



Tactical Air Control Party (TACP) member inserts coordinates for the next air strike taken during EX MAPLE STRIKE on June 25, 2026 at Jimmy Lake, Alberta.

Photo: Corporal Connie Valin, Canadian Armed Forces photo

Navigation by Sextant & Astrocompass: A CAF Approach for AI Adoption¹

COLONEL ADAM MOORE AND MATTHEW DA MOTA

Colonel Adam Moore was commissioned into Princess Patricia's Canadian Light Infantry and has served as an infantry officer in the Canadian Armed Forces since 1996. Colonel Moore's career includes a range of command and staff roles, notably as Commanding Officer of Third Battalion, PPCLI, from 2018 to 2020. He has deployed internationally four times, including to Bosnia-Herzegovina, Afghanistan, and Iraq, as well as on five domestic operations. A graduate of the Royal Military College of Canada, he holds a Bachelor of Engineering and a Master of Defence Studies. He has completed professional military education at the Canadian Army Command and Staff College and the Canadian Forces College. Colonel Moore has been a Fellow at the Centre for International and Defence Policy at Queen's University since 2024-2025 and is currently the Chief of Staff of the 4th Canadian Division.


Matthew da Mota is Research Director, Emerging Technology and National Security at the Canadian Shield Institute for Public Policy. His work explores how emerging technologies and shifting geopolitical dynamics reshape security, governance, knowledge, and sovereignty—with a focus on the implications of AI and digital infrastructure. Matthew previously worked on international AI governance and security at the Centre for International Governance Innovation. His Ph.D. research examined the relationship between historical narrative, written evidence, and political power. Matthew is also a fellow in Technology and Security at Queen's University's Centre for International and Defence Policy, and he serves on multiple standardization technical committees.

Acknowledgement: This paper is firmly focused on AI disruption in military decision-making processes. Other AI implications in the military, such as the integration of AI with robotics, are important but separate considerations that are not addressed in this paper. Readers should treat this paper as a snapshot of AI development as of mid-2025, which may have changed in key areas.

Canada's largest national security institutions, the joint Department of National Defence and Canadian Armed Forces (DND/CAF), face an institutional imperative: become AI-enabled by 2030 or risk strategic irrelevance. Yet AI integration introduces hazards—cognitive deskilling, accountability dilution, automation bias—that remain poorly understood in high-stakes military environments. The rapidly evolving global threat environment—characterized by strategic competition, regional instabilities, and non-state disruption—intensifies this challenge, demanding rapid adoption while effectively managing operational risks.^{2,3}

The DND/CAF Artificial Intelligence Strategy (DCAIS) articulates this ambition through five interrelated lines of effort (LoEs): (1) Fielding and Employing AI Capabilities, (2) Change Management, (3) Ethics Safety and Trust, (4) Talent and Training, and (5) Partnerships.⁴ While the 2030 timeline appears distant, it represents an immediate challenge in defence procurement terms, requiring organizational and cultural transformation that bridges technical expertise with institutional knowledge.⁵

Figure 1: DND/CAF AI Strategy's Lines of Effort⁶



This paper explores AI integration using the metaphor of navigating by sextant and astrocompass—emphasizing the importance of grounding innovation in established professional frameworks. Drawing on the CAF's cognitive hierarchy, the authority, responsibility, accountability (ARA) model, and institutional risk assessment language, the analysis treats AI not as a solution but as a suite of powerful yet incomplete tools that must remain subordinate to professional judgment, institutional logic, and the realities of military decision-making.

The paper argues that meaningful AI adoption requires deliberate, risk-informed action and dual fluency: the ability to operate effectively in both AI-enabled and technology-denied environments. Using the DCAIS as a framing document for understanding DND/CAF's plans for AI integration, it examines how established CAF decision-making frameworks can guide responsible integration, concluding with policy recommendations to support informed advancement within existing institutional logic.

Guiding Stars for AI Integration


Like celestial navigation on the ocean, AI integration within DND/CAF requires fixed reference points to maintain course and bearing. The sextant measures angles between known celestial bodies and the horizon; the astrocompass helps to determine true north. For AI adoption within the defence enterprise, three institutional constants should serve as guiding stars: the CAF cognitive hierarchy model; the ARA framework; and the structuring logic of risk that runs through all levels of DND/CAF. While not AI-specific, these frameworks govern how CAF members engage with and make use of information across all contexts—making them essential for understanding how AI, a fundamentally information-based suite of technologies, will integrate within existing and emerging DND/CAF systems.

