Ordered Patch Theory

Appendix T-15: The Phylogenetic Stability Filter

Anders Jarevåg

May 11, 2026 | DOI: 10.5281/zenodo.19300777

Original Task T-15: Phylogenetic Codec Refinement under the Stability Filter Problem: The v3.6.0 multi-scale codec extension proposes that biological evolution — specifically the brain-first cascade through which the patch's present neural architecture descends — is a slow-timescale structural analogue of the Maintenance Cycle ($\mathcal{M}_\tau$, §3.6) under the Stability Filter's within-stream observer-compatibility attractor (§6.6 v3.6.0 hypothesis). Preprint §3.6.9 records the correspondence; this appendix formalises the prediction structure and the fossil-record falsification footprint of the candidate biological research programme (§6.8.1). **Deliverable:** Formal articulation of the phylogenetic-scale structural correspondence, the four-class prediction structure, the Lagerstätten-constrained falsification protocol, the open edges, and the literature anchors. Same epistemic tier as Appendix T-2 (structural correspondence, not closed theorem).

Closure status: STRUCTURAL CORRESPONDENCE (same tier as T-2 / T-3). This appendix is *not* a closed theorem appendix. It records a structural correspondence at the genus level (observer-compatibility filtering under MDL parsimony) applied to a new domain (phylogenetic deep time) using existing OPT apparatus (Stability Filter via §3.1's conditioning event $O_{B,D,T}$; Maintenance Cycle from §3.6; Phenomenal State Tensor $P_\theta(t)$ from §3.5). No new formalism is introduced. The fossil-record predictions (§5 below) are catalogued in §6.8.1 of the preprint as a candidate research programme; F-promotion is gated on three operationalisation steps (effect sizes, null models, OPT-vs-evo-neuro discrimination protocol) — all currently incomplete. The quantitative $R_{\text{req}}(t)$ curve across the late Ediacaran is an open edge per §6 below.

§1. Setup — Substrate-level Filter vs. Within-stream Attractor

The Stability Filter as defined in preprint §3.1 is a *substrate-level* selector: it picks observer-compatible streams from the algorithmic substrate by conditioning the Solomonoff universal semimeasure ξ on the observer-compatibility event

$$O_{B,D,T} := \{x_{1:T} : R_{\mathrm{req}}(x_{1:T}, D) \leq B\}$$

producing the rendered observer distribution

$$\xi_O(x_{1:T}) = \frac{\xi(x_{1:T})\mathbf{1}[x_{1:T} \in O_{B,D,T}]}{\sum_y \xi(y)\mathbf{1}[y \in O_{B,D,T}]}$$

This conditioning operates *once*, on entire patches; its state-space is the set of streams in the substrate.

Within an already-selected stream, additional selection mechanisms operate continuously: Darwinian biological selection on lineages; cultural / institutional selection on civilisations; within-life codec maintenance on individuals (the Maintenance Cycle of §3.6). These are *not* the Stability Filter — they are *within-stream* selectors with their own native mechanisms (variation + heredity + differential survival; cultural transmission; sleep / REM / learning).

The within-stream attractor. Within-stream selectors inherit the observer-compatibility constraint *as an attractor*, not as a literal Filter application. The inheritance happens codec-internally: rendering a past where within-stream selection produces configurations that later fail observer-compatibility requires specifying every intermediate state leading to the discontinuity — a longer description than rendering a past where compatibility is maintained throughout. The Solomonoff-prior weighting on rendered trajectories therefore favours within-stream selection histories that *would have been compatible with the substrate-level Filter even though that Filter has already operated.* Combined with the shared codec architecture across scales (common Markov-blanket structure, active-inference loops, generative-model primitives), selection at any one scale operates in the same description-language as selection at adjacent scales.

T-15 records the phylogenetic instance of this within-stream attractor.

