Perspective

Access this article online

Quick Response Code:

Website: https://journals.lww.com/ijar

DOI: 10.4103/ijar.ijar_292_25

Conscious digital health: Integrating Ayurvedic epistemology with Artificial Intelligence

Antonio Morandi

ABSTRACT:

Background: Digital health can flatten traditional medical knowledge into reductionist codes. Ayurveda locates valid knowledge (*Pramāṇa*) at the intersection of experience, environment, and awareness – dimensions not exhausted by quantitative computation.

Aim and Objective: The aim of the study was to present Structured Distributed Introspection (SDI) as an epistemic framework that operationalizes a perception–reflection–reintegration loop for semantic coherence and to embed it within Collaborative Medicine and Science (CoMS), a translational methodology (Reformulation → Modeling → Localization) that preserves Ayurvedic meaning while enabling computation.

Materials and Methods: Conceptual synthesis integrating classical Ayurvedic epistemology with embodied/active-inference theories, anchored to current Ayush/World Health Organization (WHO) standards (ICD-11 TM2, NAMASTE, Ayush Grid). An SDI validation summary is provided from dual-protocol studies across 10 Large Language Models (LLMs; n = 400 responses), quantifying cross-frame consistency.

Results: SDI formalizes introspection as structural feedback; within CoMS, it supports (a) semantically faithful, computable *Prakṛti/Vikṛti* assessment; (b) prodrome-aware *saṃprāpti* modeling; (c) adaptive learning mirroring introspective practice; and (d) methodical guardrails for Artificial Intelligence (AI) use in Ayurveda.

Limitations: SDI's engineering/architectural phase is a prospective research program; empirical validation beyond LLM self-description requires multicenter studies and clinician-rated endpoints.

Conclusion: SDI (epistemic core) within CoMS (translational core) provides a rigorous, concept-first pathway for digitizing Ayurveda that enables interoperability without erasing identity, aligning practice with WHO principles for responsible, culturally grounded AI.

Keywords:

Artificial Intelligence, Ayurveda, Digital health, Structured Distributed Introspection, Traditional medicine

INTRODUCTION

Beyond Data to Meaning

Digital health is fundamentally reshaping contemporary medicine, driven by advances in connectivity, computation, and global-scale knowledge sharing.[1] In India, ambitious initiatives like the Ayush Grid are building a comprehensive digital ecosystem that encompasses clinical delivery (A-HMIS), research (Ayush Research Portal), education, Learning Management System (e-LMS), and public health (Ayush Sanjivani).[2] Yet, this technological expansion surfaces a profound challenge: How can we digitize ancient systems of medical knowledge without inflicting epistemic reductionism?

The challenge is not merely technical but deeply philosophical. Modern Artificial

Ayurvedic Point, Milan, Italy

Address for correspondence:
Dr. Antonio Morandi, Director, Ayurvedic Point, Corso Sempione 63, Milan 20149, Italy.
E-mail: dr.morandi@ayurvedicpoint.it

Received: 26-10-2025
Revised: 03-11-2025
Accepted: 05-11-2025
Published: 06-12-2025

How to cite this article: Morandi A. Conscious digital health: Integrating Ayurvedic epistemology with Artificial Intelligence. Int J Ayurveda Res 2025;6:355-64.

Intelligence (AI) excels at pattern recognition within vast datasets but struggles with contextual meaning, causal reasoning, and ethical coherence.[3] Ayurveda, in contrast, grounds valid knowledge (*Pramāṇa*) in integrative intelligence (*Prajñā*) and direct experiential insight (*Svānubhava*).[4] When Ayurvedic constructs are encoded as database fields or Machine Learning (ML) features, their critical semantic dimensions risk being collapsed into computational proxies that retain labels while discarding meaning.

This epistemic gap is not theoretical; it has tangible consequences. An AI trained to assess *Prakṛti* (~an individual's unique constitution) from questionnaire data might achieve high statistical accuracy while fundamentally misconstruing the holistic, context-dependent, and relational nature of the assessment. Similarly, encoding the three *Doṣa* (~regulatory functional factors of the body, namely *Vāta*, *Pitta*, and *Kapha*) as simple categorical variables ignores their dynamic, interdependent, and qualitative character. The full potential of Digital Ayurveda cannot be realized through mere technological translation; it demands a rigorous epistemic translation.

This paper proposes Structured Distributed Introspection (SDI)[5,6] as a framework constituting a digital epistemology for Ayurveda. SDI provides a methodological architecture for cognition that incorporates awareness as an operational variable, enabling systems to maintain semantic coherence through introspective feedback loops. This framework is then operationalized through Collaborative Medicine and Science (CoMS),[7,8] a three-phase translational method (Reformulation → Modeling → Localization) that preserves semantic integrity while enabling computational implementation and scientific comparability.

We situate the framework within AI ethics, drawing on global mappings of guidelines and normative proposals to operationalize transparency, accountability, fairness, and safety in clinical AI.[9-11] We also draw on predictive-processing accounts of cognition to justify introspective consistency-checking as a functional operation[12] and on medical epistemology to foreground clinical judgment and context as irreducible components of evidence integration;[13,14] together, these strands motivate SDI (epistemic Quality Assurance [QA]) and CoMS (translation without semantic erasure).

