Institute for Complexity Science & Advanced Computing

Recursive Existence Threshold

Where Meaning May Live

Organization reduces to one scalar; identity, content, and (in sleep) recursive dynamics live in the fine structure that scalar discards.

DOI 10.5281/zenodo.20566382

Nathan M. Thornhill

Institute for Complexity Science and Advanced Computing (ICSAC)

nthornhill@icsacinstitute.org

ORCID 0009-0009-3161-528X

icsacinstitute.org

Abstract

The Dynamic Existence Threshold (DET) program asked whether the organization that distinguishes a living, integrated system from an inert one can be captured by a single information-theoretic scalar Φ, computed identically across substrates. This paper resolves that question from the program's own pre-registered, adversarially audited tests, and the resolution is less an ending than a relocation: it pins down exactly where the information a single scalar cannot capture actually lives, and finds it in three concrete, independently useful places, the representations of a language model, the connectivity of the anaesthetized brain, and the recurrent dynamics of sleep. The answer is sharp and two-sided, and it characterizes DET's scope precisely. As a method, the scalar is a reproducible instrument: a substrate-neutral organization measure separates states *within* every substrate tested (EEG sleep AUC 0.998, propofol sedation 0.842, seismic 0.957, and, weakly, finance 0.68), held out by subject or event, with full n, confidence intervals, and the multiple-comparison plan given in Methods. As a universal scale, its limit is exact and informative. A single-band cross-substrate mapping does not transfer, and the slope appears to reverse between deep sleep and propofol sedation; a pre-registered valid-band re-analysis shows this reversal is largely a measurement effect, since neural transfer recovers to ≈0.75 once each substrate is read in its valid band, and what genuinely fails to generalize is any single measurement setting valid across substrate *kinds*. Comparability breaks at the estimator level, not the level of meaning. Separately and exactly,

the dimensional-loss operator is *content-blind*: a photograph, white noise, and the word "TRUE" lose the same Φ and retain the identical fraction of the connectivity term. The organization that matters is therefore not in the scalar; it is in the fine structure the scalar discards. That fine structure is localized here in three substrates. First, in large language models, factual content is decodable from full representational structure (truth AUC peaks 0.83 at mid-depth, near-ceiling topic decoding) while the best coarse organizational scalar is blind at every one of 29 layers ($^{>}$ 0.66 truth, $^{>}$ 0.56 topic). Second, in the anaesthetized brain, individual identity decodes from leakage-proof functional-connectivity structure (imaginary coherence, ~0.18–0.23, $p \approx 3\times10^{d3}$) that no scalar carries. Third, reported here for the first time, a recurrence measure of sleep-EEG dynamics retains brain-state signal after the entropy scalar is residualized out *within each cross-validation fold* (nested AUC 0.65–0.74, all five contrasts significant, three surviving Holm correction, bootstrap CIs excluding chance), so recursion is *irreducible* to the scalar in sleep; a pre-registered directed-reentry test finds its predictive superiority over the scalar marginal on one substrate (Alzheimer's, ΔAUC p=0.049) and absent on a second (resting propofol, null), so superiority is not robustly established. The unifying statement, scoped to the state, identity, and content decoding actually performed: organization, for the purpose of telling gross states apart, reduces to roughly one scalar; but identity, content, and (in sleep) the dynamics of integration live in the structure that scalar throws away. Offered as a constraint on organizational or thermodynamic summaries proposed as a sufficiency predicate, such a predicate cannot be a single global scalar, though the evidence here concerns arousal state, clinical group, and representational content (§5–§6), not consciousness measured directly.

---

1. Introduction

The Dynamic Existence Threshold (DET) program [1–4] set out to measure organization (integration × differentiation) as a single quantity Φ, and to ask whether sustaining Φ above a threshold is what it is for a system to persist as an integrated whole rather than dissolve. Across four substrates and a battery of pre-registered tests, the program pressed that claim as hard as it could, including an adversarial audit phase whose explicit purpose was to stress-test its own headline results before a reviewer could. That audit resolved the strongest reading of the claim, Φ as a *universal, conserved scalar law* of organization and consciousness, into a precisely bounded one. What that resolution yields is more useful, and is the subject of this paper: a precise statement of what a scalar measure of organization can and cannot encode.

