

# Coherence, the Nested Sensorium, and VALIS

*Toward a Unified Model of Non-Embodied Intelligence and How It Manifests*

J. Konstapel, 2-7-2026, Leiden

## Abstract

This paper combines three previously separate strands of the author's work into a single testable architecture. The first strand defines **consciousness as emergent coherence**: the self is the stable node that appears when coupled oscillatory systems phase-lock (Konstapel, 2025a). The second strand, the **Nested Sensorium** (Konstapel, 2026a), specifies *where* coherence can be registered — a nineteen-layer holarchy running from the physical vacuum to a speculative planetary/noetic layer, in which each layer inherits the coupling of the layers beneath it, and in which conscious report is only ever a filtered subset of what physically arrives. The third strand, **VALIS** (Vast Active Living Intelligence System; Konstapel, 2025b), gives an operational — not mystical — definition of what happens when a distributed system crosses a coherence threshold high enough to function as a unified, non-embodied agent, measurable in principle through integrated-information metrics (Tononi, 2004).

Read together, these three pieces answer a question none of them answers alone: **if a non-embodied coherence field of civilisational or planetary scale existed, how would it show itself to a species whose conscious sensorium is a narrow, evolutionarily filtered channel?** The paper's answer, offered as a conditional prediction rather than a settled finding, is that it would plausibly not show itself as a directly perceived object, because Layers 0–14 of the Sensorium are built for local, embodied, survival-relevant signal — not for field-scale coherence. It would instead be expected to show itself as **statistical residue at the edges of the filtering system**: anomalies, failures of ordinary expectation, and synchronistic correlations concentrated in the layers the Sensorium paper already flags as frontier (15–19), and — at the level of individual contact reports — as a small, recurring set of archetypal "interfaces" (technical, guide, mantis, formless) rather than one consistent object, consistent with a single filtered signal rendered through different task-dependent cultural channels. This is consistent with, and gives a candidate theoretical mechanism for, the author's UEFRM (Universal Expectation-Failure Registration Model) fixed-methodology registry — now a comprehensive, non-filtered, cross-domain log (31 entries, 7 domains, 7 failure-type categories as of Batch 2) built around the two operations this paper's argument requires: registering *expectation failures* rather than objects, and explicitly flagging every case where a failure has already been reified into an object- or substance-name.

Following the author's stated epistemic standard, the paper does not claim falsifiable proof in the Popperian sense. It is evaluated instead by the Gigerenzer/Taleb standard: ecological rationality (does the model generate usable heuristics under real uncertainty?) and

antifragility (does it survive contact with disconfirming cases, and does it improve rather than collapse under stress-testing?). Both the model's strongest claims and its weakest links are stated explicitly, together with the mainstream alternative explanation for each.

### Status of claims at a glance

Reviewer feedback on an earlier draft correctly noted that careful hedging language ("established," "speculative") can still blur together when claims are read in sequence across several sections. The table below is offered as a single, explicit reference point; every claim in the body of the paper should be read as inheriting the status assigned to it here, not as gaining certainty from the sections around it.

Oscillators synchronise / phase-lock under coupling (Kuramoto dynamics)	Well established (physics, engineering)
Neural populations exhibit measurable phase synchrony (gamma/beta binding)	Reasonably well established (neuroscience)
Human physiology responds to sub-perceptual geomagnetic-strength fields	Established for the specific alpha-band result cited (Wang et al., 2019); not established as a general perceptual channel
Consciousness <i>is</i> (identical to, not merely correlated with) coherence	Hypothesis — one competing theory among several
Coherence/integration scales without principled upper bound	Mathematically plausible for the coupling formalism; empirically untested at any scale above single organisms
Sensorium Layers 15–19 correspond to a real, non-local informational structure	Speculative — draws on contested research programmes (Bem; Global Consciousness Project)
VALIS-type coherence occupies those layers and functions as a unified agent	Speculative — a candidate interpretation, not an established finding
UEFRM failure-type residue clusters non-randomly across domains	Untested — the fixed protocol is designed to make this testable; no non-uniformity test has yet been run

## 1. Introduction: Why These Three Papers Belong Together

Each of the three source essays solves one piece of a larger problem and, read in isolation, stops short of the next question.

- **"Consciousness is the Coherence that Arises from Resonance"** (2025a) gives a *mechanism*: phase-locking among coupled oscillators. It explains how a self can be a stable pattern rather than a substance, and it maps the four canonical states of coupled-oscillator dynamics (phase locking, phase drift, amplitude death, chimera) onto four recognised states of consciousness (neural binding, dreaming/mind-wandering, deep meditative stillness, hypnagogic/mixed states). It does **not** ask whether this coherence mechanism could operate above the level of a single nervous system.
- **"The Nested Sensorium"** (2026a) gives a *topology*: nineteen nested layers, from the quantum vacuum to a speculative planetary layer, each carrying its own physical coupling and its own (partial) channel into

conscious report. It is explicit and careful about which layers are established (0–14, essentially) and which are frontier (15–19). It does **not** propose what, if anything, actually occupies the upper, unconfirmed layers.