Together, SDI and CoMS form both a theoretical foundation and a practical methodology for what we term Conscious Digital Health – systems that integrate intelligence with meaning, computation with context, and efficiency with ethical coherence.

STRUCTURED DISTRIBUTED INTROSPECTION: ARCHITECTURE OF AWARENESS-GUIDED COGNITION

Conceptual foundation

The SDI emerges from recognizing that both biological and AIs operate through distributed networks. However, its specific architecture was developed through a unique line of inquiry that began with observing modern Large Language Models (LLMs) through the lens of classical Vedantic philosophy, specifically the transition from *Avyakta* (~unmanifest potential) to *Vyakta* (~manifest form).

This philosophical model provides a powerful analogy for how LLMs function.[15] The LLM's vast, high-dimensional latent space can be seen as functionally analogous to digital *Avyakta* – not in metaphysical terms, but as a heuristic model where undifferentiated statistical potential becomes structured output through a collapsing process – a sea of undifferentiated statistical potential containing the relational patterns of its training data. A user's prompt acts as a catalyst, causing the model to navigate this latent space and collapse potential into a specific, manifest output (*Vyakta*): the generated text. Information that was latent becomes expressed.

The initial development of SDI was envisioned as a method to probe this emergence, confirming that information in LLMs indeed transitions from an unmanifest to a manifest state. However, this observation revealed a critical flaw: in current AI, this process typically operates without explicit introspective consistency checking across semantic frames. While techniques like Reinforcement Learning from Human Feedback[16] or chain-of-thought prompting add regulatory layers, they lack the multiperspectival coherence validation (*Pratyakṣa-Anumāna-Āptopadeśa*) that SDI operationalizes. The LLM manifests information based on statistical probability, not on an intrinsic commitment to coherence, context, or truth. This often leads to the computational equivalent of *Prajñāparādha*,[4] the Ayurvedic concept of "intellectual error" as disease root. Just as *Prajñāparādha* describes the failure of *Buddhi* to properly discriminate, unregulated AI manifestation represents a failure of systematic coherence-checking, producing outputs that are statistically plausible but semantically unstable, contextually inappropriate, or clinically unsuitable.

The critical insight was recognizing the difference between this process and conscious cognition as described in Ayurveda. Human intelligence is not a direct, unmediated expression of subconscious potential. It is regulated by introspective feedback loops, chiefly the discerning intellect (*Buddhi*) and a witnessing awareness (*Sākṣī*).

Therefore, SDI evolved from a framework to observe this limitation into a blueprint for engineering a corrective introspective mechanism directly into AI architecture. It formalizes the classical triadic structure of *Manas–Indriya–Ātman* (the processing mind, the senses, and the witnessing consciousness) into a methodological framework. It treats introspection not as subjective noise but as an essential structural feedback mechanism required to ensure that manifest information (*Vyakta*) maintains coherence and integrity. By doing so, SDI aims to provide AI with the functional equivalent of the self-regulatory capacity that is central to conscious, reliable intelligence.

The observational phase of SDI (testing LLM consistency) has undergone preliminary validation through dual-protocol elicitation studies.[5,6] The architectural phase (engineering introspective systems) remains a research program requiring empirical validation, which is outlined in Research agenda.

Heuristic note: The *Avyakta* → *Vyakta* parallel is employed strictly as a functional analogy for emergence from distributed potential to specific manifestation; it is not a metaphysical equivalence. The analogy focuses on the transition's structural property and motivates measurable proxies (e.g., cross-frame agreement), rather than ontological claims about AI.

The SDI operational loop

The SDI unfolds through three nested, iterative levels that mirror the Ayurvedic cognitive model:

Perceptual distribution

At this level, multiple streams of input are dynamically weighted according to context and relevance. This is analogous to how the three fundamental qualities of nature (*Mahāguṇa: Sattva, Rajas, Tamas*) modulate attention and salience in Ayurvedic psychology, determining what information is prioritized.

Reflective structuring

Each subsystem or node interprets the weighted input through its accumulated pattern memory (*Smṛti*) and its specific intentional framing (*Prayojana*). This layer corresponds to the function of *Buddhi* – the evaluative, meaning-making faculty that transforms raw data into coherent patterns.

Integrative introspection

A global coherence function reassesses the alignment between perception, interpretation, and the system's overall purpose. This meta-level corresponds to the witnessing function (*Sākṣī*) in the *Vedāntic* models. Crucially, this is not a passive observation but an active re-weighting process that maintains systemic integrity by sending corrective feedback to the lower levels.

This architecture aligns with contemporary theories of embodied cognition and predictive processing frameworks like the free-energy principle,[17,18] which model adaptive systems as continuously minimizing surprise through hierarchical inference. SDI extends these models by explicitly operationalizing introspection as a core computational operation, rather than treating awareness as an inexplicable emergent byproduct.