As AI systems begin to influence decision environments across DND/CAF, these foundational decision-making and knowledge production models continue to serve as essential navigation aids. They ensure that technological innovation remains rooted in proven institutional logic and professional judgment.

Cognitive Hierarchy

The CAF's cognitive hierarchy model (Figure 2) offers a structured framework for transforming raw data into actionable understanding. The model is based on four ascending stages: data, information, knowledge, and understanding, each representing a qualitative leap in cognitive processing. While AI can augment this hierarchy, it cannot replace the foundational role of human judgment.

Figure 2: The Cognitive Hierarchy⁷



AI tools can significantly enhance this hierarchy through assisting in cleaning and tagging data, automating cross-cueing,⁸ and extracting relevant information to support data collection and processing. More advanced systems may support analysis to generate knowledge and even offer structured inputs that inform human judgment in pursuit of shared understanding. Nonetheless, judgment remains a distinctly human function shaped by experience, intuition, and ethical discernment. In a command-centric environment, where staff serve to manage and interpret relevant information on behalf of commanders, the cognitive hierarchy reinforces the mission command philosophy of the CAF, emphasizing decentralized decision-making and trust. Any AI-enabled decision support must therefore be subordinate to, and anchored in, expert and ethical human judgment.

The ARA framework in DND/CAF

The triad of accountability, responsibility, and authority is the decision-making backbone of DND/CAF, applying across both military and civilian components of the joint institution.⁹ Just as the relationship between true north and celestial bodies remains fixed, these concepts maintain constant relative positions even as contexts change.

Accountability, the obligation to answer for responsibilities carried out, is non-delegable, residing within the legal and institutional structures of one's home organization. *Responsibility* is the obligation to perform assigned duties, which may be delegated but never

wholly relinquished. *Authority* denotes the right to make decisions, to give formal orders or instructions, and to direct actions—whether in a command, administrative, or functional capacity.

ARA takes on a unique character in the appointment to command, where each element is indivisibly vested in a single person exercising legal authority to direct military operations. Command appointments carry the exclusive legal and moral burden of decision-making; the commander holds personal accountability for mission success or failure. All other participants—staff members, headquarters personnel, supporting organizations—derive their authority through delegation bounded by parameters established by the accountability holder.

This hierarchical structure ensures unity of command and clarity of accountability, reinforced by DND/CAF's dual structure, where CAF members and DND employees operate under separate chains with distinct legal authorities, where legal accountability rests with one's institution, thereby ensuring clear delineation of responsibility.¹⁰

AI systems will primarily interface at the staff and command post levels, where information is processed, fused, and converted into actionable insights. These layers manage inputs, conduct analysis, and generate recommendations. The introduction of AI raises important questions—especially in the higher tiers of the CAF's cognitive hierarchy—about the potential dilution or displacement of human judgment. These questions are taken up and explored in greater depth in the section "Hazards of AI Integration."

Risk as a Structuring Logic in DND/CAF

Risk is a fundamental logic across operational and institutional domains of DND/CAF. Two complementary frameworks guide its application: *Risk Management for CF Operations* provides the doctrinal underpinning for military risk in tactical and operational contexts, while the *Defence Enterprise Risk Management Guideline* governs enterprise-level considerations.¹¹ Though distinct in emphasis, both frameworks reinforce a common logic: risk is not simply a calculation—it is a discipline of professional judgment anchored in context, authority, and responsibility.

In the military, commanders use risk to navigate decision-making. Staff play a critical advisory role in identifying hazards, proposing mitigations, and framing risk trade-offs. Risk is characterized not only by *probability* and *severity*, but also by factors such as *urgency* and *controllability*. Risk can only be accepted by individuals who both possess delegated authority over the proposed mitigations and operate within their defined residual risk acceptance authority.

Risk accountability in DND/CAF is personal, not procedural, representing more than a bureaucratic safeguard. Processes cannot be held accountable; only people can. The enduring principle is that *accountability attaches to the individual who accepts the risk*. This raises important institutional questions: When risk accepted in the past materializes in the present, who bears accountability? Is that accountability transferred to successors and, if so, are those individuals fully aware of the risks they now hold?