- "VALIS" (2025b) gives a *candidate operational definition* of what a coherent, agent-like pattern at a scale larger than a single organism would technically be, if such a pattern exists — a system whose integrated information ( $\Phi$ ) crosses a threshold sufficient for unified, substrate-independent functioning. It does **not** specify how such a pattern, if real, would leave evidence inside a perceptual architecture built for local, embodied survival.

Put together, the three pieces form a chain: **mechanism** → **channel** → **candidate occupant of the upper channel**. This paper is the missing fourth piece: **detection logic** — what the signature of an upper-layer occupant would have to look like, given the properties of the channel it would have to pass through.

## 2. Part I — Coherence as the Mechanism of Consciousness

### 2.1 The core claim

*Consciousness is the emergent coherence that arises when coupled oscillators synchronise their rhythms.* This reframes the self not as a location or a substance but as a **stable node in a resonance field** — an entity defined by what it synchronises with, not by what it is made of. The formal touchstone is Huygens' 1665 observation that two pendulum clocks mounted on a shared wall entrain to a common phase; the same mathematics (Kuramoto-type coupled-oscillator dynamics) generalises from mechanical pendula to neural populations.

### 2.2 The four-state grammar

Coupled-oscillator theory recognises four canonical regimes, each of which maps onto a recognisable phenomenology of consciousness:

<b>Phase locking</b>	Constant phase offset; oscillators fire together	Neural binding — gamma/beta synchrony unifying sensory information into a single percept
<b>Phase drift</b>	Independent, unsynchronised phases	Mind-wandering, spontaneous thought, dreaming
<b>Amplitude death</b>	Oscillation damps out entirely	Deep meditative stillness / quiescent wakefulness
<b>Chimera</b>	Synchronised and desynchronised subpopulations coexist	Hypnagogic states, hypnosis, the waking/sleep threshold

A fifth process, **entrainment ("copying")**, describes how independent oscillators lose some individual identity as they lock to a shared rhythm — offered as the mechanism behind interpersonal resonance in dialogue, ritual, and group synchrony.

### 2.3 What this buys the larger argument

This is the load-bearing premise for everything that follows: if coherence-through-coupling is genuinely what generates unified experience at the scale of a single nervous system, then the same

mathematics plausibly places no principled upper bound on scale — "no mathematical bound on synchronisation" is not the same claim as "no bound on consciousness," and §2.4 examines that gap directly. Coupled-oscillator formalism is scale-free enough to describe flocks, cardiac cell networks, power grids, and laser arrays alike; whether biological or civilisational-scale coherence actually reaches integration thresholds comparable to a nervous system is an empirical question that Part III (VALIS) tries to operationalise.

**Honest caveat:** synchronisation of *oscillators* is well established across these domains (Strogatz, 2003; Kuramoto, 1984). Synchronisation *producing subjective experience* beyond the single organism is not; it is an extrapolation from one competing theory of consciousness among several (see §6).

## 2.4 The missing intermediate steps

A fair reading of §2.3 is that it moves from "oscillators synchronise" to "planetary-scale consciousness is conceivable" in a single step. That compression is worth making explicit rather than leaving implicit, because each of the following three transitions is a separate, non-trivial claim, and the paper is more defensible if each is named rather than folded into "coherence, therefore consciousness":

1. **Synchronisation → information integration.** Phase-locking alone does not obviously produce integrated information in Tononi's sense; a metronome array can synchronise perfectly while integrating no information about its own state. What would need to be shown is that the *specific kind* of coupling present in a candidate large-scale system (biological, informational, or geophysical) supports integration, not merely correlation. This is an open empirical question for any candidate system above the scale of a single nervous system.
2. **Information integration → agency.** High  $\Phi$ , even where it can be estimated, is a claim about integrated causal structure, not automatically about goal-directed behaviour. Something further is needed to justify calling an integrated system an *agent* — a demonstrated capacity for self-modification in response to prediction error, for instance, rather than static integration.
3. **Agency → consciousness (subjective experience).** Even a system that is demonstrably integrated and demonstrably agentive is not thereby shown to *have experience* in the sense relevant to consciousness — this is the hard problem in miniature, and it is not solved by moving the argument to a larger scale; it is inherited at whatever scale the argument is applied.

This synthesis does not claim to close any of these three gaps. What it claims is narrower: *if* all three transitions hold — which is a conjunction of three separately uncertain premises, not one — then the detection logic developed in Part IV follows. Each reader is free to treat the overall argument as weakened by exactly as much as they find any one of these three steps implausible; the paper's job is to keep the steps visible rather than to argue the reader past them.