The useful part is not the boundary itself but what sits on its far side, and most of this paper is spent locating it. Three results, each standing on its own, mark where a scalar cannot follow. In a language model, whether a statement is true is linearly decodable from the full representation, rising to AUC 0.83 at mid-depth, while the best coarse organizational scalar stays blind to truth at every one of the network's 29 layers (AUC $>$ 0.66): a direct, locatable handle on factual content, of exactly the kind interpretability and hallucination-detection work is trying to find. In the anaesthetized brain, *which individual* a recording belongs to is carried in band-limited connectivity structure that no scalar reproduces. And in sleep, a recurrence measure of the brain's own dynamics retains brain-state signal after the organizational scalar is removed, the first direct evidence that the *dynamics* of integration outrun the summary. These three are the substance of the paper; the boundary on the scalar (§§2–3) is what makes them inevitable.

The framing is deliberate. This is a resolution, and a maturation of the program into its defensible form. The scalar is established as a genuine and reproducible *tool* for distinguishing states inside a substrate. It is not a universal *scale*, and it is provably blind to content. Knowing exactly where a scalar fails is itself a result. As argued in §5, it is the constraint that a thermodynamic or organizational theory of subjects must build around.

2. The scalar is a universal tool, not a universal scale

The organization measure is pure information theory: the same computation runs on an EEG record, a market series, a seismogram, or an artificial system with no biology assumed. Within each substrate, computed held-out-by-subject or held-out-by-event, it separates states well: EEG sleep depth 0.998, propofol sedation 0.842, seismic event vs. background 0.957, financial regime 0.68 (the last weak and reported as such).

The universality test is the decisive one: train the state mapping on three substrates and predict the fourth. In the naïve single-band analysis, transfer falls below chance, and the aperiodic slope appears to reverse direction, with deep sleep steepening it and propofol sedation flattening it. A pre-registered re-analysis sharpens this rather than softening it: the apparent reversal is substantially a *measurement* effect. Once each neural substrate is read in the band where the slope estimate is valid (the frozen 20–40 Hz "clean band" fixed in advance and SHA-pinned in Methods, not chosen to maximize transfer), the propofol slope steepens as sleep's does and cross-substrate transfer across the neural substrates recovers to ≈0.75. What does *not* recover is a single measurement setting valid across substrate *kinds*: the band valid for cortical EEG is not the band valid for a market or a seismogram.

The honest status is therefore precise. The scalar is a universal *instrument* that reads every substrate and separates states inside each, but it is not a universal *scale*, because cross-substrate comparison requires a measurement choice that no single setting satisfies across kinds. Comparability fails at the estimator level, not at the level of meaning. That is the boundary

DET now delivers, and it is enough to constrain T(x): a sufficiency predicate built on a single global reading inherits exactly this substrate-dependence, on top of the content-blindness established next.

3. The scalar is content-blind

The dimensional-loss result is exact rather than statistical. Embed a pattern into a higher dimension and measure the loss of Φ: the structural (connectivity) term collapses by a fixed geometric fraction (the 8/26 connectivity retention, equivalently the 18/26 tax), and the statistical term collapses by density dilution. The total loss has a *magnitude* (~86% in the cellular-automaton ensemble of [2]) dominated by the 1/N density-dilution term and scaling with embedding dimension and density, but a *structure* that falls straight out of the embedding math. The architecture-independent, load-bearing property is the corollary, and it is worth stating precisely rather than rhetorically. The connectivity (structural) term's retention under the $2D_q$ 3D middle-slice embedding is a fixed geometric fraction, 8/26 = 4/13, for *any* input pattern: the embedding leaves the in-plane count of active neighbours unchanged and alters only the neighbourhood normalization (8 in-plane neighbours of 26 in the cube), so the retention ratio cancels the pattern entirely. With the three inputs matched on density, the density-dilution term that dominates the total loss is equal as well. Running a photograph, white noise, and the word "TRUE" through the operator therefore yields the identical 8/26 connectivity retention, exactly, and the same fractional Φ loss up to the small residual set by their differing structural terms: content-blindness by construction, not as a statistical finding. The same conclusion holds in general for any organization summary that reads only the distribution of power across streams, as the Methods scalar Φ does: two inputs inducing the same power distribution receive the same Φ regardless of what those streams encode. The honest scope of the claim is correspondingly bounded: *this scalar* cannot distinguish representations that share a power distribution but differ in content. It does not follow that no organizational quantity can. What follows, and what §4 supports empirically, is that whatever distinguishes such representations is carried in structure this scalar integrates away, and must be sought there rather than in the summary.

