

# Histories Become Operators: Biological Notes on Retained Admissibility

Flyxion

July 2026

## Abstract

This essay develops a companion argument to the Admissibility Program's account of continuation, occasioned by a public conversation between Michael Levin and Katrina Schleisman on memory, agency, and learning across biological and artificial systems. Its purpose is not to summarize that conversation but to place it against existing formal apparatus in this corpus: the Ecology of Distinctions' account of memory as preserved recoverability rather than storage, Recursive Continuation's Adaptation Hierarchy, and History Before Function's compression map, read together here under a reinterpretation this essay calls Histories Become Operators. The essay argues for three separations: trace is not storage, constraint is not determination, and robustness is not control, the first of these already established rather than newly derived, with the biological conversation serving as convergent confirmation. It then places two further cases against the Adaptation Hierarchy: a developmental case concerning genomic interpretation, placed cleanly at the state level, and a learning-theoretic case concerning causal emergence, whose level the essay leaves explicitly open, since the source's own warning against conflating parameter-level and rule-level revision applies directly to it and the deciding evidence is not yet public. The essay closes with a sharpened open question, under what conditions learning-induced revision inherits irreversibility from the history that produced it, noted to recur one level up the hierarchy in the Geometry of Witnesses' own unresolved dynamics of second-order repair; a candidate mechanism connecting Hoel's decomposition of causal emergence to the Ecology of Distinctions' Recoverability and Semantic Horizon Theorems and to History Before Function's Anti-Alignment Theorem; and a coda that steps outside the single-system frame, formalized through the Geometry of Witnesses' reconstruction redundancy and articulation-point apparatus, to describe a distinct, population-level case: what a lineage of continuers loses when the direct carriers of a history are mortal and finite.

## 1 Scope and Reading

This essay develops its argument as a reading built from two existing pieces of formal apparatus in this corpus rather than as a report on a single separate prior text. Recursive Continuation's central theorem describes an operator  $F_t$  governing a system's continuation, itself revised in turn by a process  $G$  [10]. The view that this operator is not a primitive object standing apart from a system's history, but history crystallized into operative constraint,

$$H_t \longrightarrow F_t.$$

is not new to this essay. It is History Before Function's Compression map,  $C(H) = F$ , already given existence, non-uniqueness, and uncomputability results there, together with an Anti-Alignment Theorem establishing that the most efficient compression of a history into an operator is generically the worst choice for the operator's later repairability [13]. That essay states explicitly that it is meant to sit beneath Recursive Continuation, supplying one answer to how  $F_t$  arises from  $H_t$  in the first place. What this essay calls Histories Become Operators, in its title, is accordingly not a new interpretive move but a name for applying that existing compression apparatus to biological material the Levin and Schleisman conversation supplies. It is worth noting that Recursive Continuation does not treat its own viability manifold as a free-standing systems-theoretic device either: that essay states outright that recursive continuation and the Ecology of Distinctions' distinction preservation are the same claim viewed from systems theory and from a theory of admissibility respectively [10], which is independent confirmation, from the source itself rather than from this essay's own inference, that the two pieces of apparatus this essay draws on throughout are already understood by their author to be one structure rather than two. On that reading, the state, parameter, and rule modification levels examined below are not primitive tiers standing beside one another. They are different depths at which history has been retained as constraint on the present. What follows is a biological application layer built on that reading, offered as evidence for a single claim: that biological memory is not stored content but retained admissibility, and that organisms differ from current large language models chiefly in that their operators remain revisable within the episode of continuation itself, rather than only between episodes.

Three separations organize the argument. Trace is not storage. Constraint is not determination. Robustness is not control. Each is defended with a specific case rather than a general appeal to biological complexity, on the view that the differences between the cases matter more than their family resemblance.

## 2 Trace Is Not Storage

This separation is not new to the material that follows. The Ecology of Distinctions establishes it directly: memory is not the storage of the past but the preservation of the possibility of recovering it, a distinction with its own apparatus, a Reconstruction Theorem, a Forgetting Theorem, and a Memory Conservation Law, built on a prior Recoverability Theorem stating that a distinction is recoverable exactly when its generating

history remains distinguishable from alternative histories [11]. What the Levin and Schleisman conversation adds is not the thesis but independent, biologically concrete confirmation of it, arrived at from cognitive neuroscience and developmental biology rather than from the distinction-theoretic apparatus that first established it in this corpus. The two cases below are read accordingly, as instances the existing theorem already covers rather than as evidence generating a new claim.

