery and development, and tested through quality control, are assessed against the relevant legal categories and criteria, so that a decision can be made whether the entity qualifies as a military objective which may be attacked lawfully.²¹⁴ However, validation does not complete the duty of target verification either. First, validation is performed at the headquarters level upon an intelligence construct, namely the target as represented in its target folder, not upon the concrete entity that will eventually be engaged. Second, validation may precede the actual engagement of the target by days, weeks or even longer. A legal determination produced by validation that is not supported by sufficiently current intelligence at the time of the decision to attack will not discharge the duty of target verification and sustain the lawfulness of the engagement.²¹⁵ As the ICRC notes, the information supporting precautionary decisions must be ‘timely and up to date, collected and analysed up to the launch of the attack’ and ‘[i]nstructions issued in advance of an attack will not, on their own, constitute sufficient precautions’.²¹⁶

Validated targets are then nominated for engagement. Nomination advances targets from the Joint Target List to the Joint Prioritized Target List, which is approved by the joint force commander and represents a formal order to component commands to engage those targets.²¹⁷ The commander’s final approval decision rests on operational and legal considerations and thus functions as a second legal gate at which the continuing validity of the earlier legal characterization is confirmed.²¹⁸ It is at this point responsibility for the adequacy of the preceding verification is assumed at command level.²¹⁹

The final step in the sequence is positive identification of the specific target immediately prior to engagement, defined in NATO doctrine as a recognition derived from the observation and analysis of target characteristics that provides reasonable certainty that a specific, functionally and geospatially defined, entity is in fact a valid military objective.²²⁰ In the deliberate cycle, positive

humanitarian law and the rules of engagement).

²¹⁴ AJP-3.9 (n 25) 1-10.

²¹⁵ Danish Military Manual (n 167) ch 8, sec 4.2.2, Examples 8.43 and 8.44; UK Manual (n 62) para 5.32.3 (target lists to be constantly reviewed in the light of fresh information and changing circumstances).

²¹⁶ ICRC Selected Issues (n 4) 17.

²¹⁷ AJP-3.9 (n 25) 1-7 (joint prioritized target list) and 1-17–1-18 (Phases 3 and 4).

²¹⁸ AJP-3.9 (n 25) 1-18.

²¹⁹ On the decision to attack as a discrete element of directing an attack, see *Ntaganda* (n 139) para 744 (‘directing’ an attack encompasses selecting the intended target and deciding on the attack).

²²⁰ AJP-3.9 (n 25) 1-8.

identification is the point at which the target verification duty is discharged in full, in the sense that it completes the last of the duty's components. Whereas every earlier stage has operated on the target as an intelligence construct, positive identification connects that construct to the concrete person or object about to be attacked, using visual recognition, electronic support measures, identification friend or foe systems and comparable techniques. Positive identification is therefore the doctrinal counterpart of individuation, being the step that ensures that the objective actually attacked is the objective that was verified.

In the dynamic targeting methodology of F2T2E2A, the find step takes the place of discovery.²²¹ The fix step performs the information-gathering limb of development,²²² while tracking continuously monitors the target until engagement so as to hold the individuation of the target valid between identification and engagement.²²³ The target step performs the legal assessment and characterization limb of the duty. It is here that collateral damage estimation and rules of engagement are satisfied and final approval for engagement is given by the person holding target engagement authority.²²⁴ The target step thus compresses into a single decision the validation and command-approval functions that the deliberate cycle distributes between the Target Validation Board and the joint force commander. Positive identification, finally, connects the verified characterization to the concrete entity about to be attacked. Because tracking maintains positive identification throughout, the individuation that the deliberate cycle prescribes as a discrete final step is in dynamic targeting sustained continuously from the fix step onward. In dynamic targeting, the same target verification functions are thus performed in immediate succession, frequently within a single cell and within a compressed decision window.

Mapping the target verification duty onto the joint targeting cycle shows that neither the deliberate nor the dynamic cycle contains a single step at which the duty of target verification is discharged. Instead, doctrine distributes the components of the duty across the process. Each stage of the cycle performs a component of the duty, but none performs all of it. The obligation to verify the target is fully discharged only when the final stage of the cycle, where positive identification connects the verified characterization with the concrete entity to be engaged, is completed.

²²¹ AJP-3.9 (n 25) 5-3.

²²² AJP-3.9 (n 25) 5-4.

²²³ AJP-3.9 (n 25) 5-4.

²²⁴ AJP-3.9 (n 25) 5-4.

4.3. *Distribution of the duty of target verification*

How the duty of target verification is shared between human operators and AWS depends on what specific target selection tasks have been delegated to the system.²²⁵ Where an autonomous system is tasked with finding and engaging members of a category of target chosen in advance by a human operator, the system may be given anything from a single, narrowly defined identification task to a broad set of selection functions performed against several categories of objects at once. In this section, we examine two configurations to explore how the verification duty might be distributed, component by component, as more target selection tasks are handed to the machine.²²⁶ The two configurations are ideal types, but drawn from two real systems, the IAI Harpy and the Rafael SPICE-250.

The first configuration is modelled on the IAI Harpy, an anti-radiation loitering munition designed for the suppression of enemy air defences. According to its manufacturer, the Harpy is programmed before launch to fly to a defined loitering area, where it orbits and searches for radar emissions, compares the signature of any emission it detects against a stored database of known hostile air-defence systems and, once a target is confirmed, dives onto the emitter and destroys it; if the radar stops emitting before impact, the attack is aborted automatically.²²⁷ Because the Harpy can be launched without any prior knowledge of where a target is located,²²⁸ it is not fired at a particular radar that an operator has already found and selected, but is sent into a designated area to actively locate, identify and select targets for itself. This distinguishes the present configuration

²²⁵ The division of labour is not new to autonomous systems: see Henderson (n 178) 239–240.

²²⁶ See section 2.4 above on the three scenarios and, in particular, on the range of intermediate configurations within the third.

