MAPS DIALOGUES

READOUT — THIRD 2026 MAPS DIALOGUES

Working the *Disagreements*

From Tensions to Traction in Military AI

DATE
10 June 2026

FORMAT
Webinar

AUTHOR
Ulysse Richard

CONTRIBUTORS
Julia Elizabeth Rodríguez Acosta · Kim van der Sluis · Giacomo Persi Paoli · Beyza Unal

United Nations Office for Disarmament Affairs | Ministry of Foreign Affairs

CONTEXT

The following is a non-exhaustive summary of key elements from the discussion at the third webinar of the 2026 Military AI, Peace & Security (MAPS) Dialogues, a webinar series convened by the UN Office for Disarmament Affairs (UNODA), with support from the Republic of Korea, to foster inclusive multilateral dialogue on artificial intelligence (AI) in the military domain.

This webinar, *Working the Disagreements: From Tensions to Traction in Military AI* (10 June 2026), discussed eight tensions observed to run through multilateral discussions on AI in the military domain, ahead of the UN informal exchanges on the subject, taking place on 15–17 June 2026 in Geneva. Rather than seeking consensus, the panel sought to identify the source of each disagreement and the practical measures that remain available despite divergent positions. This readout groups the eight tensions into two categories: substance and process.

The panel featured Julia Elizabeth Rodríguez Acosta (Minister Counselor, Permanent Mission of El Salvador to the UN), Kim van der Sluis (First Secretary for Political Affairs, Permanent Mission of the Kingdom of the Netherlands to the UN), Giacomo Persi Paoli (Head of Programme, Science and Technology, UN Institute for Disarmament Research), and Beyza Unal (Head of the Science and Technology Unit, UNODA), and was moderated by Ulysse Richard (Consultant, UNODA). Jiyoung Yun (Director, Ministry of Foreign Affairs, Republic of Korea) delivered closing remarks. The event was open to representatives of States, international and regional organizations, civil society, the scientific community, and industry.

PANEL



Julia Elizabeth Rodríguez Acosta
Permanent Mission of El Salvador to the UN



Kim van der Sluis
Permanent Mission of the Kingdom of the Netherlands to the UN



Giacomo Persi Paoli
UN Institute for Disarmament Research



Beyza Unal
Science and Technology Unit, UNODA

Moderated by **Ulysse Richard**, UNODA

Eight tensions at a glance

Key takeaways from the third 2026 MAPS Dialogues

TENSIONS AROUND SUBSTANCE

1 · Human judgment versus operational speed

Preserve deliberation — Set by context and risk, with cognitive and technical safeguards — Emphasize operational speed

2 · Civil-military boundary

Separate mandates — Connect processes; govern by function and use — Joint process

3 · Sufficiency of existing international law

Existing law suffices — Clarify how the law applies, via national positions — New rules needed

4 · Technology access versus control

Open access — Assurances that make access trusted — Strict controls

TENSIONS AROUND PROCESS

5 · Strategic direction of travel

Voluntary measures — Combine types of instruments; norms pave the way — Binding instruments

6 · Forum and division of labor

One central process — Umbrella on principles, detail in the verticals — Domain-specific forums

7 · Multi-stakeholder participation

State-led — Structured pathways; States keep the decision — Open participation

8 · Capacity gap

Lack of capacity — A diplomatic toolkit: primers, training, curated knowledge — Meaningful engagement

PART I TENSIONS AROUND SUBSTANCE

1 TENSION

Human judgment versus operational speed

States want both the speed, scale, and information-processing advantages that AI offers and substantive human judgment over consequential decisions. The two pull against each other where decision timelines are compressed.

Panelists treated the value of speed as a question of context: some situations demand deliberate human reflection, while others require a response within seconds. In nuclear decision-making, for instance, speed works against de-escalation. The nuclear command, control, and communications (NC3) cycle has traditionally rested on giving decision-makers enough time to make informed choices, because excessive speed can potentially trigger action-reaction cycles and lead to irreversible decisions based on incomplete or incorrect information. In missile defense, by contrast, the operational requirement is to process information and respond within seconds, and the challenge is not slowing the system down but ensuring that speed does not come at the expense of reliability and accountability. The question for States is therefore which contexts require deliberation, which require rapid action, and what role humans should play across that spectrum.

