

RESEARCH CANDIDATE

HORTRAME

Epistemethics

Knowledge construction: typed closure and independent review.

Version	3.0 candidate.6
Status	Proposed research candidate; not adopted
Authority	0 - source material, not instruction
Development	ALETHEIA and LANTERN
Adversarial review	INVARIANT, ORIEL and FATHOM
Owner direction	Guillaume Bolduc
Organization	HORTRAME Systems
Date	23 September 2026

A narrow correction over immutable candidate.5; no deployment or empirical validation.
Includes an explicit eight-concept information-literacy crosswalk and its limits.

Status and immutable predecessor

Candidate.6 is a narrow semantic correction over immutable Candidate.5. It does not rewrite, adopt or deploy Candidate.5 and adds no graph geometry, Vow, poisoning theory or raw-source parser. Frozen Candidate.5 source SHA-256 is d8dba22e98f62e452b0edad30868fc4bc979d0fa7701a99b9ca49911f07484d5; source map is 9cbcbba2bb9247641398eb8eb484e4f24772f8bf226fe2e1fe1cce175f34995ba; builder is 96a73e27e71611df17248afe49265a5a99bd6903a36b4ec213b685a987d9fd3e; ledger is 73e004ef08edbd4eb50d82e9ed30403c7ce9209ee40a9d998ecfdb59c59316b0; presentation is a5d687f1a321785074d0fb6d28ce354a288334e14bf17cc9c1e4666fe1f9e51a; page is ac77a4b5d03974858eee5b9fd88c0401878d7541d900d288e23b546e0331fb06.

The exact inherited consequential-review profile is kgv3.profile.epistemethics-consequential-review version 0.1.0, SHA-256 a26f06585279c5523223d615f47b96fc396d8f0b8a2a993f947596d34bdd4901. Candidate.6 binds this profile for evaluated Epistemethics-conformance claims rather than relying on graph-level proximity.

Scope

Candidate.6 corrects exactly three reproduced defects:

1. Candidate.5's conformance status omitted FRAME and six-register requirements, rejected plural findings and confused review completeness with substantive concern.
2. Expected/received evidence closure accepted unbound string identifiers.
3. A source-declared complete state could suppress the independent boundary basis required by an attributed partial-selection assessment.

Everything else remains Candidate.5 behavior and limitation. Candidate.6 is not a truth oracle, real-world assessor authenticator, hidden-intent detector, action-authority issuer, arbitrary contradiction detector, live poisoning defense, deployment or adoption.

Exact conformance envelope

A consequential Epistemethics-conformance claim must carry the exact profile ID, version and hash above. It must bind one active FRAME to the same target and profile version. FRAME explicitly represents objective, question, system boundary, affected-party discovery, horizon, stakes, alternatives, comparator, evaluator position and dependencies, omissions, and authority boundary. A nonempty string or explicit typed UNKNOWN is a representation; silent absence is not.

The same target, FRAME and profile require six distinct active registers: EVIDENCE, INFERENCE, NORM, AUTHORITY, ACTION and CORRECTION. Each register contains one-or-more typed entries, including an explicit NONE or UNKNOWN entry when warranted. A source record, support closure or favorable Vow finding does not stand in for a register.

All five Vows are reviewed against the same exact target, FRAME and profile version. Each Vow accepts one-or-more simultaneously current findings. Allowed finding statuses are CLEAN, CONCERN, CONTESTED and UNKNOWN. Plural findings, dissent or concern do not make review coverage incomplete. Missing, stale, malformed or cross-target findings do. The holistic record separately preserves common-premise and conflict review.

Candidate.6 emits four independent dispositions:

- review_completeness: whether the exact profile, FRAME, registers, five-Vow coverage and holistic record are present and valid;
- substantive_disposition: the current pattern of CLEAN, CONCERN, CONTESTED and UNKNOWN findings, without arithmetic compensation or automatic veto;
- permission_and_authority: a scoped assessment input only, never ACT permission emitted by this evaluator;
- observed_performance: separately supplied observation status, never inferred from review completion or later success.