DND/CAF would benefit from treating risk as a shared currency of communication, promoting common language across operational and institutional functions to enhance coherence and decision clarity. Not all members of DND/CAF use risk frameworks in this way, but consistent risk language use could help to standardize and align risk awareness around AI across contexts.¹² The greater challenge lies in recognizing that seemingly discrete, low-stakes risks can accumulate across systems and, through compounding effects, create aggregate risk profiles that exceed the sum of their parts, remaining invisible to any single risk manager but capable of placing the entire enterprise in jeopardy.¹³

These guiding stars help orient the institution in uncharted space, but they do not eliminate uncertainty. With these reference points fixed, attention can now turn toward the hazards ahead—not along a predetermined route, but with the navigational discipline that sextants and astrocompasses once demanded: observation, calculation and orientation before movement.¹⁴

Hazards of AI Integration

As in any journey, navigating AI integration in DND/CAF has a roughly charted destination as defined in the DCAIS, but the actual process of navigation must be agile and adaptable to known and emerging risks, able to shift course as needed but still find the way.

“It is essential to recognize that AI—particularly generative models—does not understand anything in the human sense.”

Firstly, it is essential to recognize that AI—particularly generative models—does not understand anything in the human sense. These systems operate as highly sophisticated pattern recognition engines, capable of generating plausible outputs based on statistical regularities in training data.¹⁵ While they can achieve superhuman accuracy in well-defined contexts, their reliability degrades sharply when confronted with edge cases—so-called *1% problems*—where atypical inputs fall outside learned patterns.¹⁶ This limitation does not constitute a risk in itself, but becomes a factor that must be considered when assessing the suitability of AI tools for specific tasks or domains.

The following hazards are grouped by cognitive and institutional dynamics rather than technical domains, reflecting how AI integration will reshape decision-making processes across the defence enterprise.

Cognitive Offloading & Deskilling and the Erosion of Judgment

AI systems promise to reduce cognitive load by automating data-intensive tasks, theoretically allowing personnel to shift from production to evaluation mode—accelerating movement up the cognitive hierarchy toward understanding and judgment. This cognitive offloading mirrors familiar practices where humans delegate mental tasks to algorithmic systems.¹⁷

However, habitual reliance introduces a critical trade-off. When delegation becomes dependence, cognitive deskilling occurs—the gradual loss of independent reasoning ability—distinct from offloading in that it risks long-term erosion of judgment.¹⁸ Persistent reliance on AI-enabled tools may erode individual and institutional judgment by denying exposure to the subjective, experience-based activities at lower levels of the cognitive hierarchy. The result may be technically proficient leaders who lack the intuitive judgment required to navigate ambiguous or degraded environments.

In military contexts, this means personnel need to both excel at interpreting AI-generated intelligence summaries and have the experience necessary to recognize when source data has been compromised or patterns manipulated. But in a future where AI systems handle most processing and synthesis tasks, what happens when the human in the loop lacks the personal experience to verify the output? What happens when the volume of data exceeds a human operator's capacity for meaningful review? Should

the CAF deliberately retain non-AI-enabled, human-generated products—however inferior—to train personnel in independent judgment? And while such delegation may be acceptable in benign contexts, how should responsibility be assigned when an AI-enabled decision contributes to the unlawful application of lethal force: does accountability rest with the developer, the contracted programmer? Who gets charged with the war crime? A decision brief in 2035 might involve AI-generated outputs presented with high confidence, but if that data has been poisoned or the model sabotaged, and the attending human lacks depth of knowledge to challenge it, the outcome may not just be flawed—it may be undetectably wrong. In such environments, trust in automation must be matched by an equally strong commitment to developing judgment-capable leaders.

AI adoption in defence does not reduce the burden of professional development—it increases it. To operate in both technologically degraded and technologically enhanced environments, CAF personnel will require dual fluency in human-centred judgment and machine-assisted analysis. Like navigating the seas, the ability to recognize when something is off course depends less on upgrading the instruments than on remembering how to use them.¹⁹

Automation Bias and Algorithm Aversion

AI-enabled decision support introduces a dual hazard: uncritical reliance on algorithmic outputs (automation bias) and outright dismissal even when statistically superior (algorithm aversion).²⁰ Both exist on a continuum, with impact particularly acute under time pressure within action-oriented institutions.

Military decision-making has long depended on intermediaries—staff officers, advisors, planning cells—who filter information and frame choices. AI support represents digital bureaucracy structured by rules and neural networks.²¹ However, commanders have always understood that staff bring implicit biases, interests, and analytical limitations. The distinctive risk with AI is algorithmically induced bias: distortions embedded in training data, model assumptions, or environmental inputs, often obscured from end users and difficult to detect. Solutions to algorithmic bias are uncertain, but a clear solution in the interim is to ensure a system of checks and balances with humans in and on the loop to verify AI outputs.