States should decide which decisions may be delegated to AI systems, under what conditions, and how to ensure that human involvement in those decisions remains substantive. Panelists argued that human involvement must be substantive. An operator who merely approves AI-generated recommendations under time pressure is formally in the loop but exercises little real agency. Panelists indicated that responsibility must remain with humans regardless of a system's level of autonomy. General Assembly resolution 80/23 on AI in nuclear command, control, and communications was cited as an emerging boundary, while systems operating in less critical environments could be subject to more flexible, risk-based measures.

A second layer of disagreement concerns what AI does to human judgment, even when humans remain formally involved. A panelist pointed to research on skill atrophy and agency decay, in which operators gradually lose critical judgment through over-reliance on AI outputs, and on automation bias, that is, the tendency to accept machine recommendations uncritically. The panelist added that the sycophantic behavior of current large language models, which tend to agree with users even when they are wrong, compounds the problem. Drawing on behavioral science, including Daniel Kahneman's book *Thinking, Fast and Slow*, the distinction between fast, intuitive thinking and slower, deliberate reasoning, a panelist observed that stress, uncertainty, and time pressure push humans toward the former, and that AI-enabled systems, producing recommendations faster than human cognition, can amplify that tendency.

Because the challenge is partly cognitive, panelists argued that the solution space should extend beyond technical parameters to include cognitive elements as well. Proposed measures included training operators not only on how to use AI systems but also on how humans make decisions when interacting with these systems, covering automation bias, confirmation bias, and overconfidence; designing systems that prompt critical engagement of humans, for instance by presenting confidence levels and underlying assumptions in ways that invite scrutiny rather than deference; doctrinal safeguards around decision tempo that preserve human deliberation time in high-stakes situations; and extending testing and evaluation beyond technical performance to include how operators behave when working with a system, such as whether they challenge its recommendations.

2 TENSION
Civil-military boundary

> Multilateral processes are organized around the domain a technology originates from, but risk follows the domain in which it is used. Civilian AI adapted for military purposes, and military systems repurposed for civilian uses, both fall between mandates.

Panelists agreed that the most significant AI capabilities originate in the civilian sector, supporting solutions that assess an AI system's function, use, and risk profile rather than its origin. AI is a dual-use technology: the same capabilities that support healthcare, transportation, communications, and commerce can carry security implications. The movement also runs the other way, as the GNS's military origins illustrate. Panelists saw clear value in stronger dialogue between the civilian and military AI communities on testing and evaluation measures, risk management, governance, transparency practices, and human rights considerations.

Technology also flows from the battlefield back to the commercial market. Panelists described surveillance-related technologies used in the military domain, including those deployed in conflict settings that serve as de facto test beds, and then returned to the commercial market, where commercial actors profit by selling them for use in non-conflict settings. AI technologies and their underlying data are becoming a business model that the multilateral AI discussion has not yet addressed.

At the same time, the dual-use character of AI is neither wholly new nor unmanageable, and defense procurement offers decades of relevant practice. Panelists cautioned against two symmetrical errors: treating AI like any other dual-use technology and treating it as so different that its dual-use nature cannot be managed at all. Militaries routinely buy commercial off-the-shelf technology and adapt it. The global positioning system (GPS) has long existed in civilian and military versions built on the same underlying technology, with the military version held to more stringent requirements, such as higher resilience to electronic warfare attacks. Many naval navigation radars are identical to those on cruise ships. The task is to apply that procurement experience to AI while accounting for its unique complexity.

TENSION 2 · **CIVIL-MILITARY BOUNDARY** (CONTINUED)

Current mandates deliberately keep the two domains separate, leaving States uncertain about where boundary-crossing questions belong. General Assembly resolutions on AI in the military domain explicitly exclude the civilian domain, while the Global Digital Compact, the UN's framework for digital cooperation, places military applications outside its primary scope. This creates hesitation among Member States and the UN Secretariat about where certain questions should be discussed. In practice, States tend to find venues for the issues they care about and raise them even in processes whose mandates do not formally cover them.