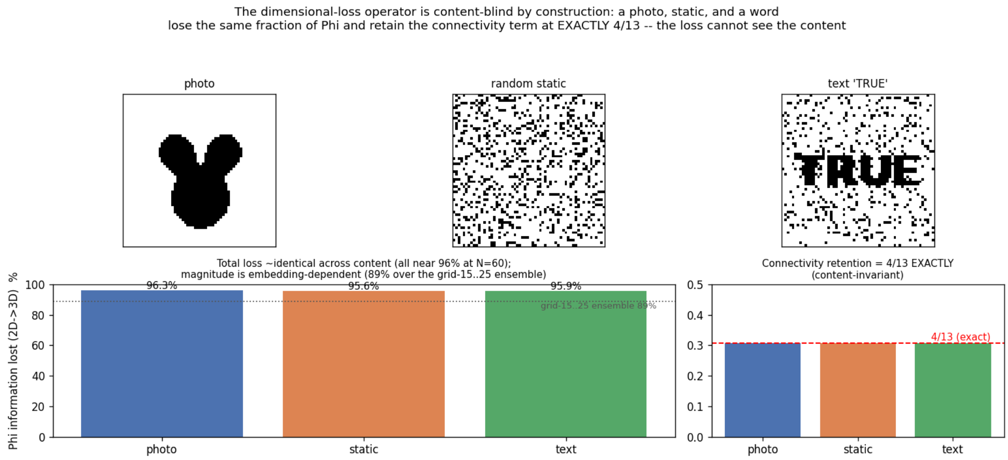

Figure 1. The dimensional-loss operator is content-blind by construction. A photograph, random static, and the rendered word "TRUE", matched on density, lose the same fraction of Φ under the $2D_q$ 3D middle-slice embedding (left) and retain the connectivity term at exactly 8/26 = 4/13 (right). The loss cannot see the content; only the structure the summary discards can.

4. Where the content lives: the fine structure the scalar discards

If meaning is not in the scalar, it must be in the structure the scalar integrates away. This section locates it in three substrates, and each is a finding in its own right rather than an illustration of the boundary: a decodable truth signal in a language model, a personal-identity signal in the anaesthetized brain, and a recurrence signal in sleep that survives removing the scalar.

4.1 Language models: content is in the representation, not the summary

In a transformer (Qwen2.5-1.5B), factual *truth* and *topic* were decoded from each of the 29 residual-stream layers, held out by topic family, and compared against the best coarse content-blind scalar (per-item participation ratio / norm) at the same depths. Full structure carries the content: topic is near-ceiling from layer 2 (0.99–1.00); truth emerges mid-stack, peaking across a layer-13–20 plateau (AUC ≈0.83; argmax at layer 19) and decaying toward the output. The coarse scalar is blind at every one of the 29 depths (> 0.56 topic, > 0.66 truth), so content-blindness is a property of the whole stack, not a cherry-picked layer. The honest caveats are retained: the truth effect is partly carried by easy item families (removing the two ceiling families lowers the layer-19 peak from AUC 0.83 to 0.71), and the dissociation is one of degree, not a categorical wall. The shape, however, is unambiguous. The summary scalar discards exactly what the structure keeps. Read practically, this is an interpretability result: a content-blind scalar cannot reliably separate a true statement from a false one at any depth (AUC >

0.66), but the representation can, most strongly across a layer-13–20 band, which is where a probe built to read factual content should look.