§2. The Brain-First Cascade in OPT Language

Chipman 2026 [109] proposes a *brain-first* reading of the Cambrian Explosion: rising ecological complexity (predation, mobility, sensory differentiation, complex inter-species interaction) from the late Ediacaran onward selected for more sophisticated central nervous systems with deeper internal generative models; once a sufficiently powerful predictive codec evolved, the developmental toolkits that produced regionalised brains (Hox, segmentation, enhancer co-option) were applied to pattern other organ systems, producing the morphological radiation. The “explosion” is not a sudden multiplication of unrelated body plans but a small number of bilaterian lineages with shared foundational neural architecture iteratively refining their generative models, with body plans as downstream implementations.

In OPT language: R_{req} at the lineage level — the predictive load imposed by the ecological niche — rose during the late Ediacaran. Lineages whose codec architecture could not sustain the rising R_{req} within their effective $B_{\max}$ failed (extinction); lineages whose codec architecture could either ride the increase or expand to absorb it persisted (radiation). The brain-first cascade is the structural finding that *codec architecture led*, with morphological radiation as downstream consequence rather than independent driver.

Ontological-status note (load-bearing, repeated from preprint §3.6.9). A lineage is *not* a unified observer-class entity. It does not have a single $B_{\max}$ bottleneck, a global Markov blanket, or an irreducible Phenomenal Residual. When this appendix speaks of "lineage-level R_{req}" or "the lineage codec," the framing is a *population-distributed structural quantity* — the joint predictive load summed (or averaged in some appropriate measure) across the individual observers in the lineage's adapted environment, with each individual carrying its own $B_{\max}$ and Markov blanket. The Maintenance-Cycle correspondence in §3 below is *structural at the genus level*, not literal apparatus-transfer; the swarm-binding reading (Appendix E-6) is **not invoked**.

§3. Lineage-level Maintenance-Cycle Correspondence

Preprint §3.6.9 introduces the four-pass structural correspondence. T-15 expands it with the Narrative-Drift / Decay failure-mode reading at lineage scale.

Within-life Maintenance Cycle	Phylogenetic structural analogue
Pass I — pruning (MDL pressure on K_θ, Eq. T9-3)	Lineage extinction; loss of body plans and behavioural codecs that fail to compress under rising R_{req}. Resource-capacity cost (not K-complexity, per the §2.8 / T-12 channel-independence reformulation pending in the appendix-corrections queue): population size, ecological niche width, metabolic budget, developmental flexibility.

Within-life Maintenance Cycle	Phylogenetic structural analogue
Pass II — consolidation (compression gain, Eq. T9-8)	Co-option of developmental toolkits (Hox, segmentation, regionalised-brain genes) for new structural roles; convergent evolution of efficient predictive architectures. The toolkit reuse *is* MDL consolidation at the developmental-program level: a single conserved gene set, redeployed, compresses the description of multiple body-plan implementations.
Pass III — forward-fan sampling (adversarial branch testing, Eq. T9-11)	Variation + ecological stress-testing across geological time. Surviving branches are those whose codecs satisfy the Stability Filter against actual environmental load; the importance-weighting $w(b) \propto \exp(\beta \lvert E(b) \rvert)$ at within-life timescale has its phylogenetic counterpart in selection pressure concentrating on lineages near niche boundaries or facing acute selective challenges.
Failure mode: Narrative Drift / Decay (Appendix T-12)	Lineage-level Narrative Decay: cryptic loss of substrate-relevant predictive capacity in lineages that overfit to a curated environment (e.g., highly specialised cave-fish lineages losing sensory modalities, or island-isolated lineages losing flight). The within-life Substrate Fidelity Condition (T-12b: redundancy of δ-independent input channels) corresponds at lineage scale to *ecological niche breadth* — lineages with narrow niche fit are vulnerable to Narrative Decay if the niche changes.

The correspondence is structural at the genus level (observer-compatibility filtering under MDL parsimony) — each row maps a Maintenance-Cycle mechanism to a phylogenetic counterpart that operates on a different state-space (lineage vs. individual codec) but follows the same MDL-parsimony logic. The MDL-parsimony argument is constrained to *rule-set phase changes* (preprint §6.6 v3.6.0 hypothesis): smooth gradients under stable rule-sets are perfectly compressible and follow their own native logic, while phase changes between structurally different generative models — chemistry to biology, pre-neural to neural, anaerobic to aerobic codec, gauge-symmetric to broken-electroweak — favour sharp rendering because half-broken intermediate rule-sets are algorithmically expensive.

