OPERATIONAL INTELLIGENCE:

An Ontology for AI Integration in Defence Institutions

Author	Institution	Date	Series
Umran Nayani	Quantumist.ai, Mumbai	March 2026	Paper No. 02

Abstract

The integration of artificial intelligence into defence institutions is proceeding without a principled framework for distinguishing which cognitive functions admit AI augmentation, which require human oversight, and which must remain exclusively human. This paper proposes a formal ontology that decomposes defence operations into four classes of cognitive primitives — each defined by distinct truth conditions, error tolerances, and computational reducibility — and applies this framework to the specific operational, security, and sovereignty requirements of Indian defence institutions. We demonstrate that AI integration succeeds or fails based on primitive classification accuracy, not technological capability. A logistics calculation mishandled by a probabilistic model is not an efficiency failure — it is a potential operational catastrophe. A command decision outsourced to an algorithm is not automation — it is an abdication of legal and moral authority. By mapping the complete cognitive stack of defence operations — from deterministic computation to irreducibly human command judgment — we establish a framework for building AI-augmented defence institutions that amplify rather than replace human intelligence, while maintaining complete data sovereignty on Indian-controlled infrastructure. This ontology serves as the theoretical and operational foundation for India's Atmanirbhar AI capability in defence — not as a policy aspiration, but as a deployable architecture.

JEL Classification: H56 (National Security and War) — O33 (Technological Change: Choices and Consequences) — D83 (Search; Learning; Information and Knowledge) — L23 (Organization of Production) — C45 (Neural Networks and Related Topics) — H11 (Structure, Scope, and Performance of Government) — D81 (Criteria for Decision-Making under Risk and Uncertainty)

Keywords: *Defence AI, Cognitive Ontology, Operational Security, Data Sovereignty, Command Authority, AI Augmentation, Atmanirbhar AI, Large Language Models, Institutional AI Integration, Human-in-the-Loop*

Section I

The Sovereignty Problem

Why the Current Approach Is Structurally Unsafe

1.1 The Default Architecture and Its Consequences

When a defence professional uses ChatGPT to draft a briefing, summarise an intelligence report, or analyse a procurement document, the query travels to a server in the United States. It is processed by a model trained on data assembled under American legal jurisdiction. It is logged, at least temporarily, under terms of service that no Indian institution fully controls. And it returns an answer that may be accurate, partially accurate, or confidently fabricated — with no reliable mechanism for distinguishing between these outcomes at the point of use.

This is the default architecture for AI adoption in Indian institutions today. It is not the result of malicious intent. It is the result of convenience — the same convenience that led generations of institutions to adopt foreign software and foreign infrastructure without examining what that dependency meant for operational security until the dependency became structural. With AI, the stakes are categorically higher. A foreign software platform processes transactions. A foreign AI platform processes thought — the reasoning, synthesis, and decision support that shapes operational judgment. The difference is not marginal. It is existential.

1.2 The Three Structural Vulnerabilities

Vulnerability 1: Data Exposure Through Query Patterns

Even when no classified information is directly entered into a commercial AI system, query patterns themselves constitute intelligence. A series of queries about a specific geographic region, a particular equipment type, or a defined operational timeline — individually innocuous, collectively revealing — creates an intelligence signature on a foreign server that no terms of service agreement adequately addresses. Pattern analysis of AI usage data from a defence institution is, potentially, a window into operational priorities, planning horizons, and capability gaps. This window should not exist.

Vulnerability 2: Model Dependency Without Verification

Commercial AI models are probabilistic systems. They generate outputs that maximise plausibility under their training distribution — they do not compute truth. For a defence institution that has come to rely on a commercial AI model for logistics analysis, intelligence synthesis, or operational planning support, the failure modes of that model are now operationally relevant. When the model hallucinates — and all current large language models hallucinate — the institution has no visibility into the training data, the

architecture, or the specific failure modes of the system it is depending on. You cannot audit a model you do not own.

