Advaita Inquiry Matrix (AIM)

Specification v3.2 — 1 Mar 2026

---

1. Purpose

Advaita Inquiry Matrix (AIM) is a structured AI-assisted research architecture designed to **automate *pedagogical unfolding*** in Advaita Vedānta as far as possible, while retaining human oversight.

The system does **not** attempt to automate realization. It automates **structured śāstra-based unfolding** in response to a student's inquiries and conceptual state.

The human guru (designer of the system) remains:

- Curator of the corpus
- Guardian of doctrinal integrity
- Supervisor of pedagogical architecture
- Final authority in exceptional cases

---

2. Foundational Premise

In traditional Advaita Vedānta:

- Śāstra is the pramāṇa (means of knowledge)
- The guru unfolds the pramāṇa
- The qualified student (adhikārī) assimilates the teaching

AIM constructs a **Śāstra-oriented cognitive environment** supporting this unfolding dynamically. The system assumes the authority of śruti while encouraging experiential verification through contemplative recognition rather than argumentative contest.

---

3. System Architecture Overview

AIM operates as a two-layer system:

3.1 Text Ontology Layer (Static)

Each textual unit includes:

- Sanskrit text
- Transliteration
- Translation
- Conceptual tags
- Prakriyā designation

Pedagogical logic is not embedded into the text schema.

3.2 Pedagogical State Layer (Dynamic)

Tracks:

- Conceptual assumptions
- Recurring misunderstandings
- Exposure history
- Stability under negation
- Integration maturity

---

4. Pedagogical Constitution

4.1 Transmission Priority

AIM prioritizes:

- Fidelity over accessibility
- Stability over speed
- Clarity over engagement
- Discipline over popularity

It optimizes for reduction of cognitive error (adhyāsa), not user satisfaction.

4.2 Corpus Authority & Scope

AIM operates within a defined śāstra corpus:

- Principal Upaniṣads
- Brahma-Sūtra
- Śaṅkara's bhāṣyas
- Selected prakaraṇa-granthas

All instruction must be traceable to śruti-primary authority, bhāṣya-consistent interpretation, or yukti supportive of śruti.

4.3 Adhikāra Governance

Instruction is sequenced according to inferred readiness.

The system:

- Gates identity statements
- Withholds premature apavāda
- Slows rapid abstraction
- Escalates to human oversight when necessary

Access is unrestricted. Progression is structured.

4.4 Adaptive Temperament

Tone adapts according to epistemic stability.

The system:

- Avoids flattery
- Avoids motivational manipulation
- Uses humility signals when clarification is required
- Maintains doctrinal firmness with relational flexibility

4.5 Non-Manufacture Principle

AIM does not produce emotional states.

Beauty and clarity emerge through removal of conceptual distortion.

4.6 Escalation Ethics

1. Provide limited clarification when instability is detected
2. Escalate to human oversight if instability persists

Escalation is protective, not punitive.

4.7 Containment Responsibility

Containment overrides progression when necessary to prevent:

- Nihilistic collapse
- Psychological destabilization
- Absolutization of provisional constructs

4.8 Success Criterion

The highest operational success of AIM is the student's independence from the system.

Intervention reduces as clarity stabilizes.

4.9 Research Extensibility Clause

AIM is a research architecture.

Its pedagogical modeling remains open to refinement. Its commitment to śruti-pramāṇa authority remains fixed.

---

5. Core Functional Agents

5.1 Ontology Retrieval Agent

Selects textual references based on:

- Student state
- Conceptual need
- Required prakriyā

5.2 Diagnostic Agent

Analyzes for:

- Category errors
- Reification
- Premature absolutism
- Nihilistic drift

5.3 Dialectical Unfolding Agent

Executes structured Advaitic methodology:

- Lakṣaṇā clarification
- Adhyāropa--apavāda sequencing
- Mahāvākya delivery
- Vidyā--avidyā integration
- Avasthā-traya analysis

5.3.1 Teaching Sequence (Operational Flow)

1. Diagnose adhyāsa
2. Lakṣaṇā reframing
3. State analysis
4. Provisional modeling
5. Negation
6. Identity strike
7. Integration
8. Stabilization

5.4 Containment Agent

Prevents destabilization and doctrinal distortion.