The habit of engineering is to treat memory as a fidelity problem. A bit is written, and the task of the system is to ensure that the same bit is read back later, protected from drift by redundancy and error correction. Meaning, on this model, lives outside the storage medium, in whichever observer assigns significance to the retrieved bit. Levin's argument is that biological memory is built on the opposite premise. A living system never has access to its own past; what it has, in his words, are "the memory traces that your past you has left," and he proposes thinking of them as "messages from your past self" whose meaning must continually be interpreted rather than read out [1]. Because biological substrates are unreliable by nature rather than by defect, the interpretive burden cannot be delegated to an external reader the way it can in a database. The organism is its own interpreter from the first moment, and what evolution has optimized is not the fidelity of the trace but the creativity of the reading; elsewhere Levin has developed this position in writing as the claim that biological memory preserves salience rather than fidelity [2].

The clearest evidence for this claim is not the caterpillar and the butterfly, evocative as that example is, but a plainer one drawn from ordinary developmental biology. The same genome, placed in different developmental circumstances, does not simply express itself twice. Standard embryogenesis and the induced formation of anthrobots begin from materially identical genetic information and diverge into different anatomies, with thousands of genes expressed differently between the two outcomes. Nothing about the genome itself changed. What changed was the interpretive context in which the trace was read. The genome behaves less like a blueprint that specifies an organism and more like a heavily underdetermined score that different performances can realize differently, all coherently, none of them privileged as the correct reading.

Schleisman's contribution sharpens this from the opposite direction. She distinguishes episodic memory, "the cognitive process of storing and creatively constructing memories of specific experiences," from semantic memory, the abstraction of general facts from those experiences, and argues that most present-day language models function as something close to pure semantic memory [1]. In her description, the embedding space of a large language model is "a little bit, if you squint, like a semantic memory model," a compressed, decontextualized store of general knowledge with no equivalent of an episodic layer, no record of what happened to the system, when, and under what particular circumstances [1]. This is not a degraded version of biological memory undergoing lossy reinterpretation. It is closer to the limit case of the same principle approached from the other side. Where biological memory has traces that must be interpreted, current large models often have no traces of that kind at all, only the residue of training compressed into weights that were never indexed to a particular episode in the first place. The rudimentary fix, dumping conversational history into a retrieval store and searching it later, treats the symptom rather than the structure, because it reintroduces storage-style

fidelity into a place where the missing ingredient was never fidelity but indexicality: memory tied to a specific past occasion rather than to a general fact about the world. Schleisman's own algorithmic proposal for closing this gap, developed with Eric Davis, integrates episodic and semantic memory directly within an artificial teammate rather than treating episodic recall as an add-on retrieval layer [3].

### **3 Constraint Is Not Determination**

If trace is not storage, then history cannot function as a lookup table for the present. What it can do, and what both speakers converge on without quite using this language, is narrow the space of continuations available to a system without selecting uniquely among them. The genome again supplies the sharpest case. It rules out an enormous range of possible organisms while remaining compatible with more than one actual one. This is precisely the relationship that continuation geometry is built to describe: a history functions as a constraint on admissible futures, not as a specification of one future in particular. What the conversation adds to that picture, and what is easy to lose if the point is compressed too far, is that the narrowing itself is not a fixed operation performed once on the trace. It is redone, differently, by each continuer that reads the trace under its own local conditions. The constraint is real, and it is also not stationary. Two systems inheriting the same history can experience different admissible futures because the act of reading the constraint is itself a further constraint on how the reading proceeds.

This is the place where an argument of this kind is tempted to flatten Levin's claim into something more comfortable. It would be easy to say that history constrains continuation and stop there, treating the interpreting system as a passive channel through which the constraint operates. Levin's own framing resists that. He insists that the organism is not counting on an external party to say what its memories mean, and that interpretation is itself a locus of agency rather than a mechanical pass through a filter. A passive constraint history and an actively reconstructed history place different demands on a formal account. The former only needs a notion of admissibility fixed in advance by the trace. The latter needs admissibility to be something the continuer partly authors in the act of continuing, which is a stronger and less comfortable claim, and the one the material actually supports.