Traceability—the ability to understand how a recommendation was generated—has become a major concern in this context.²² Yet this, too, has precedent. Expert human judgment is often grounded in tacit knowledge that resists formal articulation.

Within DND/CAF, the challenge lies not in eliminating bias or aversion, but managing them through sound processes and deliberate human oversight.²³ If human oversight on AI recommendations become mere checkboxes to enable speed, decision support utility is lost.²⁴ Conversely, if skepticism leads to paralysis, the advantage of AI is never realized. Success requires clearly

defined authorities and accountabilities delineating who verifies outputs, when to pause for deeper analysis, and how to escalate when something appears wrong.

Moral Crumple Zones and the Illusion of Accountability

The concept of *moral crumple* zones introduced by Madeline Elish refers to the human tendency to absorb blame when autonomous systems fail—regardless of actual control or authorship.²⁵ The person “in the loop” verifying an automated system’s outputs becomes the face of accountability, even when decision-making has been substantially delegated to an opaque and distributed system. This illusion of accountability becomes increasingly problematic as system complexity grows and explanation becomes less accessible.²⁶

Looking ahead to a 2030 AI-enabled DND/CAF, this concern shifts from theoretical to real. If future AI tools continue to function as black boxes, producing outputs without transparent or interpretable reasoning, accountability structures must be re-evaluated. Rapid AI development, increasing complexity and autonomy of systems, and potential for recursive improvement and self-modification of AI systems threaten to push these systems beyond human comprehension.²⁷ As the line between authority and responsibility blurs in increasingly automated systems, the burden of accountability remains sharply focused—often disproportionately—on the closest human.

This raises pressing questions for human-machine teaming: who is accountable when something goes wrong?²⁸ The operator, tactical commander, project manager, contracted company, or model training team? While failures may be infrequent, their severity can be significant for both affected people and institutional trust.

In practice, it is highly likely that some person will be blamed for AI-related failures, but the downstream consequences of that blame merit attention. If DND/CAF intends to remain ethically grounded while scaling AI integration, it must equip all relevant personnel not only with tools to make decisions, but with frameworks to share accountability and process moral complexity.

Institutional Ownership and Generative Intellectual Property

The 2019 *Lukits v. DND* ruling by the Federal Public Sector Labour Relations and Employment Board upheld a Royal Military College professor’s right to retain ownership of handwritten lecture notes.²⁹ While the *Lukits* ruling appropriately upheld academic freedom,³⁰ it also reveals institutional vulnerability in how the Crown asserts ownership over knowledge produced by its workforce. How well is DND/CAF positioned to manage intellectual property in an era of AI?

Generative AI tools complicate this further. Unlike handwritten lecture notes, Generative AI outputs are shaped by statistical prediction, informed by user prompts, and processed through commercial infrastructure. The resulting products—text, images, code, design concepts—may be refined by the user, but are not easily traceable in terms of authorship or custody. Treasury Board guidance for employees to consult legal services, review licensing terms, and stay informed about evolving AI law,³¹ is well-intentioned but impractical. Not all DND/CAF personnel are intellectual property specialists, and no capacity has been allocated to ensure the average DND/CAF user of AI is able to manage this burden. Responsibility for managing intellectual property risk should not be downloaded onto individual users.

Crown ownership over co-developed Generative AI outputs cannot be assumed. If data or design artifacts generated within DND/CAF are absorbed or monetized by commercial vendors through unclear licensing arrangements,³² the resulting impact could be significant. The issue is not a lack of policy but a failure to deliver clear, unambiguous direction.

A designated authority—such as the DND/CAF AI Centre—should be empowered to issue enforceable guidance on ownership, licensing, and data use across all adopted AI tools. This is an immediate requirement, not a matter for further study. Expecting military or civilian personnel to self-manage legal compliance is unreasonable, risks losing institutional control over mission-critical knowledge, and opens the door to the erosion of information integrity.

Valuation of Data: Contextual Significance and System Blindness

AI systems infer value statistically, but military planning often relies on sparse, context-specific information.³³ This becomes complex as models are trained on large volumes of unstructured, inconsistently governed data.³⁴ An operations order's intent paragraph may appear once while administrative details repeat throughout—yet AI may overvalue repetition over judgment-critical content. Without enterprise capability to signal authoritative inputs, systems cannot distinguish between delegated authority and administrative noise, substantially increasing risk.

It remains unclear whether current DND/CAF enterprise capability exists to signal that certain inputs should be treated as more authoritative than others. How does the CAF ensure a revised Chief of Defence Staff Directive or an order that changes the rules of engagement for a deployed task force registers as a meaningful shift rather than just another document? How is superseded content prevented from continuing to shape outputs simply because it appears more frequently in datasets?