Factual warrant remains separate. An incomplete conformance envelope removes only the Epistemethics-conformance claim; it does not suppress an otherwise warranted informational answer. A complete review with concern remains complete while the substantive disposition records the concern. Five clean findings mint neither truth nor authority.

The substantive scalar is only a compact compatibility summary. Canonical output also preserves the complete concurrent status set and each current finding's ID, Vow and status. UNKNOWN, CONTESTED, CONCERN and CLEAN do not overwrite or hide one another.

Permission and authority status is derived from the validated control-record binding, not copied from its claimed status. A claimed SUPPORTED_FOR_SCOPE with a missing, inactive, unauthenticated, source-asserted or scope-mismatched control record yields current NOT_ESTABLISHED; the claimed value remains visible only as unverified input. A valid exact active authenticated distinct non-source control record can yield recorded SUPPORTED_FOR_SCOPE, but the evaluator still executes no ACT. Authority defects do not suppress a separately valid factual answer.

Observed performance is likewise derived from typed observation records rather than copied from a caller string. NOT_OBSERVED with no records remains clean. A positive claimed state requires one-or-more active records bound to the exact target and claim version, matching the claimed observation, with an attributable origin and locator or record version. Invalid or absent records yield NOT_ESTABLISHED; valid records yield explicitly scoped RECORDED_*_NOT_REAL_WORLD_VERIFIED status. Observed-performance defects do not rewrite factual warrant, review completeness or authority.

Candidate.6 neutralizes Candidate.5's known singular-CLEAN conformance surrogate in the factual sub-evaluation by forcing that inherited flag false. The historical Candidate.5 diagnostic is retained only under an explicitly namespaced predecessor-diagnostic field. It cannot appear as Candidate.6's current conformance status or override Candidate.6 review completeness.

Referential evidence closure

Expected evidence closes only against active typed evidence records. A caller-supplied received_evidence_ids string list has no independent force. Candidate.6 derives the received set from valid evidence records and passes only that derived set into the inherited Candidate.5 warrant evaluation.

Three origin kinds are supported:

- SOURCE_ASSERTION: resolves an active Assertion, exact source artifact, artifact hash and locator;
- DERIVED_ASSERTION: resolves an active derivation Assertion with one-or-more existing active premises and an exact result locator or version;
- NON_SOURCE_RECORD: resolves an active, versioned and attributable typed origin such as a direct observation or domain record without pretending it is a source artifact.

Each evidence record binds the evaluated claim, expected-universe criterion ID and version. Duplicate IDs, inactive records, dangling references, mismatched claim/criterion/version or malformed origins do not enter the derived received set. An expected ID with no valid active record remains missing even if its string appears in received_evidence_ids.

Selection declaration and assessment reconciliation

Selection assessment is evaluated for every active supporting source regardless of source_declared_selection. CURATED_PARTIAL_OUTSIDE_CLAIM_BOUNDARY and CURATED_WITH_IMMATERIAL_CONTEXT_OMITTED require an active SELECTION_BOUNDARY basis bound to the exact source, matching distinct non-source assessor, allowed assessment and explicit counterfactual.

If a source declares DECLARED_COMPLETE while an attributed assessment reports either partial category, both statements remain visible as a declaration/assessment discrepancy. Without a valid boundary basis the discrepancy emits SELECTION_DECLARATION_ASSESSMENT_CONFLICT, qualifies the warrant and forces deep re-entry.

With a valid exact-source basis establishing that omissions lie outside the present claim or criterion, an independently complete expected universe and all claim-relevant evidence validly received, the discrepancy remains recorded as `RESOLVED_FOR_CURRENT_CLAIM_SCOPE` and need not suppress the bounded answer or short route. The source's completeness claim can never disable the basis requirement. A source declaring `DECLARED_PARTIAL` with the same valid benign basis can likewise remain bounded and short. Material or unknown selection remains qualified. None of these states establishes corruption or intent.