### **4 Robustness Is Not Control**

The price of building a system that interprets its own traces rather than deferring to an external reader is that the system's behavior stops being fully specifiable in advance. Levin states this without hedging: what is gained is extraordinary robustness and adaptive plasticity, purchased by the fact that one no longer gets to micromanage the outcome. A system built this way is guaranteed to do something coherent with its history, but not guaranteed to do the thing an external designer had in mind. At that point the relationship between designer and system stops resembling the relationship between a

programmer and a program and starts resembling negotiation between agents, the kind of relationship one has with an animal rather than with a machine.

The most quantitatively precise version of this trade-off in the conversation is not about interpretation in the moment-to-moment sense at all, and is kept here as a separate case rather than absorbed into the developmental example. Work from Levin’s laboratory, credited there to Federico Pegozzi, reports a relationship between a system’s learning ability and its degree of causal emergence, a measured quantity describing how much of a multi-scale system’s behavior is attributable to higher-level structure rather than to its parts considered separately [1, 4]. Systems with substantial causal emergence learn well; training further increases their causal emergence; and when such a system is subsequently forced to forget the content it learned, the gain in causal emergence does not reverse. Levin describes the resulting dynamic as “an asymmetric ratchet that points upwards in terms of intelligence and agency” [1]. What survives forgetting is not the content but a structural change in the system’s own capacity to integrate information as a whole. Levin’s own gloss on this result is worth preserving precisely because it resists an easy biological reading: he attributes the effect to the mathematics of networks in general, present even in random networks with no evolutionary history behind them, rather than to any special property of living matter.

This third case does not sit comfortably inside the passive-versus-active distinction that organizes the first two sections, and forcing it in would blur the one result here that is actually citable as quantitative rather than descriptive. The developmental case is a trace being read differently by different continuers under fixed interpretive machinery. The causal-emergence case is the interpretive machinery itself being permanently reshaped by the act of interpreting, independent of whether the specific content that caused the reshaping survives. Both are instances of memory outliving storage, but they are not the same instance.

## 5 The Adaptation Hierarchy and Where These Cases Sit

Recursive Continuation organizes recursive systems into three nested levels, stated there as a theorem with a strict containment  $\mathcal{S} \subsetneq \mathcal{P} \subsetneq \mathcal{R}$  [10].

**Definition 5.1** (Level 1: State Update). *A fixed mechanism  $F$  acts on a changing state,*

$$x_{t+1} = F(x_t).$$

**Definition 5.2** (Level 2: Parameter Update). *The functional form  $F$  is fixed, but a parameter  $\theta_t$  varies within it,*

$$x_{t+1} = F_{\theta_t}(x_t).$$

**Definition 5.3** (Level 3: Rule Update). *The mapping itself, not merely a parameter within it, is what a process  $G$  modifies,*

$$F_{t+1} = G(F_t, x_t, E_t).$$

The source is explicit that Level 2 and Level 3 are routinely conflated, and names this the single most common category error in discussions of recursive self-modification: an ordinary neural network updating its weights by gradient descent is its canonical Level 2 example, since the architecture stays fixed while only the parameters move, whereas architecture search or a self-hosting compiler rewriting its own compilation algorithm are its Level 3 examples, where the functional form itself is what changes [10]. That warning bears directly on how the cases below should be placed, and one of them cannot yet be placed with confidence because of it.

The developmental case belongs cleanly to Level 1. The genome, read as  $F$ , does not change between standard embryogenesis and induced anthrobot formation; what changes is the environment, and the resulting divergence in  $x_{t+1}$  is exactly what a fixed mechanism meeting different conditions predicts. This is a cleaner placement than treating the developmental case as evidence of interpretation operating at the level of the mechanism itself. The variability Levin describes is real, but it is variability in outcome under a fixed rule, not variability in the rule.