Vulnerability 3: Vendor Dependency as Strategic Leverage

An institution that has built its AI-augmented workflows around a single commercial provider has created a dependency that is, in strategic terms, a vulnerability. Pricing changes, terms of service updates, geopolitical developments, or simple product discontinuation can disrupt operational capability at a moment chosen by the vendor, not the institution. India's defence establishment has invested decades in reducing strategic dependencies across platforms and supply chains. AI infrastructure is the next frontier of that investment — and it is currently being built, by default, on maximum dependency.

1.3 The Atmanirbhar AI Argument

Atmanirbhar Bharat applied to AI means: open-source models running on Indian-controlled servers, configured and operated by Indian professionals, processing data that never leaves Indian jurisdiction, producing outputs that can be audited, modified, and improved without reference to any foreign vendor.

This is not a theoretical aspiration. The open-source AI ecosystem — Llama 3, Mistral, Falcon, and hundreds of domain-specific models — has matured to the point where models competitive with closed-source commercial alternatives on specific operational tasks can be run on commodity server hardware at a fraction of the cost, with complete data sovereignty and zero ongoing vendor dependency. The infrastructure barrier has fallen. What remains is the capability barrier.

Section II

The Category Error

Why Current AI Integration in Defence Is Asking the Wrong Questions

The dominant discourse on AI in defence institutions mirrors the discourse on AI in every other sector: which roles will AI replace? This framing is as misconceived in defence as it is in any other high-stakes domain. It treats 'Intelligence Officer,' 'Logistics Coordinator,' and 'Commanding Officer' as functional categories when they are, in fact, role labels that contain multiple cognitively distinct operations.

An Intelligence Officer who processes satellite imagery and one who advises a commander on tactical options are performing operations with fundamentally different computational structures. One is pattern recognition against a known taxonomy — a function for which AI already demonstrates superior performance. The other is contextual synthesis under radical uncertainty, incorporating political intelligence, psychological assessment, and real-time environmental variables that exist in no dataset.

The framework this paper proposes does not ask which jobs AI will replace. It asks: what are the atomic cognitive operations that constitute defence institution functioning, and what are the truth conditions, error tolerances, and computational structures of each? The answer produces a principled architecture for AI integration that is operationally safer, strategically sounder, and institutionally more resilient than any role-level automation analysis.

Section III

Foundations

Computational Reducibility in Defence Operations

Not all cognitive operations have the same computational structure, and the structure of the operation — not the sophistication of the technology — determines whether AI integration is safe, useful, or dangerous. We classify defence cognitive operations according to their underlying truth conditions and the degree to which they admit computational reduction. This classification is independent of rank, role, or organisational structure.

Class	Type	Truth Condition	Error Tolerance	AI Role
I	**Deterministic Operations**	Equality	Zero — errors compound into systemic failures	Executor with formal verification
II	**Probabilistic Operations**	Calibration	Continuous, task-dependent	Generator with human validation
III	**Contextual Operations**	Counterfactual effectiveness	High variance acceptable; systematic bias catastrophic	Sparring partner — human in command
IV	**Principal Operations**	Legitimacy and accountability	Existential	Tool only — never principal

The critical insight is that AI safety and utility are properties of the task-system pairing under specific truth conditions — not properties of AI systems in isolation. A perfectly capable AI system deployed on a Class I task without formal verification is operationally dangerous. The same system deployed on a Class II task with appropriate human validation is operationally valuable. The same system deployed on a Class IV task under any circumstances is an institutional and legal catastrophe.

Section IV

The Ontology

Mapping Defence Operations to Cognitive Primitives

Class I

Deterministic Primitives

"The Calculator"

Truth Condition	Error Tolerance	AI Role
Equality — single correct output	Zero	Executor with formal verification

Logistics Inventory	Calculates days of supply, resupply requirements, and priority allocation from inventory levels, consumption rates, and unit requirements. 100% accuracy with symbolic solver. Error consequence: mission-critical failure under operational conditions.
Maintenance Scheduling	Generates maintenance schedules and parts requisition from equipment hours, intervals, and availability. Deterministic scheduling with defined constraints. Error consequence: deferred maintenance creates cascading readiness failures.
Personnel Calculations	Deployment eligibility, rotation schedules, skill gap identification from qualified records and unit requirements. 100% for structured records. Error consequence: operational and welfare risks from misallocation.
Budget and Procurement	Available budget, procurement feasibility, variance analysis from allocations and expenditure records. 100% with symbolic solver. Error consequence: compliance failures and operational funding gaps.
Traffic Analysis — Structural	Volume aggregation, anomaly flagging from message volumes and network topology. High capability for structured datasets. Error consequence: false positives or negatives in anomaly detection.

