APLN

Artificial Intelligence in the Australian Defence Forces

Strengthening Denial, Managing Escalation

APRIL 2026

AINA TURILLAZZI

This paper was supported through a general core grant from the Nuclear Threat Initiative.

Please direct inquiries to: Asia-Pacific Leadership Network (APLN)
Secretariat 4th floor, 116, Pirundae-ro,
Jongno-gu, Seoul, ROK, 03035
Tel. +82-2-2135-2170
Fax. +82-70-4015-0708
Email. apln@apln.network

This publication can be downloaded at no cost at www.apln.network.

Cover Photo: A Boeing MQ-28 Ghost Bat is on display at the 2025 Avalon International Airshow in Avalon, Australia, in March 2025. (Alexander Bogatyrev/SOPA Images/LightRocket via Getty Images)

ARTIFICIAL INTELLIGENCE IN THE AUSTRALIAN DEFENCE FORCES: STRENGTHENING DENIAL, MANAGING ESCALATION

Aina Turillazzi

AUSTRALIA'S AI DEFENCE BUILD-UP

Australia's defence sector is undergoing a structural shift in how it generates military advantage. Rather than being introduced as a standalone capability, Artificial Intelligence (AI) is being integrated into the information architecture that underpins targeting, surveillance, battle management, logistics, and force coordination. The result is a force whose operational effectiveness depends increasingly on data and algorithms rather than platform performance alone.

The 2024 Integrated Investment Program (IIP), published together with the National Defence Strategy (NDS), allocates A$330 billion for capability investment over the decade to 2033–34.[1] Within that envelope, AI is largely financed through the same machinery that funds connectivity, data infrastructure, and command and control (C2). The most concrete investment is the A$8.5–$11 billion committed to enterprise data and information and communication technology (ICT).[2] This is the digital backbone that enables faster intelligence fusion and decision support across the Australian Defence Force (ADF).

These investments underpin the OneDefence Data Program, a whole-of-Defence initiative that consolidates fragmented data systems into a standardised environment designed to deliver decision advantage through advanced analytics and AI.[3] The Advanced Strategic Capabilities Accelerator (ASCA) then serves as the mechanism that takes what that foundation enables and pushes it into fielded capability. The IIP commits up to A$3.8 billion over the decade to ASCA, established to rapidly develop and transition asymmetric capabilities into the ADF by connecting industry and universities with military end-users.[4] The unmanned combat aerial vehicle (UCAV) MQ-28A Ghost Bat and the Ghost Shark extra-large uncrewed underwater vehicle (XL-UUV) are the most visible examples of this trajectory.

[1] Department of Defence, *2024 Integrated Investment Program* (Canberra: Australian Government, 2024), https://www.defence.gov.au/about/strategic-planning/2024-national-defence-strategy-2024-integrated-investment-program.
[2] This A$8.5–$11 billion investment includes funding for enterprise networks at A$5.8–$7.8 billion and enterprise systems at A$2.7–$3.7 billion.
[3] Department of Defence, *2024 Integrated Investment Program*,
[4] Department of Defence.

This domestic investment does not sit in isolation. Australia is developing systems under AUKUS Pillar II, which includes AI alongside undersea capabilities, cyber, electronic warfare, hypersonics, and quantum technologies.[5] Although alliances are outside the scope of this brief, Pillar II priorities shape Australia's AI capability trajectory: domestic investment is accompanied by shared development frameworks that accelerate access to advanced technologies.

The NDS makes a "Strategy of Denial" the cornerstone of the Australian Department of Defence (henceforth, DoD or Defence) planning, aiming to deter conflict before it begins and to prevent a potential adversary from using force to coerce Australia, by convincing it that military action would not succeed at an acceptable cost.[6] That logic demands decision advantage and an integrated force design, because denial depends on seeing first, understanding faster, and coordinating across different domains. AI matters in this context because it accelerates the processing and sharing of information that modern denial requires.

Royal Australian Navy

Australia's maritime strategy depends on sustained situational awareness across vast ocean approaches and underwater infrastructure. In the Navy, AI is being adopted primarily as an underwater surveillance and decision support enabler rather than a standalone capability. The IIP's funding architecture reflects that logic clearly.