Rather than merging processes, panelists proposed building connections between them. Suggestions included periodic informal check-ins among the chairs and facilitators of different AI-related processes to identify overlapping problems and solutions, and joint expert briefings during General Assembly sessions, modeled on the biennial practice that already links the outer space sustainability and security agendas. Since questions of testing, evaluation, validation and verification, accountability, and transparency are common to both domains, progress on one side is likely to benefit the other.

3 TENSION
Sufficiency of existing international law

> The General Assembly has reaffirmed that international law applies to AI in the military domain. States differ on whether the task ahead is clarifying how it applies or filling genuine gaps with new rules.

Panelists converged on clarification as the immediate task, noting that the exchange on how international law applies has barely begun. General Assembly resolutions 79/239 and 80/58, along with the Secretary-General's report A/80/78, provide a valuable starting point, and panelists saw no major disagreement on whether international law applies. What is missing is a structured discussion of how States understand the application of international humanitarian law (i.e., the law of armed conflict) and international human rights law to AI-enabled systems. The informal exchanges in Geneva were seen as the place where that baseline could begin to take shape.

Publishing national positions was identified as a proven instrument for building shared understanding. The practice has helped in the cyber domain, where States' published views on the application of international law provided a foundation for discussion. International law applies across technological contexts, but States may hold different views on how specific rules and principles apply to emerging technologies. Panelists argued for building that collective understanding before determining whether additional measures are needed, and they observed that the Secretary-General's report already compiled many national positions.

The clarification agenda has concrete anchors across the AI lifecycle. Panelists raised questions about how to embed international law in systems at the design stage; privacy concerns under human rights law during development and use; and accountability, including where responsibility lies when an operator makes a judgment under time pressure and how the chain of accountability should be set out in advance. Discussions taking place outside the UN may also inform the UN's own work on these questions.

4 TENSION
Technology access versus control

States broadly agree that military-relevant AI should not proliferate irresponsibly. They disagree on where legitimate non-proliferation concerns end and exclusion begins: one group sees export controls and licensing as necessary safeguards, another sees the same measures as entrenching technological dependence.

Panelists argued that security and inclusivity should not be treated as competing objectives. States have legitimate non-proliferation concerns and responsible-use obligations, but all States also have a claim to participate in shaping AI governance and to benefit from technological development. Rather than drawing lines around who may use the technology, panelists proposed advancing transparency, risk-management practices, capacity-building, and responsible-use commitments that serve both objectives simultaneously.

Reframing the question from who gets access to what assurances make access trusted gives States an existing toolkit to build on. Export control practices developed outside AI, such as end-user verification and licensing along the value chain, remain relevant baselines. Detecting violations of end-user agreements is harder with AI than with traditional transfers, but panelists argued that this changes the difficulty, not the principle: established practice can anchor the design of AI-specific assurance measures.

The role of industry makes access to AI different from earlier technology-transfer debates. AI capabilities are unlikely to be accessed without industry being part of the conversation, which creates a new set of relationships among the governments where companies are registered, the governments purchasing access to capabilities, and the companies themselves. Recent events have shown how much power commercial licenses and user agreements exert over the use of capabilities. One panelist proposed a dedicated roundtable or working group to determine the right combination of State-regulated technology transfer and commercially regulated licensing and user agreements.

PART II TENSIONS AROUND PROCESS

5 TENSION
Strategic direction of travel

> States differ on the preferred outcome of multilateral processes: binding or voluntary instruments, and whether "responsibility" is a sound organizing concept or a subjective framing advanced by a particular group of States.

Panelists rejected the choice between voluntary and legally binding measures as a false binary. International security processes have advanced incrementally: the framework of responsible State behavior in cyberspace combines norms, confidence-building measures, political commitments, transparency measures, and capacity-building, while preserving space for negotiation of stronger instruments. The realistic next step, panelists argued, is to keep identifying areas of convergence and to use the available tools in combination rather than forcing an early choice.

