

Israel – Hamas 2024 Symposium – The Gospel, Lavender, and the Law of Armed Conflict

lieber.westpoint.edu/gospel-lavender-law-armed-conflict

June 28, 2024

by

Michael N. Schmitt

| Jun 28, 2024



In November 2023, *+972 Magazine* and *Local Call* published a joint [report](#) on Israel Defense Force (IDF) targeting operations, labeling them a “mass assassination factory.” It correctly stated that the IDF was using “Gospel,” a system characterized as relying upon artificial intelligence (AI). The report alleged the system “generates, among other things, automatic recommendations for attacking private residences where persons suspected of being Hamas or Islamic Jihad operatives live. Israel then carries out large-scale assassination operations through heavy shelling of these residential homes.” Other media sources, such as the [Guardian](#), quickly picked up the story, sparking a firestorm of condemnation.

In April, the focus of attention shifted to the IDF’s use of “Lavender,” again based on a [+972 report](#) that was [echoed](#) in other media (including the [Guardian](#)). Lavender was described as identifying Hamas and Islamic Jihad operatives, with *+972’s* sources claiming that its “influence on the military’s operations was such that they essentially treated the outputs of

the AI machine ‘as if it were a human decision.’” The report further charged that “the army gave sweeping approval for officers to adopt Lavender’s kill lists, with no requirement to thoroughly check why the machine made those choices or to examine the raw intelligence data on which they were based.” Reportedly, “human personnel often served only as a ‘rubber stamp’ for the machine’s decisions,” sometimes approving them in 20 seconds. The +972/*Local Call* accounts, taken at face value, are exceptionally troubling. They essentially depict a mechanistic AI “targeting machine” that is subject to negligible verification.

Many in the international law community have weighed in. One report asserts that “[s]trikes against [Gospel] targets are often authorized without further oversight.” Other commentary suggested that “the supervision protocol before targeting suspected militants involves confirming the AI-selected target’s gender, with the assumption that female targets are erroneous and male targets are appropriate.” Some have opined that “AI is increasing the tempo of operations and expanding the pool of potential targets, making target verification and other precautionary obligations much harder to fulfil, increasing the risk that civilians will be misidentified and mistakenly targeted.” Yet, other legal commentators vigorously defended the systems (see, e.g., here).

The IDF had previously acknowledged the use of Gospel (and other systems like “Alchemist” and “Depth of Wisdom”) during Operation Guardian of the Walls in 2021. It also responded to the *Guardian* report on Lavender, asserting that “[s]ome of the claims . . . are baseless in fact, while others reflect a flawed understanding of IDF directives and international law.” It described a human-controlled process of identifying targets in which Gospel or Lavender focuses a human analyst on information that is but a starting point in a complex targeting process. Nevertheless, these explanations seem to have gained little traction in the discourse over the employment of these and other such systems. On 18 June, the IDF published another explanation, offering greater detail.

In this post, I will attempt to provide operational context to the use of the Gospel and Lavender systems. I then assess their relationship to the law of armed conflict (LOAC) rules governing targeting. The piece concludes with thoughts on their use in practice.

Before turning to the systems, one critical caveat is necessary. *Every* weapon and weapon system, and much of the intelligence infrastructure that supports attacks, can be used in a manner that violates LOAC, *either systemically or by individual soldiers*. In that regard, I have no first-hand, on-the-ground knowledge of how the IDF is using Gospel and Lavender, especially in relation to individual strikes. Therefore, my analysis is based solely on media reports, explanations of how the systems are being used, my own experience regarding targeting in general, and extensive discussions with targeting experts. Accordingly, this post deals with the use of the two systems in the abstract, not individual instances of such use.

The Systems

Gospel

According to the IDF, Gospel is not used to identify human targets. Instead, it is limited to pointing intelligence analysts to information on “objects” (e.g., buildings and other structures) that might qualify as military objectives.

Gospel compiles, fuses, and cross-references layers of information from different datasets to generate suggestions for intelligence analysts regarding objects that have the potential to qualify as a military objective. The analyst is required to assess them by reference to a “standard operating procedure” (SOP) that applies the criteria for military objectives found in Article 52(2) of the 1977 Additional Protocol I to the 1949 Geneva Conventions. Israel is not a party to Additional Protocol I, but like the United States, it recognizes the provision as an accurate restatement of customary international law.

While the SOP (and other relevant guidance) is not publicly available, an example of a requirement drawn from U.S. practice is that identification as a potential military objective sometimes requires multiple independent sources of verification. Additionally, from experience, the SOP probably mandates specified forms of verification for some targets as a matter of policy. This is particularly likely for those that are sensitive because they have humanitarian, strategic, or diplomatic significance or enjoy special protection under LOAC.