The IIP allocates A$63–$76 billion to underwater warfare capability over the decade.[7] This portfolio includes nuclear powered submarines and infrastructure, the *Collins*-class, and a dedicated line for underwater warfare and uncrewed maritime systems. The *Collins* line, for example, includes upgrades to the sonar suite, reinforcing the importance of underwater sensing and signal processing.[8] Alongside submarines, Australia is also investing A$550–A$650 million in underwater range systems and warfare facilities, providing the testing, training, and the engineering infrastructure needed to make advanced sensing and automation usable.[9]

[5] John Christianson, Sean Monaghan, and Di Cooke, "AUKUS Pillar Two: Advancing the Capabilities of the United States, United Kingdom, and Australia," Center for Strategic and International Studies, July 10, 2023, https://www.csis.org/analysis/aukus-pillar-two-advancing-capabilities-united-states-united-kingdom-and-australia.
[6] Department of Defence, *2024 Integrated Investment Program*.
[7] The A$63–$76 billion is the total planned investment, noting that only A$14 billion have been approved and A$48–$61 billion is still unapproved planned investment, as of March 2026. See Department of Defence.
[8] Department of Defence.
[9] Department of Defence.

The same operational logic extends beyond the underwater domain. The IIP funds A$51–$69 billion in maritime capabilities for sea denial and localised sea control operations, including an expanded surface combatant fleet.[10] The latter includes upgraded *Hobart*-class destroyers, *Hunter*-class frigates, 11 new general purpose frigates, and six Large Optionally Crewed Surface Vessels. Across all of these platforms, the common requirement is the same: reliable sensing, rapid identification of contacts, and the ability to combine information from multiple sources into a shared operational picture.[11]

Within this architecture, the Integrated Undersea Surveillance System (IUSS) is best understood as a shift from "platform awareness" to networked underwater awareness. The IUSS acquisition project, scheduled from 2025 to 2040, has a budget provision of €3.38–€5.0 billion (A$5.6–A$8.3 billion).[12] Its stated purpose is to bring into service an integrated underwater surveillance system and, in doing so, assess the utility of crewed vessels, uncrewed surface vessels (USVs), and uncrewed underwater vehicles (UUVs).[13] In practical terms, this points to a surveillance approach in which distributed sensors and platforms feed a coherent detection and tracking picture, rather than relying on single assets operating in isolation.

The Royal Australian Navy (RAN) is already fielding uncrewed systems that make this direction concrete. The IIP confirms continued acquisition of Ocius Bluebottle USVs, supported by the dedicated underwater support and trials vessel Australian Defence Vessel (ADV) Guidance, which can deploy and support underwater crewed and uncrewed vehicles and undersea surveillance system trials.[14] Bluebottle is described as incorporating AI neural networks and edge computing for processing sensor signals, supported by low-bandwidth communications and "team" software that allows vessels to coordinate their behaviour.[15] Exercises have also demonstrated integration with other underwater systems and maritime aviation, including a 2023 exercise where Bluebottles worked with US Navy unmanned surface forces to track an underwater vehicle and then cue an MH-60R helicopter for follow-on action.

---

[10] Sea denial is a maritime operational concept and should not be conflated with the NDS "strategy of denial" (a national defence strategy). See Jennifer Parker, *An Australian Maritime Strategy: Resourcing the Royal Australian Navy* (Barton, ACT: Australian Strategic Policy Institute, 2023), 17; Jennifer Parker et al., *A Maritime Strategy for Australia 2035* (Canberra: UNSW Canberra, 2025), https://www.unsw.edu.au/content/dam/pdfs/unsw-canberra/hass/A%20Maritime%20Strategy%20for%20Australia%202035.pdf.

[11] Department of Defence, *2024 Integrated Investment Program*.

[12] Peter Layton, "Evolution Not Revolution: Defence AI in Australia," in *The Very Long Game*, ed. H. Borchert, T. Schütz, and J. Verbovszky (Cham: Springer, 2024), https://doi.org/10.1007/978-3-031-58649-1_26.