²²⁷ Israel Aerospace Industries, ‘Harpy: Autonomous Weapon for All Weather’ (*LAI*) <<https://www.iai.co.il/p/harpy>> accessed 19 August 2026; Automated Decision Research, ‘Israel Aerospace Systems HARPY loitering munition’ (*Automated Decision Research*) <<https://automatedresearch.org/weapon/israel-aerospace-systems-harpy-loitering-munition/>> accessed 19 August 2026 (the Harpy performs ‘autonomous flight to a pre-defined “Loitering Area”, in which they orbit while searching for radiating targets’, the signature of a detected target being ‘compared with a preprogrammed database of known air defense systems’ and the attack ‘automatically aborted’ if the radar ceases to emit).

²²⁸ Palladyne AI and Israel Aerospace Industries, ‘HARPY Anti-Radiation Loitering Munition’ (*Palladyne AI*) <<https://www.palladyneai.com/products/iai/harpy-anti-radiation-loitering-munition/>> accessed 19 August 2026 (the Harpy ‘can be launched without prior intelligence on the target’s location’); see also Riley Simmons-Edler and others, ‘AI-Powered Autonomous Weapons Risk Geopolitical Instability and Threaten AI Research’ (2024) arXiv:2405.01859 (noting that the Harpy ‘can autonomously detect and destroy radar targets in an area whose precise locations were not known before launch’).

from the second scenario considered earlier,²²⁹ where an anti-radiation missile is fired at a specific, preselected radar and uses its seeker only to sharpen its accuracy against a target the operator has already chosen.²³⁰ It also separates the Harpy from a mine.²³¹ Whereas a mine is passively actuated by a person or object,²³² the Harpy actively searches its environment to find and identify previously unknown targets and select those for engagement which match its onboard library of target profiles.

Even so, systems such as the Harpy occupy the conservative end of the automation range. The selection tasks delegated to the system are limited to the sensory recognition of candidate entities, matching them to a stored target profile and deciding to engage. The determination that an active hostile air-defence radar constitutes a category of military objectives, the assembly of the library of emitter signatures that expresses that category in a form the system can act upon and the level of certainty required for the system to identify candidate targets and confirm them for engagement are all made in advance. The verification duty is divided in line with this narrow delegation. Of the four components of the duty identified above, the legal characterization is completed by the operator before the system is launched, when the category of hostile radars is defined and the level of certainty fixed. The other three fall to the system, in whole or in part. The system gathers some of the information on which verification depends, sensing the emissions present within its assigned area; it assesses the reliability of that information, matching a detected emission against its library to the required level of confidence; and it individuates the target, performing the analogue of positive identification, when it determines that a particular emitter is one of the hostile radars the operator has characterized as a lawful target. Two features keep the machine's share of the duty modest. The system confirms its target by the very feature that defines the category, the hostile emission, so that its individuation of the target is a direct application of the operator's legal characterization and calls for no separate judgment about the object's identity. Because the system strikes upon confirmation, with no interval between identifying the emitter and engaging it, it need not hold the individuation current over time.

The second configuration is modelled on the Rafael SPICE-250, a precision-guided

²²⁹ See section 2.4.

²³⁰ Kelsey Atherton, 'Loitering Munitions Preview the Autonomous Future of Warfare' (*Brookings*, 4 August 2021) <<https://www.brookings.edu/articles/loitering-munitions-preview-the-autonomous-future-of-warfare/>> accessed 19 August 2026 (contrasting the HARM anti-radiation missile, which is 'human-directed and then ha[s] a short flight time in which to find a target', with the Harpy, which can 'search for 2.5 hours' and travel more than 300 miles after launch, looking for 'both known and unknown missile installations').

²³¹ The two are sometimes assimilated, but this overlooks the crucial differences between them.

²³² See section 2.4 above; Amended Protocol II, arts 2(1) and 2(4).

munition fitted with automatic target recognition. According to its manufacturer, the weapon uses AI and deep-learning techniques to learn the characteristics of a target type before the strike. As it approaches the target area, it detects and recognizes targets by matching what its electro-optical sensor sees against those learned representations, homing in on its assigned target either autonomously or under human supervision.²³³ The weapon reportedly can be directed against several types of target. In its standard use, the operator selects the target type and assigns a particular target to each weapon before release. Whether the weapon may engage another object where it cannot find the one assigned to it is, on the manufacturer's own account, a matter for the user to settle through its own policy and rules of engagement, and the automatic selection of a secondary target is contemplated as a possible future addition to the munition.²³⁴ The boundary between operator and machine in the selection of individual targets is therefore not permanently fixed by the design of the weapon, but is one which the operator is in a position to move. Systems such as the SPICE-250 accordingly illustrate a far more extensive delegation of target selection activities and a corresponding reduction in human involvement. Assume that such a system is tasked to search a designated area for members of more than one category of military objective, for example mobile air-defence systems and self-propelled artillery pieces, whose individual positions are not known when the system is launched. The system searches the area, detects candidate entities, classifies each of them against the profiles of the target categories and, where it finds several qualifying objects, ranks them and distributes its attack among them, tracking a

²³³ Rafael Advanced Defense Systems, 'Rafael Unveils New AI and Deep-Learning Technologies in SPICE-250 to Enable Automatic Target Recognition' (press release, 10 June 2019), reproduced in 'Rafael adds Automatic Target Recognition to SPICE-250' (*Army Recognition*, 12 June 2019) <<https://www.armyrecognition.com/archives/archives-aerospace-defense/aerospace-defense-2019/rafael-adds-automatic-target-recognition-to-spice-250>> accessed 19 August 2026 (the ATR feature enables the weapon to 'learn the specific target characteristics ahead of the strike, using advanced AI and deep-learning technologies'; '[w]hen approaching the target area, the weapons use the ATR mode for detection and recognition of the targets', each homing on its target 'either autonomously or with a human-in-the-loop'). For the manufacturer's own account of the limits of the capability, see Seth J Frantzman and Kelsey D Atherton, 'Israel's Rafael integrates artificial intelligence into Spice bombs' (*C4ISRNET*, 17 June 2019) <<https://www.c4isrnet.com/artificial-intelligence/2019/06/17/israels-rafael-integrates-artificial-intelligence-into-spice-bombs/>> accessed 19 August 2026 (reporting Rafael's position that whether the weapon may engage a secondary target where it cannot find the target assigned to it is a matter for user-defined policy and rules of engagement, and that the automatic selection of a secondary target might be added to the munition's upgrade profile).