The IDF describes Gospel as scanning datasets containing information that intelligence analysts have entered from a wide array of sources. Those datasets are presumably drawn from sources like satellite imagery, drone footage, cyber intelligence, phone intercepts, human intelligence, open-source information, and reports from ongoing operations. Reportedly, these data are continuously updated as new information becomes available. It is a process similar to that in which a human analyst would engage. Gospel, however, filters information much more quickly and with greater certainty that all available information on a potential military objective has been factored into the assessment. Think of it as a system that “connects the dots” for the intelligence analyst.

Apparently, targeting analysts can access the raw intelligence upon which a Gospel-generated suggestion is based. Analysts are thereby able to independently assess the quality of the information and the accuracy of determinations such as the object’s location and nature. They may consider other relevant information not found in the system as well.

The IDF has emphasized that Gospel does not definitively determine an object to be a military objective, nor does it pick targets; Gospel’s suggestions alone are not a sufficient basis for concluding that an object is lawfully targetable. Moreover, the system does not assess potential collateral damage for a proportionality analysis or identify viable precautions in attack. Compliance with those LOAC rules is implemented separately during the targeting

process stages that follow target identification. Instead, Gospel simply generates suggestions as to objects that may qualify as military objectives for further assessment by an intelligence analyst.

Analysts may reject the suggestions. But, according to the IDF, if the analyst determines that the object qualifies as a military objective, the determination will be evaluated by at least one higher-level intelligence officer. If approved, the potential target will be passed to a “mission planning cell” (operators, planners, intelligence analysts, legal advisors, etc.) or other body tasked with planning and approving attacks. The cell then determines whether and how to attack the object. Its decisions will be based on a wide array of factors, including the extent to which the attack is time-sensitive, its significance in ongoing operations, the availability of weapon systems capable of striking it, any opportunity cost of attacking the target, whether the attack as planned complies with the rule of proportionality, and whether any precautions in attack are required by law or otherwise appropriate in the circumstances to avoid or minimize collateral damage.

There is no indication that every IDF *pre-planned* targeting operation relies on suggestions generated by Gospel; indeed, I would be surprised if they did. For various operational and other reasons, many targets are surely identified in the traditional manner. Furthermore, in *dynamic* targeting scenarios such as “troops in contact” (TIC) receiving fire from a particular building, Gospel is not needed to determine that the building is a military objective by use. And it may not be operationally feasible to employ the system when attacking targets of opportunity, such as so-called “*fleeting*” targets.

Lavender

Whereas Gospel is a *system* that fuses and cross-references layers of information to identify potential military objectives, Lavender is, technically speaking, a *smart database* that fuses and sorts information on individuals who might be members of Hamas or other organized armed groups. Lavender’s database appears to be vast, which is unsurprising given that Hamas’s *al Qassam* brigades numbered 30-40,000 members when the most recent hostilities broke out. Islamic Jihad adds another few thousand fighters.

Perhaps most importantly, Lavender does not independently identify individuals as members of organized armed groups or direct participants in hostilities who are subject to attack under LOAC. In other words, it does not produce suggested “kill lists” or otherwise serve as a confirmed and approved list of targetable individuals. Further, it has no “predictive” function; it does not “anticipate” who might become a member of any of the Palestinian fighting groups.

Like Gospel, Lavender reportedly accesses an array of datasets that have already been gathered, assessed, and categorized by analysts and which are regularly updated. To illustrate, consider a simple case of an intercept regarding a future rocket attack from the

mobile phone of an individual suspected of being a Hamas operative. An intelligence analyst assessing (per the applicable SOP) whether that person is a targetable member of an organized armed group or direct participant can query Lavender, which will, in turn, provide all the data points collected on that individual. It eases the intelligence analyst's job, thereby accelerating the process while contributing to its comprehensiveness.

Apparently, the same subsequent process as described for Gospel applies. In particular, a higher-level intelligence analyst vets conclusions drawn by more junior analysts based on Lavender information before passing it to those responsible for planning and executing attacks. Only the latter decide whether and how the attack is conducted.

As a general database, Lavender presumably could be used for various purposes. Some are unrelated to targeting, as in the case of identifying individuals who should be detained or questioned. It could also be helpful when planning an operation in an area in order to identify organized armed group members or other persons of interest or to get a sense of the enemy's density in it.