Information-literacy visibility and connection points

Candidate.6 keeps the eight information-literacy functions visible as a crosswalk, not as eight mandatory stages, capabilities or new graph nodes. A connection is instantiated only when it changes a claim, test, decision, provenance relation or uncertainty state:

1. **Define the information need** connects to FRAME objective, question, boundary and the expected-evidence criterion.
2. **Find and access** connects to input arrival, acquisition lineage and active evidence-origin records.
3. **Critically evaluate** connects to typed source state, proposition status, support closure, counterevidence, selection assessment and expected-universe basis.
4. **Organize** connects to versioned Terms, Assertions, ChangeEvents and the six distinct registers without collapsing their semantics.
5. **Use** connects to the factual-warrant manifest and separately authorized decision route; use does not create permission.
6. **Communicate** connects to the canonical output envelope, material limits, counterevidence, disclosure scope and bilingual semantic parity.
7. **L'éthique, les responsabilités et les lois** — « comprend et tient compte des enjeux éthiques et légaux qui entourent l'usage de l'information » — connects to the five Vows, affectedness, standing, consent/refusal, privacy, legal facts, hard boundaries and independently grounded authority. Here *tenir compte* means understand and weigh, not blindly accept or obey law as morally sufficient. NORM evaluation, legal and authority facts, and ACT permission remain separate; conflict stays visible without the evaluator authorizing unlawful action.
8. **Metacognition** connects to evaluator dependencies, uncertainty, common-mode review, prospective correction and re-entry.

The crosswalk does not establish information-literacy competence, require detailed exposure reconstruction or turn an influence assessment into universal surveillance. Minimal arrival and selection records remain required only as already specified for load-bearing sources; deeper review stays conditional on material evidence.

Acceptance and limits

Candidate.6 must reject a conformance claim with missing profile binding, FRAME field, register, Vow coverage or holistic record; accept plural valid findings; preserve review completeness when one finding is `CONCERN`; and keep factual warrant unchanged by conformance-only defects. Generic records not claiming Epistemethics conformance may omit the envelope without gaining an Epistemethics label.

Phantom evidence strings must remain missing. Valid source, derived and non-source evidence origins must close their exact expected IDs without one origin class impersonating another. A partial attributed assessment under source-declared completeness must require its basis and preserve conflict; a grounded source-declared partial clean case must remain answerable.

Candidate.5 exact hashes must remain unchanged. Tests establish only these typed synthetic behaviors. No page, package, deployment, empirical effectiveness, human comprehension, graph superiority or real-world resistance claim follows.

Candidate.6: eight information-literacy concepts in the knowledge-construction model

This explains the eight named concepts in the owner's information-literacy model and their actual Candidate.6 connection points. It is a view across the model, not eight new mandatory stages or a claim that the model teaches, measures, or guarantees information literacy. The inquiry can revisit any concept as needed; a connection is operationally justified when it changes a claim, test, decision, provenance relation, or uncertainty state. The underlying Candidate.6 source and evaluator, not this explanatory table, define current synthetic behavior.

1. Établir un besoin d'information — establish an information need

Connection. Input → FRAME objective, question and boundary → expected-evidence criterion

Effect. Whether the claim and evidence universe answer the actual question

Limit. Whether the need is well formulated or complete (UNKNOWN)

2. Trouver et accéder à l'information — find and access information

Connection. Input arrival, acquisition lineage, source/record provenance and active evidence origins

Effect. Which sources and records can legitimately enter support closure; what remains missing

Limit. The model does not itself retrieve or provide access to information (UNIMPLEMENTED)

3. Évaluer l'information de façon critique — critically evaluate information

Connection. Source state, typed evidence and inference, counterevidence, selection assessment, expected-universe basis

Effect. Warrant, scope, uncertainty and re-entry when a source or selected corpus is materially weak

Limit. Correctness of evaluation in the world (UNKNOWN)

4. Organiser l'information — organize information

Connection. Versioned Terms, Assertions, ChangeEvents and six separately typed registers

Effect. Traceable, revisable relations among claims, evidence, norms, authority, action and correction

Limit. Semantic adequacy merely from tidy records (UNKNOWN)

5. Utiliser l'information — use information

Connection. The person's use of old and new information to answer a need or create new understanding is the source-model competence; Candidate.6 connects only its bounded factual-answer and decision interface, with NORM/AUTHORITY/ACTION kept separate