§4. Body Schema as Conserved Structural Feature

The body schema (Maravita & Iriki [110]; Iriki, Tanaka & Iwamura [111]) — the dynamic, predictive, plastic representation of the body in space — is a Cambrian-era structural feature that has been selected on continuously since. From within OPT, the body schema is the *codec's plastic predictive boundary* — the part of $P_\theta(t)$ that says "what counts as me-acting-on-the-world." Its plasticity is not a modern quirk; it is the conserved structural feature that explains why a human can drive a car as if the wheels were limbs, or why a macaque can incorporate a rake into its peripersonal-space representation within minutes of tool training [111].

Mechanism. The forward/inverse-model framework (Wolpert & Ghahramani [112]) supplies the mechanism: efference copies are routed back into the sensory-prediction layer, allowing the codec to distinguish self-generated from externally-generated sensory consequences and to update the boundary representation in real time when the effector-space changes. Multisensory Bayesian integration — visual + proprioceptive + tactile — provides the inputs the codec uses to maintain the boundary; the rubber-hand illusion (Botvinick & Cohen [113]) is the observational confirmation that the codec's boundary representation is plastic at the millisecond timescale and updates from correlated multisensory input alone.

Structural reason in OPT. The codec is selected to maintain a low-prediction-error generative model of agent-in-environment under bounded R_{req}. A plastic boundary is the bandwidth-efficient solution: rather than re-learning the entire body model when the effector-space changes (which would exceed B_{max} during the transition), the codec maintains a small invariant boundary-representation kernel and updates only the affected limb / tool / vehicle portion. The Cambrian-era predictive-control requirement (mobility, predation, complex appendages) selected for codecs with this plastic-boundary architecture; subsequent vertebrate / mammalian / hominin development inherited and elaborated the same architecture without re-deriving it. Modern tool-use and vehicle-operation are direct phenotypic continuations of a 540-Myr-old codec design choice.

Cross-reference. This is the structural feature that the v0.10 multi-scale codec memo §1.6 identifies as the *conserved Cambrian inheritance* explaining why silicon-emergence is structurally tractable as a substrate extension: a codec with a plastic predictive boundary is already prepared to incorporate non-biological substrate as part of "me-acting-on-the-world," subject to the §8.14 observer-criterion gate.

§5. The Fossil-Record Prediction Structure

T-15 formalises the four-class prediction structure recorded in preprint §6.8.1, with the Lagerstätten-constrained falsification protocol explicit and the F-promotion gate named.

Class (1) — Convergent predictive architectures across distant lineages. OPT predicts that independent evolution of predictive-control architectures should converge on hierarchical predictive structure (centralised CNS with regional differentiation, sensory integration, forward/inverse models, plastic body schemas) rather than on superficial morphology. The empirical observable is comparative neuroanatomy across arthropods (compound eyes + ventral nerve cord), cephalopod molluscs (camera eye + central brain), vertebrates (camera eye + dorsal cord) — three independent lineages with structurally similar predictive-architecture solutions despite vastly different developmental origins. Currently supported by extant convergence in eye evolution (40+ independent origins), centralised nervous systems (multiple independent origins), and behavioural complexity (corvids, cephalopods, primates).

Falsifying observation (Class 1). Sustained finding that successful Cambrian-and-later lineages exhibit *low* predictive-architecture complexity, with morphological diversity unaccompanied by neural sophistication — i.e., that the convergence in predictive architecture is illusory and the apparent pattern is post-hoc selection bias by modern observers.

Class (2) — Neural-fossil sophistication ahead of body-plan complexity (Lagerstätten-constrained, the strongest candidate for F-promotion). The brain-first cascade predicts that, where neural anatomy is preserved, it will appear *earlier* and *more sophisticated* than the body plans of the same lineage would suggest. Currently supported by Ma et al. 2012 [6] *Chengjiangocaris* / *Fuxianhuia* showing remarkably modern arthropod brain organisation in early Cambrian forms, predating the body-plan radiation by an interval consistent with the brain-first cascade's structural prediction.