²³⁴ Rafael (n 233) ('[d]uring flight, the pilot selects the target type to be attacked and allocates a target to each weapon'). On the treatment of secondary target selection as a matter of user-defined policy and rules of engagement, see Frantzman and Atherton (n 233).

chosen object until the moment it strikes and confirming the match before releasing its effect.

In this configuration, humans make category-level choices, including the approval of the classes of object that may be attacked, the development and characterization of target categories, their validation as military objectives, the construction of target profiles, the setting of the level of certainty a match must reach before an object is treated as a target and the criteria by which the system is to prioritize if it finds more than one qualifying entity.²³⁵ Meanwhile, the AWS makes choices about individual objects. These are the search of the area and the detection of candidate entities, the classification of each detected entity as a member or non-member of a certain target category, the ranking of several qualifying entities and the allocation of the attack among them, the tracking of a chosen target and the final confirmation to engage.

The verification duty is divided in the same way, but the wider delegation of selection allocates a wider share of the duty to the machine. As with the Harpy, the operator completes the legal characterization in advance, fixing the categories that qualify as military objectives, the parameters of membership and the level of certainty a match must satisfy, while the system gathers some of the information on which verification depends, assesses its reliability and individuates the target before it strikes. Three differences nonetheless enlarge the system's contribution. First, the system does not confirm its target by the single feature that defines the category, as the Harpy does, but classifies each candidate against learned proxy features standing in for the legally relevant ones, so that its individuation of the target is a less direct application of the operator's characterization whose reliability turns on how well the proxy tracks the legal predicate.²³⁶ Second, because an interval separates detection from engagement, the system's tracking of the chosen entity holds the individuation current across that interval. Third, the system is said to learn the characteristics of its target types. In so far as this implies that the system autonomously adjusts the features that define target profiles or the thresholds of certainty required for the identification of specific instances of a target category, the line between applying a predefined legal characterization and (re)making one becomes blurred.

4.4. *Implications*

The duty to verify a target under Article 57(2)(a)(i) AP I requires an individuated classification of

²³⁵ For these category-level acts of selection and their place in the targeting cycle, see section 3.3 above; on target selection standards and the commander's guidance fixing the categories and the operational envelope, see section 4.2 above.

²³⁶ For the underlying technique and its vocabulary, see Afina (n 159) 167.

the specific entity selected for attack and therefore remains engaged until this requirement has been discharged by those who plan or decide upon the attack. When target selection is delegated to an AWS, some of the activities that implement the duty of target verification are also delegated to the autonomous system. At a minimum, this includes determining that a specific entity identified by the system qualifies as a member of the category of military objectives it was programmed to engage, otherwise the system would not be selecting its target autonomously and therefore be excluded from the characterization of AWS by subparagraph 1A. This has several implications for the duty of target verification and beyond.

The first consequence concerns other precautionary obligations, some of which are engaged or can be discharged only after the specific entity to be attacked has been selected. This includes the proportionality rule, which forbids an attack which may be expected to cause incidental loss of civilian life, injury to civilians or damage to civilian objects that would be excessive in relation to the concrete and direct military advantage anticipated.²³⁷ Neither the incidental civilian harm expected from an attack, nor the concrete and direct military advantage anticipated, can be assessed without knowing the nature, function, location and immediate environment of the actual target of the attack.²³⁸ More fundamentally, the proportionality rule applies to attacks which are directed against a specific military objective. The rule is inapplicable to a military operation which does not, or does not yet, involve directing violence against a specific military objective. The rule is engaged only when the actual target of the attack has been chosen.²³⁹ If this task has been delegated to an AWS, it follows that the proportionality rule is engaged only after the system has been activated.²⁴⁰ The same is true for the obligation to take all feasible precautions in the choice of means and methods of attack with a view to avoiding, and in any

²³⁷ Additional Protocol I, arts 51(5)(b) and 57(2)(a)(iii).

²³⁸ ICRC Selected Issues (n 4) 15 (describing the assessment as a ‘prospective (ex-ante)’ one in which those launching the attack ‘must determine the incidental harm to civilians and civilian objects that is expected, and compare it to the anticipated, concrete and direct military advantage’).

²³⁹ This point is often overlooked: eg Boothby, ‘AI Warfare and the Law’ (n 53) 185 (suggesting that where an autonomous weapon is autonomously undertaking attacks, art 57(2)(a)(iii) applies ‘during the process that precedes the determination by the AI system that the object or persons that the sensors are observing is a military objective’).

²⁴⁰ Daniele Amoroso, *Autonomous Weapons Systems and International Law: A Study on Human-Machine Interactions in Ethically and Legally Sensitive Domains* (Edizioni Scientifiche Italiane 2020) 96 (predetermining the proportionality calculation at the planning stage ‘provides an adequate legal solution only when AWS without loitering capabilities are programmed to execute deliberate targeting missions’). See also Jonathan Kwik, *Lanfully Using Autonomous Weapon Technologies* (TMC Asser Press 2024) 247 (recognizing the problem, but not the full implications).

event to minimizing, incidental civilian harm. Since the incidental civilian harm to be expected cannot be assessed before a specific military objective has been selected for attack, the duty to avoid or minimize this harm cannot be discharged in full by human operators before that point in time either. Where a system does not simply engage the first qualifying entity it detects, but is programmed to choose between several possible targets, the precautionary obligation under Article 57(3) AP I may also be applicable. Once again, this obligation cannot be fully discharged by human operators before the AWS has prioritized the actual targets.