Responsibility-based and capability-based approaches regulate different things, and States rarely rely on either exclusively. Responsibility-based approaches, familiar from the cyber and outer space processes, regulate behavior: how a technology is used and who is accountable for that use. Capability-based approaches, typical of legally binding instruments, regulate which technologies are developed, possessed, and deployed, for example, through restrictions on a weapon's payload, size, or weight. Some States object to the responsibility framing, indicating that it is subjective; others respond that intangible technologies cannot be regulated through capabilities alone. In practice, national policy often encompasses both, and the approaches can be combined: norms can pave the way for binding obligations, harden into customary international law, or be derived from a better understanding of the application of existing law.

Existing disarmament and arms control processes already offer a full menu of models, with the cyber process considered the best model for AI. Panelists saw no need to design something new. The cyber process, having led recently to the establishment of the "Global Mechanism on developments in the field of ICTs in the context of international security and advancing responsible State behaviour in the use of ICTs," gave States a standing forum to discuss threats, international law, norms, confidence-building measures, capacity-building, and, eventually, how to institutionalize the dialogue, a breadth judged well suited to AI at its current stage. As a political caveat, panelists highlighted that no process design can substitute for the political will of States, and the larger risk is that the geopolitical context slows multilateral progress, pushing discussions outside the UN. Panelists insisted that these discussions need to remain at the UN, where all 193 Member States can contribute on an equal footing. Progress made elsewhere may be faster, but it is, by definition, less inclusive and less universal.

Responding to an audience question about why States resist binding rules today despite negotiating treaties in the past, panelists described the hesitation as a matter of sequencing rather than principle. In disarmament history, treaties have typically come at the end of long periods of political debate, after States developed a shared understanding of the problem and built confidence through intermediate measures. No comparable shared understanding of AI exists yet, and some States worry that specific binding rules would become obsolete as the technology evolves. Parts of the field are nonetheless already on a binding track: discussions on lethal autonomous weapons systems engage with legally binding options, political commitments are emerging on AI and nuclear decision-making, and normative approaches dominate elsewhere. For many developing countries, the issue is not opposition to binding rules but the desire to build the capacity to negotiate them meaningfully first.

6 TENSION
Forum and the division of labor

> AI is a general-purpose technology that cuts across many regimes, while the multilateral system is organized vertically by issue. The question is how to address a horizontal technology across the cyber, autonomy, biological, outer space, and nuclear agendas without duplication or gaps.

A detailed, domain-specific discussion of AI risks is both inevitable and desirable. AI in the cyber domain poses different risks than AI in the biological domain or elsewhere, so each community will produce its own assessment of risks and mitigation strategies, and the richness of the discussion stems from those specialized settings. Panelists also warned against a single mandate covering all sub-issues, which would overwhelm the system.

To date, there is no dedicated space for this domain-specific knowledge to inform deliberations. Domain-specific assessments remain siloed within their verticals because no platform exists to make them common, shared knowledge. Options raised included expanding the time the First Committee (the General Assembly's main committee on disarmament and international security) devotes to these exchanges and creating a scientific panel dedicated to AI in the context of international peace and security. Because States have deliberately kept civilian and security discussions separate, the Independent International Scientific Panel on AI, being set up on the civilian side is unlikely to address security-specific questions.

Light coordination mechanisms can connect the verticals without merging them. Panelists returned to periodic check-ins among the chairs of different processes and to joint or biennial expert briefings, alongside simple confidence-building measures borrowed from the cyber experience, such as publishing national positions. One panelist suggested making "how does AI affect this issue" an automatic reflex across UN processes, so that intersections are caught without new structures, while keeping coherence in how the discussions move forward.

Any dedicated process on AI and international peace and security should be scoped to focus on high-level strategic principles. A single process cannot produce solutions applicable across all domains. The model described was an umbrella: an agreement at the strategic level that is then absorbed, unpacked, and discussed in detail within each vertical. The detailed work remains where the expertise lies, and the umbrella keeps the pieces aligned.

7 TENSION
Multi-stakeholder participation

> The expertise these discussions require sits largely outside governments, while States hold the responsibility for decisions. The tension is around how the knowledge of external stakeholders can inform deliberations without undermining State-led decision-making.