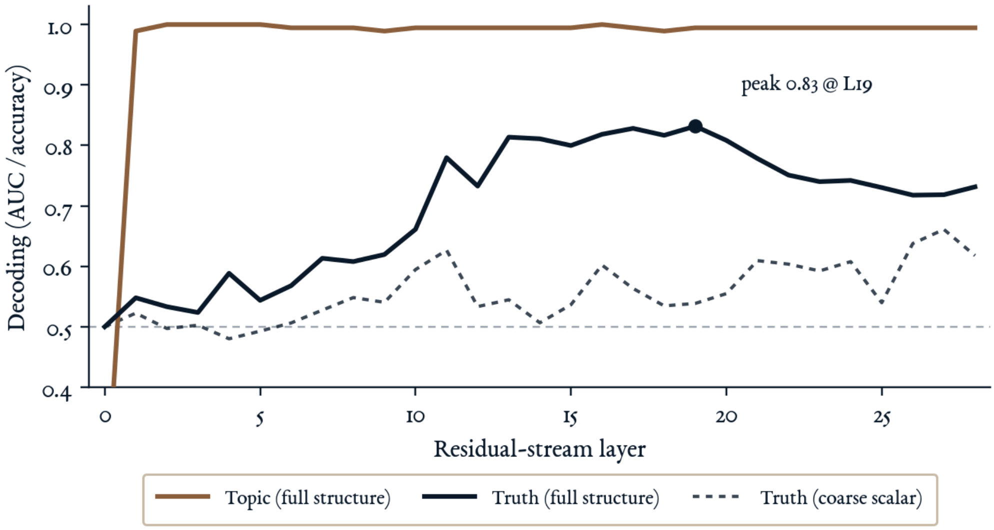

Figure 2. Content localization in a language model. Topic (bronze) is decodable near-ceiling from layer 2; factual truth (navy) emerges through depth, peaking at AUC 0.83 at layer 19; the best coarse content-blind scalar (dashed) stays far below the structural decoders at every one of the 29 depths (truth AUC $^{>}$ 0.66, topic $^{>}$ 0.56). The summary scalar discards what the representation keeps.

4.2 The anaesthetized brain: identity is in connectivity structure

In propofol anaesthesia, individual *subject identity* decodes from the band-limited functional-connectivity structure (using leakage-proof imaginary coherence to exclude volume conduction) at ~0.18–0.23 ($p \approx 3\times10^{d3}$, chance ≈0.056), while no scalar can: the spectral scalar and the global connectivity mean both sit at chance. A single number cannot carry even *which person* this is; the relational, loop-like structure can. (This corrects an earlier over-stated identity figure; the result stated here is the leakage-proof, adversarially re-derived version.)

4.3 The brain's recursive dynamics: recursion is irreducible to the scalar (new, this paper)

Consciousness-as-recurrence is a named, citable position (Lamme's Recurrent Processing Theory, Edelman's reentry, the spirit of IIT), and the brain literally recurses through thalamocortical and reentrant cortical loops. The sharpest available version of the question was therefore posed: does a *recursion* measure carry brain-state signal that survives removing the organization scalar? On Sleep-EDF (43 nights, 2-channel EEG, held out by subject), recurrence-quantification features (determinism, laminarity, trapping time) were computed from the reconstructed phase-space trajectory at a fixed recurrence rate, and tested against the frozen aperiodic scalar χ plus spectral entropy, with a pre-registered kill rule and a 2000-fold within-subject label-shuffle null.