This temporal and functional relationship between the duty of target verification and other precautionary duties has implications for the design, development and operation of AWS. That certain duties are engaged only after activation, and cannot therefore be discharged in full before the system is deployed, in no way absolves those who plan or decide upon the attack from complying with them.²⁴¹ In essence, they have three options to do so. The first is to rely on humans to carry out the assessments required by the proportionality rule and other precautionary duties. This is not impossible, but the AWS would have to share at least some of the information gathered through its sensors with human operators to enable them to make these contextual assessments. However, bringing humans back into the targeting cycle at this point will almost certainly slow it down. This is suboptimal from an operational perspective. The second option is to delegate these tasks to the AWS, automating the implementation of the proportionality rule and other precautionary obligations.²⁴² This too runs into difficulties. There is near universal agreement that proportionality is among the most challenging rules to apply and that AI systems, at their current level of development, lack the sophistication to make the contextual and value-based judgments it requires.²⁴³ The third option is to approach the matter from a legal risk perspective and programme

²⁴¹ Rolling Text, 5 June 2026 (n 6) paras 12(A) and 18. Paragraph 12(A) requires that LAWS be operated under a responsible chain of command and control, ‘including ensuring assessment of legal obligations by a human, in particular with regard to the effects of the identification, selection, and engagement functions’. Paragraph 18 provides that States, parties to armed conflict and individuals ‘remain at all times responsible under applicable international law and accountable for their decisions with regard to LAWS since responsibility cannot be transferred to machines’, and that this ‘should be considered across the entire life cycle of the weapon system’.

²⁴² Boothby, ‘AI Warfare and the Law’ (n 53) 44 (where an autonomous weapon determines what is to be attacked and no human has been in a position to ensure compliance in advance, ‘it must of necessity be for the autonomous system itself to make those determinations’). Not everyone accepts that autonomous systems may make legal determinations: see Kwik (n 240) 7–8.

²⁴³ eg Boothby, *Weapons and the Law* (n 23) 249 (autonomous attack decision-making being unlikely for the foreseeable future to be capable of use consistently with the targeting rules save in unusual battle spaces from which civilians are absent); Camilla G Cooper, ‘Ensuring Lawful Use of Autonomous Weapons: An Operational Perspective’ (2024) 15 *Journal of International Humanitarian Legal Studies* 317, 340; McFarland (n 2) 123–124;

AWS to disengage as soon as they identify collateral concerns.²⁴⁴ This deliberately avoids engagements requiring a proportionality assessment and so imposes lower demands on the technology, but at the cost of severely limiting the operational utility of these systems.²⁴⁵ Moreover, while it may be possible in certain environments to avoid engagements risking incidental harm to civilians, it will be more difficult to avoid all engagements risking incidental harm to civilian objects. Whichever approach is adopted, autonomous systems must be equipped with the capabilities required to perform the tasks delegated to them and procedures must be put in place to provide for human input to complement machine functions, as necessary.

For the work of the GGE, the preceding analysis reinforces the conclusion that the tripartite enumeration of identification, selection and engagement should give way to selection and engagement, with selection understood in a broad sense. Just as identification and selection are not two self-contained functions that can be fully assigned to a human or to a machine, the duty of target verification too involves activities that are distributed across the targeting cycle in a manner that the tripartite enumeration tends to obscure. The exceptions introduced in subparagraph 1A should be retained for the same reasons. The duty of target verification requires classifying the specific entity selected for attack. Human operators cannot carry out this classification without the actual target having been selected first. If humans were to classify the actual target, they would also have to preselect it, meaning that the AWS would no longer be autonomously selecting its target.

Beyond this, the rolling text offers only limited guidance on how the duty of target verification is to be discharged when an AWS is used to select and engage targets. Paragraph 11 provides that such systems must only be used in accordance with applicable international humanitarian law, including the principles and requirements of distinction, proportionality and precautions in attack.²⁴⁶ The proposition is of fundamental importance, as it confirms that those

Pablo Kalmanovitz, 'Judgment, Liability and the Risks of Riskless Warfare' in Nehal Bhuta and others (eds), *Autonomous Weapons Systems: Law, Ethics, Policy* (CUP 2016) 145, 151.

²⁴⁴ See Seixas-Nunes (n 2) 34 (precautions 'can only be taken ... when those in charge have a clear knowledge of the situation in which an AWS is to be deployed', so that such systems 'will only be used for missions where no risk of violating IHL is envisaged', and it is therefore the mission 'that will govern the degree according to which AWS can be deployed and be autonomous'); Natalia Jevglevskaia, *International Law and Weapons Review: Emerging Military Technology under the Law of Armed Conflict* (CUP 2022) 230–231; William H Boothby, 'How Far Will the Law Allow Unmanned Targeting to Go?' in Dan Saxon (ed), *International Humanitarian Law and the Changing Technology of War* (Martinus Nijhoff 2013) 45 (hereinafter Boothby, 'How Far Will the Law Allow') 57–58.

²⁴⁵ Amoroso (n 240) 92.

²⁴⁶ Rolling Text, 5 June 2026 (n 6) para 11. The obligation was strengthened during the first 2026 session from 'may

who plan or decide upon an attack remain bound in full by their legal obligations when they employ an autonomous system, no less than when they employ any other means of warfare. Building on this, paragraph 12 offers some guidance on how those obligations are to be met once the activities that implement them are divided between the operator and the machine. It stipulates the overarching principle that human judgment and control with regard to such systems are needed in order to uphold compliance with international law, but adds that this does not imply constant direct human control. Instead, what is required is an appropriate level of human judgment and control, which must be determined by reference to a range of factors, including the operational context and the characteristics and capabilities of the weapon system as a whole.²⁴⁷

That the requisite degree of human involvement must vary with the circumstances is a sound principle reflecting the context-dependent nature of the law of targeting. The rolling text says nothing, however, about what that variation consists of. It does not indicate which target selection functions may be delegated to the system, which of the activities implementing the duty of target verification may be performed by it or how the resulting division of labour bears on the level of human involvement required. The notion has been criticized within the Group on precisely this ground, several delegations objecting that it is understood differently across delegations, that it is too vague and that it does not clearly reflect an existing rule of law.²⁴⁸

Two readings of the context-appropriate human involvement standard suggest themselves. Since the overriding goal is to secure compliance with the applicable law, the first holds that any distribution of tasks between the operator and the system which ensures such compliance satisfies the standard. The level and nature of human involvement must reflect whatever compliance with the applicable rules requires in the prevailing circumstances, the precise balance between human involvement and machine automation, and the identity of the tasks delegated, being secondary matters. The rolling text lends this reading some support. The Chairperson has described the requirement of human judgment and control as a notion specific to autonomous weapons systems which is required in order to ensure compliance rather than as an obligation arising under existing

only be used' to 'must only be used', a change supported by all delegations as reflecting a requirement already established under international humanitarian law: CCW/GGE.1/2026/WP.2 (n 86) para 49.