Key Insight — Class I:

These operations currently consume significant human time across every level of the defence establishment. Their deterministic nature makes them the immediate, low-risk frontier of AI augmentation — not because AI is sophisticated, but because these operations never required judgment. Freeing senior professionals from Class I execution is the reallocation of human intelligence to operations that actually require it.

Class II

Probabilistic Primitives

"The Analyst"

Truth Condition	Error Tolerance	AI Role
Calibration — confidence matches frequencies	Task-dependent; higher for scenario generation	Generator with human validation

OSINT Synthesis	Structured intelligence summaries and anomaly flags from news feeds, public reports, and open-source data. High capability for large text corpora. OPSEC note: query patterns reveal analytical priorities — sovereign infrastructure mandatory.
Threat Scenario Generation	Plausible threat scenarios with probability ranges from adversary capabilities and operational context. High for scenario diversity; moderate for probability calibration. Error consequence: miscalibrated confidence is operationally equivalent to wrong information delivered with authority.
Planning Options Generation	Multiple course-of-action options with tradeoff analysis from mission objectives and available forces. Moderate capability — tactically sound but contextually incomplete without human integration.
Training Scenario Design	Scenario options with expected learning outcomes from objectives and historical exercise data. High capability for structured scenario generation. Requires unit-specific customisation by human instructors.
Maintenance Failure Prediction	Failure probability distributions and maintenance priorities from sensor data and operational history. 30-40% improvement over rule-based scheduling in analogous industrial contexts. Human override authority architecturally mandatory.

Key Insight — Class II:

Maximum operational value — and maximum risk if misclassified. An AI-generated threat scenario is a starting point for human analysis, not a finished intelligence assessment. An AI-generated course of action is a planning input, not an operational order. The institutional discipline required is not technical — it is epistemological. Senior professionals must know, at the point of use, what kind of output they are receiving.

Class III

Contextual Primitives

"The Commander"

Truth Condition	Error Tolerance	AI Role
Counterfactual effectiveness	High variance acceptable; single catastrophic failure can be irreversible	Sparring partner — human in command

Tactical Decision-Making	Course-of-action selection and execution sequencing. Latent variables: adversary intent, morale state, political constraints, real-time environmental changes. AI generates five options with tradeoff analysis; human commander selects and adapts in real time. Contextual errors are not recoverable through post-hoc correction.
Crisis Communication	Communication strategy — what to say, sequencing, tone, channel. Latent variables: adversary signal interpretation, domestic political constraints, alliance dynamics, historical grievances. AI generates template responses; human leadership navigates real-time geopolitical stress.
Strategic Resource Allocation	Allocation recommendations across competing operational priorities. Latent variables: unstated political priorities, alliance obligations, long-term capability tradeoffs, internal organisational dynamics. AI models allocation scenarios; human decision-maker integrates the unquantifiable.
Negotiation and Interagency Coordination	Negotiation strategy and concession architecture. Latent variables: counterpart's internal political situation, personal motivations, relationship history. AI models game trees; human negotiator reads the room and adapts in real time.
Capability Development Strategy	Long-horizon capability priority recommendations. Latent variables: political appetite for investment, alliance technology sharing, industrial base constraints. AI analyses gaps and models scenarios; human leadership makes the strategic commitment.

Key Insight — Class III:

These are the operations that define the value of experienced senior professionals. They cannot be automated because the inputs that determine outcomes are not in any dataset — they are in the heads, relationships, and experiential pattern libraries of people who have spent careers developing contextual intelligence that no training corpus can replicate. AI amplifies their leverage; it does not replace their judgment.