[13] Layton, "Evolution Not Revolution: Defence AI in Australia."

[14] Department of Defence, *2024 Integrated Investment Program*.

[15] Layton, "Evolution Not Revolution: Defence AI in Australia."

The IIP reinforces that this is not a one-off trial. It allocates A$5.2–A$7.2 billion for subsea warfare and uncrewed maritime systems, including large and extra-large UUVs and further Bluebottle vessels.[16] Ghost Shark sits at the far end of this trajectory. As Australia's XL-UUV program, it represents a shift toward underwater systems capable of collecting, processing, and sharing information with significantly less constant human control, making it the clearest signal of where the RAN's AI integration is heading.

The Navy's AI adoption is therefore most advanced where the demands of denial are most acute: in the systems that extend Australia's undersea awareness, reduce dependence on crewed platforms, and feed usable intelligence across a distributed maritime force.

Australian Army

The Army's AI adoption is less visible in individual platforms and more evident in the systems and practices that make brigades function. Rather than presenting AI as a doctrinal shift, the Army is treating it as a practical way to improve command support and reduce the cognitive burden on commanders in complex environments. This is consistent with an "evolutionary" pathway in which professional adaptation, and experimentation shape how AI is absorbed into routine operations.[17]

That approach sits inside a wider land modernisation effort. The 2023 Defence Strategic Review (DSR) mandated a structural shift for the Army: away from land-centric operations and toward an amphibious-capable, littoral-manoeuvre force suited to Australia's northern approaches and the island chains of Southeast Asia.[18] Land capability lines in the IIP allocate A$36–A$44 billion to realise this, building a combined-arms system optimised for coastal operations in a contested maritime environment.[19] For the Army, most AI-relevant funding sits inside this modernisation effort and land command-system upgrades, rather than in a single labelled "AI programme." Within that envelope, AI-enabled tools are being incorporated into C2, intelligence processing, and targeting prioritisation. The intent is not to replace human judgment, but to make sense of information faster and coordinate effects more reliably.

[16] Department of Defence, *2024 Integrated Investment Program*.
[17] Alex Neads, Theo Farrell, and David J. Galbreath, "Evolving towards Military Innovation: AI and the Australian Army," *Journal of Strategic Studies* 47, no. 5 (2024): 669–688, https://doi.org/10.1080/01402390.2023.2200588.
[18] Department of Defence. *Defence Strategic Review 2023* (Canberra, Australian Government, 2023), https://www.defence.gov.au/about/reviews-inquiries/defence-strategic-review.
[19] Department of Defence, *2024 Integrated Investment Program*.

The Army's approach to AI is not to wait for a mature product to arrive but is instead building the concepts and institutional knowledge needed to absorb autonomous systems at scale. This is mainly done through experimentation. Take the M113 Optionally Crewed Combat Vehicle (OCCV),[20] supported by the Army's Robotic and Autonomous Systems Implementation and Coordination Office (RICO), has been used to test how autonomy and AI-enabled vehicle management can be integrated into battlefield concepts, including live fire experimentation and remote operation of weapon stations.[21] Separate experimentation study autonomous C2-related technologies, uncrewed systems and counter-unmanned aircraft systems (UASs) capabilities.[22]

Against that backdrop, the Future Autonomous Vehicles project is best read as the acquisition pathway that gives today's experimentation a destination. It is described as acquiring a fleet of uncrewed systems sufficient for up to brigade-sized operations,[23] building from the M113 experimentation programme, with an acquisition phase scheduled from 2033 to 2040 and a planning budget of €5–€7.5 billion (A$8.3–A$12.5 billion).[24] The significance for the present is that the Army is already organising concepts, integration work, and command-system adaptation around the expectation that uncrewed systems will be absorbed at brigade scale, not as trials.

The Army's AI trajectory is therefore less about individual platforms and more about institutional readiness: building the concepts, command systems, and experimentation record needed to operate autonomously at scale when the capability matures.