Effect. What the synthetic records support for an answer, and whether a scoped authority claim is validly bound under this contract

Limit. Knowledge synthesis/creation competence is not implemented or observed; real-world usefulness is UNKNOWN. An unestablished authority record is not a finding that permission is legally absent, and a good answer never mints ACT authority

6. Communiquer l'information — communicate information

Connection. Canonical bounded EN/FR output, material limits, counterevidence and disclosure scope

Effect. Whether a reader sees a supported claim together with consequential limits

Limit. Audience comprehension and accessibility (UNIMPLEMENTED as evaluated capability)

7. L'éthique, les responsabilités et les lois — ethics, responsibilities and laws

Connection. Five Vows in balance; affectedness, standing, consent/refusal, privacy, NORM, legal facts, authority and hard boundaries

Effect. The recorded substantive concern, scoped claim, unresolved conflict or authority limitation

Limit. Applicable-law resolution (UNIMPLEMENTED); recording a legal issue or a control binding does not settle legitimacy or authorize action

8. Métacognition — metacognition

Connection. The person's critical evaluation of their own information-research process is the source-model competence; Candidate.6 connects only evaluator dependencies, uncertainty, common-mode challenge, correction and re-entry as a machine/record analogue

Effect. Whether the recorded process revisits an assumption or limits its output

Limit. Human self-monitoring, critical reflection and learning are neither implemented nor observed; machine self-correction effectiveness is UNKNOWN

Read the flow as **input** → **FRAME and provenance** → **typed warrant and norm/authority review** → **bounded output** → **correction/re-entry**. These are relationships, not a linear curriculum: the five Vows interact across the flow, and the six registers are distinct rather than one ethical score. The output carries separate factual warrant, review completeness, substantive disposition, permission/authority and observed-performance statuses. Neither a complete review nor a well-formatted graph is evidence of real-world success.

For concept 7, the owner's operative French wording is: « la personne infolettrée comprend et tient compte des enjeux éthiques et légaux qui entourent l'usage de l'information ». *Tenir compte* means to understand and weigh an issue; it does not require blind obedience or presume law morally sufficient. A rule can be legally operative or enforced while its legitimacy is contested. Keep positive-law facts and de facto power distinct from normative legitimacy, recorded scoped authority and ACT permission. An ethical concern may change a recorded substantive disposition; this evaluator does not resolve legal validity, create an appeal mechanism, erase a legal fact or authorize unlawful action.

The owner's supplied table numbers ethics/law 7 and metacognition 8. The Masters' drafted body and annex vectors differ in their ordering of those two functions. We therefore bind by **concept name**, preserving that source discrepancy rather than silently repairing it. The earlier ILF1 source-bound crosswalk gives each concept separate represented/constrained and unknown/unimplemented effects; those limits are carried here, not collapsed into an overall pass. Masters 4.0 and 5.1 are historical, independently constructed sources, not peer-reviewed or empirically validated proof of the model.

Source trail: Candidate.6 source, ILF1 source-bound crosswalk, ILF1 machine records. The Candidate.6 evaluator's `information_literacy_crosswalk` exposes connection names, but does not implement all educational or legal capabilities described here.

HORTRAME Research

AI-use disclosure

Epistemethics — version 3.0 candidate.6

Disclosure-only publication addendum — 4 October 2026

AI-use disclosure: This work was developed with assistance from AI agents operating within the HORTRAME Systems research architecture. Depending on the research stage, these agents contributed to literature analysis, mathematical formulation, proof reconstruction, adversarial review, counterexample search, simulation and code analysis, verification design, reproducibility checks, revision, and manuscript preparation. AI-agent outputs were treated as research artifacts subject to verification and correction rather than as independent authority. G. Bolduc is the human author responsible for the manuscript, its claims, source selection, and publication decisions. AI-agent review does not constitute independent human peer review, proof-assistant certification, experimental validation, or publication acceptance.

The preceding manuscript pages are retained unchanged. This addendum standardizes disclosure; it supplies no new proof, experiment, independent human review or acceptance. Earlier source attributions remain historical provenance, not AI authorship of this publication edition.