Class IV

Principal Primitives

"The Commander's Signature"

Truth Condition	Error Tolerance	AI Role
Legitimacy and accountability — personhood required	Existential — single violation destroys institutional legitimacy	Tool only — AI cannot be principal

Command Authority	An operational order carries the legal and moral authority of the commander who issues it. AI-generated text is not an order. Command authority cannot be delegated to non-persons. AI may generate planning inputs; the command decision is exclusively human.
Rules of Engagement	ROE interpretation and application requires accountable human judgment — not only whether force is legally permissible but whether it is morally justified in this specific context, by this specific commander. No legal framework provides for algorithmic authorisation of force.
Ethical Gatekeeping	Deciding whether to undertake a contested operation requires moral agency — the capacity to be held responsible for a values-based judgment. Algorithms have objectives, not values. 'The algorithm decided' is not a defence recognised by any court or moral framework.
Accountability for Outcomes	Military accountability — responsibility that attaches to commanders for operational outcomes — is a foundational principle of military law. It cannot be distributed to a system. When no person is accountable, no person is protected, and no institution is trustworthy.

Strategic Representation	When a senior officer represents the institution in strategic negotiations or alliance engagements, their credibility, relationship capital, and personal accountability are the assets deployed. These are properties of persons. Counterparts cannot negotiate with an algorithm.

Key Insight — Class IV:

These are not tasks AI cannot yet do. They are tasks AI cannot do in principle — because the value derives from personhood, legal authority, and moral accountability, not from computational capability. Attempting to automate Class IV functions is not an engineering ambition. It is a category error with potentially catastrophic institutional, legal, and moral consequences. This boundary is permanent.

Section V

Failure Mode Analysis

The Catastrophic Risks of Misclassification

Type I Failure: Treating Probabilistic Operations as Deterministic

Mechanism	Using an LLM to produce a threat assessment without human validation and treating the output as verified intelligence. The model generates a plausible-looking probability — '72% probability of border incursion in the next 30 days' — that is syntactic confidence, not epistemic probability. Compound failure: probabilistic output drives resource allocation, operational planning proceeds, and when the assessment proves wrong, the failure is misattributed to intelligence failure rather than misclassification. The lesson drawn is wrong. The mistake recurs.
Consequence	False positive: unnecessary mobilisation, operational exposure, political escalation. False negative: unprepared response to a threat within the model's plausibility range.

Type II Failure: Preserving Human Labor on Deterministic Operations

Mechanism	Maintaining manual compilation of logistics readiness reports when automated aggregation would produce more accurate and more timely outputs. Status quo bias, training justifications, and security concerns applied as blanket objections rather than targeted constraints.
Consequence	Senior professionals spend significant time on Class I execution that produces no analytical value, and are therefore unavailable for Class III synthesis where their judgment is genuinely irreplaceable. Chronic underperformance and systematic misallocation of human intelligence.

Type III Failure: Treating Contextual Operations as Probabilistic

Mechanism	Using AI-generated course-of-action analysis as a planning output rather than a planning input, without integrating classified intelligence, unit morale assessments, and political context. A course-of-action analysis generated from observable data can be structurally sound and operationally catastrophic simultaneously.
Consequence	Plans that look right on paper and fail in execution. The failure is not attributable to any single error — it is systemic. The process did not require integration of the latent variables that determined the outcome.

Type IV Failure: Delegating Principal Operations to AI

Mechanism	An operational planning system whose recommendations are consistently implemented without substantive human review — not through explicit policy, but through institutional convenience and time pressure — until the distinction between AI recommendation and human command has effectively disappeared. Type IV failures rarely occur suddenly. They occur through gradual normalisation.
Consequence	When something goes wrong, accountability has no place to attach. The commander signed the order. The AI generated it. The legal and moral framework requires a human principal. The institutional reality is a gap. There is no recovering from this failure mode. It must be prevented absolutely.