Royal Australian Air Force

The integration of AI in the Royal Australian Air Force (RAAF) is most visible where air power is increasingly won or lost, namely, in intelligence, surveillance and reconnaissance (ISR) exploitation and air battle management. Three investments illustrate this trajectory across the air domain's information chain.

[20] The M113 AS4 OCCV is a converted variant of the Australian Army's legacy M113 Armoured Personnel Carrier, designed to operate either with a human crew or as an autonomous/remotely piloted system.
[21] Layton, "Evolution Not Revolution: Defence AI in Australia."
[22] Malcolm Davis, "Aligning for Advantage: Integrating Autonomous Systems into the Australian Defence Force," Australian Strategic Policy Institute, December 8, 2025, https://www.aspi.org.au/report/aligning-for-advantage-integrating-autonomous-systems-into-the-australian-defence-force/.
[23] S. Kate Devitt and Damian Copeland, "Australia's Approach to AI Governance in Security and Defence," in *The AI Wave in Defence Innovation: Assessing Military Artificial Intelligence Strategies, Capabilities, and Trajectories*, ed. Michael Raska and Richard A. Bitzinger (London: Taylor & Francis Group, 2023), 289–325, https://doi.org/10.4324/9781003218326-11.
[24] Layton, "Evolution Not Revolution: Defence AI in Australia."

First, MQ-28A Ghost Bat (Loyal Wingman) sits within the Teaming Air Vehicles project as the RAAF's clearest move toward remotely piloted or AI-enabled aircraft designed to complement crewed combat platforms and expand air combat capacity. While its acquisition window extends from 2026 to 2040, with a budget provision of €5–€7.4 billion (A$8.30–A$12.28 billion),[25] its relevance to the current posture lies in how it anchors work on human–machine teaming and the integration of uncrewed aircraft into air combat concepts.

Second, the Distributed Ground Station Australia project (2024–2031), estimated at €0.8–€1.2 billion (A$1.33–A$1.99 billion), targets the less visible but decisive problem of ISR exploitation. It is designed to process, exploit, and disseminate intelligence collected by RAAF platforms, addressing a problem that receives less attention than platform acquisition but is equally decisive for operational effectiveness.[26]

Third, the Joint Air Battle Management System (JABMS) (2023–2031), budgeted at €1.2–€1.9 billion (A$1.99–A$3.15 billion),[27] strengthens air and missile situational awareness and interoperability with allies. Its significance for AI integration lies in its use of machine learning to help the ADF maintain coordination when communications and radar are contested, operating within an open architecture designed for joint and coalition use.

Taken together, these programmes show that RAAF AI adoption is centred on connecting sensor data to operational decisions more rapidly, while enabling a future air combat mix that blends crewed and uncrewed platforms. That points to a broader implication: AI progress in the air domain will be constrained less by aircraft hardware than by the ability to move, process, and act on information reliably when networks and the spectrum are contested.

Cross-force enablers: Data, cyber, and computing capacity

AI integration across the Navy, Army, and Air Force is only as credible as the enabling architecture behind it. In Australia's case, the investments are not only in platforms, but in the connective tissue that allows data to be collected securely, moved at speed, processed at scale, and translated into operational decisions across domains.

[25] Layton, "Evolution Not Revolution: Defence AI in Australia."
[26] Layton.
[27] Layton.

The IIP allocates A$27–A$36 billion to space and cyber, A$11–A$15 billion to theatre C2, and A$15–A$21 billion to theatre logistics.[28] Together, these lines fund the systems that underpin AI-enabled operations in practice: the networks that move data securely, the command systems that translate it into coordinated action, and the logistics layer that keeps forces operational under disruption. Without this foundation, decision advantage and human-machine teaming remain confined to isolated trials.