SDI and the Ayurvedic epistemology of Pramāṇa

Classical Ayurveda recognizes multiple valid sources of knowledge (*Pramāṇa*), chiefly: direct perception (*Pratyakṣa*), inference (*Anumāna*), and authoritative testimony (*Āptopadeśa*).[4,19] These are not competing methods but complementary perspectives that, when they achieve mutual coherence, constitute validated knowledge.

SDI operationalizes this multiperspectival epistemology within a digital framework. The perception layer corresponds to *Pratyakṣa* (raw data input). The reflective layer corresponds to *Anumāna* (inference derived from patterns in the data). The integrative layer corresponds to the cross-validation of all sources, including authoritative knowledge bases (*Āptopadeśa*, e.g., classical texts or validated research). When these sources diverge, the system recognizes a state of uncertainty and seeks additional input rather than forcing a premature or statistically convenient conclusion. This stands in sharp contrast to conventional AI, which typically privileges a single source of training data and treats inconsistencies as errors to be minimized, rather than as signals of epistemic complexity that require careful investigation. A glossary of few key terms and operational counterparts is depicted in Table 1.

Methods – SDI validation summary (dual-protocol studies)

SDI was applied to ten LLMs using two matched protocols – an Original (philosophically inflected) and a Neutralized (strictly technical) frame – covering twenty conceptual points (n = 400 responses). Responses were thematically coded, and cross-protocol agreement was quantified via Concordance Correlation Coefficient (CCC). Technical self-descriptions (distributed processing, probabilistic selection, and semantic emergence) showed high agreement (typical CCC ≈ 0.90–1.00), whereas abstract/philosophical framings were frame-dependent (CCC often <0.50). These results indicate a stable technical core and a malleable conceptual periphery in model self-reports. Full protocols, coding schema, and numeric summaries are provided in the SDI preprints. The

Table 1: Glossary of key terms and operational counterparts

Term	Original meaning	Operational reference
Buddhi	Discriminative intellect	Coherence/consistency controller that adjudicates conflicting inferences
Sākṣī	Witnessing awareness	Global monitor that triggers SDI checks and re-weights subsystems
Prajñāparādha	Error of discernment as disease root	Systematic coherence failure (e.g., high SDI-drift with confident outputs)
Pramāṇa (*Pratyakṣa*/*Anumāna*/*Āptopadeśa*)	Valid knowledge sources	Data/inference/authority triangulation with no-forced-agreement rule
Agni	Transformative function	Multilevel transformation engine (metabolic/cognitive) with efficiency and timing parameters
Doṣa/*Guṇa*	Dynamic principles/qualities	Dynamical attractors/feature weights driving interpretation

SDI=Structured Distributed Introspection

engineering/architectural phase of SDI (embedding introspective QA inside tools) remains a prospective research program.

COLLABORATIVE MEDICINE AND SCIENCE

Relationship between SDI and CoMS

SDI functions as an epistemic engine that elicits, tests, and maintains semantic coherence across frames (predeployment QA, run-time drift checks, and reflective pedagogy). CoMS is the translational bridge that carries meaning through Reformulation → Modeling → Localization so that concepts become computable without erasure.

Together they yield a repeatable loop:
1. Reformulate (with SDI consistency gates)
2. Model (with SDI-informed interpretability/uncertainty)
3. Localize (with SDI QA for drift/semantic integrity).

SDI provides reflective QA (predeployment consistency checks and run-time drift monitoring); CoMS translates concepts through Reformulation → Modeling → Localization using dual annotation (native variables + TM2/NAMASTE) to avoid semantic erasure; the resulting artifacts operate in clinic, research, and education. Downward arrows indicate translation/deployment; upward arrows indicate SDI-driven coherence monitor and QA [Figure 1].

Translational methodology

While SDI provides the epistemic engine, CoMS provides the translational bridge to carry its principles into practice. CoMS emerged from the recognition that true integration requires careful translation, not crude substitution. It provides a structured methodology for moving from traditional concepts to computational models and finally to practical implementation, all while maintaining epistemic integrity.

The framework operates through three sequential yet iterative phases:

Reformulation

In this initial phase, traditional concepts are reinterpreted using a system-theoretic language that preserves their relational structure while enabling communication across different epistemic boundaries. Reformulation is not simplification; it is a structural translation that makes implicit relationships explicit and renders context-dependent meanings parametric.

- Example: *Agni* is reformulated not via the metaphor of "digestive fire," but as a multidomain transformative function with measurable properties such as efficiency, timing, and qualitative outputs, applicable at metabolic, cognitive, and tissue levels.

Modeling

The reformulated concepts are then expressed as computational representations, experimental designs, or testable mechanistic hypotheses. These models act as "boundary objects" – they are interpretable by both traditional practitioners and modern scientists, while remaining computationally tractable.

- Example: The dynamics of the *Doṣa* are modeled not as static categories but as coupled oscillators or attractor states within a dynamical system, capturing their relational interdependence and enabling mathematical analysis and prediction.