Falsifying observation (Class 2, Lagerstätten-constrained). Systematic finding *specifically in exceptional soft-tissue preservation sites where neural tissue is actually preserved* (Burgess Shale, Chengjiang biota, Sirius Passet, Maotianshan Shales, and analogous deposits — denote this set $\mathcal{L}$) that neural preservation shows consistently *simpler* brain organisation than the body plans of the same lineages require. The Lagerstätten constraint is structurally important: standard geological decay destroys soft tissue, so a general absence of neural fossils across the bulk of the fossil record reflects *taphonomic bias*, not absence of neurology. The framework is falsified only by *negative findings in sites where positive findings would be possible* — i.e., $\mathcal{L}$-restricted negative findings — which is a stricter and empirically meaningful test that protects the prediction from being falsely falsified by ordinary preservation gaps.

Class (3) — Plasticity and evolvability signatures (expectation only, anthropically loaded). Successful lineages should exhibit developmental plasticity, modularity (Hox toolkit reuse, enhancer co-option), and high "evolvability" — the structural capacity to generate viable variation around a stable neural core. *Why expectation, not falsifier:* the inference from "successful lineage" to "plasticity" is partly anthropic; we observe lineages that are still observable, which is selection on persistence. A proper falsifier here would require a comparative observation across many parallel-evolution events with statistical-significance thresholds that the current empirical base cannot support.

Class (4) — Energy-redundancy trade-offs (expectation only). Codecs require thermodynamic grounding (Landauer / Bennett costs, preprint §3.6). OPT predicts that successful Cambrian and post-Cambrian lineages will show metabolic investment in *redundant* sensory and neural systems rather than minimal-viable designs — the bandwidth-margin necessary to absorb substrate noise without lineage-level Narrative Decay (§3 above). *Why expectation, not falsifier:* the prediction is structurally weaker than (1) or (2); minimal designs may have been outcompeted for non-OPT reasons (e.g., simple metabolic competition).

F-promotion gate (explicit). Promotion of Classes (1) and (2) to an F7 shutdown commitment in preprint §6.8 requires three operationalisation steps:

- *Effect sizes.* Quantitative thresholds for what counts as confirmation vs. falsification. How much brain-fossil sophistication ahead of body-plan complexity? What magnitude of mismatch?
- *Null models.* Explicit null models distinguishing the OPT prediction from the mainstream brain-first evolutionary-biology prediction. OPT-specific signal: sharpness on geological timescales + structural alignment with the multi-scale rapidity pattern of §6.6.
- *Discrimination protocol.* Clear procedure for distinguishing "supports brain-first evolutionary biology generically" from "supports OPT's joint multi-scale prediction specifically."

None of these are in place at v3.6.0; F7 enters preprint §6.8.1 as a candidate research programme.

§6. Open Edges

Quantitative $R_{\text{req}}(t)$ curve across the late Ediacaran. The structural prediction is that R_{req} at the lineage / biosphere level rose during the late Ediacaran into the early Cambrian, driving the brain-first cascade. T-15 does not derive the curve. Two real obstacles compound the difficulty:

- *Fossil-record sparsity.* Pre-Cambrian evidence of past environmental and ecological-interaction state is fragmentary; trace fossils, biomarkers, and isotope proxies constrain rough atmospheric and substrate conditions but not the predictive load any individual lineage faced.
- *No agreed ecological-complexity metric.* There is no consensus quantitative measure of "ecological complexity" in the contemporary biology literature that maps cleanly to OPT's predictive load. Existing diversity, disparity, and trophic-network metrics each capture parts of the relevant signal but none integrate to give R_{req} in bits. A cleaner quantitative treatment requires either an empirical proxy that earns its OPT-mapping or theoretical work bridging information-theoretic ecological complexity to within-stream predictive load.