²⁴⁷ Rolling Text, 5 June 2026 (n 6) para 12. In the rolling text of 18 December 2025 the requirement was framed in terms of 'context-appropriate human judgement and control' and appeared as Box III, para 6, the implementing measures now found in sub-paragraphs 12A to 12D appearing as Box III, para 7. The Chairperson relocated the requirement from Box II to Box III in order to distinguish what the law prohibits from what it requires and from the manner in which compliance is to be secured: CCW/GGE.1/2026/WP.2 (n 86) paras 25 and 32.

²⁴⁸ CCW/GGE.1/2026/WP.2 (n 86) para 35. See also para 53 of the same summary (recording the persistence of the objection that the notion does not constitute a rule of international law and risks introducing legal ambiguity).

law, and the Group has accepted that such compliance does not necessarily call for direct, continuous or real-time human control at every stage, but rather for appropriate forms of human involvement at different stages of the life cycle.²⁴⁹

On the second reading, certain tasks must be performed by human beings whatever the circumstances, so that the level of human involvement is context dependent only within these limits. There is a possible hint of this position in subparagraph 12A, which provides that operating such systems under a responsible chain of command and control includes ‘ensuring assessment of legal obligations by a human’, in particular with regard to the effects of the identification, selection and engagement functions.²⁵⁰ The words are capable of bearing either meaning. They may mean no more than that the implementation of the applicable legal obligations, including the duty of target verification, is itself subject to the requirement of appropriate human involvement, so that a human must be answerable for the assessment without necessarily performing every element of it. Alternatively, they may mean that the assessment of legal obligations is a matter reserved to human operators which may never be delegated to a machine. The Chair’s commentary offers no assistance in resolving the ambiguity.

The ICRC is one of the key proponents of the second position. Its starting point is that the rules governing the conduct of hostilities presuppose context-specific human judgment, that it is human beings who must determine the lawfulness of the attacks they plan, decide upon or execute and that therefore these determinations ‘cannot be delegated to machine processes’.²⁵¹ The

²⁴⁹ CCW/GGE.1/2026/WP.2 (n 86) para 25 (the notion is ‘specific to lethal autonomous weapons systems required to ensure compliance rather than an obligation under existing international humanitarian law’) and para 52 (recording broad recognition of the essential role of the human element while acknowledging that compliance does not necessarily require direct, continuous or real-time human control at every stage, appropriate forms of human involvement being exercised at different stages instead).

²⁵⁰ Rolling Text, 5 June 2026 (n 6) sub-para 12A.

²⁵¹ ICRC Selected Issues (n 4) 2. The same proposition is developed at 8. The proposition has a longer pedigree. See Richard Moyes, ‘Meaningful Human Control over Individual Attacks’ in International Committee of the Red Cross, *Autonomous Weapons Systems: Implications of Increasing Autonomy in the Critical Functions of Weapons* (Expert Meeting, Versoix, 15–16 March 2016) 46, 48–49 (because ‘those who plan or decide upon an attack’ refers to humans, a machine cannot identify and attack a military objective ‘without human legal judgement and control being applied in connection with an attack on that specific military objective at that time’ and programming that capacity into a machine is ‘an abrogation of human agency with respect to the law’); Dustin A Lewis, ‘On “Responsible AI” in War: Exploring Preconditions for Respecting International Law in Armed Conflict’ in Silja Voenekey and others (eds), *The Cambridge Handbook of Responsible Artificial Intelligence* (CUP 2022) 488, 497–498 (positing as preconditions of lawful employment that the exercise and implementation of international law is ‘reserved solely to humans’ and may not be ‘reposed in non-human (artificial) agents’); Boulanin, Bruun and

ICRC accepts that certain technical tasks may be carried out by machine, but it distinguishes these from the determination of the lawfulness of an attack, including whether an object constitutes a military objective, which in its view calls for context-specific human assessment and judgment for which human beings remain accountable.²⁵² The need for such judgment is said to be particularly evident in relation to the status of persons and objects, whose classification depends on contextual and values-based assessments that resist reduction to a quantitative formula and cannot be captured in the generalized target profile on which an autonomous system relies.²⁵³

From this premise the ICRC draws a conclusion about timing. Since an autonomous weapon system is by definition one that selects and engages targets without human intervention after activation, the assessments that the law reserves to human beings can only be made before the system is activated.²⁵⁴ The ICRC accordingly treats each instance in which such a system selects and engages a specific person or object as an attack commencing at the moment of activation, and it requires those responsible to anticipate and assess the lawfulness of the attack in advance, that is to say at the time of activation, in respect of all possible strikes throughout the entire area and for the whole period in which the system is active.²⁵⁵ Applied to the duty of target verification, the requirement is that the user, when activating the system, must ensure that whatever entity will be

Goussac (n 12) 20 (recording the view of some States that, because the rule addresses ‘those who plan or decide upon an attack’, the evaluation it demands ‘cannot be delegated to an AWS or substituted with preprogrammed precautionary (e.g. fail-safe and self-destruct) mechanisms’); Sadjad Soltanzadeh, ‘Normative Reasoning and Military AI: DILEMA’s Dilemmas’ in Bérénice Boutin, Taylor Kate Woodcock and Sadjad Soltanzadeh (eds), *Legal, Ethical, and Technical Dilemmas in Military Artificial Intelligence* (T.M.C. Asser Press 2026) 31, 40–41 (the precautionary principle is ‘inherently reliant on continuous human judgement’ and respect for legal norms is an inalienable obligation which cannot, as a matter of logic rather than of epistemic capacity, be delegated to a machine).

²⁵² ICRC Selected Issues (n 4) 8.