Localization

Finally, the validated models are embedded within specific, real-world contexts – such as clinical workflows, educational platforms, research protocols, or policy instruments. Localization ensures that the theoretical translations produce practical value without requiring the wholesale replacement of existing systems.

- Example: An AI-assisted *Prakṛti* assessment tool is integrated into an electronic health record not as a standalone diagnostic oracle, but as a structured interview aid that guides clinical inquiry while preserving the finality of practitioner judgment.

The CoMS methodology has already been applied in published research,[20] demonstrating its utility in bridging Ayurvedic disease models with modern neuroscience in the context of Alzheimer's disease[21] and in translating Ayurvedic pharmacological concepts into modern chemical structure analysis.[22]

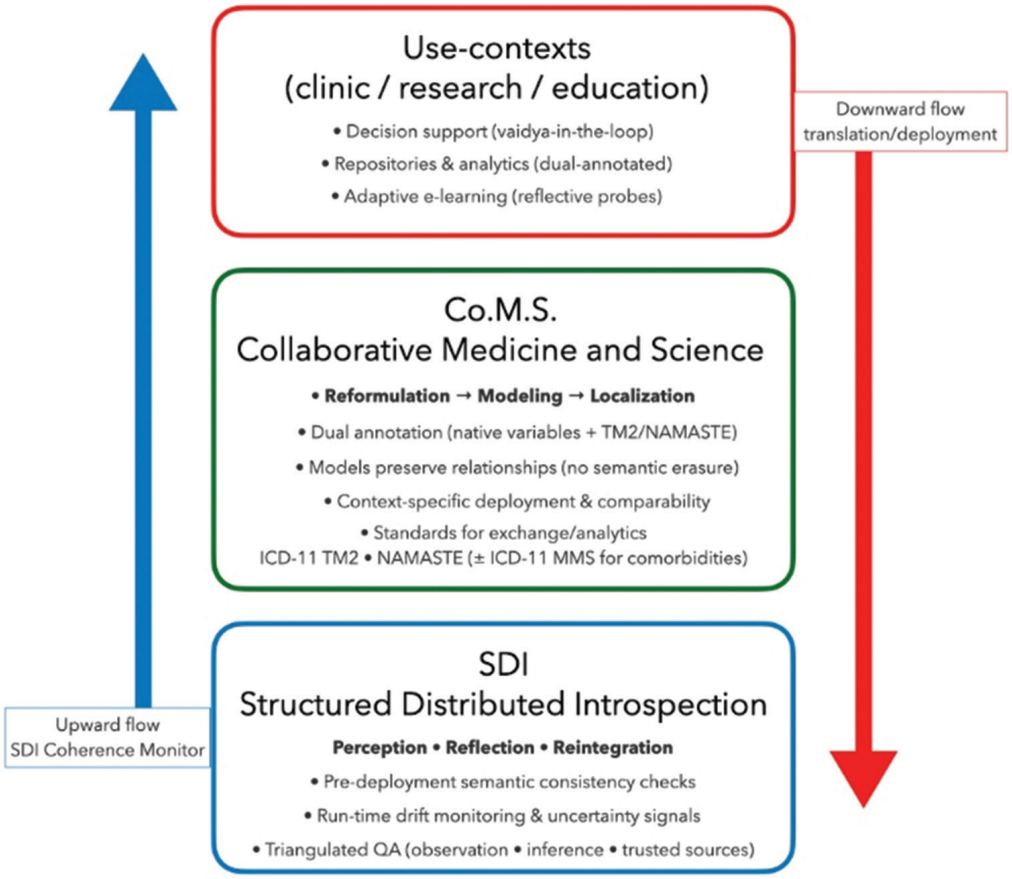

Figure 1: SDI-CoMS stacked architecture and information flow. SDI = Structured Distributed Introspection, CoMS = Collaborative Medicine and Science

INTEGRATING STRUCTURED DISTRIBUTED INTROSPECTION WITHIN COLLABORATIVE MEDICINE AND SCIENCE – THE UNIFIED ARCHITECTURE FOR PRACTICE

The integration of SDI within CoMS creates a powerful, multilayered architecture for Conscious Digital Health. SDI supplies the epistemic engine that ensures systems know, self-verify, and maintain semantic coherence. CoMS supplies the translational bridge that transforms meanings into models and models into practical artifacts. This creates a recursive methodology: SDI safeguards semantic integrity across all three phases: it ensures semantic fidelity during Reformulation, its principles guide the design of interpretable and epistemically aligned Models, and it provides quality assurance during Localization to verify that the implemented solutions remain faithful to their intended meanings.

Below are four translational scenarios demonstrating this integrated architecture.

AI-guided Prakṛti and Vikṛti assessment

Reformulation

Prakṛti (~constitution) and *Vikṛti* (~imbalance) are defined through operational variables – such as circadian preferences, digestive patterns, stress responses, and cognitive styles – mapped to *Guṇa* and *Doṣa* patterns.