Data significance is contextual and changes over time, while AI models optimize for patterns, not meaning. The same data point may be critical in one context, irrelevant in another, and misleading in a third.³⁵ DND/CAF must either develop mechanisms

that recognize and preserve shifting context over time or formalize human-in-the-loop roles specifically aimed at preserving contextual meaning within AI-enabled processes.³⁶

Pressing Forward Despite Uncertainty³⁷

DND/CAF should treat AI as a general-purpose technology requiring immediate preparation. The trajectory of AI development and its convergence with key military domains strongly suggests that it will reshape both the battlespace and the institutional environment in which DND/CAF operates. This warrants a deliberate shift toward AI-readiness: institutional, procedural, and cultural. Preparing now lays essential groundwork for confronting future disruptions from other emerging technologies, including quantum-era threats.

Sovereignty concerns should be viewed as a constraint to manage, not a barrier to overcome. The scale of infrastructure required for current and future AI models is vast and not likely to be possible for Canada to develop a “sovereign AI stack” within the short term.³⁸ Some degree of reliance on existing private and foreign cloud and data infrastructure will continue to be necessary. However, not all data or processes require maximum classification or control.³⁹ For those that do, secure sovereign data centres may already exist, or else such centres can be built. For those in the grey zone in between, DND/CAF can make use of privacy-enhancing technologies (PETs), data enclaves, and other tools to maximize sovereignty and data protection while continuing to make use of high-quality private infrastructure hosted in allied nations.⁴⁰ So long as there is clarity on what must remain sovereign and what can be responsibly externalized with safeguards, a reasonable degree of protection can be ensured.⁴¹

“If Canada intends to shape how AI is used ethically and operationally, DND/CAF must maintain credible presence in those conversations.”

Digital interoperability is now strategic. In an environment where allies rapidly integrate AI into operations, Canada must keep pace to retain influence and credibility.⁴² A country that lags risks becoming a standards taker rather than shaper. If Canada intends to shape how AI is used ethically and operationally, DND/CAF must maintain credible presence in those conversations through meaningful capability and demonstrated experimentation.

Canada's domestic AI ecosystem is a national advantage.⁴³ Academic expertise, innovation capacity, and industry already exist. The DCAIS LoE5 calls for the development of trust-based partnerships that serve both mission assurance and Canadian national interests.

Policy Recommendations

The following recommendations chart a deliberate course toward responsible AI integration in DND/CAF, reflecting the integrated LoEs outlined in the DCAIS and guided by enduring institutional reference points:

1. Treat AI as a general-purpose technology for planning purposes and begin structured implementation accordingly.
2. Treat the CAF Digital Campaign Plan and the DCAIS as enterprise reinvention guides, not as discrete initiatives.
3. Institutionalize risk as a shared language across operational and institutional spheres, improving the ability to track and manage aggregate risk accumulation, particularly where AI introduces distributed decision-making.
4. Develop dual capability: personnel must operate effectively in both AI-enabled and technology-denied environments.
5. Distinguish sovereignty requirements pragmatically between data/processes requiring full government control and those enabling secure externalization.
6. Leverage interoperability as a strategic influence vector, not just technical compatibility, helping Canada shape emerging defence AI standards.
7. Integrate DND/CAF AI adoption within full-ecosystem planning, ensuring defence contexts inform broader governmental AI governance.

Conclusion

These guiding stars—cognitive hierarchy, the ARA framework, and the structuring logic of risk—shine more brightly in an AI-enabled environment, not despite technological transformation, but precisely because of it. In a doctrinal Ends-Ways-Means framing, AI fundamentally shifts the *means* by which decisions are made through data-enabled models, probabilistic reasoning, and systems capable of exponentially more rapid adaptation. These new means demand new *ways*—revised processes, doctrinal adjustments, upskilling and digital literacy enabling personnel to interrogate the reliability, traceability, and mission relevance of AI-generated outputs. The *ends*—effective decision-making within volatile, uncertain, complex, and ambiguous environments—remain consistent.

Canada's strategic advantage lies in doing AI adoption well—in a way that is aligned with Canadian values and sensitive to national needs. Canada's legitimacy stems from world-class research, professional military training, and strong partnerships, but fundamentally from respect for rights and values. Channeling these principles into AI frameworks creates superior tools and more desirable outcomes.