Although Lavender seems useful in determining whether a particular individual is targetable, situations undoubtedly exist, as with objects, in which targetability is obvious. For instance, intelligence analysts would not need to resort to Lavender to assess whether a known *a/-Qassam* Brigades senior leader is targetable as an organized armed group member. Yet, I imagine that Lavender might still be queried to see whether there is recent information about their location (e.g., last whereabouts), habits, etc.

The Law

The law of armed conflict rules governing targeting are relatively straightforward in the abstract, although their application on the battlefield can be complex. Three rules lie at the heart of the law of targeting. First, only military objectives, combatants, members of organized armed groups, and direct participants in hostilities may be attacked. Second, an attacker must take feasible precautions in attack to minimize harm to civilian objects or injury or death to civilians (collateral damage). Third, the expected collateral damage may not be excessive relative to the attack's anticipated concrete and direct military advantage (proportionality).

Gospel and Lavender are directly related to compliance with the first rule. In principle, and if used in accordance with reasonable procedures, they both should enhance the ability of an attacking force to reliably identify legally targetable objects and persons. On the modern battlefield, especially for an armed force with advanced intelligence, surveillance, and reconnaissance (ISR) capabilities like the IDF, the sheer volume of information can overwhelm human capacity to sort through it to identify points that bear on whether a target qualifies as a military objective or targetable person. Overlooked or misrouted intelligence can be a significant contributing factor in targeting errors. Systems such as Gospel and

Lavender can provide more information to analysts or highlight relevant intelligence that might otherwise be missed. They also free up time to assess the object or person more thoroughly.

It might be argued that these tools narrow the information that reaches the analyst. That is true, but they do so by weeding out apparently irrelevant information, allowing the analyst to give greater attention to that which is most relevant, especially in light of requirements regarding independent sources. By separating the wheat from the chaff, Gospel and Lavender ease the information-sifting burden, thereby heightening the reliability of assessments. They can also refresh information faster than human analysts, which lessens the risk of reliance on “stale” data. A corresponding reduction in errors should result. Of course, this benefit depends on the accuracy and reliability of the process generating the Gospel or Lavender output.

The use of Gospel and Lavender also affects the obligation to take feasible precautions in attack. Of particular significance in this regard is the customary law obligation to take feasible steps to verify that the targets to be attacked are not civilians, civilian objects, or other protected persons or objects. The Gospel system and Lavender database facilitate compliance with this obligation by ensuring that all available information regarding the status of a potential target is immediately available. This is a considerable benefit during fast-paced combat operations, especially given the number of human and technological sources feeding information to the IDF. It reduces the likelihood that relevant information might be inadvertently missed. Indeed, it is at least arguable that a failure to use these systems would, in certain circumstances, qualify as a violation of the obligation to take feasible precautions.

Lastly, neither Gospel nor Lavender generates collateral damage estimates to be factored into the proportionality analysis. Nevertheless, the more granular the information the attacker has about the target, the greater the opportunity to plan and execute an attack in a manner that avoids or mitigates potential collateral damage. For instance, knowing what part of a building is being used for military purposes, as distinct from simply knowing the building is being so used, may allow planners to employ tactics and weapons that can avoid collateral damage while still achieving the desired effect on enemy operations.

Gospel and Lavender in Practice

There is nothing inherently unlawful about the use of Gospel or Lavender. On the contrary, they are decision-support systems that can significantly enhance LOAC compliance. But like all weapons and much of the technological infrastructure used to support them, they can be misused in violation of law, policy, and good sense.

Some of the criticism leveled against Gospel and Lavender appears unfounded. For instance, neither system “picks targets” in the sense that they make the final (or any) “decision” regarding target identification, let alone whether to attack the target. Gospel

highlights potential military objectives for human analysts by applying the same criteria they would use if operating without the systems. Indeed, an analyst could perform the same task manually, albeit in most cases without comparable speed and comprehensiveness. Lavender does not appear to go even that far; it is a database available to intelligence analysts. In both systems, humans make the critical decisions, in particular, the determination that an object qualifies as a military objective or a person qualifies as a legally targetable individual, whether to conduct an attack, and how the attack will be conducted

Likewise, criticism that analysts using Lavender sometimes verify targets in “20 seconds” strikes me as somewhat isolated from the broader targeting context, which, for most professional armed forces, comprises a linear series of steps for identifying and verifying targets. Review of Lavender’s output is but one step in that process. A balanced assessment of the database’s use would require an understanding of how its output is generated and an examination of the subsequent steps that are taken before determining that there is sufficient reason to conclude that a person qualifies as a military objective.