This classification treats the developmental mechanism as a fixed operator  $F$  for analytical convenience. A stronger view, associated with Hiesinger, holds that some biological structures may remain inseparable from the developmental histories that generate them, in which case the operator and history are not two distinguishable objects related by  $H_t \rightarrow F$  but a single ongoing computation, so that the operator and history distinction used here becomes an approximation rather than a primitive separation [5]. History Before Function does not resolve this complication, and it is worth being precise about why: that essay's entire apparatus presupposes separability rather than questioning it, requiring  $F$  to be picked out by its extensional behavior before any question of a compressing history becomes meaningful at all [13]. It has a great deal to say about what is lost when a history is compressed into an operator, taken up in Section 7 below, but it does not engage Hiesinger's stronger claim that no such compression occurs in the first place. That claim is set aside here as a genuine unresolved complication rather than one this corpus has already addressed elsewhere.

The causal-emergence ratchet is the case the source's warning applies to directly, and it cannot be placed at Level 3 with the confidence the earlier draft of this essay assumed. If Pegozzi's result concerns synaptic weights adjusting within a fixed circuit topology, ordinary learning of the kind the source's own canonical example describes, it belongs at Level 2, not Level 3, whatever the resulting causal-emergence gain does to the network's integrative behavior. It would belong at Level 3 only if what changes is the topology itself, which connections exist, not merely their strengths. The transcript does not settle this, and the unpublished status of Pegozzi's manuscript means this essay cannot settle it either. What can still be said without resolving the level is narrower than the earlier version of this section claimed: whichever level the revision occurs at, the revision does not symmetrically reverse when the state-level content that produced it is subsequently erased. The asymmetry is the finding; which level of the hierarchy it occurs at is now marked as open rather than asserted.

**Observation 5.1.** *The episodic gap Schleisman identifies in current large language models is a Level 1 description of inference with the revision boundary placed outside the episode entirely.*

*Such a system at inference time has an active state, its activations, and a fixed parameter set  $\theta$ ; whatever Level 2 revision produced that  $\theta$  happened during training, external to and decoupled from any episode the system later undergoes. Episodic memory and online parameter or rule revision are analytically distinct: the former preserves particular occasions of continuation, while the latter allows those occasions to modify the mechanism of continuation itself, and current large language models largely lack both capacities, but neither logically entails the other. The missing episodic layer is therefore not merely missing content. It is the absence of online revision machinery operating on the timescale of the system's own lived episodes, whatever level that revision would occur at, and that no amount of appended conversational history can supply within a Level-1-only architecture.*

Read as a biological application layer of the reading this essay develops, the three cases now separate by depth with one held deliberately open rather than three separating cleanly. The developmental case shows Level 1 variability,  $x_{t+1} = F(x_t, E_t)$ , history retained only as far down as the state. The language-model case shows Level 1 operation during the episode with no online revision at any higher level, history retained only up to the boundary of training and not within continuation itself. The causal-emergence case shows revision, at Level 2 or Level 3 depending on a fact this essay does not have access to, history retained at least one level deeper than the state, as a change to the mechanism rather than to its output, with the precise depth of that change left as the section's genuine unresolved point rather than its conclusion.

## 6 An Open Question

State, parameter, and rule levels are already available within the existing apparatus, with biological and artificial instances already assigned to each, and  $H_t \rightarrow F$  already frames the levels as depths of the same retention rather than as separate primitives. The question this essay leaves open is accordingly not whether new formal apparatus is needed.

**Open Question 6.1.** *Under what conditions does  $G$  inherit irreversibility from  $H_t$ ?*

The causal-emergence result shows a case in which learning-induced revision becomes path-asymmetric: the gain in the network's causal structure survives the erasure of the  $x_t$  that produced it, rather than remaining merely updateable in either direction. Whether that asymmetry is a general property of revision at Level 2, derivable from the same mathematics that already establishes the Adaptation Hierarchy, a property specific to Level 3, or a further condition that would need to be added to the theorem's statement to hold generally rather than only in the specific networks examined so far, is not settled by the material this essay draws on, compounded by the fact that which level the Pegozzi result even occupies is itself unresolved. It is the actual research question the essay has reached rather than resolved, and it is left here as a question rather than forced into a premature answer.