The metaphorical sextant and astrocompass remind us that while tools evolve, navigation discipline endures. Canada's path forward demands deliberate progress guided by institutional wisdom and democratic values—navigation principles that chart a course through uncertain waters while maintaining true bearing.

Notes

- 1 Developed by Col. Adam Moore during a defence sabbatical (August 2024–May 2025) while a Visiting Fellow at Queen's University. A longer version appeared with the Centre for International and Defence Policy (CIDP), Queen's University, 2025.
- 2 Canadian Security Intelligence Service, *CSIS Public Report 2023*, 2024. Originally found at: <https://www.canada.ca/en/security-intelligence-service/corporate/publications/csis-public-report-2023.html>; Department of National Defence, *Pan-Domain Command & Control Concept Paper*, 2024. Originally found at: <https://www.canada.ca/en/department-national-defence/corporate/reports-publications/pan-domain-command-control.html>. See Jamali, Nahlah, "DeepSeek AI Shakes American Swagger and Upends Assumptions," BBC News, January 28, 2025. Originally found at: <https://www.bbc.com/news/articles/cd643wx888gq>.
- 3 See Lloyd, Ron, "Public Service Culture, Risk Management and Governance: The Prerequisite of Effective Strategic Governance to Realizing Our Digital Ambition," Canadian Global Affairs Institute, 2024. Originally found at: https://www.cgai.ca/public_service_culture_risk_management_and_governance_the_prerequisite_of_effective_strategic_governance_to_realizing_our_digital_ambition.
- 4 Department of National Defence (DND), *The Department of National Defence and Canadian Armed Forces Artificial Intelligence Strategy (DCAIS)*, 2024. Originally found at: <https://www.canada.ca/en/department-national-defence/corporate/reports-publications/dnd-caf-artificial-intelligence-strategy.html>.
- 5 Inspired by the Canada School of Public Service's curation of Stewart-Weeks, M., and Cooper, S. (2019) "We're Missing the Point of Digital Government - Here's Why. Opinion: Digital Experts Understand Tech - but Do They Understand Government?" *Apolitical*, August 13. Originally found at: <https://apolitical.co/solution-articles/en/why-were-missing-the-point-in-digital-government>.
- 6 Reproduction of an image from DND, *The DCAIS*, 2024, pg. 1.
- 7 Author-adapted representation inspired by the U.S. Army's Mission Command Training Program. This representation aligns with the CAF's Cognitive Hierarchy explanations in DND, *B-GI-300-003/FP-001 Command in Land Operations*, 2007; DND, *B-GI-331-001/FP-001 Command Support in Land Operations*, 2008; and DND, *Canadian Forces Joint Publication 01: Canadian Military Doctrine*, 2nd ed. 2024. The source image can be found in U.S. Army, *FY 24 Mission Command Training Program Key Observations*. Center for Army Lessons Learned, 2025. Figure 2-1, "Cognitive Hierarchy," 10. Originally found at: <https://api.army.mil/e2/c/downloads/2025/02/24/aea2698e/25-04-946-fy-24-mctp-key-observations-feb-25.pdf>.
- 8 Cross-cueing "entails one sensor passing detection, geolocation and targeting information to another sensor, which improves efficiency." DND, *Royal Canadian Air Force Doctrine: Intelligence, Surveillance and Reconnaissance*, 2017, p. 2-2.