²⁵³ ICRC Selected Issues (n 4) 12 and 13. On the ICRC’s understanding of the generalized target profile, see 6.

²⁵⁴ ICRC Selected Issues (n 4) 11. The definition on which the inference rests is set out at 6, autonomous weapons systems being those that select and engage targets without human intervention.

²⁵⁵ ICRC Selected Issues (n 4) 11. The temporal boundaries of an attack conducted by an autonomous system are contested rather than settled. See Boulanin, Bruun and Goussac (n 12) 23–24 (direct endangerment may be located at the point when the person or object is selected for the application of force, when it is identified as a potential target or when the system is activated and begins matching data against its target profile library, while on a triggering-act approach activation is ‘probably most similar to the placement of a mine’); Kwik (n 240) 231 (attributing much of the disagreement over the legality of autonomy to divergent understandings of the scope of an attack, the narrowest reading, on which each firing is an attack, ‘essentially “preclud[ing] autonomous systems”’); Mauri (n 2) 36–37 (deriving from art 49 AP I and the Commentary’s treatment of mine-laying the view that a weapon system attacks ‘when it releases force ... against a selected target’).

selected for engagement satisfies the definition of a military objective at the time of the strike, while accounting for all reasonably foreseeable and relevant changes in circumstances.²⁵⁶ The restrictions for which the ICRC calls, on the types of targets that may be engaged, on the duration and geographical scope of operations and on the situations in which such systems may be used, follow directly from this, since they are the conditions under which an anticipatory assessment of that kind could realistically be made at all.²⁵⁷

However, the ICRC's position does not accurately reflect the character of the duty of target verification. As we have seen, target verification combines factual and legal determinations that divide into distinct activities performed at different stages of the targeting cycle, namely the gathering of information, the assessment of its reliability, the legal characterization of the entity and its individuation.²⁵⁸ A position insisting that the legal assessments involved in target verification may never be delegated to a machine must say which component it means. The ICRC accepts that certain technical tasks may be carried out by machine processes and reserves to human beings only the determination of lawfulness.²⁵⁹ However, it is far from clear how the technical is to be separated from the legal. A machine which collects information through its sensors and matches it against a stored profile to a prescribed level of confidence performs the gathering of information and the assessment of its reliability that Article 57(2)(a)(i) AP I requires. None of this is exclusively technical or exclusively legal, but mixes factual and legal determinations that are components of the duty.²⁶⁰

More importantly, the ICRC position overlooks the individuated nature of the duty of target verification. The duty attaches to the specific person or object against which force will be directed, not to the category to which that entity belongs, and it is discharged in full only when the concrete entity about to be attacked has been classified.²⁶¹ Where target selection has been delegated to an autonomous system, the operator cannot perform that classification, because at the activation stage there is no specific entity to classify. The duty therefore cannot be discharged in full at the point of activation and remains engaged after the deployment of the system.

Those who adopt the view that legal assessments may not be delegated to AWS are thus

²⁵⁶ ICRC Selected Issues (n 4) 13.

²⁵⁷ ICRC Selected Issues (n 4) 17 and 22.

²⁵⁸ See sections 4.1 and 4.2 above. Doctrine distributes the components of the duty across the targeting cycle, and the duty is discharged in full only at the point of positive identification, when the verified characterization is connected to the concrete entity about to be attacked.

²⁵⁹ ICRC Selected Issues (n 4) 8.

²⁶⁰ See section 4.2 above.

²⁶¹ See sections 4.1 and 4.2 above.

left with a choice between two positions. The first is to maintain that the legal determinations required by the targeting rules, verification among them, must be performed by human operators before activation. Since the individuated component of the duty cannot be discharged at that point, but must be performed by the system selecting its own targets, it follows that no AWS could ever be used in compliance with the law of armed conflict. The conclusion is logically coherent and some commentators embrace it,²⁶² but neither the ICRC nor the majority of States take the view that AWS are legally impossible.²⁶³ The second position is to abandon the categorical claim that legal assessments may never be delegated to a machine. Nothing in Article 57(2)(a)(i) AP I reserves the activities which implement the duty of target verification exclusively to humans.²⁶⁴ While the obligation is addressed to those who plan or decide upon an attack, this language is without prejudice to the instruments or methods they may rely on to implement it. The authentic language versions bear this out, the French text speaking of ‘ceux qui préparent ou décident une attaque’, a formulation which, unlike the English and the Spanish, is not confined to natural persons.²⁶⁵ It might be objected that other provisions of the Protocol presuppose distinctly human cognition and that the presumption laid down in Articles 50(1) and 52(3) AP I in case of doubt therefore bears on the verification duty as well.²⁶⁶ However, that rule simply prescribes what is to be done

²⁶² eg Liebllich and Benvenisti (n 7) 271 and 283 (based on the argument that automated systems cannot exercise discretion as required by the law of armed conflict).

²⁶³ ICRC Selected Issues (n 4) 10, 13 and 22 (recommending the prohibition only of autonomous weapons whose effects cannot be sufficiently understood, predicted and explained and of those designed or used to attack persons directly). Neither is it the position of States. The rolling text prohibits certain systems and requires that the rest be used in accordance with the applicable law, an architecture which would be incoherent if every system within the characterization were unlawful by definition: Rolling Text, 5 June 2026 (n 6) paras 7 to 11.

²⁶⁴ McFarland (n 2) 116–117 (‘there do not appear to be explicit requirements for steps in a targeting process ... to be completed directly by a human rather than via software’); Boothby, ‘AI Warfare and the Law’ (n 53) 177–178 (the passive mood of art 57(2) leaves open whether the precautions may be delegated to a technological entity, the provision being ‘agnostic as to whether a human being or some other non-human entity performs the precautions’); Sassòli (n 176) para 10.090 (all rules of international humanitarian law are addressed to humans, but this does not preclude planners and decision-makers ‘from being temporally and geographically removed from the attack as long as they define the parameters according to which the robot attacks and make sure that it has the information necessary to apply such parameters’); Robert Frau, ‘Neue Waffensysteme als Herausforderung für das Völkerrecht’ (2018) 1 Humanitäres Völkerrecht 5, 18 (neither general nor special international law prohibits non-human actors from taking attack decisions). See also Sati (n 55) 222 (art 57 ‘requires the fulfilment of the tasks it prescribes; it does not specify whether compliance must be achieved through humans or machines’).