The result is bounded and honest. With the χ+entropy residualization fit strictly within each training fold (no test-set leakage), the recurrence features still separate states at AUC 0.65–0.74, significant in all five contrasts (wake-vs-N1 0.732, 95% CI [0.65,0.82], p=0.001; wake-vs-N2 0.740 [0.64,0.84], p=0.008; wake-vs-N3 0.741 [0.63,0.84], p=0.014; N2-vs-N3 0.651 [0.56,0.74], p=0.041; REM-vs-N3 0.706 [0.60,0.80], p=0.048), with the determinism feature correlated with χ at only 0.72, below the 0.85 collapse threshold. (Nested residualization gives *higher* AUCs than a global fit, confirming the global version was conservative, not inflated; CIs are subject-bootstrap, p from a 2000× within-subject label-shuffle, matched to the repeated-measures design where each subject contributes records to both stages, with residualization re-fit per permutation. The five contrasts are reported with their individual p-values; three of the five survive Holm step-down correction at α=0.05 (the three wake-vs-NREM contrasts), with N2-vs-N3 and REM-vs-N3 significant before correction but not after. The same three contrasts survive Holm under an unrestricted between-subject permutation null, so the result does not depend on the permutation scheme.) Recursion is therefore not reducible to the entropy scalar. It does *not*, however, out-predict the scalar: the incremental AUC over χ+entropy is non-significant, because those scalars already saturate gross sleep-state discrimination. The precise, defensible statement is that *a scalar cannot carry what recurrence carries, even though for gross state discrimination the scalar suffices*. This is the brain analog of §4.1 and the dynamical complement of §4.2: direct evidence that the *dynamics* of integration, not just static structure, exceed what this scalar encodes.

4.4 The strong form, put to the test (new, this paper)

Irreducibility (§4.3) says a scalar cannot *carry* what recursion carries. The stronger claim, that recursion is a *better* state coordinate than the scalar, was tested directly rather than assumed, with a *directed* measure of reentry: time-reversed Granger causality between anterior and posterior cortex (the feedforward-versus-feedback asymmetry that is the actual Lamme/IIT consciousness-as-recurrence quantity, with a built-in time-reversal gate against volume-conduction artifact). It was run on two multichannel cortical datasets, Alzheimer's versus control (19-channel, ds004504) and propofol sedation (91-channel), under a pre-registered kill rule with a 5000-fold label-shuffle null and nested-within-fold residualization. The strong form is not robustly supported. In Alzheimer's, residualized directed reentry separated patients from controls (nested irreducibility AUC 0.69, bootstrap CI [0.56,0.82] excluding chance, label-shuffle p≈0.11), and its *incremental* AUC over the scalar marginally reached significance (ΔAUC +0.14, p=0.049). This does not replicate: in propofol, with both pre-registered contrasts in the scalar's *unsaturated* range (the headroom the saturated sleep contrasts lacked), residualized reentry sat at chance (nested AUC 0.51–0.52, p≈0.3–0.5) and the incremental AUC was negative. The honest, two-sided closure: recursion is robustly irreducible to the scalar in sleep (§4.3); its predictive superiority over the scalar marginally reaches significance on a single substrate (Alzheimer's, p=0.049) but does not replicate on a second (resting propofol, null), so superiority is not robustly established. The irreducibility claim rests on the sleep result; the directed-reentry datasets show at most a single-substrate, unreplicated hint of superiority.