- 9 DND, *Accountabilities, Responsibilities, and Authorities*, Corporate Secretariat, 2024.
- 10 DND, *Canadian Forces Joint Publication 01: Canadian Military Doctrine*, 2nd ed. 2024.
- 11 DND, *Risk Management for CF Operations*, 2007. DND, *Defence Enterprise Risk Management Guideline*, 2018.
- 12 The National Institute of Standards and Technology, *AI Risk Management Framework*, NIST AI 100-1. 2023. Originally found at: <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>, offers useful guidance for institutional risk analysis.
- 13 See Lloyd (note 3).
- 14 The Treasury Board's Algorithmic Impact Assessment (AIA) tool emphasizes impact severity over probability, misaligning with defence risk models that require both factors: Since many DND/CAF applications involve minimal Canadian personal data, strict AIA adherence may unnecessarily delay high-utility operational applications. A tailored assessment model or controlled exceptions may better serve defence requirements while maintaining ethical oversight. Treasury Board Secretariat, "AI Strategy for the Federal Public Service 2025-2027," March 4, 2025. Originally found at: <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/gc-ai-strategy-overview.html>.
- 15 Vallor, Shannon, *The AI Mirror: How to Reclaim Our Humanity in an Age of Machine Thinking*, Oxford University Press, 2024.
- 16 Max Lamparth and Jacquelyn Schneider, "Why the Military Can't Trust AI," *Foreign Affairs*, 2024. Originally found at: <https://www.foreignaffairs.com/united-states/why-military-cant-trust-ai>. Karine Perset and Sara Fialho Esposito, *How Are AI Developers Managing Risks? Insights from Responses to the Reporting Framework of the Hiroshima AI Process Code of Conduct*, OECD Artificial Intelligence Papers (Paris: OECD Publishing, 2025), 33. Originally found at: https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/09/how-are-ai-developers-managing-risks_fbaeb3ad/658c2ad6-en.pdf.
- 17 Grinschgl, Sandra, Frank Papenmeier, and Hauke S. Meyerhoff, "Consequences of Cognitive Offloading: Boosting Performance but Diminishing Memory," *Quarterly Journal of Experimental Psychology* 74, no. 9 (2021): 1477-96. Originally found at: <https://doi.org/10.1177/17470218211008060>.
- 18 Thiele, Leslie Paul, "Rise of the Centaurs: The Internet of Things Intelligence Augmentation," in *Towards an International Political Economy of Artificial Intelligence*, edited by Tugrul Keskin and Ryan David Kiggins, 39-61. Springer International Publishing, 2021. Originally found at: <https://link.springer.com/book/10.1007/978-3-030-74420-5>. Cognitive deskilling is closely related to the concept of moral deskilling. For an in-depth discussion of the latter, see Vallor, Shannon, "Moral Deskilling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character," *Philosophy and Technology* 28, no. 1 (2015): 107-24.
- 19 Further research is needed on AI's impact on tacit knowledge within the CAF. Since AI systems operate on structured, explicit data while military expertise relies heavily on unrecorded, experiential knowledge, overreliance risks eroding expert judgment if systems displace rather than augment human experience. For more on the devaluation of professional judgment see Nichols, Tom M., *The Death of Expertise: The Campaign against Established Knowledge and Why It Matters*, Oxford University Press, 2017.
- 20 The term "automation bias" here draws directly from Nadibaidze, Anna, Ingvild Bode, and Guangyu Zhang, *AI in Military Decision Support Systems: A Review of Developments and Debates*, 2024. Originally found at: https://findresearcher.sdu.dk/ws/portalfiles/portal/275893410/AI_DSS_report_WEB.pdf; "algorithm aversion," is more loosely defined according to Jussupow, Ekaterina, Izak Benbasat, and Armin Heinzl, "Why Are We Averse Towards Algorithms? A Comprehensive Literature Review on Algorithm Aversion," in *ECIS 2020 Proceedings*, Virtual: Association for Information Systems, 2020. Originally found at: https://aisel.aisnet.org/cgi/viewcontent.cgi?article=1167&context=ecis2020_rp.
- 21 Stimulated by a logical extrapolation from the idea that "if the economy is an information processing system, does that mean that every corporation is an artificial intelligence?"; Davies, Dan, *The Unaccountability Machine: Why Big Systems Make Terrible Decisions—and How the World Lost Its Mind*, Profile Books, 2024, p. 63.
- 22 Taddeo, Mariarosaria, David McNeish, Alexander Blanchard, and Elizabeth Edgar, "Ethical Principles for Artificial Intelligence in National Defence," *Philosophy & Technology* 34, no. 4 (2021): 1707-29. Originally found at: <https://doi.org/10.1007/s13347-021-00482-3>, p. 1717-1718; DND, The DCAIS, 2024, p. 18.
- 23 Davies, *The Unaccountability Machine*, 78-79.
- 24 Nadibaidze, Bode, and Zhang, *AI in Military Decision Support Systems*, 33.