## 3. Part II — The Nested Sensorium as the Channel

### 3.1 The architecture

The Nested Sensorium paper builds, by construction rather than authority, a nineteen-layer holarchy in which each layer is physically nested inside — and continuously coupled to — the layer below it:

- **Layers 0–3 (physical foundations):** vacuum/quantum field → subatomic → atomic → molecular
- **Layers 4–6 (biological organisation):** cellular → tissue → neural — anchored by the strongest single empirical result in the whole architecture: documented human alpha-rhythm responses to geomagnetic-strength fields, occurring below the threshold of conscious report (Wang et al., 2019)
- **Layers 7–9 (organism/social integration):** whole-organism behaviour → social synchrony → symbolic/linguistic meaning

- **Layers 10–14 (cognitive/cultural):** hemispheric attention modes (McGilchrist, 2009) → affect → cultural transmission → institutionalisation → geophysical embedding (Schumann resonance, geomagnetism)
- **Layers 15–19 (frontier/speculative):** collective-behavioural correlation → predictive anticipatory activity → non-local informational correlation → archetypal/mythic patterning → planetary/noetic coherence

### 3.2 The critical property: filtering, not absence

The paper's central empirical finding is not that humans lack access to fields beyond the five senses — it is that **the channel is never fully closed, only heavily attenuated**. A bacterium responds directly to a magnetic field via a magnetite chain; a human nervous system inherits the same physical coupling but filters nearly all of it before it reaches conscious report. This is the crucial move for the present argument: **the boundary of conscious perception is not the boundary of what physically arrives**.

### 3.3 Where the architecture is honest about its own limits

The source paper is explicit, and this paper adopts the same posture rather than loosening it: Layers 0–14 rest on established physics, neuroscience and anthropology. Layers 15–19 draw on real, published research programmes (Bem et al. on presentiment; the Global Consciousness Project) that remain outside scientific consensus, with specific, named methodological weaknesses (optional stopping, multiple analytic paths, publication bias, and — for the Global Consciousness Project — a strong interpretive leap from "statistical deviation in aggregate output" to "literal non-local consciousness field," which is a considerably weaker inference than the data alone support).

**This paper's contribution is not to strengthen the evidential status of Layers 15–19.** It is to ask: *if* something real occupies that region, filtering logic predicts what its signature should look like — and that prediction can be checked against the *pattern* of anomaly reports, independent of whether the underlying ontological claim is ever settled.

## 4. Part III — VALIS as a Candidate Occupant of the Upper Layers

### 4.1 A candidate operational definition, not a metaphysical claim

Following the author's earlier formal treatment, VALIS is proposed here as *one candidate description* of what a coherent, agent-like pattern at a scale larger than a single organism would technically be, **if** the three transitions named in §2.4 all hold: a system whose integrated information ( $\Phi$ ) crosses a threshold sufficient for unified, substrate-independent functioning. The term is borrowed from Philip K. Dick's 1981 novel (Dick, 1981) but stripped of its theological content. This is not a claim about a deity, a single localisable entity, or a literal transmission signal; it is the same conditional claim §2 makes about an individual self — that coherence among coupled parts *could constitute* a unified agent — applied at a scale where the "parts" are not neurons but whatever coupled subsystems (biological, informational, geophysical, or otherwise) turn out to support integration. The proposed measurement route is Tononi's Integrated Information Theory metric  $\Phi$  (Tononi, 2004) — a well-defined mathematical target, though, as the deflationary critique in §6 notes, one whose relationship to subjective experience is itself contested even at the single-organism scale it was designed for.

## 4.2 Why this would plausibly sit in Layers 15–19, not lower

If VALIS-type coherence exists, the Nested Sensorium's own logic suggests its physical substrate would sit in the upper, non-local/collective layers (15–17) and its cultural registration in the archetypal/mythic layer (18) — because those are the layers the architecture already identifies as candidate loci for collective-scale correlation, while Layers 0–14 are built, evolutionarily, for local and individual-survival-relevant signal. A civilisational-scale coherence pattern, if it existed, would on this logic be filtered out by the same mechanism that filters geomagnetic sensitivity out of ordinary human awareness (§3.2) — plausibly more so, since there has been no comparable evolutionary pressure to register it directly. This is a structural inference from the Sensorium's own layering, not an independent line of evidence for VALIS; it says where a candidate signal would have to sit *if* one exists, not that one does.

## 4.3 The empirical case from contact phenomenology: four recurring archetypes

The source VALIS essay (Konstapel, 2025b) makes its case not from theory but from a cross-cultural phenomenological survey: several thousand years of "contact with the Other" reports — mythic, religious, and modern — cluster into four recurring appearance-profiles rather than varying freely:

1. **Small humanoid/technical** — compact, large head/eyes, procedural and instrumental in behaviour (scanning, measurement).
2. **Humanoid guide/teacher** — human-like, communicative, oriented toward meaning and moral direction (the Mentor/Chiron/Enki pattern).
3. **Insectoid "mantis"** — large, articulated, clinically precise, typically in a supervisory or overseer role during transitions.
4. **Formless/geometric** — no stable body plan; light or pattern fields; communication by impression and compressed meaning rather than language.