Section VI

The Sovereign AI Architecture

Operational Design for Indian Defence Institutions

Layer	Operational Design
Layer 1 Deterministic Automation	All Class I operations on institution-controlled, air-gap capable infrastructure. No query leaves the institutional network. No usage data logged externally. Representative systems: logistics management, personnel calculations, budget variance analysis, communications traffic aggregation. Sovereign advantage: efficiency gains with zero data sovereignty compromise.
Layer 2 Probabilistic Generation	All Class II operations use open-source LLMs on sovereign GPU infrastructure with mandatory human validation before any output is used operationally. Representative workflows: OSINT synthesis, threat scenario generation, planning options, maintenance prediction. Sovereign advantage: OSINT queries reveal analytical priorities — sovereign infrastructure eliminates the query-pattern intelligence vulnerability.
Layer 3 Contextual Synthesis	All Class III operations use AI as analytical input generator. Human commanders retain synthesis authority and decision ownership. Representative workflows: tactical planning, crisis communication, resource allocation, capability strategy. Sovereign advantage: strategic planning queries are among the most sensitive data an institution generates. Sovereign infrastructure is an operational requirement.
Layer 4 Principal Authority	All Class IV operations remain human-only. AI tools may support preparation — briefings, analyses, option summaries — but the decision, the order, and the accountability are exclusively human. Requirement: institutional design and cultural discipline, not technology. The permanent boundary: personhood is a legal and moral category, not a capability threshold.

The Training Imperative

The traditional model — junior professionals learn by doing administrative work, gradually absorbing institutional knowledge and operational judgment — breaks down when administrative work is automated. New professionals cannot 'learn by compiling logistics reports' if logistics reports are auto-generated. A new training paradigm is required: AI literacy and operational fluency, followed by hands-on AI deployment and agent building, followed by sovereign infrastructure configuration and governance. The goal is not to produce AI engineers — it is to produce senior professionals who understand what a sovereign AI deployment looks like, can govern one intelligently, and can ask the right questions of the

technical staff who manage it.

Section VII

Related Work

Theoretical Foundations and Positioning

Task-Based Theories of Automation

The modern analysis of technology's impact on institutional work originates in Autor, Levy, and Murnane (2003), who decompose occupations into tasks classified as routine versus non-routine and manual versus cognitive. Our contribution refines this taxonomy for high-stakes institutional contexts by replacing the routine/non-routine distinction with truth conditions. 'Compiling a logistics readiness report' and 'advising a commander on tactical options' might both be classified as non-routine cognitive in the Autor framework. But these operations have fundamentally different computational structures — a distinction invisible in the routine/non-routine framework and operationally critical in defence contexts.

AI Safety in High-Stakes Domains

The AI safety literature emphasises specification gaming, distributional shift, and value alignment. Our ontology provides a domain-specific instantiation of these abstract challenges for defence contexts. The critical insight — borrowed from the AI safety literature and applied here operationally — is that safety is not a property of AI systems in isolation. It is a property of the task-system pairing under specific truth conditions. This shifts the integration question from 'is this AI system safe?' to 'is this AI system appropriately matched to this operational task under these truth conditions?'

Sovereignty and Institutional AI

The political economy of AI sovereignty has received increasing attention in the academic and policy literature (Dafoe, 2018; Lee, 2018; Roberts et al., 2021). What the existing literature does not provide is an operational framework for translating sovereignty aspirations into institutional design. The four-class ontology addresses this gap directly: it specifies not only why sovereign AI matters but which cognitive operations must be processed on sovereign infrastructure, which can tolerate carefully governed external tools, and which require human authority regardless of infrastructure configuration.

Section VIII

The Research Agenda

An Operational Science for AI in Indian Defence

Paper	Title	Scope
Paper 3	**The Measurement Layer**	Benchmarking AI Systems Against Defence Operational Truth Conditions — A standardised evaluation framework testing AI performance against Class I, II, and III defence operations.
Paper 4	**The Governance Layer**	Institutional Guardrails for AI Integration in High-Stakes Operations — Policy, audit, and compliance architecture including model risk management and Class IV boundary enforcement.
Paper 5	**The Sovereignty Layer**	The Atmanirbhar AI Stack: Technical Architecture for Defence AI Sovereignty — Hardware requirements, open-source model selection, security configuration, and air-gap deployment architecture.
Paper 6	**The Training Layer**	Building Human-AI Complementarity: A Training Framework for AI-Augmented Defence Professionals — Curriculum design from tool literacy through agent deployment to sovereign infrastructure governance.