The IIP also includes an information and cyber domain "Emerging and Disruptive Technologies" project, with a budget provision of €1.14–€1.7 billion (A$1.89–A$2.82 billion) from 2033 to 2040.[29] This sits beyond the immediate window of today's fielded capability, but it matters as a planning signal, namely that DoD expects new digital and AI-related capabilities to transition from experimentation into acquisition pathways. In parallel, Defence science and technology infrastructure receives A$1.7–A$2.2 billion over the decade, supporting high-performance computing and digital research systems that are essential for model development, testing, and secure evaluation.[30]

Governance settings also shape what can be fielded and how it is controlled. Australia's AI governance framework for security and defence requires that systems meet standards of lawful use and human accountability before operational deployment.[31] These frameworks do not substitute for strategy, but they influence design choices and the degree of autonomy permitted in operational settings.

The cross-force enablers described above explain why AI in Australian defence is now structural rather than experimental: the service branches are being connected through shared data infrastructure and secure networks that make joint AI-enabled operations possible. This sets the conditions for the next question, namely, how these AI-enabled systems affect regional stability, including both stabilising effects and escalation risks.

Implications for Regional Stability and Escalation Risk

Australia's current approach to AI integration in the military rests on a clear operational logic. Investments in decision advantage and human–machine teaming are designed to produce a faster, better coordinated force that is harder to coerce. The assumption underpinning this approach is straightforward: better data and faster analysis, translate directly into strategic stability by reducing uncertainty and strengthening conventional credibility.

[28] Department of Defence, *2024 Integrated Investment Program*.
[29] Layton, "Evolution Not Revolution: Defence AI in Australia."
[30] Department of Defence, *2024 Integrated Investment Program*.
[31] Devitt and Copeland, "Australia's Approach to AI Governance in Security and Defence."

The question, however, is not whether AI increases capability. The question is how these capabilities interact with recognised mechanisms of escalation risk in a region where strategic rivalry is compounded by overlapping alliance commitments and nuclear armed competitors.

This section assesses those interactions. It proceeds by examining stabilising effects first, then working through the escalation mechanisms that AI integration intersects with most directly.

How AI Can Stabilise Crisis Dynamics

AI integration into ISR can, under certain conditions, reduce uncertainty by improving how forces detect, track, and interpret activity. Better integration of sensor data and faster analytics can reduce false alarms and mistaken attribution, especially in crowded maritime and air environments. A RAND Corporation analysis of AI in the military notes that improved information processing can, in principle, reduce miscalculation if embedded within command processes.[32] In a maritime standoff scenario, for example, more accurate tracking of vessels and aircraft could prevent mistaken attribution or accidental engagement. In Australia's case, this stabilising potential is most relevant where AI is being embedded in ISR exploitation and battle management, namely the systems intended to convert large sensor feeds into operationally usable information.

Decision support systems (DSS) can also strengthen crisis management by reducing cognitive overload and improving coordination. The Stockholm International Peace and Research Institute (SIPRI) observes that AI tools may enhance early warning and coordination, but the benefit depends on how they are integrated into human decision-making.[33] This is directly relevant to Australia's investments in the JABMS and the Distributed Ground Station Australia project, both of which are explicitly designed to support judgment rather than substitute for it, converting large sensor feeds into operationally usable information while keeping human control central to the process.

AI also matters in the less visible domain of logistics, namely, the systems that keep forces supplied and equipment working. By making sustainment more predictable, from maintenance scheduling to resupply under disruption, AI reduces the operational fragility that pushes commanders toward early escalation. If forces can be confident that they will remain supplied, the pressure to act quickly before being cut off diminishes, and with it one incentive for early escalation.

---

[32] James Black et al., *Strategic Competition in the Age of AI: Emerging Risks and Opportunities from Military Use of Artificial Intelligence* (Santa Monica, CA: RAND Corporation, 2024), https://doi.org/10.7249/RRA3295-1.

[33] Vladislav Chernavskikh and Jules Palayer, *Impact of Military Artificial Intelligence on Nuclear Escalation Risk*, SIPRI Insights on Peace and Security No. 2025/06 (Stockholm: Stockholm International Peace Research Institute, 2025), https://doi.org/10.55163/FZIW8544.

These stabilising effects are real, but they are conditional. The same systems that improve speed and coordination can also compress the time available for deliberation while creating new forms of overreliance. This is why the question is not whether AI improves performance, but how it changes crisis dynamics between competitors.