More generally, execution of the targeting process is always highly situational. Some targets can be verified very quickly, as with a rocket launcher, whereas others, such as a potential direct participant in hostilities who is not presently engaged in a firefight, demand far greater effort before verifying status. The operational tempo of a battle and the time-sensitivity of potential targets also exert a powerful effect on the process and inform the reasonableness of its execution. Indeed, an array of factors bears on whether the LOAC standard of reasonable grounds to believe an object or person is a lawful target has been satisfied, a subject I have examined elsewhere with a colleague at the Naval War College.

This is not to say that there are no risks. Reports that some intelligence analysts, and perhaps even operational decision-makers, accept whatever the systems generate without further assessment are especially troubling. This might occur because of laziness, dismissive attitudes towards the importance of accurate target verification, undue bias towards machine-generated products, or incompetence. Modern warfare is replete with examples of such failures regardless of the use of AI systems.

But equally, anyone familiar with military operations knows that humans frequently trust their instincts and conclusions more than machine-generated ones. Indeed, there is a risk that intelligence analysts will discount information from the systems and over-confidently rely upon their training and experience. Operational commanders are particularly unlikely to defer entirely to machines, for they often harbor a great deal of confidence in their own abilities. They do not want to lose control of decision-making, for doing so can, as experience teaches, sometimes backfire. The point is that while there is a risk of undue deference to the systems, there is also a tendency by humans in combat to approach machine-generated information with a degree of skepticism (which itself can be a risk to the civilian population).

Concluding Thoughts

There is much debate about the term “artificial intelligence.” If these systems qualify as such, they lie at the lower end of the sophistication continuum. Gospel and Lavender are better characterized as decision-support tools humans use to improve the identification of military objectives and targetable individuals.

But neither relieves intelligence analysts or others involved in the targeting process of their responsibility to make reasonable determinations and decisions in accordance with LOAC and policy. This demands that, as with any weapon, weapon system, or associated infrastructure, those using Gospel and Lavender are well-trained to leverage the systems’ capabilities as a tool, not as a replacement for sound intelligence analysis. And the IDF is responsible in law for ensuring that both systems are employed lawfully and responsibly; it must remain alert to the possibility of abuse. Indeed, Israel should remember that State responsibility attaches even when members of the armed forces act in an *ultra vires* manner (Articles on State Responsibility, art. 7).

Michael N. Schmitt is the G. Norman Lieber Distinguished Scholar at the United States Military Academy at West Point. He is also Professor of Public International Law at the University of Reading and Professor Emeritus and Charles H. Stockton Distinguished Scholar-in-Residence at the United States Naval War College.

Photo credit: Unsplash

SUBSCRIBE

RELATED POSTS

The Legal Context of Operations Al-Aqsa Flood and Swords of Iron

by Michael N. Schmitt

October 10, 2023

—

Hostage-Taking and the Law of Armed Conflict

by John C. Tramazzo, Kevin S. Coble, Michael N. Schmitt

October 12, 2023

–

Siege Law and Military Necessity

by Geoff Corn, Sean Watts

October 13, 2023

–

The Evacuation of Northern Gaza: Practical and Legal Aspects

by Michael N. Schmitt

October 15, 2023

–

A “Complete Siege” of Gaza in Accordance with International Humanitarian Law

by Rosa-Lena Lauterbach

October 16, 2023

–

The ICRC’s Statement on the Israel-Hamas Hostilities and Violence: Discerning the Legal Intricacies

by Ori Pomson

October 16, 2023

–

Beyond the Pale: IHRL and the Hamas Attack on Israel

by Yuval Shany, Amichai Cohen, Tamar Hostovsky Brandes

October 17, 2023

–

Strategy and Self-Defence: Israel and its War with Iran

by Ken Watkin

October 18, 2023

–

The Circle of Suffering and the Role of IHL

by Helen Durham, Ben Saul

October 19, 2023

–

Facts Matter: Assessing the Al-Ahli Hospital Incident

by Aurel Sari

October 19, 2023

–

Iran's Responsibility for the Attack on Israel

by Jennifer Maddocks

October 20, 2023

–

Inside IDF Targeting

by John Merriam

October 20, 2023

–

A Moment of Truth: International Humanitarian Law and the Gaza War

by Amichai Cohen

October 23, 2023

–

White Phosphorus and International Law

by Kevin S. Coble, John C. Tramazzo

October 25, 2023

–

After the Battlefield: Transnational Criminal Law, Hamas, and Seeking Justice – Part I

by Dan E. Stigall

October 26, 2023

–

The IDF, Hamas, and the Duty to Warn

by Michael N. Schmitt

October 27, 2023

–

After the Battlefield: Transnational Criminal Law, Hamas, and Seeking Justice – Part II

by Dan E. Stigall

October 30, 2023

–

Assessing the Conduct of Hostilities in Gaza – Difficulties and Possible Solutions

by Marco Sassòli

October 30, 2023

–

Participation in Hostilities during Belligerent Occupation

by Ioannis Bamnios

November 3, 2023

–

What is and is not Human Shielding?