²⁶⁵ Frau (n 264) 10.

²⁶⁶ Damian Copeland, *A Functional Approach to the Legal Review of Autonomous Weapons Systems* (Brill Nijhoff 2025)

where the available information falls short of the required standard. It does not prescribe who or what must apply that standard. An automated system that withholds engagement wherever its confidence in a match falls below a certain level gives effect to the presumption in the manner the rule requires, by declining to treat the entity as a military objective. Whether the level fixed is one a reasonable commander would accept in the circumstances is a question about the adequacy of the system and its programming, not about the identity of the agent applying the rule. The question is therefore not whether legal determinations and assessments may be delegated to a machine at all, but whether a machine can perform any of the constituent elements of the duty of target verification in a manner that meets the requirements imposed by that duty on humans. The answer depends, ultimately, on the technical sophistication of these systems and the operational circumstances.²⁶⁷

This is not to suggest that those who plan or decide upon an attack have unfettered discretion in delegating target verification tasks to autonomous systems. Recall that they must do everything feasible to verify the objectives to be attacked.²⁶⁸ Where it is feasible for human operators to carry out some or all target verification tasks themselves and do so at higher levels of confidence compared to a machine, the duty to do everything feasible demands their involvement in the process, all other things being equal.²⁶⁹ The duty does not, however, mandate human involvement under all circumstances, contrary to what has sometimes been suggested.²⁷⁰ What is feasible falls to be determined by weighing humanitarian and military considerations. While human involvement may enhance humanitarian protections by reducing or even eliminating machine

207–208 (the text of some rules ‘clearly contemplate the application of distinctly human cognition such as “recognise” and “doubt”, so that a State may require, as a matter of law, that the functions governed by such rules be performed or directly controlled by humans, the protection of art 52(3) AP I being the clearest case).

²⁶⁷ To similar effect, Tomasz Zurek and Tom van Engers, ‘The Computational Representation of Rules of International Humanitarian Law in Military Autonomous Devices: Issues, Problems, and Controversies’ in Bérénice Boutin, Taylor Kate Woodcock and Sadjad Soltanzadeh (eds), *Legal, Ethical, and Technical Dilemmas in Military Artificial Intelligence* (T.M.C. Asser Press 2026) 239, 258 (‘the issue is not the delegation of the evaluation process to the machine, but rather the quality of the decision-making process itself’); Mauri (n 2) 147–148 (the argument that autonomous systems cannot lawfully be used because the classification of status requires a context-related assessment ‘proves too much’, being ‘[u]ltimately ... a technological objection’ which advances in machine learning may overtake’).

²⁶⁸ Additional Protocol I, art 57(2)(a)(i).

²⁶⁹ See also Boothby, *Weapons and the Law* (n 23) 254–255 (if the obligation under art 57(2)(a)(i) can be fulfilled using a manned mission, but not using an autonomous one, the manned mission should be employed).

²⁷⁰ Amoroso (n 240) 106–113.

errors,²⁷¹ it may also slow down the targeting process to a pace where the military advantage essentially evaporates. In some circumstances, human involvement in verifying a target could impose an operational penalty that manifestly outweighs any protective benefit it would offer. In such cases, nothing in Article 57(2)(a)(i) AP I prevents the four constituent elements of the duty of target verification from being delegated to an autonomous system in their entirety, provided that the system is capable of verifying targets above the threshold of doubt and the military benefit of relying on the machine outweighs any likely adverse humanitarian consequences.²⁷²

5. CONCLUSION

Even though target selection features in virtually all definitions of autonomous weapons systems, what exactly selecting a target involves has received little systematic attention in the literature and in international debates.²⁷³ Target selection, as understood in military doctrine and operations, is not a single act, but a structured process in which a pool of candidate entities is progressively narrowed through the repeated application of political, operational, legal and other criteria. Target identification, as the epistemic component of target selection, is likewise a family of activities dispersed across the targeting cycle.²⁷⁴ Even in a genuinely autonomous system, neither function is wholly automated. Some of the constituent tasks are performed by human operators before the system is activated, while others fall to be performed by the system once it is deployed.

The duty of target verification tracks this division of labour. The obligation imposed by Article 57(2)(a)(i) AP I breaks down into four components, the gathering of information about candidate entities, the assessment of the reliability of this information, the legal characterization of an entity and the individuation of the target.²⁷⁵ The duty requires the specific person or object selected for attack to be identified and it remains engaged until that step has been completed. It follows that the duty cannot be discharged in full at the point when an autonomous system is activated. Whereas the operator completes the legal characterization of the target category in advance, the individuation of the target has, at that point, no object. In any system that

²⁷¹ Amoroso (n 240) 108.

²⁷² See Sassòli (n 176) para 10.089 (emphasizing that such assessments must be made on a case-by-case basis).

²⁷³ The principal exception is Merel A C Ekelhof, 'Complications of a Common Language: Why It Is So Hard to Talk about Autonomous Weapons' (2017) 22 *Journal of Conflict and Security Law* 311, 326–327 (distinguishing target selection understood as target recognition from target selection understood as the deliberative process of target planning and development).

²⁷⁴ See sections 2 and 3 above.

²⁷⁵ See sections 4.1 and 4.2 above.

autonomously selects its own target, this final individuation remains to be performed by the system itself.²⁷⁶

The multistage and individuated nature of the duty of target verification means that system activation is the wrong legal moment for implementing most of the precautionary obligations engaged by autonomous targeting. Assessing the proportionality of an attack, taking precautions to avoid or minimize collateral harm and selecting the objective which may be expected to cause the least danger to civilians and civilian objects all presuppose a specific target. None of these obligations can be completed before an autonomous system tasked with selecting a target has done so. The insistence that precautionary duties must be completed at the point of activation derives from the assumption that all legal assessments must be made by humans, not from anything in the law of targeting itself.²⁷⁷ Once that assumption is set aside, the precautionary obligations fall to be discharged where they are engaged and by whatever combination of human and machine action is capable of discharging them.