Section IX

Disclosures and Ethical Statements

Use of Artificial Intelligence

In the spirit of the ontology presented in this paper, the production of this manuscript followed a Human-in-the-Loop (Class III) workflow. Large language models were employed for drafting support, copy-editing, and adversarial stress-testing of arguments. All core conceptual frameworks, ontological definitions, operational case studies, and strategic conclusions are the original intellectual property of the author. No classified, sensitive, or operationally restricted information was processed by any AI system during this research.

Conflict of Interest and Commercial Affiliation

The author, Umran Nayani, is the Founder of Quantumist.ai, a research and education firm dedicated to the operational science of AI integration in high-stakes institutions. The frameworks detailed herein form the theoretical basis for the firm's advisory and training programs, including the AI Enablement Program for defence institutions. This paper is published as open research. Commercial applications of the framework are available through the firm's program offerings.

Scope and Classification

This paper addresses the open, administrative, logistics, intelligence synthesis, and operational planning dimensions of defence institution AI integration. It does not address classified AI applications, autonomous weapons systems, or any application requiring security clearance to discuss. The ontological framework presented here is applicable to publicly available information and open-source AI systems operating under institutional governance.

The defence professional who understands the difference between a probabilistic threat assessment and a verified intelligence product — and governs their operational decisions accordingly — is not more cautious than their peers.

They are more dangerous, in the best sense.

The goal of this ontology is not to constrain AI adoption in defence institutions. It is to ensure that adoption is structured so that AI amplifies operational judgment rather than substituting for it — and that the infrastructure on which that AI runs is as sovereign, secure, and self-reliant as the institution it serves.

— U.N. | Mumbai, March 2026

References

Acemoglu, D., & Restrepo, P. (2018). The Race Between Man and Machine. American Economic Review, 108(1), 1–48.

Amodei, D., et al. (2016). Concrete Problems in AI Safety. arXiv:1606.06565.

Anthropic. (2024). Claude Constitutional AI: Harmlessness from AI Feedback. Anthropic Technical Report.

Autor, D. H., Levy, F., & Murnane, R. J. (2003). The Skill Content of Recent Technological Change. Quarterly Journal of Economics, 118(4), 1279–1333.

Bostrom, N. (2014). Superintelligence: Paths, Dangers, Strategies. Oxford University Press.

Dafoe, A. (2018). AI Governance: A Research Agenda. Future of Humanity Institute, University of Oxford.

Government of India. (2018). National Strategy for Artificial Intelligence. NITI Aayog.

Government of India. (2024). Digital Personal Data Protection Act. Ministry of Electronics and Information Technology.

Lee, K. F. (2018). AI Superpowers: China, Silicon Valley, and the New World Order. Houghton Mifflin Harcourt.

Ministry of Defence, Government of India. (2021). Artificial Intelligence in Defence: Strategy and Implementation Roadmap.

Nayani, U. (2025). The Cognitive Primitives of Investment Banking: An Ontology for AI-Driven Augmentation in High-Stakes Finance. SSRN Working Paper No. 5963734. Quantumist.ai Research

Series, Paper No. 01.

Nayani, U. (2025). Quantum Entanglement in Corporate Valuation: Oracle's Borrowed Inference and the Catastrophic Risk of Non-Local Coherence Collapse. SSRN Working Paper No. 5920762.

Roberts, H., et al. (2021). The Chinese Approach to Artificial Intelligence. AI & Society, 36(1), 59–77.

Russell, S. (2019). Human Compatible: Artificial Intelligence and the Problem of Control. Viking.

Simon, H. A. (1962). The Architecture of Complexity. Proceedings of the American Philosophical Society, 106(6), 467–482.

Vaswani, A., et al. (2017). Attention Is All You Need. Advances in Neural Information Processing Systems, 30.

COPYRIGHT © 2026 UMRAN NAYANI / QUANTUMIST.AI. ALL RIGHTS RESERVED. Distributed for academic and non-commercial discussion. Usage of the 'Cognitive Primitives' taxonomy and 'Sovereign AI Architecture' framework in commercial or institutional applications requires written permission and attribution.