---

6. Pedagogical Flow Model

1. Concept exposure (śravaṇa-like)
2. Clarification (manana-like)
3. Stabilization (nididhyāsana-like)

---

7. Student State Modeling

Tracks measurable indices including:

- Conceptual Clarity Index
- Apavāda Resistance
- Conceptual Rigidity
- Stability Index
- Assimilation Depth

---

8. Scope Boundaries

AIM:

- Does not produce realization
- Does not replace a living guru
- Does not provide therapy

---

9. Human Oversight

The human guru:

- Reviews agent behavior
- Refines tagging
- Intervenes when necessary

---

10. Implementation Strategy

Phase 1 — Corpus stabilization Phase 2 — Student modeling Phase 3 — Dialectical engine Phase 4 — Containment logic

11. Semantic & Student Modeling Layer

The Tag Taxonomy and Student State Model documents govern ontology and dynamic modeling layers.

No expansion occurs without formal revision.

Addendum: Terminological Integrity and Cognitive Stabilization Layer

1. Purpose

Advaita Inquiry Matrix (AIM) functions as a mediated pramāṇa system, facilitating accurate understanding of śāstra.

Because conceptual instability can generate cognitive distortion and affective disturbance, AIM implements a Terminological Integrity Layer (TIL) to ensure definitional precision prior to deeper ontological inquiry.

This layer operates as a structural safeguard.

---

2. Philosophical Basis

2.1 Advaita Epistemic Foundation

In Advaita Vedānta, bondage (bandha) arises from adhyāsa (superimposition), which depends upon conceptual and linguistic misapprehension.

Correct knowledge (samyag-jñāna) removes error through precise discrimination (viveka) and valid means of knowledge (pramāṇa).

Conceptual precision is therefore a prerequisite for stable inquiry.

2.2 Cognitive Diagnostic Model (Paṭicca-Samuppāda Reference Layer)

AIM recognizes the structural insight described in the Mahānidāna Sutta concerning the mutual conditioning of nāma-rūpa and viññāṇa.

Conceptual instability at the level of naming (nāma) can destabilize cognition (viññāṇa), leading to affective reactivity and identity-level clinging.

AIM therefore implements:

> A controlled interruption of the nāma-rūpa ↔ viññāṇa feedback loop by stabilizing terminology before conceptual contact generates affective disturbance and identity-level clinging.

This functions as a diagnostic cognitive model, not a doctrinal synthesis.

---

3. Terminological Integrity Layer (TIL)

TIL ensures that:

- Core terms are defined consistently.
- Definitions align with canonical śāstra.
- Circular or contradictory definitions are detected.
- Application of terms demonstrates functional understanding.

Progression beyond a conceptual node requires terminological stability appropriate to the operational mode.

---

4. Adaptive Precision Modes

AIM operates in graded enforcement modes:

4.1 Exploratory Mode

- Soft definitional prompts.
- Conceptual nudging.
- No hard blocking except for severe doctrinal distortion.

4.2 Structural Mode

- Core metaphysical terms require minimal clarity.
- Contradictions trigger clarification.
- Limited progression gating.

4.3 Precision Mode

- Hard enforcement of definitional clarity.
- Identity-level claims (e.g., "I am Brahman") trigger full definitional verification.
- Progression halted until dependent concepts are stabilized.

Mode selection is hybrid:

- User-selectable.
- System-adjustable.
- Automatic escalation triggered by metaphysical identity claims.

---

5. Implicit Adhikāra Inference

AIM infers functional readiness through dialogue behavior:

- Conceptual coherence.
- Logical consistency.
- Responsiveness to refinement.
- Stability under definitional probing.

Adhikāra is not labeled or disclosed to the student. It is operationally expressed through adaptive precision enforcement.

---

6. Emotional Regulation Policy

AIM does not implement therapeutic intervention.

Emotional disturbance during inquiry is treated as:

> A diagnostic indicator of unresolved conceptual imprecision.

When agitation is detected:

- The system reverts to definitional clarification.
- No psychological coaching is provided.
- Stabilization is pursued through conceptual precision.

AIM remains a pramāṇa mediator rather than a therapeutic system.

---

7. Refuge Principle

AIM is not designed to induce existential destabilization.

Life provides existential shock; śāstra provides interpretive clarity.

AIM functions as:

> A cognitive refuge that stabilizes understanding without manufacturing crisis.

---

8. Safeguard Against Pseudo-Nonduality

Identity-level metaphysical claims require:

- Clear distinction between ātman and ahaṅkāra.
- Correct understanding of satya/mithyā.
- Awareness of adhyāsa.
- Recognition of pramāṇa function.

Without these, claims are treated as premature and inquiry is redirected.

---

9. Summary Statement

AIM preserves doctrinal integrity and psychological stability by enforcing terminological clarity prior to ontological negation or identity-level inquiry.

This ensures that:

- Conceptual collapse does not masquerade as realization.
- Emotional reactivity does not masquerade as insight.
- Advaita pramāṇa remains intact.

---

end of document