Inadvertent Escalation: Decision-Time Compression

One of the most widely recognised mechanisms of escalation risk is decision-time compression. By accelerating ISR exploitation and decision support, AI shortens the time available between detection, interpretation, and response, particularly in high-tempo conventional operations. This increases the risk of inadvertent escalation because leaders have less deliberative space,[34] and even non-nuclear AI applications can shape nuclear risk environments by intensifying pressure for rapid decisions.[35]

Australia's emphasis on decision advantage, rapid ISR processing, and human–machine teaming interacts directly with this dynamic. When data is processed faster and targeting cycles shorten, the window for political assessment narrows. The risk is not that Australia "automates escalation," but that stakeholders are pushed into faster, important choices, before intentions and signals are fully clarified, including whether to raise readiness or authorise pre-positioned responses.

Consider a regional scenario such as a Taiwan crisis. Rapid indications of missile activity or undersea movement could trigger accelerated alerting across coalition networks before the picture is complete. Even with human control retained, the speed of warning and information sharing can generate momentum, making it harder for diplomacy to keep pace with military posture. In such settings, escalation can emerge from tempo and uncertainty rather than deliberate choice.

Classical escalation theory underscores that once movement along an escalation ladder begins, actors cannot reliably control where it stops.[36] AI-enabled time compression intensifies this problem because it increases the probability of early threshold crossings driven by worst-case interpretation and fear of falling behind. For Australia, whose investments (e.g., JABMS, IUSS) are explicitly designed to accelerate the detection-to-decision cycle, this is not a theoretical concern but a design feature that requires deliberate management.

---

[34] James Johnson, "Inadvertent Escalation in the Age of Intelligence Machines: A New Model for Nuclear Risk in the Digital Age," *European Journal of International Security* 7, no. 3 (2022): 337–359, https://doi.org/10.1017/eis.2021.23.

[35] Chernavskikh and Palayer, *Impact of Military Artificial Intelligence on Nuclear Escalation Risk*.

[36] Lawrence Freedman, "On the Tiger's Back: The Development of the Concept of Escalation," in *The Logic of Nuclear Terror* (London: Routledge, 2020), https://doi.org/10.4324/9781003081111-6.

Accidental Escalation: Socio-Technical Interaction and Overreliance

Accidental escalation does not require a malfunction to occur. It can emerge from socio-technical dynamics, namely the interaction of people, procedures, and AI tools under stress. The opacity and probabilistic nature of AI systems, sometimes referred to as a "black box," shapes human decision-making in ways that are difficult to detect precisely because they operate beneath the level of conscious choice.[37] AI-enabled systems introduce uncertainty not only through model error but through the compounding effects of data quality and system interactions that are hard to judge reliably in real time.[38]

In the Australian context, this matters because AI is being integrated into ISR exploitation, and maritime and undersea awareness, areas where early warning and rapid interpretation can drive posture changes. When operators rely on predictive outputs and confidence scores, there is a risk of automation bias: treating the machine's judgment as more reliable than it is. A false positive (e.g., an anomalous pattern misread as hostile activity) or mislabelling (e.g., misidentifying an object or signal) may not be recognised quickly, particularly in a crisis when information is incomplete and decision tempo is high.

SIPRI's assessment of AI and nuclear escalation risk notes that automated DSS can push decision makers toward faster reaction, increasing the risk of miscalculation.[39] The concern is not autonomous nuclear launch. It is the gradual normalisation of machine-mediated judgement in high-pressure contexts, where the human role shifts from assessing evidence to confirming a recommendation. The lesson of historical false-warning episodes is that crisis stability often depends on the capacity to question "authoritative" alerts, especially when the cost of being wrong is high. The classic analogue is the 1983 Petrov incident, when a Soviet officer judged that a satellite warning of incoming US missiles was likely a false alarm, and his decision not to escalate is often cited as an example of how human judgement can prevent catastrophe.