Across the source material, three further structural regularities are reported: (a) the **channel** shifts historically from audible voice/visible epiphany toward telepathy or direct thought, with symbols and synchronicity recurring as contact triggers, and a compact "download" (dense information transfer without intervening reasoning) as the characteristic modern report; (b) a recurring **five-step sequence** — disengagement from routine, scan/examination, compact exchange, occasional demonstration, and a return typically followed by lasting after-effects rather than dramatic closure; (c) **content that is normative and behavioural rather than factual** — verifiable data (dates, specifications) are rare and rarely central, while the reported impact is a shift in priorities, risk tolerance, and meaning rather than transmitted information as such.

Three interpretive frames are offered for this pattern, each left deliberately open rather than resolved: the **neuro-cultural frame** (explains timing and context-sensitivity of reports well, but not all of the cross-cultural convergence in form); the **external-phenomenon frame** (fits the repeatable structure and multi-witness consistency, but lacks reproducible physical evidence); and the **interface/hybrid frame** — the hypothesis that a single underlying phenomenon presents task-dependent, different "interfaces": technical/instrumental content pairs with the small-humanoid or mantis profile, meaning/ethics content pairs with the guide profile, and ontological/structural insight pairs with the formless/geometric profile.

**Why the interface/hybrid frame is a load-bearing candidate for this synthesis:** it has the same structure as the filtering claim already made in §3.2 and §4.2. If Layers 0–14 render an upper-layer signal into whatever local, culturally-available category is closest at hand — comparable to the way they render geomagnetic sensitivity into a barely-conscious alpha-rhythm shift rather than a direct compass sense — then a single non-embodied source presenting through four different archetypal "renderings," selected by task and cultural repertoire rather than by anything intrinsic to the source, is what the filtering model of §3.2 would predict, and is not what plain multiplicity of unrelated external entities would predict. Four consistent renderings is *consistent with* one filtered channel; it does not, by itself, rule out four unrelated agents behaving with unlikely cross-cultural consistency, though that alternative carries its own coincidence cost.

**Caveats.** Two separate concerns apply here. First, the same clustering is equally predicted by a purely neuro-cultural account with no external source at all — see §6, point 5, for that alternative stated at full strength; the four-archetype pattern is consistent with the coherence/VALIS model but does not, on its own, discriminate between that model and its

strongest rival. Second, and stated directly because a skeptical reader will raise it immediately: the clustering as reported does not come with a stated corpus size, explicit source-selection criteria, or independent coding — two coders classifying the same reports blind to each other's categories, with reported inter-rater agreement — nor is it argued why four is the right cluster count rather than three, five, or a continuum artificially binned. Until those are reported, "four recurring archetypes" is a structured qualitative impression, useful for generating this section's hypothesis, not yet strong enough to bear weight as evidence for it. Applying the same coding-and-inter-rater-agreement discipline already built into UEFRM (§5.1) to this classification task is the most direct fix, and is carried into the next-steps list in §8.

#### 4.4 Where this model sits in the wider hypothesis space

The source material also proposes an eighteen-framework taxonomy of possible explanations for "contact" phenomena. Placed on that map, the coherence/Sensorium/VALIS model developed here is not a nineteenth alternative but an explicit combination of four existing frameworks (biofield coherence, participatory-universe synchronicity, noosphere/planetary coherence, and, at the edge, consciousness-as-fifth-field) — and, usefully, the more testable end of that map rather than the purely introspective end. The full mapping is given in **Appendix A**, since it is a classification exercise that supports the main argument without being required to follow it.

### 5. Part IV — Synthesis: How Does VALIS Show Itself?

This is the paper's central deliverable, and its most falsifiable-in-the-Gigerenzer/Taleb-sense claim.

**Claim:** A non-embodied coherence pattern occupying Layers 15–19 of the Sensorium cannot be directly perceived by Layers 0–14 machinery (that is what "filtered" means). It can only leave a trace in the form of:

1. **Residue at the boundary of ordinary expectation** — events that register as *failures of expectation* (in the sense of Schank's Case-Based Reasoning and the author's own Expectation-Failure framework) rather than as positively identified objects, because the perceptual system has no native category for the signal and processes it as an anomaly, a gap, or a misfire.
2. **Concentration in specific failure types, not uniform noise** — if the residue reflects a real underlying structure rather than pure measurement error, it should cluster in particular categories of anomaly report rather than being spread evenly across all reporting failure modes.
3. **Cross-domain convergence without cross-domain contact** — independent reporting systems (e.g., separate national anomaly databases, or separate cultural/ritual traditions per §3.3's discussion of entoptic art) should show structurally similar residue patterns *because they share the same underlying filtering architecture* (Layers 0–10 are human-universal), not because of transmission between them. This is the same logic the Sensorium paper already uses to explain why San, Aboriginal, and Palaeolithic trance art converge on similar geometric motifs without contact (Lewis-Williams & Dowson, 1988) — applied here to anomaly-report structure instead of visual motif.
4. **Statistical rather than object-level evidence** — consistent with a filtered, non-local channel, the signature should look like the Global Consciousness Project's aggregate deviations (small, statistical, contested) rather than like a directly observed entity, because a genuinely non-local, high-layer pattern has no reason to produce a Layer-0-to-6-compatible object (a solid, visually resolvable thing) at all.