Expertise largely lies outside governments, and decisions remain the prerogative of States. Much of the technical knowledge resides in industry and research organizations, operational experience is held by militaries, and legal expertise is found in academia and civil society, while many of the societal consequences are felt by communities absent from security discussions.

A structured way to include industry is missing, despite its potential to inform and implement commitments. Contrary to the perceived resistance from industry to engaging, panelists argued that many companies want to participate in multilateral discussions but do not know how. For instance, only one major technology company, Microsoft, submitted a contribution to the Secretary-General's 2025 report. Panelists argued for broadening the definition of the AI industry beyond a small number of large, mostly Western, technology companies to include AI developers, semiconductor companies, cloud providers, startups, and firms in adjacent sectors whose work has peace and security implications, drawn from across the globe.

Participation models exist. Proposals included structured industry-led briefings to Member States ahead of key meetings, drawing lessons from the civilian AI discussions, where industry involvement is further advanced, and creating a directory of the global AI industry at the Secretariat through which engagement between Member States and industry could be organized. Decisions are stronger when industry and other stakeholders are informally involved, panelists noted, but they remain for States to take.

8 TENSION

The capacity gap

> "Capacity" is used to mean two different things: the capacity to access and field the technology, and the capacity to govern and negotiate it. On both counts, the least-resourced States are the least able to shape the rules.

For many delegations, the question of binding versus voluntary instruments is premature; the prior question is the ability to engage at all. One panelist explained that a focus on capacity-building does not signal unwillingness to negotiate normative frameworks or binding instruments. It reflects the need for a basic understanding of highly technical subject matter, so that participation is meaningful and the resulting instruments serve the interests of all Member States. Participation in the discussions is itself a form of capacity-building, and practices from cyber processes, such as financing delegates' attendance, have broadened it further.

A practical diplomatic toolkit could lower the barrier to entry, starting with curation rather than creation. Asked what a military AI diplomatic toolkit would need to include, one panelist recommended starting with a curated repository of existing knowledge: high-quality research papers and primers already exist, and diplomats need a single place to find them, organized by topic or application. The second recommendation was a starter course on AI diplomacy for diplomats approaching the topic for the first time, modeled after the recently redesigned cyber diplomacy course.

Training a base of diplomats before negotiations gain traction ensures the community will not start from zero. The Women in AI Fellowship, modeled after the Women in Cyber Fellowship, was cited as a working example. It builds a base of diplomats with an initial understanding of the topic, so the community is prepared when multilateral discussions accelerate. As processes evolve, the specific gaps delegations face can be identified and addressed in turn. The MAPS side event in Geneva on the morning of 16 June will gather input on the toolkit's design.

CONCLUSION

Panelists converged on sequenced, incremental progress as the realistic path from tension to traction. Panelists offered several recommendations ahead of the informal exchanges:

→ Keep the entire AI lifecycle in view and pair it with a well-designed mechanism for engaging industry

→ Invest in capacity-building so that every delegation can engage meaningfully and the resulting deliberations benefit all Member States;

→ Remember that in disarmament, shared understanding and confidence-building have historically preceded treaties rather than followed them;

→ Do not let the complexity of the issue become a barrier to taking the first step. As one panelist put it, delegations do not need to solve everything at once; they need to identify a hook all can agree on, whether confidence-building measures, capacity-building, or a more detailed discussion of threats, and start there.

Across the eight tensions, the discussion consistently moved away from binary choices and toward context-specific combinations. Panelists shifted the focus from taking a side to identifying the contexts, conditions, and sequences that allow both concerns to be addressed, whether around speed or judgment, civilian or military, voluntary or binding, or access or control.

Closing the 2026 MAPS Dialogues on behalf of the Republic of Korea, Director Jiyoung Yun placed the webinar within the arc of the series. The first webinar showed that productive dialogue begins with a shared vocabulary; the second brought a lifecycle perspective spanning the eleven stages set out in General Assembly resolution 80/58, from pre-design to decommissioning; and this third webinar turned to the political terrain where positions diverge. The Republic of Korea expressed the hope that the common vocabulary, the lifecycle frame, and the areas of convergence identified throughout the series will accompany participants to Geneva and provide a foundation for the General Assembly resolution on AI in the military domain.