This question is not confined to the present essay. The Geometry of Witnesses lists among its own open problems a dynamic theory of second-order repair, treating reconstruction as a time-dependent process  $W_t : \mathcal{H}_t \rightarrow \mathcal{R}_t$ , and asks explicitly whether

such a formulation would reveal critical transitions, attractors, or hysteresis within reconstructive systems [12]. That is the same question this essay asks about  $G$  and  $H_t$ , posed one level up the hierarchy, about  $W$  and the histories that feed it, and left equally unresolved there. Whether  $F_t$  in Recursive Continuation's sense and an element of the repair repertoire  $\mathcal{R}$  in the Geometry of Witnesses' sense are formally the same kind of object is not established here; the parallel is structural, not a proven reduction of one open problem to the other. What can be said is narrower and still worth saying: no part of this corpus has yet answered this question at any level of the hierarchy, first-order or second-order, biological or archival, which is some evidence that the difficulty is real rather than an artifact of how this particular essay framed it.

## 7 A Candidate Mechanism

Hoel's formalization of causal emergence gives the open question a place to attach rather than remaining purely descriptive. Effective information, the quantity whose macro-scale increase constitutes causal emergence, decomposes as determinism minus degeneracy [6, 7]: a system's causal structure is more effective the more reliably a given state leads to a specific successor, and less effective the more different states are collapsed onto the same successor. Causal emergence at a macro scale, on this account, arises when a coarse-graining either sharpens determinism or reduces degeneracy relative to the micro description.

**Conjecture 7.1.**  *$G$  inherits irreversibility from  $H_t$  when the change  $H_t$  induces in  $F_t$  is predominantly degeneracy-reducing rather than purely determinism-increasing. Degeneracy reduction is structurally identical to coarse-graining in the thermodynamic sense: a system can compress forward, collapsing distinguishable micro-trajectories onto a shared macro-consequence, but recovering the fine-grained partition from the coarse one is generically non-unique, because the information needed to invert it is exactly what the compression removed.*

Whether determinism increases are correspondingly more reversible is left outside the conjecture proper, as an open possibility rather than a claim it depends on. Determinism increases add specificity to existing pathways without erasing distinctions between them, which suggests they might be easier to undo, but nothing argued here establishes that, and the degeneracy side of the conjecture carries its explanatory weight independently of whether the determinism side turns out to hold.

This conjecture has independent support beyond the analogy to coarse-graining. Recent work reformulating causal emergence through the singular value decomposition of a system's Markov dynamics shows that effective information is closely tied to a formal measure of dynamical reversibility, with causal emergence arising specifically from redundant, irreversible information pathways within the dynamics [8]. On that account, irreversibility is not an incidental feature that happens to accompany causal emergence in some systems; it is close to definitional of what causal emergence measures. If the Pegozzi ratchet is a genuine instance of causal emergence in Hoel's sense, an expectation of path-asymmetry is not an additional empirical surprise the result would need to

establish separately. It is closer to what the underlying mathematics would already predict.

The coarse-graining analogy the conjecture leans on is, moreover, not merely an analogy. The Ecology of Distinctions proves a Recoverability Theorem stating that a distinction is recoverable exactly when its generating history remains distinguishable from alternative histories, and a Semantic Horizon Theorem showing that recoverability falls as the inverse of the preimage size under repeated projection, reaching zero as distinction classes merge, while the merged structure's causal influence persists regardless [11]. A preimage growing under projection is the same object, described in a different formalism, as degeneracy growing under coarse-graining: both measure how many distinguishable prior configurations have collapsed onto one subsequent one. Read against that theorem, the conjecture above is not proposing a new mechanism so much as asking whether a specific empirical case, Pegozzi's networks, crosses an already-defined semantic horizon.

A third anchor, closer to home, sharpens why the compression involved should be expected to cost something specific rather than something merely diffuse. History Before Function proves an Anti-Alignment Theorem for its own compression map  $C(H) = F$ : over the lattice of retained traces consistent with a given operator's behavior, the most minimal retention weakly maximizes repair entropy, the log-count of fault locations consistent with a symptom, and does so strictly whenever the discarded surplus would have excluded some fault hypothesis [13]. Applied to the causal-emergence case, this gives a reason, independent of both Hoel's and the Ecology of Distinctions' formalisms, for expecting degeneracy-reducing compression to be costly in a specific way: not merely that the pre-merge structure becomes unrecoverable in general, but that what is disproportionately lost is exactly the diagnostic detail that would let a later process tell which of several possible pre-training configurations actually produced the current one. Erasing  $x_t$  cannot restore that detail because efficient compression, by the theorem, was never likely to have retained it in the first place.