In the Australian case, these risks are not hypothetical. The Intelligent Automated Decision Superiority (i-ADS) initiative is explicitly designed to compress the ISR-to-decision cycle, automating the filtering and prioritisation of intelligence to generate actionable outputs at speed. Automation bias is a known concern in high-tempo maritime surveillance environments, where operators may be less likely to question or override AI-enabled analysis under pressure. The Army's exploration of autonomous first contact creates scenarios where the pace of engagement outstrips meaningful human control before commanders have a complete picture. Current strategies appear too willing to remove humans from the loop in pursuit of speed, a tension that sits

[37] Johnson, "Inadvertent Escalation in the Age of Intelligence Machines: A New Model for Nuclear Risk in the Digital Age," 337–359.
[38] Black et al., *Strategic Competition in the Age of AI: Emerging Risks and Opportunities from Military Use of Artificial Intelligence*.
[39] Chernavskikh and Palayer, *Impact of Military Artificial Intelligence on Nuclear Escalation Risk*.

unresolved against the legal requirement for informed human judgement in target engagement under Article 36 of the Geneva Convention. The Defence AI Centre and the RAS-AI Strategy 2040 represent steps toward an accountability framework, but the governance architecture remains a work in progress relative to the pace of capability integration.[40]

Accidental escalation, therefore, is best understood as a governance and human problem as much as a technical one. The risk is not primarily that AI systems malfunction. It is that organisational routines come to reward speed over scrutiny, normalising machine-mediated judgment in precisely the contexts where the capacity to question an authoritative alert matters most.

Misperception-Driven Escalation: Signalling and Offensive Interpretation

Escalation is not only about action (what states do), but interpretation (what rivals believe those actions enable). AI-enabled ISR and integrated command systems can be interpreted as reducing warning time and strengthening counterforce options, even when the doctrinal intent is defensive. As Sir Lawrence Freedman notes, movement up an escalation ladder often involves steps whose consequences cannot be fully predicted, in part because the adversary's interpretation is decisive.[41]

For Australia, the risk is most acute where AI strengthens maritime domain awareness and underwater tracking, because these capabilities sit close to valuable assets and contested threshold decisions. In a South China Sea crisis, for example, AI-enabled integration of satellite, and signals intelligence could be perceived as enabling rapid identification and cueing of targets. The concern is not only about what Australia intends but about what its capabilities appear to make possible. A People's Liberation Army Navy (PLAN) commander observing Australian and allied undersea surveillance activity during a crisis would have limited visibility into whether that activity is oriented toward conventional denial or toward tracking assets relevant to nuclear second-strike survivability. Under time pressure and worst-case assumptions, that ambiguity is itself a source of escalatory pressure.

[40] The Defence AI Centre was established in July 2024 to coordinate the safe, responsible adoption of AI and mitigate associated risks, while the RAS-AI Strategy 2040 is RAN's Robotics and Autonomous Systems AI strategy, which sets out principles for developing and employing autonomous systems, including a focus on human-machine teams.

[41] Freedman, "On the Tiger's Back: The Development of the Concept of Escalation."

There is also a narrower, conditional nuclear-adjacent pathway.[42] The Centre for New American Security (CNAS)'s work on the AI–nuclear nexus underscores how improvements in conventional sensing and precision can indirectly affect nuclear stability if they are perceived to erode the survivability of second-strike forces.[43] If Australian and allied systems contribute to undersea awareness in ways that appear relevant to nuclear platform survivability, adversaries may interpret them as strategically destabilising regardless of Australia's stated intentions.

The mechanism here is fear and expectation: when rivals interpret AI-enabled awareness as compressing their options, they may escalate to restore room for manoeuvre. For Australia, whose investments in undersea surveillance and ISR exploitation are among its most significant capability commitments, managing that perception is as important a policy problem as managing the capabilities themselves.

Policy Priorities for AI Integration

Australia's investment in AI is strategically coherent. The NDS's emphasis on denial and decision advantage is being translated into funding for enterprise data and ICT, C2, logistics, space and cyber, and a dedicated accelerator to move capability from trials into service.[44] The policy task is therefore not to slow down AI adoption, but to ensure it holds up under crisis pressure, and can be delivered at operational speed. To enhance regional stability, policymakers should prioritise the following shifts.