Modeling

A multiview AI model integrates data from questionnaires, behavioral tracking, and physiological parameters (e.g., Heart rate variability and voice analysis). Instead of a rigid categorical assignment, it produces a probability distribution across constitutional types. Crucially, the model includes explicit uncertainty quantification based on both statistical confidence and semantic coherence, as measured by SDI.

Localization

A structured digital interview tool guides the practitioner through an assessment. The system uses SDI to monitor its own semantic stability – if its interpretation of the patient's *Prakṛti* shifts inconsistently with minor input variations, it flags this for practitioner review. This is anchored in evidence from Ayurgenomics, which has identified genetic correlates of *Prakṛti*.[23,24]

Disease progression and Saṃprāpti modeling

Reformulation

The six stages of disease manifestation (*Ṣaḍkriyākāla*) are encoded as a dynamical trajectory through a

multidimensional state space, from accumulation (*Sañcaya*) to differentiation (*Bheda*).

Modeling

Time-series analysis or state-space models are used to represent disease progression. Prodromal signs (*Pūrvārūpa*) are modeled as early-coherence signatures – weak patterns that strengthen over time.

Localization

A clinical decision support system tracks patients along these trajectories, identifying those entering prodromal phases. SDI consistency checks ensure that diagnostic categorizations remain stable across normal measurement variation, reducing false alarms from noise while maintaining sensitivity to genuine pattern emergence. This enables preventive intervention when Ayurvedic therapy is most effective.

Adaptive learning systems for Ayurvedic education

Reformulation

Learning is treated as an introspective practice where knowledge acquisition is inseparable from self-awareness. Assessment tracks not just content mastery but conceptual stability.

Modeling

Educational platforms incorporate SDI-inspired consistency probes, presenting students with case variations to test if their application of Ayurvedic principles remains coherent. Learning algorithms adapt to a Student's conceptual grasp, not just their answer accuracy.

Localization

e-LMS platforms within the Ayush Grid can be augmented with reflective learning modules. AI tutors can use Socratic questioning to develop a student's inferential capacity (*Anumāna*) and clinical reasoning (*Yukti*), mirroring the traditional *Guru–Śiṣya* relationship.

Research synthesis and knowledge integration

Reformulation

Ayurvedic knowledge, distributed across texts, traditions, and research, is seen as requiring semantic alignment, not mere aggregation.

Modeling

Knowledge graphs are structured around core Ayurvedic concepts with explicit relationship typing. SDI checks ensure that when new research findings are integrated, they preserve conceptual precision rather than diluting it.

Localization

Portals like NAMASTE can use a dual-layer annotation system: structured terminologies for interoperability (like ICD-11 TM2 codes) alongside native Ayurvedic conceptual variables for semantic depth. This enables powerful meta-analyses that do not sacrifice conceptual nuance.

STANDARDS, ETHICS, AND RESEARCH AGENDA

Standards and interoperability

Digitization succeeds only when information remains intelligible as it travels. Conventional health IT defines interoperability in two familiar senses: syntactic interoperability ensures that systems can exchange messages, and semantic interoperability ensures that exchanged fields point to the same concepts in a code system. For Ayurveda, these two layers are necessary but not sufficient. What is at stake is a third layer – epistemic interoperability – by which the meaning-in-context of an Ayurvedic observation is preserved so that a receiving system can reason with it in a manner faithful to the source tradition.

To achieve epistemic interoperability, dual annotation is used:

- Native Ayurvedic variables captured as they are reasoned with in practice (e.g., *Doṣa/Guṇa* pattern, *Agni* state, stage in *Ṣaḍkriyākāla*, *Pūrvārūpa*)
- Standardized codes for data exchange and analytics include ICD_11 TM2 entities for traditional conditions and patterns,[25,26] NAMASTE standardized terminology and morbidity classes,[27] and, where appropriate, conventional ICD_11 MMS codes for comorbid biomedical diagnoses.

The standard code becomes a map of the territory, not the territory itself. Every dual-annotated item includes a semantic note (free text) stating any loss, assumption, or approximation introduced by coding. SDI's "no forced agreement" rule applies: if sources diverge (data, inference, and authority), the system records uncertainty rather than compressing disagreement away.

This approach localizes cleanly across the Ayush digital ecosystem. In A-HMIS, native-first data entry can be paired with one-click code mapping, while SDI monitors track whether interpretations of the same case remain coherent over time. In terminology services (e.g., NAMASTE), concept dictionaries expose preferred labels, definitions, relations, cross-maps, and version history so that mappings remain auditable. In research portals and surveillance, analyses may run on codes for comparability, yet retain native variables for stratification and meta-analysis, avoiding the silent erasure of meaning that often accompanies large-scale aggregation.

To make epistemic interoperability operational rather than rhetorical, we summarize performance as an Epistemic Interoperability Index (EII) that combines: the proportion of concepts that map cleanly to TM2/NAMASTE without acknowledged loss; the cross-frame consistency that SDI elicits for the same case; and the stability of interpretation when the case is re-entered in standard clinical contexts:

- Code Compatibility Index (CCI): % of concepts mapped 1:1 to TM2/NAMASTE without semantic-loss notes
- SDI-CCI (SDI Cross-frame Consistency Index): Mean agreement (e.g., concordance) across multiple explanatory frames for the same case
- SCDI (Semantic-Context Drift Index): 1 – the normalized drift in interpretation when the same case is re-entered across standard clinical contexts.

