

ARIP Defence Decision Assurance Module: Reconstructing Human-AI Decisions for After-Action Review and Institutional Learning

Executive Summary

Artificial intelligence is increasingly used to compress military decision cycles by fusing data, identifying patterns, generating assessments, and supporting courses of action. The Department of National Defence and Canadian Armed Forces (DND/CAF) have stated that AI should augment rather than replace human decision-making, while decisions that augment human judgment should be traceable, explainable and accountable. IDEaS challenges issued in 2022 and 2026 have also sought trusted human-autonomy teaming, after-action analysis, audit of AI-enabled decisions, and explainable multi-modal fusion for operational decision-making. [1][2][4][5]

These requirements create a problem that is different from building a faster sensor-fusion model or a better course-of-action optimizer: after an operation or exercise, can the organization reconstruct what information was available, how AI and human reasoning interacted, why a particular decision was reached, and whether the reasoning was sound given what was reasonably knowable at the time?

The AI-Augmented Research & Insight Platform (ARIP) can be adapted as a governed post-decision assurance environment. The proposed ARIP Defence Decision Assurance Module would not attempt to replace live Command and Control systems, multi-modal fusion engines, or tactical decision-intelligence platforms. Its initial role would be retrospective: preserve or reconstruct the decision-time evidence state, evaluate how the evidence was interpreted, compare the original analytical signals with facts revealed after the outcome, examine contradictions and alternatives, and create a structured record of lessons that can be reused in later training, experimentation and capability evaluation.

The core distinction is between **decision-time truth** and **outcome-time knowledge**. ARIP should never rewrite what was known at the time simply because the outcome is now known. Instead, it preserves the original epistemic state and adds a separate post-outcome validation layer. A defence deployment could therefore support a structured sequence:

1. **CAPTURE / RECONSTRUCT** — preserve the decision event, alternatives, AI outputs, source artifacts, warnings, timestamps and the actual User decision.
2. **DEFINE** — express the actual decision problem and material courses of action as explicit review hypotheses or candidates.
3. **MEASURE** — reconstruct the CTQs and decision-time evidence set, while separately recording what post-outcome facts reveal about the accuracy, timeliness, independence and relevance of those inputs.
4. **ANALYZE** — reproduce the decision-time analytical state using only contemporaneous information, then compare CTQ-level analytical signals with realized facts for calibration and learning.
5. **IMPROVE / AFTER-ACTION REASONING REVIEW** — examine contradictions, missed alternatives, AI reliance, assumptions, counterfactuals and whether better use of information already available could have improved the decision.
6. **CONTROL / CLOSURE** — record the review conclusion, distinguish decision quality from outcome quality, and preserve lessons, follow-up actions and future monitoring requirements.
7. **ULA LEARNING** — observe repeated post-decision records to identify recurring evidence failures, reasoning weaknesses, trust-calibration patterns and transferable lessons without mutating historical decisions.

The result is a new form of decision provenance. Data provenance answers where information came from, how that information became an accepted belief, how it was challenged, what the human decided, and what the outcome later revealed.

The shift: The after-action question changes from “Did the mission succeed?” to “Given what was knowable at the time, was the reasoning sound, what caused the outcome, and what can the human-AI team learn before the next decision?”