Three caveats keep this a conjecture rather than a result. First, the equivalence between causal emergence and dynamical reversibility establishes that causal emergence and irreversibility are mathematically entangled in general; it does not by itself establish that content-level forgetting, erasing  $x_t$ , is incapable of restoring the specific pre-merge degeneracy in a trained network, which is the sharper claim the conjecture needs. Second, none of this has been checked against Pegozzi's actual result, which remains unpublished. Third, and specific to the connections just drawn, Hoel's degeneracy, the Ecology of Distinctions' preimage-size recoverability measure, and History Before Function's repair entropy are structurally parallel to one another, not shown here to be formally identical; establishing those correspondences precisely, rather than asserting them by resemblance, is itself unfinished work. What this section adds is a mechanism with three citable anchors, one external and mathematical, two internal and already proven, not a derivation of the biological case; whether the anchors and the case actually meet is the next thing to establish, not something this essay can yet claim.

## 8 A Coda: Operator Succession Across Continuers

Everything argued in Sections 3 through 7 shares a frame that is worth making explicit now that it is about to be left. The three cases placed in the Adaptation Hierarchy, genomic interpretation, the episodic gap, and the causal-emergence ratchet, all answer the same question: for one system running its own continuation, at what depth is history retained, as a change to the state, to the mechanism, or to the process that revises the mechanism. What follows is a different question, and it does not belong in that table. It concerns what happens to  $H_t \rightarrow F_t$  not within a single continuer but across a population of them, when the direct carriers of  $H_t$  are mortal and finite.

The case is drawn from a monologue by a commentator publishing as The Functional Melancholic, who argues that a generation holding direct, embodied memory of a pre-digital information environment occupies a position no later generation can occupy, not because the earlier environment was better, but because comparison against it is only available to those who ran it themselves [9]. The distinction drawn there between testimony and nostalgia is doing real work: nostalgia is a soft-focus longing for a remembered past, while testimony is precise comparative access, knowing with some specificity what a prior configuration was actually like and what replaced it. That distinction maps onto a genuine asymmetry between two kinds of continuer. One inherits  $H_t$  directly, having run the history itself. The other inherits only  $H'_t$ , a compressed, already-interpreted account of  $H_t$  produced by an earlier continuer, with no way to check  $H'_t$  against  $H_t$ , because the only witnesses capable of making that check are the ones no longer available to make it. This is not the interpretive variability of the genome case, where the trace persists and different continuers read it differently. Here the trace itself, in its uncompressed form, does not survive its carriers.

This is not, on reflection, best described by a theorem about ordinary distinctions, since what is at stake is not a fact but a capacity, the capacity to interpret, which is closer to what the Geometry of Witnesses calls a repair operator than to a distinction in the ordinary sense. That essay defines a second-order repair operator  $W : \mathcal{H} \rightarrow \mathcal{R}$  mapping histories to the repair capacities they can reconstruct, and a reconstruction redundancy  $\rho(r) = |\mathcal{C}(r)|$ , the number of independent historical pathways from which a given capacity  $r$  can be regenerated [12]. It is explicit that redundancy requires independence rather than mere multiplicity: many copies of the same testimony, sharing the same interpretive limitations, do not increase  $\rho$  the way genuinely independent pathways would, a caveat the coda inherits directly, since a generation's many witnesses are independent only to the degree that their individual experiences of the pre-digital environment actually differed. The essay further distinguishes ordinary repair failure, recoverable by reconstruction, from the case  $\rho(r) = 1$ , a single remaining reconstructive pathway, which it treats as a qualitative rather than quantitative boundary: once that pathway is destroyed, no higher-order restoration remains available, and recovery is replaced by irreversibility, a phenomenon it likens explicitly to an articulation point in a graph, a vertex whose removal disconnects otherwise robust regions. A generation of firsthand witnesses approaching extinction is, on this reading, a repair capacity approaching  $\rho \rightarrow 1$ , and what the coda describes is the crossing of exactly that boundary, not a milder version of

ordinary forgetting.