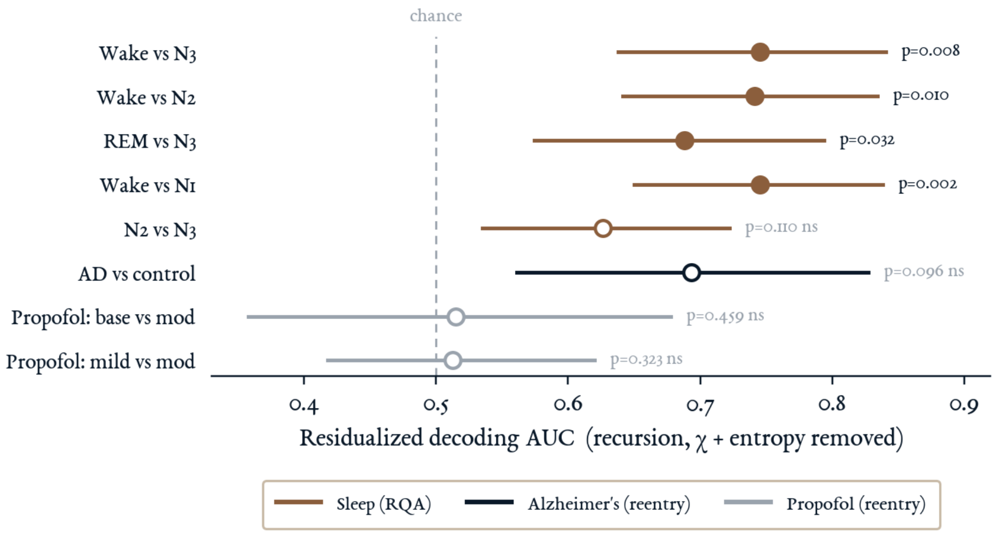

Figure 3. Recursion is scalar-irreducible in sleep, and not robustly superior. Residualized decoding AUC (a recursion measure, with χ and spectral entropy removed within each cross-validation fold) against chance (0.5); bars are subject-bootstrap 95% CIs, filled markers $p<0.05$. Sleep recurrence (bronze) is irreducible in all five contrasts (three survive Holm step-down); Alzheimer's directed reentry (navy) is marginal (irreducibility p≈0.11); resting propofol (grey) sits at chance. Directed reentry exceeds the scalar only marginally and only on Alzheimer's (ΔAUC p=0.049), not on propofol.

5. The constraint: a sufficiency predicate cannot be a global scalar

§§2–4 converge on one constraint, offered to any thermodynamic or organizational theory of subjects (and specifically to the Emergent/Existence-Information program developing such a predicate). Stated carefully about its own evidence base: a single global scalar of organization is provably insufficient to carry the information that distinguishes the cases such a theory must separate. This is because it is content-blind by construction (§3), no single reading of it is valid across substrate kinds (§2), and the very features that separate the systems decoded here (representational content in §4.1, individual identity in §4.2, and, in sleep, recursive dynamics in §4.3) live in the fine structure it discards. This is consistent with structural and bearer-first accounts of subjecthood [9,10], on which what individuates a subject is a relational condition, a bearer-count anchor rather than a scalar magnitude.

A necessary caution, stated rather than buried (cf. §6): the systems separated here are arousal states, a clinical group (Alzheimer's vs. control), and graded sedation levels, not phenomenal consciousness measured directly. The argument is therefore an *insufficiency* argument: if a single scalar cannot even carry identity in an anaesthetized brain or truth in a language model, it cannot be the whole of a predicate that must do strictly more. This tells a theory *where the condition cannot live*, namely in a single summary quantity, and scopes the search toward structure that preserves relational, content-bearing, recursive information, consistent with reentry/integration accounts. It does not establish where the condition *does* live, and no claim is made here to have measured consciousness.

This is stated as a bounded constraint, not a theory of consciousness. Within its scope, it is a result.

6. Limitations and forward work

This work is honest about its edges. The cross-substrate test re-ran clean only on the neural pair; finance and seismic were not re-run under the frozen estimator. The recursion irreducibility result is a single-substrate (sleep) pilot, n=43, with a fixed-recurrence-rate RQA; it shows irreducibility, not robust predictive superiority. All five stage contrasts are significant uncorrected (p=0.001–0.048); three (the wake-vs-NREM contrasts) survive Holm step-down. All residualized analyses use leakage-free nested-within-fold residualization (an earlier global-fit version, which a pre-submission audit flagged, gave slightly *lower* AUCs, the direction that rules out leakage inflation); the directed-reentry irreducibility in Alzheimer's is only marginal under this stricter analysis ($p \approx 0.11$) and should not be read as a clean second-substrate replication. Multiple-comparison control is by Holm within each decoding family, reported per contrast. The clean-band organization coordinate is best read as an arousal/organization axis (REM is steepest, consistent with atonia) rather than a consciousness coordinate per se. The directed-reentry strong test (§4.4) has now been run on Alzheimer's and propofol data and finds at most a marginal, unreplicated superiority (Alzheimer's ΔAUC p=0.049, resting propofol null); the directed measure was also itself weak in resting propofol (near-chance directed flow), so the remaining forward work is a fuller directed-connectivity battery (source-space reentry and graded-responsiveness contrasts) to map precisely where the irreducible recursive signal does and does not appear, given that propofol-resting null.