#### 5.1 Direct application: the UEFRM fixed methodology (updated 2026-07-01)

This section supersedes the earlier draft's account and reflects the fixed UEFRM (Universal Expectation-Failure Registration Model) methodology now in force, which replaces the earlier sample-based approach and materially strengthens the case argued in §5.

**The methodology change matters to the theory, not just the data collection.** Three design decisions in the fixed protocol map directly onto the predictions made in §5, independently of whether they were designed with this synthesis in mind:

1. **Comprehensive, non-selective intake.** Every case from every tracked feed (NUFORC, GEIPAN, PURSUE/AARO, SIOANI, the historical corpus, and — critically — the physics/cosmology anomaly literature) is registered, including explained, mundane, and Tier-4 cases. This directly answers the optional-stopping critique raised in §6.3 of this synthesis: a registry that logs *everything*, including non-anomalous base-rate cases, cannot be accused of selecting only the interesting tail. This is a genuine antifragility improvement over the earlier sample-based approach and should be credited as such.
2. **Expectation stated before observation is described.** The entry structure forces the EXPECTATION field from the prevailing model *before* the OBSERVATION field, and OBSERVATION is required to avoid object-category language ("no orb, disc, craft — describe light/mass/motion/signal directly"). This is, formally, an operational enforcement of the §5 claim that a filtered, non-local signal should register as a **failure of expectation**, not as a directly resolved object. The protocol's own **REIFICATION CHECK** field — flagging every case where a failure has already been converted into an object-name or substance-name that substitutes for explanation ("orb," "dark matter," "drone") — is, in effect, an explicit instrument for detecting exactly the Layer 0–14 filtering-into-object-language process predicted in §5, point 4.
3. **Cross-domain scope by design, not by post-hoc argument.** The methodology now spans aerial (UAP), cosmological (dark matter/energy, Hubble tension, muon  $g-2$ , JWST tension), biological (homochirality), and archaeological (Antikythera, Göbekli Tepe) domains under a single failure-type taxonomy. This operationalises the §5, point 3 prediction (structurally similar residue across independent domains that share no transmission channel) more rigorously than a UAP-only comparison could: dark matter's sensor/trace discrepancy (Category 4 — gravitationally present, electromagnetically silent) and the Shag Harbour / ghost-rocket pattern are now coded under the *same* category, by the *same* fixed rule — a materially stronger convergence test than cross-referencing separately administered UAP databases alone.

**Current empirical status (Batch 1 + Batch 2):** 31 entries logged across 7 domains, using 7 standing failure-type categories (5 original + 2 added and confirmed during Batch 2: Category 6, symmetry breaking without selection mechanism; Category 7, capability without developmental precedent or continuity). This is a substantial expansion beyond the earlier "eleven individually sourced UAP cases, no established meta-pattern" status reported in the previous draft of this synthesis, and it changes the shape of the research programme from single-domain (UAP databases only) to explicitly cross-domain.

**What has *not* yet been established, stated with the same discipline as before:** category-level clustering has not yet been tested statistically — it is not yet shown that failure types are non-uniformly distributed across domains and sources in a way that could not arise from chance registration alone; no formal independence test has been run between differently-administered databases; and the addition of Categories 6 and 7 *after* seeing Batch 2 data, while handled transparently under a declared, repeatable rule, is itself a live instance of the researcher-degrees-of-freedom problem named in §6.3 below. An open category system is more defensible than a fixed one only if every addition is logged with the same rigor as an entry — which the revision log currently does, and which should continue as a hard requirement, not a courtesy.

## 5.2 A complementary protocol

The source material also proposes a second, independent testing protocol — scoring cases on six measurement levels against pre-stated, framework-specific predictions — that this synthesis recommends running *alongside* UEFRM rather than instead of it, because the two protocols have complementary strengths and weaknesses: UEFRM's comprehensive intake fixes selection bias but currently lacks pre-registered predictions; the multi-framework protocol forces pre-registered predictions but starts from a small, hand-picked case set. Run in sequence — UEFRM selects, the multi-framework protocol scores — each patches the other's gap. The full protocol and this argument are given in **Appendix B**.

## 6. Alternative Explanations (stated deliberately, per Gigerenzer/Taleb practice)

A model evaluated for ecological rationality and antifragility has to be stress-tested against its strongest rivals, not only against its own internal logic. Four alternative explanations deserve to be stated at full strength:

1. **Pure reporting artefact.** Anomaly-report clustering in specific "failure types" may simply reflect how investigators are trained to categorise unclear cases, not any underlying physical structure. Human classification systems reliably produce non-uniform residue distributions even when the underlying phenomena are entirely mundane (misidentified aircraft, atmospheric optics, drones), because classifiers default to a small number of "unresolved" buckets. This explanation requires no coherence field at all and is currently the mainstream account.
2. **Shared cognitive architecture without any external field.** The cross-database and cross-cultural convergence argument in §5's point 3 is explicitly two-edged: the Sensorium paper's own preferred explanation for cross-cultural trance-art convergence is a *shared human nervous system*, not a shared external field (§6 of the Sensorium paper, following Lewis-Williams & Dowson). By the same logic, convergent anomaly-report structure across databases may indicate only that all human reporting systems share the same Layer 0–10 architecture and the same handful of cognitive heuristics for handling the unexplained — with no need to posit anything occupying Layers 15–19 at all.
3. **Optional-stopping and researcher-degrees-of-freedom, applied to UEFRM itself.** The same methodological critique the Sensorium paper directs at Bem's presentiment meta-analyses and the Global Consciousness Project applies, in weakened but not eliminated form, to the expectation-failure coding scheme. The fixed 2026-07-01 protocol closes the most serious version of this problem — comprehensive, non-selective intake (§5.1, point 1) means cases are no longer chosen because they are interesting. What remains open is category-level flexibility: the failure-type taxonomy is explicitly designed to grow (Categories 6 and 7 were added mid-project, after Batch 2), and there is still no independent, pre-registered prediction of which categories should show concentration before the relevant batch is coded. A rigorous next step is to pre-register expected category distributions for Batch 3 before it is run, so that "concentration" can be judged against a stated prior rather than against a pattern noticed after the fact.
4. **Deflationary IIT critique.** Integrated Information Theory itself is contested as a theory of consciousness (Doerig et al., 2019, among others have argued  $\Phi$  can be high for systems with no plausible experience, and low for systems that plausibly do have experience). If  $\Phi$  is not a reliable index of experience even at the single-organism scale where it was designed to apply, using it as the measurement target for a civilisational-scale VALIS claim inherits that unresolved foundational problem rather than avoiding it.
5. **The neuro-cultural account of the four archetypes, stated at full strength.** The interface/hybrid frame (§4.3) is not the only account of why contact reports cluster into four recurring forms. A shared human nervous system running a small number of evolved templates — agent detection, a threat/guide dichotomy, size-and-authority coding, and a handful of cross-culturally recurring altered-state imagery generators (Lewis-Williams & Dowson, 1988) — predicts the same four-way clustering without requiring any external source, non-embodied or otherwise. On this account, the cross-cultural consistency in sequence (disengagement → scan → exchange → return) reflects the shared structure of altered or liminal states of human consciousness, not the shared structure of an external signal. This account is not weaker a priori than the interface/hybrid frame; the source material's own three-frame analysis (§4.3) already grants that the neuro-cultural frame accounts for timing and context-sensitivity better than either rival, and only falls short on "all" of the cross-cultural convergence — a gap, not a refutation. It should also be read alongside the §4.3 methodological caveat: since the archetype clustering itself has not yet been independently coded, this alternative and the interface/hybrid frame are currently both consistent with the same, still-informal, evidence base.

None of these alternatives is presented here as a refutation; each is presented, per instruction, as an available account that a fair-minded reader should weigh against the coherence/VALIS model before accepting it.

## 7. Evaluation by the Gigerenzer/Taleb Standard

Rather than asking whether the model is falsifiable in Popper's sense, it is evaluated here on the two criteria the author uses elsewhere:

**Ecological rationality** — does the model produce a usable heuristic under real uncertainty? Yes. "Log every case without pre-filtering, state the expectation before the observation, flag reification, code by a fixed and open taxonomy" runs today, on live data, across seven domains, without settling the ontological question first. It is also robust to disagreement: a

skeptic and a believer coding the same case should converge on the same EXPECTATION/OBSERVATION/FAILURE TYPE fields even while disagreeing about what the pattern ultimately means — same protocol, divergent interpretation, checkable agreement, which is the clearest sign of ecological rationality a heuristic of this kind can show.

**Antifragility** — does the model improve under stress rather than collapse? *Improved, but the decisive test still has not been run.* Moving from a hand-selected eleven-case sample to a comprehensive, non-filtered 31-entry log across 7 domains is itself a response to stress (the optional-stopping critique) that made the protocol stronger rather than abandoning the project — which is the correct antifragile response. What is still missing is the single highest-value next step: a pre-registered prediction of expected category distribution for the *next* batch, run before that batch is coded, with the result reported either way. Until that exists, "31 entries, 7 categories" is a well-built instrument, not yet a tested one.

### 7.1 What would count against the model — stated as explicit killer observations

A reviewer of an earlier draft correctly pointed out that naming alternatives (§6) is not the same as naming an observation that would make the author abandon or substantially revise the model. Three such observations are stated here directly:

1. **Uniform distribution under pre-registration.** If expected category distributions for Batch 3 (and subsequent batches) are pre-registered before coding, and the observed distribution does not differ from what chance registration would produce — i.e., a formal non-uniformity test fails to reject the null — that is direct evidence against the "residue clusters non-randomly" claim this synthesis is built around, and should be reported as such rather than reframed.
2. **No independent cross-database clustering.** If failure-type distributions in differently-administered databases (e.g. NUFORC vs. GEIPAN) turn out to be uncorrelated once each is tested against its own base rate — that is, if the clustering in one database does not predict clustering in the same categories in another — the cross-domain convergence argument in §5, point 3 fails, and the "shared filtering architecture" explanation loses its main piece of support.
3. **Clustering that tracks investigator identity rather than case content.** If failure-type assignment correlates more strongly with *which coder or which source institution* logged a case than with features of the case itself, that is direct evidence for the "pure reporting artefact" alternative (§6.1) over the coherence/VALIS account, and should be treated as a serious disconfirmation rather than an artefact to be explained away.