by Michael N. Schmitt

November 3, 2023

–

The Obligation to Allow and Facilitate Humanitarian Relief

by [Ori Pomson](#)

November 7, 2023

–

[Attacks and Misuse of Ambulances during Armed Conflict](#)

by [Luke Moffett](#)

November 8, 2023

–

[Distinction and Humanitarian Aid in the Gaza Conflict](#)

by [Jeffrey Lovitky](#)

November 13, 2023

–

[Targeting Gaza's Tunnels](#)

by [David A. Wallace](#), [Shane Reeves](#)

November 14, 2023

–

[Refugee Law](#)

by [Jane McAdam](#), [Guy S. Goodwin-Gill](#)

November 17, 2023

–

[After the Conflict: A UN Transitional Administration in Gaza?](#)

by [Rob McLaughlin](#)

November 17, 2023

–

[The Law of Truce](#)

by [Dan Maurer](#)

November 21, 2023

–

International Law “Made in Israel” v. International Law “Made for Israel”

by Yuval Shany, Amichai Cohen

November 22, 2023

–

Cyberspace – the Hidden Aspect of the Conflict

by Tal Mimran

November 30, 2023

–

Israel’s Right to Self-Defence against Hamas

by Nicholas Tsagourias

December 1, 2023

–

Time for the Arab League and EU to Step Up on Gaza Security.

by Michael Kelly

December 4, 2023

–

Attacking Hamas – Part I, The Context

by Michael N. Schmitt

December 6, 2023

–

Attacking Hamas – Part II, The Rules

by Michael N. Schmitt

December 7, 2023

–

Flooding Hamas Tunnels: A Legal Assessment

by Aurel Sari

December 12, 2023

–

Damage to UN Premises in Armed Conflict: IHL and Beyond

by Ori Pomson

December 12, 2023

–

Applicability of Article 23 of the Fourth Geneva Convention to Gaza

by Jeffrey Lovitky

December 13, 2023

–

Delivery of Humanitarian Aid from the Sea

by Martin Fink

December 13, 2023

–

The Question of Whether Gaza Is Occupied Territory

by Michael W. Meier

December 15, 2023

–

Sexual Violence on October 7

by Noëlle Quénivet

December 19, 2023

–

Hostage Rescue Operations and the Law of Armed Conflict

by Kevin S. Coble, John C. Tramazzo

December 20, 2023

Qassam Rockets, Weapon Reviews, and Collective Terror as a Targeting Strategy

by Arthur van Coller

January 17, 2024

A Gaza Ceasefire: The Intersection of War, Law, and Politics

by Marika Sosnowski

January 18, 2024

Information Warfare and the Protection of Civilians in the Gaza Conflict

by Tamer Morris

January 23, 2024

Algorithms of War: Military AI and the War in Gaza

by Omar Yousef Shehabi, Asaf Lubin

January 24, 2024

The Ibn Sina Hospital Raid and International Humanitarian Law

by Michael N. Schmitt

February 1, 2024

Beyond the Headlines: Combat Deployment of Military AI-Based Systems by the IDF

by [Tal Mimran](#), [Magdalena Pacholska](#), [Gal Dahan](#), [Lena Trabucco](#)

February 2, 2024

–

[Ruminations on the Legal, Policy, and Moral Aspects of Proportionality](#)

by [Peter C. Combe](#)

February 9, 2024

–

[Israel's Declaration of War on Hamas: A Modern Invocation of Recognized Belligerency?](#)

by [Benjamin R. Farley](#)

March 5, 2024

–

[Reflections on the Invocation of Common Article 1](#)

by [Tal Mimran](#)

March 13, 2024

–

[Press Access during Armed Conflict](#)

by [Emily E. Bobenrieth](#)

March 18, 2024

–

[Civilian Protection as an Element of Military Advantage in Determining Proportionality](#)

by [Jeffrey Lovitky](#)

April 2, 2024

–

[Targeted Sanctions against West Bank Settlers](#)

by [Larissa van den Herik](#)

May 1, 2024

–

What Happens If the ICC Issues Warrants for Senior Hamas and Israeli Leaders?

by Douglas Guilfoyle

May 2, 2024

–

AI-Based Targeting in Gaza: Surveying Expert Responses and Refining the Debate

by Klaudia Klonowska

June 7, 2024