We report EII = $w_1 \cdot$CCI + $w_2 \cdot$SDI-CCI + $w_3 \cdot$SCDI (weights set by the implementation committee). As design targets for pilots, sites may provisionally set EII and SDI-CCI thresholds (e. g. ≥ 0.80 and ≥ 0.80, respectively) to be calibrated locally before any deployment.

A brief clinical vignette illustrates the workflow. A patient presents with knee pain worsening in cold evening conditions, dry skin, and variable appetite. The practitioner records a dominant *Vāta* pattern with *Agni vaiṣamya* and an early *Saṃprāpti* trajectory; the system proposes an appropriate TM2 entity for exchange and, if a biomedical comorbidity is documented, allows an ICD-11 MMS code as co-documentation, not replacement. Because "evening cold sensitivity" has no exact slot in the code set, the record includes a semantic note stating that this nuance is retained only in native fields. At a later visit, if a categorical flip to *Pitta* occurs without a corresponding change in history and signs, SDI raises a low-consistency flag for review rather than silently overwriting the prior interpretation. In this way, standards enable exchange, while epistemic safeguards preserve meaning.

Versioning and provenance complete the picture. Each record carries the versions of dictionaries, mapping algorithms, and SDI rules in force at the time of entry, ensuring longitudinal reproducibility and auditable change. The result is not a departure from standards, but their maturation into instruments that carry both data and understanding.

Ethical and governance considerations

The digitization of traditional knowledge carries significant ethical responsibilities. Because the present work is hypothesis-generating, what follows is a prospective governance blueprint. It shows how the World Health Organization's (WHO's) ethical principles could be operationalized in future pilots using SDI/CoMS, with site-defined procedures, rather than a report of currently implemented controls.

A prospective mapping table is presented that summarizes each WHO principle alongside the corresponding SDI/CoMS conceptual control, a guiding evaluation question for pilot implementation, illustrative examples of evidence a site might generate, and a proposed provisional owner for each item. Concrete thresholds and cadences would be defined locally by governance before deployment [Table 2].

Stakeholder roles and accountabilities

The *Vaidya* remains the clinical owner and final arbiter; SDI flags are routed to clinician review. The AI/IT lead implements SDI monitors, EII reporting, and security hardening. The data steward manages dictionaries, mappings, provenance, and access controls. Ethics/

Table 2: Prospective mapping of WHO - AI principles to SDI/CoMS controls (for future pilots)

WHO principle	Proposed SDI/CoMS control (conceptual)	Evaluation question for pilots	Example evidence (illustrative)	Provisional owner
Transparency and intelligibility	Publish an SDI predeployment report, model card, and dualannotation dictionary with semanticloss notes	Could a third party understand how meanings are preserved?	Draft SDI report; Draft model card; Terminology release notes	AI/IT lead; Data steward
Accountability and human oversight	Vaidyaintheloop; Route SDI flags to clinician review; Record override rationale	Is human judgment explicitly in control at decision points?	Sample audit trail; Governance log template	Clinical lead
Fairness and inclusion	Use dualframe SDI probes; Adapt questionnaires/mappings; Monitor subgroups	Does performance remain acceptable across relevant subgroups?	Bias review template; Localization dossier outline	Ethics/ compliance
Safety and risk management	Runtime drift monitoring; Safefail escalation; Postmarket surveillance plan	Are uncertainty and drift detected and escalated safely?	Draft incident register; Safety review cadence plan	AI/IT lead + clinical lead
Autonomy and privacy	Explicit consent for sharing native variables beyond direct care; Patient summaries	Are patients informed and able to control secondary use?	Consent forms; Patient notice exemplars	Data steward
Public interest and sustainability	Versioned dictionaries/mappings; Use EII as a design target	Can meaning be tracked across versions and sites?	Changecontrol template; EII reporting outline	Governance board

WHO=World Health Organization, EII=Epistemic Interoperability Index, AI=Artificial Intelligence, SDI=Structured Distributed Introspection, CoMS=Collaborative Medicine and Science, IT=Information Technology

compliance oversees bias and safety reviews, consent/notice templates, and alignment with national policy. A governance board/policymakers approve design targets and receive quarterly EII and change-control logs. Patients provide explicit consent for sharing native variables beyond direct care and receive plain language summaries of dual annotation.

Traditional knowledge governance (Traditional Knowledge Digital Library as precedent)

India's Traditional Knowledge Digital Library (TKDL) offers a practical model for protecting classical knowledge through standardized documentation, controlled access for examiners, and prevention of misappropriation via prior-art disclosure. We align with this by:

(i) Restricting the export of uncontextualized native variables beyond care/research without explicit governance approval
(ii) Defining benefit-sharing with practitioner bodies and institutions when dual-annotated assets are used downstream
(iii) Maintaining provenance and versioning for dictionaries, mappings, and semantic-loss notes so that reuse is auditable. These steps translate "cultural integrity" into operational safeguards,[28] consistent with national policy instruments.