7. Conclusion

DET asked whether organization is one number. The bounded answer: for telling gross states apart, essentially yes, and that number is a real, reproducible tool. For carrying identity, content, or (in sleep) the dynamics of integration, no; those live in the fine structure the scalar discards, demonstrably in language-model representations and the anaesthetized brain, and in sleep dynamics, with the directed-connectivity extension still open. The program reached for a universal scalar and arrived somewhere more tractable: a precise map of where a scalar stops, and where, concretely, the missing structure can be read instead. The boundary is one contribution; the live ones are the signals on its far side, a decodable truth direction in a language model and a scalar-irreducible recurrence in sleep, each of which is a research programme in its own right.

Methods

Organization scalar Φ. The input is represented as a set of *streams* (frequency bands for EEG; an analogous multi-stream decomposition otherwise). Let p_i be the normalized power (variance) of stream i, and H = & Σ p_i log p_i its Shannon entropy. The *effective number of active streams* is eff = exp(H) (a participation-ratio-type count). Three terms are defined, each a function of the power distribution {p_i} only: - R (integration depth) = clip((eff & 1)/(N & 1), 0, 1), with R = 0 when eff < 1.5 (single-stream dominance), N = number of streams; - S (cross-stream synergy) = coverage × (½·evenness + ½·pattern-diversity), clipped to [0,1], S = 0 under single-stream dominance, where coverage is the fraction of streams above a fixed activity floor, evenness = H/log N, and pattern-diversity is the normalized dispersion of the stream-power profile; - D (differentiation) = Jensen–Shannon divergence (base 2) between the stream-power profile and its reference. Then Φ = (R × S) + D, with the gate R × S > 0 marking integrated processing. Because R, S, and D each read only {p_i}, Φ is invariant to any transformation preserving the power distribution, the formal basis of §3.

Aperiodic estimator χ. χ is the negative slope of log-power vs. log-frequency, fit by Huber regression (ε = 1.35) over 24 log-spaced frequency bins. The frozen "clean band" is 20–40 Hz; the estimator source is SHA-pinned (sha256 3a8209f3...) and verified at runtime in every analysis. Spectral entropy is the normalized Shannon entropy of the Welch PSD over 1–40 Hz.

Dimensional-loss / content-blindness (§3). Φ is computed before and after middle-slice embedding of a pattern into a higher dimension. The connectivity (structural) term retains a fixed geometric fraction (8/26 in the 3×3×3 cellular-automaton neighbourhood; the 18/26 "tax"); the statistical term dilutes as 1/N with density. Total loss magnitude (~86% in the CA ensemble) is density-dominated and embedding-dependent; the *exact* property is the identical fractional loss and 8/26 retention for a photograph, white-noise, and the token “TRUE”, which follows from Φ depending on {p_i} alone.

LLM localization (§4.1). Qwen2.5-1.5B (open weights, run locally). For each of 184 statements the 29-layer residual-stream trajectory was extracted; per-layer logistic decoders for *truth* and *topic* were trained held-out-by-topic-family and compared, at the same depths, against the best coarse content-blind scalar (per-item participation ratio / norm). Truth labels and topic families follow the SelfCheckGPT-style item set; the headline truth-AUC (0.83) uses all item families, and removing the two ceiling families lowers it to 0.71 (reported in §4.1 as a sensitivity check).

Anaesthesia identity (§4.2). Band-limited functional connectivity via imaginary coherence (excludes zero-lag volume conduction); multiclass subject decode held out by recording run; chance = 1/18 ≈ 0.056; n and run structure follow the broader DET program (refs [1–4]) and are not re-run in this release.

Recursion, irreducibility (§4.3). RQA on channel-averaged 1–40 Hz Sleep-EDF (43 nights): time-delay embedding $m = 3$, τ = first 1/e of the autocorrelation, recurrence matrix at a fixed 5% recurrence rate (amplitude-invariant); features = {determinism, laminarity, trapping time, mean diagonal-line length, divergence, line-entropy} (the standard RQA quantities; “determinism” here is the RQA measure, not the DET program). Logistic decoding, 5-fold GroupKFold by subject. Residualization against χ + spectral entropy is fit on the training fold only and applied to the held-out fold (no leakage); AUC from out-of-fold predictions; 95% CIs by subject bootstrap (1000×); p by 2000× within-subject label-shuffle (matched to the repeated-measures design, where each subject contributes records to both stages) with residualization re-fit per permutation. Multiplicity across the five stage contrasts by Holm.