This has practical implications for how autonomous systems are designed and operated. If key aspects of precautionary obligations must be performed after activation, they must either be carried out by the system itself, which demands an appropriate level of technical sophistication, or by a human to whom the system returns control together with the information needed to act, which requires clear lines of responsibility. Alternatively, the system could be programmed to disengage whenever its actions would engage legal obligations that exceed its compliance capabilities. In each case, compliance is the product of the socio-technical whole, comprising the autonomous system, the operator who deploys it and the wider decision-making process in which the system is embedded.

For the regulatory debate, several conclusions follow. First, the GGE should resist the tripartite enumeration of identification, selection and engagement. The underlying assumption that these are self-contained functions which may be allocated wholesale to one side or the other in the human-machine divide does not reflect the nature of the targeting process. Instead, States should characterize AWS as systems that select and engage targets autonomously, defining selection in a broad sense. Second, the Group should retain the exclusions of subparagraph 1A, which draw the boundary of the characterization exactly where the verification analysis places it, namely at the determination of the specific target. Third, the notion of context-appropriate human involvement is of critical importance, but it will remain an empty shell until States clarify which tasks may be

²⁷⁶ Rolling Text, 5 June 2026 (n 6) sub-para 1A.

²⁷⁷ ICRC Selected Issues (n 4) 11. On the two standpoints and the point of principle that divides them, see Boulanin, Bruun and Goussac (n 12) 20.

delegated to an automated system, in what circumstances and subject to what conditions. The categorical position that legal assessments may never be delegated to a machine cannot be sustained. Pressed consistently, it implies that no autonomous weapons system could ever be used in accordance with the law of armed conflict, a conclusion that neither the ICRC nor States accept.²⁷⁸ Developing the notion of context-appropriate human involvement requires asking which aspects of precautionary duties a system is capable of performing in a manner that meets the legal obligations imposed on those who plan or decide upon an attack.

Much of the wider debate about military autonomy and AI has been conducted at a high level of abstraction, revolving around notions such as meaningful human control and judgment. Although essential for framing the analysis, most of these concepts offer limited guidance as to what exactly compliance with the law of armed conflict entails in specific circumstances. One strand of the literature has therefore long urged that the debate should pay more attention to how targeting is actually practised.²⁷⁹ The analysis developed here shows that target selection as a military enterprise and target verification as a precautionary obligation are both multistage processes involving a series of steps that do not lend themselves to any simple apportionment between humans and machines. In the majority of use-cases, component activities will interlock and reinforce one another. The division of labour between human operators and automated systems therefore results in a system of dynamic and mutually dependent interaction, rather than distinct decision-making stovepipes, where target selection and verification are co-produced.²⁸⁰

²⁷⁸ See section 5 above; ICRC Selected Issues (n 4) 14, 16 and 17. For authors who do press the position to that conclusion, see Lieblich and Benvenisti (n 7) 283; Kwik (n 240) 243.

²⁷⁹ eg Merel AC Ekelhof, 'Lifting the Fog of Targeting: "Autonomous Weapons" and Human Control through the Lens of Military Targeting' (2018) 71 *Naval War College Review* 61; Merel Ekelhof, 'Moving Beyond Semantics on Autonomous Weapons: Meaningful Human Control in Operation' (2019) 10 *Global Policy* 343; Mark Roorda, 'NATO's Targeting Process: Ensuring Human Control Over (and Lawful Use of) "Autonomous" Weapons' in Andrew P Williams and Paul D Scharre (eds), *Autonomous Systems: Issues for Defence Policymakers* (NATO Allied Command Transformation 2015) 152; Alan L Schuller, 'At the Crossroads of Control: The Intersection of Artificial Intelligence in Autonomous Weapons Systems with International Humanitarian Law' (2017) 8 *Harvard National Security Journal* 379.

²⁸⁰ Ingvild Bode, *Human-Machine Interaction and Human Agency in the Military Domain* (Centre for International Governance Innovation, Policy Brief No 193, 2025) 5–6, developing the concept of distributed agency: agency is located neither in the human nor in the machine but 'in the relations that exist between them' and in the practices they perform together, so that it cannot be understood by treating humans and machines as separate entities in isolation. See also Ekelhof (n 46) 163 (compliance with the targeting rules is 'typically shared or distributed' among many personnel and does not succeed or fail with the decision of one person); Amoroso (n 240) 109; Klaudia Klonowska and Taylor Kate Woodcock, 'Rhetoric and Regulation: The (Limits of) Human/AI

These findings carry broader lessons. The line that currently divides human from machine contributions is not fixed. As the underlying technologies mature, components of the target selection process and the verification duty that remain with operators for now may increasingly be delegated to autonomous systems. The regulatory frameworks and approaches developed today therefore may age quickly. However, what remains constant amidst rapid technological and organizational change is the fact that the law of targeting imposes obligations on humans. Regardless of what tasks are delegated to autonomous systems, the enduring question is whether machines are capable of performing them in a manner that reflects and respects the legal obligations of their human operators. To make that happen, compliance with the law must be embedded into the design and development of autonomous systems. The rolling text points in this direction when it provides that the formulations adopted by the Group are to be taken into account in the design, development and use of such systems and that their critical functions should be adequately predictable, reliable, traceable and explainable.²⁸¹ The direction of travel is therefore set, but substantial work remains to be done to translate high-level principles into concrete standards and performance requirements.

Comparison in Legal Debates on Military AI' in Bérénice Boutin, Taylor Kate Woodcock and Sadjad Soltanzadeh (eds), *Legal, Ethical, and Technical Dilemmas in Military Artificial Intelligence* (T.M.C. Asser Press 2026) 51, 55–57 (identifying the 'replacement fallacy' underlying attempts to allocate tasks wholesale between humans and machines and its genealogy in the literature on which tasks machines and men are better at).

²⁸¹ Rolling Text, 5 June 2026 (n 6) paras 13 and 15C.