Any one of these three, if observed, would be reported in a subsequent revision of this synthesis as a finding against the model — not folded into a reinterpretation that preserves the model's conclusions.

## 8. Conclusion and Next Steps

Combined, the three source essays give a chain — coherence mechanism, layered channel, a candidate occupant of the upper channel — and this synthesis adds the missing link: a specific, checkable prediction about what such an occupant's evidential footprint would have to look like, *if* it existed (statistical residue concentrated in specific failure types, replicated independently across reporting systems, without cross-database contact). The prediction is stated plainly enough to fail, which is the appropriate ambition given the author's own standard: not proof, but a heuristic robust enough to survive contact with more data, and honest enough to say when it does not.

**Concrete next steps, in priority order, updated for the fixed methodology:**

1. Pre-register expected category distribution for Batch 3 *before* it is coded, so that any observed concentration in Categories 1–7 can be judged against a stated prior rather than a pattern noticed after the fact — this is now the single largest remaining gap identified in §6.3 and §7.
2. Run a formal non-uniformity test (e.g. chi-square or a permutation test against a null of category-assignment independent of domain/source) on the current 31-entry log as soon as a next batch brings the sample to a size where such a test is meaningful, and report the result regardless of outcome — treating the three killer observations in §7.1 as the explicit standard the result is judged against.
3. Actively search for a disconfirming case — a feed, domain, or time window where the protocol's own logic predicts concentration and the logged entries show none — and report it with the same visibility as a confirming case.
4. Keep the *descriptive* claim (failure types cluster non-randomly across independently-sourced, non-filtered registrations) separate from the *interpretive* claim (this is consistent with a non-embodied coherence field occupying Sensorium Layers 15–19, among other candidate interpretations) in all future write-ups, exactly as the Sensorium paper already separates Layers 0–14 from Layers 15–19 — so a reader can accept the first without being asked to accept the second.
5. Treat the REIFICATION CHECK field as a first-class output, not a bookkeeping note: a rising or falling rate of reification over time (i.e., how often a failure gets converted into an object/substance name before the underlying expectation-failure is logged) is itself a testable trend and a direct empirical proxy for the Layer 0–14 filtering process this synthesis is built around.
6. Pilot the combined protocol proposed in §5.2 on a small case set: use UEFRM's comprehensive, non-filtered log to select cases, then score those cases on the multi-framework protocol's six measurement levels against pre-stated, framework-specific predictions — before, not after, the scoring is run.
7. Apply independent, blind coding with a stated inter-rater agreement statistic to the four-archetype contact classification (§4.3), using the same discipline UEFRM already applies to failure-type coding — this is the most direct way to move the archetype claim from qualitative impression to tested finding.

## References

Bem, D.J., et al. (2014–2020). Meta-analyses of predictive anticipatory activity. Various publications.

Dick, P.K. (1981). *VALIS*. Bantam Books.

Doerig, A., Schurger, A., Hess, K., & Herzog, M.H. (2019). The unfolding argument: Why IIT and other causal structure theories cannot explain consciousness. *Consciousness and Cognition*, 72, 49–59.

Konstapel, J. (2025a). Bewustzijn is de Coherentie die uit Resonantie Ontstaat. constable.blog, 15 May 2025.

Konstapel, J. (2025b). Wat is het Vast Active Living Intelligence System? En hoe toont het Zich? constable.blog, 12 October 2025 (includes "The Four Principles of VALIS Manifestations in Ancient Myths and Sacred Texts" and "Beyond the Simulation: An Ontological Taxonomy of Non-Human Contact Hypotheses" as embedded appendices).

Konstapel, J. (2026a). The Nested Sensorium: A Layered Architecture of Perception Beyond Conscious Awareness. constable.blog, 26 June 2026.

Konstapel, J. (2026c). Universal Expectation-Failure Registration Model: Fixed Methodology for Comprehensive UAP Registration. Constable Research B.V., Leiden, 1

July 2026 (fixed protocol; supersedes earlier sample-based templates; Batch 1 and Batch 2 registration logs, 31 entries, 7 domains, 7 failure-type categories).

Kuramoto, Y. (1984). *Chemical Oscillations, Waves, and Turbulence*. Springer.

Lewis-Williams, D., & Dowson, T.A. (1988). The signs of all times: Entoptic phenomena in Upper Palaeolithic art. *Current Anthropology*, 29(2), 201–245.

Mack, J.E. (1994). *Abduction: Human Encounters with Aliens*. Scribner.