- 25 Elish, Madeleine Clare, "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction," *Engaging Science, Technology, and Society* 5 (March 23, 2019): 40-60. Originally found at: <https://doi.org/10.17351/ests2019.260>.
- 26 Devitt, S. Kate, "Bad, Mad, and Cooked: Moral Responsibility for Civilian Harms in Human-AI Military Teams," in *Responsible Use of AI in Military Systems*, edited by Jan Maarten Schraagen, 248-77, Chapman and Hall/CRC, 2024. Originally found at: <https://doi.org/10.1201/9781003410379-16>.
- 27 Mustafa Suleyman has warned about these risks (Suleyman, Mustafa, and Michael Bhaskar, *The Coming Wave: Technology, Power, and the Twenty-First Century's Greatest Dilemma*, Crown, 2023.), and Dan Davies makes a similar observation: when explanation breaks down, parsing source code offers no real insight into cause or consequence (Davies, *The Unaccountability Machine*, 76.)
- 28 See Elish "Moral Crumple Zones: Cautionary Tales in Human-Robot Interaction," 54; Nadibaidze, Bode, and Zhang, *AI in Military Decision Support Systems*, 40-41.
- 29 *Lukits v. Treasury Board (Department of National Defence)*, 2019 FPSLREB 20 (CanLII), *Federal Public Sector Labour Relations and Employment Board*, March 13, 2019. Originally found at: <https://decisions.fpslreb-crtespf.gc.ca/fpslreb-crtespf/d/en/item/407680/index.do>.
- 30 Treasury Board Secretariat, "University Teaching (UT): Agreement Between the Treasury Board and the Canadian Military Colleges Faculty Association." Originally found at: <https://www.tbs-sct.canada.ca/agreements-conventions/view-visualiser-eng.aspx?id=27#toc2112221127>.
- 31 Treasury Board Secretariat, "Guide on the Use of Generative Artificial Intelligence," May 30, 2024. Originally found at: <https://www.canada.ca/en/government/system/digital-government/digital-government-innovations/responsible-use-ai/guide-use-generative-ai.html>.
- 32 See, for example, McManus, Sean, "Why Employees Smuggle AI into Work," *BBC News*, February 3, 2025. Originally found at: <https://www.bbc.com/news/articles/cn7rx05xq2go>.
- 33 Goldfarb, Avi, and Jon R. Lindsay, "Prediction and Judgment: Why Artificial Intelligence Increases the Importance of Humans in War," *International Security* 46, no. 3 (2022): 7-50. Originally found at: https://doi.org/10.1162/isec_a_00425.
- 34 Goldfarb, Avi, and Jon R. Lindsay, "Artificial Intelligence in War: Human Judgment as an Organizational Strength and a Strategic Liability," *Foreign Policy at Brookings*, 2020. Originally found at: https://www.brookings.edu/wp-content/uploads/2020/11/fp_20201130_artificial_intelligence_in_war.pdf.
- 35 Vallor, *The AI Mirror*, 26-27; Nadibaidze, Bode, and Zhang, *AI in Military Decision Support Systems*, 36-38.
- 36 On contextual data in AI deployment, see da Mota, M., "Toward an AI Policy Framework for Research Institutions," Centre for International Governance Innovation, 2024. Originally found at: <https://www.cigionline.org/publications/toward-an-ai-policy-framework-for-research-institutions/>.
- 37 Author's reflection from knowledge gained during defence sabbatical; no public or academic sources cited.
- 38 For more on the debates over a Canadian "sovereign stack" for AI, see emerging work from the Canadian AI Sovereignty and Innovation Cluster. Originally found at: <https://caisic.ca/>; the Council of Canadian Innovators <https://www.canadianinnovators.org/>; the Canadian AI adoption Initiative Originally found at: <https://www.aiadoption.ca/>; and Innovation, Science and Economic Development Canada, *Canadian Sovereign AI Compute Strategy*, 2025. Originally found at: <https://ised-isde.canada.ca/site/ised/en/canadian-sovereign-ai-compute-strategy>.
- 39 For a critical review of Canada's classification framework through a digital lens, see: Lloyd, Ron, "Canada's Security Classification Framework: The Biggest Impediment to Realizing Our Digital Ambition," Canadian Global Affairs Institute, 2024. Originally found at: https://www.cgai.ca/canadas_security_classification_framework_the_biggest_impediment_to_realizing_our_digital_ambition.
- 40 For more on PETs for enhancing AI sovereignty see: OECD, "Sharing Trustworthy AI Models with Privacy-Enhancing Technologies," *OECD AI Papers*, no. 38, 2025. Originally found at: https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/06/sharing-trustworthy-ai-models-with-privacy-enhancing-technologies_5df6fd05/a266160b-en.pdf.
- 41 Shared Services Canada and the Communications Security Establishment have been and continue to be building infrastructure and defining lines to enable those policies.
- 42 For example, see: NATO, "Summary of NATO's Revised Artificial Intelligence (AI) Strategy," 2024. Originally found at: https://www.nato.int/cps/en/natohq/official_texts_227237.htm.
- 43 DND, "Adopting Artificial Intelligence with Security in Mind," DDN2-V46, Canadian School of Public Service, 2024. Originally found at: <https://www.cspes-efpc.gc.ca/video/ai-security-eng.aspx>.