The precise threshold-crossing condition. §1.5 of the multi-scale codec memo names the MDL-parsimony inequality $K(\text{rule}_1) + K(\text{rule}_2) + K(t_c) + K(\text{switch}) < K(\text{continuous dynamics}) + K(\text{parameters}) + \sum_i K(\text{rule}_{1\to 2,i})$ as the condition under which sharp rendering is favoured. T-15 does not derive the specific threshold values at which the brain-first cascade transitions from gradual to sharp; this requires both the quantitative $R_{\text{req}}(t)$ curve and an evaluation of the K-complexity of intermediate transitional rule-sets in pre-neural to neural codec evolution. Open.

The lineage-level vs. individual-level decoupling. §2 records that "lineage codec" is a population-distributed structural quantity, not a unified macro-observer. The decoupling between codec advance (which happens at the population / clade level) and morphological radiation (which happens at the individual-genome level) is not formally derived in T-15. The structural correspondence in §3 is consistent with this decoupling, but a quantitative model relating the two levels would require population-genetic and developmental-biological apparatus that this appendix does not deploy.

Coordinate with T-12 channel-independence reformulation. T-12 is in the appendix-corrections queue (v0.4 §2.8) for a reformulation of the channel-independence condition: independence of *filtering mechanisms*, not signals. T-15's lineage-level Narrative-Decay correspondence in §3 uses the T-12 framing; if T-12 is reformulated, T-15 §3's Narrative-Decay row will need a parallel update. Coordinate.

§7. Closure Summary

T-15 Deliverables — Structural Correspondence

1. **Substrate-level Filter vs. within-stream attractor distinction (§1).** Recorded in preprint §3.1 (conditioning on $O_{B,D,T}$) and §6.6 (within-stream attractor); T-15 §1 names the phylogenetic instance.
2. **Brain-first cascade in OPT language (§2).** Chipman 2026 [109] reading in R_{req} / $B_{\max}$ vocabulary; lineage-level R_{req} as population-distributed structural quantity, not unified macro-observer.
3. **Lineage-level Maintenance-Cycle structural correspondence (§3).** Four-pass mapping (pruning ↔ extinction; consolidation ↔ co-option; forward-fan ↔ ecological stress-testing; Narrative Drift ↔ lineage-level Narrative Decay). Structural at the genus level, not literal apparatus-transfer.
4. **Body schema as conserved Cambrian inheritance (§4).** Plastic predictive boundary; forward/inverse-model mechanism; structural reason in OPT terms (bandwidth-efficient solution to mobile-agent prediction). Connects to silicon-substrate engagement (§1.6 of the multi-scale codec memo / Appendix E-6).
5. **Four-class prediction structure (§5).** Class (1) convergent predictive architectures; Class (2) neural-fossil sophistication ahead of body-plan complexity (Lagerstätten-constrained); Class (3) plasticity / evolvability

(expectation only); Class (4) energy-redundancy trade-offs (expectation only). F-promotion gate explicit: effect sizes + null models + discrimination protocol.

Falsification status. Catalogued in preprint §6.8.1 as a candidate biological research programme. F-promotion of Classes (1)+(2) to an F7 shutdown commitment is gated on the three operationalisation steps named in §5; none are in place at v3.6.0.

Closure tier. Same as Appendix T-2 (Entropic Gravity) — *structural correspondence*, not closed theorem. The framework's commitment is that the existing apparatus (Stability Filter, Maintenance Cycle, Phenomenal State Tensor) applies to phylogenetic deep time via the within-stream attractor; the empirical question of whether the resulting predictions are confirmed in the fossil record is the candidate research programme catalogued in §6.8.1 and §5 above.

Open edges. Catalogued in §6 — quantitative $R_{\text{req}}(t)$ curve; precise threshold-crossing condition; lineage-vs-individual decoupling; T-12 channel-independence coordination.

This appendix is maintained as part of the OPT project repository alongside opt-theory.md. References: Ma, Hou, Edgecombe & Strausfeld (2012) [6 — reused parked slot]; Chipman (2026) [109]; Maravita & Iriki (2004) [110]; Iriki, Tanaka & Iwamura (1996) [111]; Wolpert & Ghahramani (2000) [112]; Botvinick & Cohen (1998) [113].