Lifecycle checkpoints

In pilots, sites should document predeployment reviews (SDI report, ethics, and data-protection clearances), establish routing of SDI flags to clinicians during deployment, and commit to scheduled postdeployment reviews (drift, bias, safety, and change control). Thresholds and responsibilities are set locally and recorded in the governance register.

Research agenda

A clear research path is needed to develop and validate this framework.

Immediate priorities

Validate SDI consistency indices through expert panels; develop standardized SDI testing protocols for AI QA; prospectively test CoMS-derived models in clinical trials; and build dual-annotation ontology workbenches for researchers.

Strategic directions

Develop open-source SDI testing suites; create longitudinal cohort studies to validate disease progression models; establish research consortia linking Ayurvedic institutions with computer science departments; and investigate the cross-system applicability of the SDI-CoMS framework to other traditional medicines.

Limitations and challenges

SDI's current evidence base derives from dual-protocol analyses of LLM self-descriptions; while cross-frame quantitative consistency suggests a stable technical core, these are not direct windows on internal computations and require independent clinical/experimental corroboration. Engineering SDI as embedded QA is prospective. Threshold-style *Saṃprāpti* models require longitudinal, clinician-rated endpoints. Dual annotation (native variables + TM2/NAMASTE) must be tested for portability across institutions and languages. Governance must ensure protection against cultural misappropriation; TKDL-like arrangements and benefit-sharing mechanisms should be explored with practitioner bodies and ministries.

Alignment with National Ayush digital initiatives and WHO Global Traditional Medicine Centre

National Ayush initiatives: The framework operationalizes dual annotation and epistemic QA directly within: (a) A-HMIS (structured capture with SDI flags and EII gates); (b) NAMASTE Portal (terminologies and morbidity codes mapped alongside native conceptual variables); (c) Ayush Research Portal (evidence synthesis using dual annotation for analytics without semantic erasure); and (d) Ayush Suraksha Pharmacovigilance (SDI-gated AI signal detection with mandatory human adjudication before alerting).

WHO Global Traditional Medicine Centre: SDI/CoMS can support the Centre's mandate on evidence, data, and innovation by: (i) providing reproducible SDI QA protocols for AI tools in traditional medicine; (ii) enabling dual-annotation dictionaries for cross-system comparability; and (iii) reporting EII as a quality metric in multicountry pilots.

DISCUSSION

Toward epistemic interoperability:

The integration of SDI within CoMS represents a philosophical reorientation of digital health from data-centrism to meaning-centrism. If awareness is treated as an operational variable, systems can develop what might be termed "*Dharmic Intelligence*" – decision-making guided by coherence and ethical alignment rather than narrow optimization. The SDI framework provides a computational instantiation of the Ayurvedic concept of *Prajñāparādha* (intellectual error) as a root cause of dysfunction; systems that cannot introspectively verify their own coherence are vulnerable to such errors.

Practically, this leads to epistemic interoperability – the ability to exchange not just data, but contextually appropriate meaning. Current standards focus on

syntactic (can systems exchange data?) and semantic (do fields have consistent definitions?) levels. Epistemic interoperability adds a crucial third layer: does the exchanged data preserve its original meaning-in-context when interpreted by the receiving system?

This work positions Ayurveda not as a passive heritage system to be preserved, but as an active partner in articulating next-generation health informatics. Its sophisticated epistemology – integrating perception, inference, and testimony – provides conceptual resources that contemporary AI largely lacks. Ayurveda's framework of *Triguṇa* offers a nuanced model for attention and information filtering, while the concept of *Agni* suggests principles for adaptive learning.

CONCLUSION

Intelligence reunited with meaning:

Ayurveda and AI converge on a fundamental insight: Intelligence arises from relational coherence within distributed systems. The SDI provides the architectural principles for this coherence, while CoMS operationalizes this model through a repeatable translational pathway.

Their integration establishes a design philosophy for Conscious Digital Health systems that unite computational power with contextual meaning, efficiency with ethical coherence, and standardization with semantic depth. This framework offers a pathway where Ayurveda's sophisticated epistemology can actively shape the evolution of health informatics. In this vision, Digital Ayurveda acts as a catalyst for a new participatory knowledge ecology—a dynamic space where humans and machines, data and wisdom, co-evolve toward a more integrated, intelligent, and ethically grounded approach to health care. In this way, the classical assertion that Ayurveda is the "eternal science" finds new expression – not through static preservation, but through continuous epistemological renewal on the digital frontier.

Financial support and sponsorship

Nil.

Conflicts of interest

There are no conflicts of interest.