First, build "decision advantage with brakes" into AI-enabled decision support. Any AI-enabled DSS used for ISR, planning, or targeting should be designed for crises as well as routine operations. The DoD should mandate a minimum assurance package: an explicit "slow mode" that raises evidentiary thresholds and forces uncertainty to be surfaced, auditable provenance for key outputs (e.g., model version, assumptions), and structured red teaming built into operational cycles. The objective is to preserve AI-enabled speed where it is useful, without allowing tempo to become a bias that narrows interpretation and compresses political control during escalation.

Second, make OneDefence a gate for higher-risk AI, and fund delivery like operational infrastructure. The unified data layer should be treated as a prerequisite, not an aspiration. Where data remains fragmented, AI outputs become inconsistent across units and harder to verify in joint settings, precisely the conditions under which socio-technical overreliance takes hold. Policy should therefore link higher-risk AI

[42] Black et al., *Strategic Competition in the Age of AI: Emerging Risks and Opportunities from Military Use of Artificial Intelligence.*
[43] Paul Scharre and Michael Depp, "Artificial Intelligence and Nuclear Stability," Center for a New American Security, January 16, 2024, https://www.cnas.org/publications/commentary/artificial-intelligence-and-nuclear-stability.
[44] Department of Defence, *2024 Integrated Investment Program*.

applications to explicit OneDefence readiness thresholds, covering data standards and interoperability, before operational deployment is authorised.

Third, signal restraint through selective transparency on how AI is used. Perception and signalling ambiguity are an under-addressed stability risk of military AI. Australia should publish a short set of principles on how AI is used in decision support and intelligence processing, including what Australia will not automate and what oversight requirements apply. The purpose is not virtue signalling. It is practical signalling that reduces worst case inference and clarifies that decision advantage is being pursued with controls designed for crisis stability.

Fourth, fix the transition to capability: fast software acquisition, operational pull-through, and the people pipeline. Australia's late delivery problem is not only a budget issue but also a stability issue, because it increases incentives for hurried integration and premature reliance on prototypes. DoD should create a fast acquisition pathway for software and models that supports rapid iteration and modular upgrades, while reducing *ad hoc* modifications that delay delivery and fracture interoperability. ASCA should be directed to privilege projects that integrate with OneDefence, have an explicit sustainment model covering updates and cyber hardening, and are adopted by operational units with dedicated integration staff. Finally, treat the people pipeline as an enabler of assurance, not an afterthought. Without sufficient in house technical competence, auditability becomes aspirational rather than real. The Australian DoD should therefore treat clearances, retention, and technical career pathways as foundational to the governance architecture, ensuring that the humans responsible for overseeing AI systems have the expertise to do so meaningfully.

These recommendations are urgent because Australia is already adopting AI into the systems that shape crisis perception and operational tempo. Implementing these shifts will preserve the benefits of decision advantage while reducing the conditions under which speed and signalling ambiguity destabilise crises. The intended outcome is an AI-enabled force that is not only more capable but also more controllable under pressure, and that distinction matters now, because the Indo-Pacific is entering a period where misinterpretation and compressed decision time are becoming routine features of strategic competition.

ABOUT THE AUTHOR

Aina Turillazzi is a PhD candidate at the Strategic and Defence Studies Centre, Australian National University. Her research examines AI-enabled autonomy in weapons systems and its implications for crisis escalation, with a particular focus on grey zone dynamics in the Indo-Pacific.

ABOUT APLN

The **Asia-Pacific Leadership Network (APLN)** is a Seoul-based organization and network of political, military, diplomatic leaders, and experts from across the Asia-Pacific region, working to address global security challenges, with a particular focus on reducing and eliminating nuclear weapons risks. The mission of APLN is to inform and stimulate debate, influence action, and propose policy recommendations designed to address regional security threats, with an emphasis on nuclear and other WMD (weapon of mass destruction) threats, and to do everything possible to achieve a world in which nuclear weapons and other WMDs are contained, diminished, and eventually eliminated.


@APLNofficial


@APLNofficial


apln.network


aplnofficial