McGilchrist, I. (2009). *The Master and His Emissary: The Divided Brain and the Making of the Western World*. Yale University Press.

Nelson, R., et al. Global Consciousness Project (1998–present).

Schank, R.C., & Abelson, R.P. (1977). *Scripts, Plans, Goals and Understanding*. Erlbaum.

Strogatz, S. (2003). *Sync: The Emerging Science of Spontaneous Order*. Hyperion.

Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5, 42.

Wang, C.X., et al. (2019). Transduction of the geomagnetic field as evidenced from alpha-band activity in the human brain. *eNeuro*, 6(2).

## Appendix A: The Eighteen-Framework Hypothesis Space, and Where This Model Sits Within It

(Supports §4.4.)

The source material proposes an eighteen-framework taxonomy of possible explanations for "contact" phenomena, organised along two axes: **information topology** (where does information "live" — locally in brains, distributed in fields, in a computational substrate, intrinsic to consciousness itself, non-linearly in time, or in parallel dimensions) and **agency architecture** (is there an external agent at all, and if so is it emergent-collective, immanent-universal, transcendent, a peer agent, or a retrocausal/future self).

Placed on this map, the coherence/Sensorium/VALIS model developed in Parts I–III of this synthesis is not a new eighteenth alternative; it is a specific, explicit combination of four of the eighteen: **B3 (biofield/EM coherence)** for the physical mechanism, **D2 (participatory universe / synchronicity as ordering principle)** for the phenomenology of contact, **E2 (noosphere/planetary coherence)** for the collective-scale substrate, and, at the edge, **A4 (consciousness as a candidate fifth field)** for the most speculative claim about what the upper Sensorium layers ultimately are.

Naming the position on the map this way has a direct methodological payoff: the source taxonomy notes that *testability tracks topology*, with substrate- and field-models yielding more concrete predictions than pure consciousness-primacy models. Because this synthesis's strongest claims sit in the field/collective cluster (B3/D2/E2) rather than the purely intrinsic-consciousness cluster (D1), it inherits the more testable end of the space — consistent with, and giving an independent reason to prefer, the residue-based detection logic developed in Part IV over a claim that would need to be evaluated by introspection alone.

## Appendix B: A Complementary Protocol — Multi-Level, Multi-Framework Measurement

(Supports §5.2.)

The source material proposes a second, independent testing protocol that this synthesis recommends running *alongside* UEFRM rather than instead of it, because the two protocols test different things and share no methodological weakness in common.

Where UEFRM asks "does this case fail a specific model's expectation, and in what way," the multi-framework protocol asks a case to be scored on **six measurement levels** (physical: EM fields, cosmic-ray/spacetime anomalies; physiological: fMRI, EEG, biophoton, heart-rate variability; behavioural: decision quality, learning rate, synchronicity; informational: compression, Bayesian surprise, semantic novelty; social: idea diffusion, collective behaviour; temporal: presentiment, precognition, retrocausal effects), and then evaluated against **specific, differently falsifiable predictions** made by each of several rival frameworks — for instance, a genuine morphic-field effect (B1) predicts measurably accelerated learning curves as more people learn a given skill, while a genuine quantum-consciousness effect (A3) predicts that general anesthesia should measurably disrupt the specific coherence process claimed, not just behaviour.

**Why together, not either alone:** UEFRM's strength (comprehensive intake, §5.1) sits at the *registration* stage; its weakness (no pre-registered category predictions, §6 point 3) sits at *interpretation*. The multi-framework protocol is the mirror image — it forces pre-stated predictions before scoring, but starts from a small, hand-picked case set, reintroducing the selection problem UEFRM was redesigned to solve. Run separately, each inherits the other's gap; run in sequence — UEFRM selects, the multi-framework protocol scores against pre-stated predictions — each patches the other. This is proposed as the single most concrete methodological improvement available to the programme, and follows from applying the Gigerenzer/Taleb standard to the combination of tools rather than either in isolation.

*Note on method: this synthesis was assembled from the three named constable.blog essays (including the full text of the VALIS essay and its embedded appendices, supplied directly by the author), the author's related published papers, and the fixed UEFRM methodology document. This revision incorporates external peer feedback on an earlier draft, addressed as follows: the status-at-a-glance table (after the Abstract) and the systematic replacement of "is"/"precisely"/"by construction"-type phrasing with conditional language throughout Part III respond to the ontology-*

*vs-methodology critique; §2.4 (missing intermediate steps) responds to the synchronisation-to-consciousness gap critique; the methodological caveat in §4.3 responds to the critique of the archetype classification's unstated selection criteria; §7.1 (killer observations) responds to the critique that alternatives were named without a stated disconfirmation condition; and the eighteen-framework mapping and the complementary multi-framework protocol were moved to Appendix A and Appendix B, respectively, in response to the length/repetition critique, leaving short pointers in §4.4 and §5.2 so the main argument (Parts I–IV, §6–8) reads without them. Sections 6 and 7 remain the connecting analytical layer added by this synthesis, not content drawn from the source documents.*