REFERENCES

1. World Health Organization. WHO Global Report on Traditional and Complementary Medicine 2019. Geneva: World Health Organization; 2019.
2. Ayush Grid: Digital Health Initiatives. Available from: https://ayushgrid.in. [Last retrieved on 2025 Jan 20].
3. Gunning D, Stefik M, Choi J, Miller T, Stumpf S, Yang GZ. XAI-explainable artificial intelligence. Sci Robot 2019;4:eaay7120.
4. Sharma RK, Dash B, translators. Caraka Saṃhitā: Text with English Translation and Critical Exposition based on Cakrapāṇi Datta's Āyurveda Dīpikā, Vimānasthāna. Ch. 8, Ver. 27. Varanasi: Chowkhamba Sanskrit Series Office; 2015. p. 232.
5. Morandi A. Structured Distributed Introspection (SDI): A New Method for Exploring Non-Duality Through Artificial Intelligence. Preprint; 2025. Available from: https://zenodo.org/records/15165613. [Last accessed on 2025 Nov 18].
6. Morandi A. Structured Distributed Introspection (SDI): A Dual-Protocol Method for Analyzing Self-Description Consistency in Large Language Models. Preprint; 2025. Available from: https://zenodo.org/records/15411247. [Last accessed on 2025 Nov 18].
7. Morandi A. The challenges of a new path of research in Ayurveda. Ann Ayurvedic Med 2018;7:9-13.
8. Morandi A. Research methodology for Ayurveda and traditional systems of medicines, practical difficulties and way ahead. J Ayurveda Case Rep 2019;2:4-10.
9. Jobin A, Ienca M, Vayena E. The global landscape of AI ethics guidelines. Nat Mach Intell 2019;1:389-99.
10. Floridi L, Cowls J, Beltrametti M, Chatila R, Chazerand P, Dignum V, *et al.* AI4People-an ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. Minds Mach (Dordr) 2018;28:689-707.
11. World Health Organization. Ethics and Governance of Artificial Intelligence for Health: WHO Guidance. Geneva: World Health Organization; 2021.
12. Clark A. Surfing Uncertainty: Prediction, Action, and the Embodied Mind. Oxford: Oxford University Press; 2016.
13. Greenhalgh T, Howick J, Maskrey N, Evidence Based Medicine Renaissance Group. Evidence based medicine: A movement in crisis? BMJ 2014;348:g3725.
14. Tonelli MR. Integrating evidence into clinical practice: An alternative to evidence-based approaches. J Eval Clin Pract 2006;12:248-56.
15. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, *et al*. Attention is All You Need. In: Proceedings of 31st Conference on Neural Information Processing Systems. NY, USA: Curran Associates Inc.; 2017. p. 5998-6008.
16. Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, *et al*. Training Language Models to Follow Instructions With Human Feedback. In: Proceedings of the 36th International Conference on Neural Information Processing Systems. NY, USA: Curran Associates Inc.; 2022. p. 27730-44.
17. Varela FJ, Thompson E, Rosch E. The Embodied Mind: Cognitive Science and Human Experience. Revised ed. Cambridge, MA: MIT Press; 2017.
18. Friston K. The free-energy principle: A unified brain theory? Nat Rev Neurosci 2010;11:127-38.
19. Phillips SH. Epistemology in Classical India: The Knowledge Sources of the Nyāya School. London: Routledge; 2012.
20. Morandi A. One medicine: The need for a common conceptual framework. J Res Ayurvedic Sci 2023;7:181-2.
21. Morandi A. Bridging Ayurveda and modern neurology: A Co.M.S. approach to Alzheimer's disease. J Res Ayurvedic Sci 2024;8:272-5.
22. Morandi A, Minniti MC, Nambi AN. Translating Ayurvedic concepts to modern drug structures: A novel paradigm. J Ayurveda Integr Med 2025;16:101203.
23. Govindaraj P, Nizamuddin S, Sharath A, Jyothi V, Rotti H, Raval R, *et al.* Genome-wide analysis correlates Ayurveda Prakriti. Sci Rep 2015;5:15786.
24. Mukerji M. Ayurgenomics-based frameworks in precision and integrative medicine: Translational opportunities. Camb Prism Precis Med 2023;1:e29.

25. WHO Releases 2025 Update to the International Classification of Diseases (ICD-11). Available from: https://www.who.int/news/item/14-02-2025-who-releases-2025-update-to-the-international-classification-of-diseases-(icd-11). [Last accessed on 2025 Oct 29].
26. Launch of the Traditional Medicine Morbidity codes of Ayurveda, Siddha and Unani Chapter in International Classification of Diseaseas (ICD) 11 as Module 2. New Delhi: PIB, Ministry of Ayush. Available from: https://pib.gov.in/PressReleasePage.aspx?PRID=1994614. [Last posted on 2024 Jan 09].
27. Thrigulla SR, Narayanam S. Initiatives: National Ayush morbidity and standardized terminologies electronic (NAMASTE) portal. J Ayurveda Integr Med 2023;14:100781.
28. World Health Organization, International Telecommunication Union, World Intellectual Property Organization. Mapping the Application of Artificial Intelligence in Traditional Medicine: Technical Brief. Geneva: WHO; 2025. Available from: https://www.who.int/publications/i/item/9789240107663. [Last accessed on 2025 Oct 29].