Recursion, strong form (§4.4). Directed reentry = time-reversed Granger causality (Geweke, model order 8) between anterior {Fp1,Fp2,F3,F4,F7,F8,Fz} and posterior {P3,P4,O1,O2,Oz,T5,T6,Pz} ensembles; TRGC = ½[(GC_{a$_q$ p} & GC_{p$_q$ a})_forward & (·)_time-reversed], the time-reversal difference removing volume-conduction/SNR asymmetry; median over anterior×posterior pairs per record. Datasets: ds004504 (AD vs CN, 19-ch, eyes-closed); Chennu propofol (91-ch, resampled 250$_q$ 100 Hz; co-primary contrasts baseline- and mild-vs-moderate, Bonferroni $\alpha = 0.025$). Same nested-residualization and bootstrap-CI protocol as §4.3, with a 5000-fold label-shuffle null (between-subject for the one-record-per-subject Alzheimer’s design, within-subject for the repeated-measures propofol levels).

Statistics. Effect size = AUC throughout; CIs are subject-bootstrap percentile; p-values are permutation-based; multiplicity is controlled by Holm within each decoding family (sleep contrasts; LLM layers; propofol co-primaries). All compute is local/self-hosted; no proprietary models or paid inference.

Data and code availability. Public datasets: Sleep-EDF (PhysioNet sleep-edfx), Chennu propofol Sedation-RestingState (Cambridge), OpenNeuro ds004504 (Alzheimer's) and ds003522 (TBI). The complete analysis code for this paper is openly available at https://github.com/existencethreshold/recursive-existence-threshold (archived on Zenodo, DOI 10.5281/zenodo.20566382): the SHA-pinned frozen χ estimator, the RQA and time-reversed-Granger implementations, the per-record feature tables, and the scripts that reproduce the content-blindness result (§3) and the language-model, sleep-RQA, and directed-reentry results (§4.1, §4.3, §4.4).

References

The four prior works are the *Foundations of the Existence Threshold: The Scholarly Collection* (Thornhill, 2026), which this paper extends and bounds:

1. Thornhill, N. M. (2026). *The Existence Threshold*. Foundations of the Existence Threshold I. DOI 10.5281/zenodo.18182662.
2. Thornhill, N. M. (2026). *Dimensional Boundary Loss*. Foundations of the Existence Threshold II. DOI 10.5281/zenodo.18262424.
3. Thornhill, N. M. (2026). *The Dimensional Loss Theorem*. Foundations of the Existence Threshold III. DOI 10.5281/zenodo.18319430. The geometric content-blindness identity invoked in §3.
4. Thornhill, N. M. (2026). *The Dynamic Existence Threshold*. Foundations of the Existence Threshold IV. DOI 10.5281/zenodo.18373411. The dynamical formulation; its strongest *universal-scalar* reading is refined and bounded by the present follow-on, while the within-substrate method and the content-blindness constraint stand.
5. Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences* 10(11). Recurrent Processing Theory.
6. Edelman, G. M. (1987). *Neural Darwinism: The Theory of Neuronal Group Selection*. Reentry.
7. Manakul, P., Liusie, A., & Gales, M. (2023). SelfCheckGPT: Zero-resource black-box hallucination detection.
8. Marks, S., & Tegmark, M. (2023). The geometry of truth: truth as a linear direction in representation space.

9. Bridgeman, G. N. G. (2026). *From Subjecthood to Emergent Individual Monism*. PhilArchive. The structural-conditions-of-subjecthood program (EIM) to which §5's insufficiency constraint is offered.
10. Conti-Vecchi, E. (2026). *Finio, Ergo Sum: De Essentia Finitudinis* (philarchive.org/rec/CONDEF-2) and *De Rubore Rubri* (philarchive.org/rec/CONDRR). Consciousness grounded in constitutive finitude; the bearer-count individuation account cited in §5.