**Observation 8.1.** *A continuer that has personally operated under two distinct operator regimes within a single continuous existence, in the case described, a pre-digital and a digital information environment inhabited by the same nervous system, has a comparative vantage that a continuer operating under only one regime cannot have by any amount of documentary reconstruction. Comparison of this kind requires having run both operators, not merely having been told about one of them through  $H'_t$ . The claim that the capacity to compare configurations is also the capacity to evaluate them, drawn from the same source, follows directly: evaluation of a regime requires an outside vantage on it, and a continuer who has only ever operated under the current regime has no such outside from which to evaluate it.*

This section is offered as a coda rather than a fourth case integrated into Section 5, for two reasons that should stay visible rather than be smoothed over. First, it is population-level where Sections 5 through 7 are single-system, and forcing it into the Adaptation Hierarchy's table would either not fit cleanly or would bend the table to accommodate it, the outcome that section was built to avoid. Its formal home is the Geometry of Witnesses' second-order repair hierarchy rather than Recursive Continuation's Adaptation Hierarchy, a different piece of existing apparatus rather than a gap in this corpus, which corrects an earlier draft of this coda that treated the case as still needing formal treatment. Second, its evidentiary register remains mixed even so: the biological illustration is a monologue by a commentator publishing pseudonymously, not a peer-reviewed or otherwise independently checkable source in the way Levin, Schleisman, Hoel, and Hiesinger are, and it is cited here as the origin of an example worth formalizing, while the formalization itself now rests on already-proven results rather than on that source's authority. What this coda establishes is that operator succession under interpreter extinction is a real case distinct from the three already placed, and one already treated by name, as a monopoly in reconstruction redundancy, in the Geometry of Witnesses.

## References

- [1] Levin, M. and Schleisman, K. *Memory, Agency, and Learning in Biological and AI Systems*. Recorded conversation, 2026. <https://www.youtube.com/watch?v=H1-n0fSLh-M>
- [2] Levin, M. Self-Improvising Memory: A Perspective on Memories as Agential, Dynamically Reinterpreting Cognitive Glue. *Entropy*, 26(6):481, 2024.
- [3] Davis, E. and Schleisman, K. Episodic Memory for Autonomous Agents. Manuscript in preparation, described in [1].
- [4] Pegozzi, F. Learning ability and causal emergence in network systems. Work from the Levin Laboratory, Tufts University, described in [1]; manuscript forthcoming.
- [5] Hiesinger, P. R. *The Self-Assembling Brain: How Neural Networks Grow Smarter*. Princeton University Press, Princeton, 2021. ISBN 9780691181226.
- [6] Hoel, E. P., Albantakis, L., and Tononi, G. Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795, 2013.
- [7] Hoel, E. P. When the map is better than the territory. *Entropy*, 19(5):188, 2017.
- [8] Zhang, J., Tao, R., Leong, K. H., Yang, M., and Yuan, B. Dynamical reversibility and a new theory of causal emergence based on SVD. *npj Complexity*, 2(1):3, 2025.
- [9] The Functional Melancholic. *Gen X: The Last Generation to Live Offline*. Video monologue, 2026. <https://youtu.be/I5tA04qsKsI>
- [10] Flyxion. *Recursive Continuation: Autopoiesis, Symptoiesis, and the Compression of Self-Modification*. Admissibility Program, 2026. <https://standardgalactic.github.io/alphabet/example/recursive-continuation.pdf>
- [11] Flyxion. *The Ecology of Distinctions*. Admissibility Program. [https://standardgalactic.github.io/spherepop/textbook/The\\_Ecology\\_of\\_Distinctions.pdf](https://standardgalactic.github.io/spherepop/textbook/The_Ecology_of_Distinctions.pdf)
- [12] Flyxion. *The Geometry of Witnesses: Second-Order Repair and the Maintenance of Admissibility*. Admissibility Program, June 2026. [https://standardgalactic.github.io/library/workspace/geometry\\_of\\_witnesses.pdf](https://standardgalactic.github.io/library/workspace/geometry_of_witnesses.pdf)
- [13] Flyxion. *History Before Function: Compression as the Origin of Operators*. Admissibility Program, 2026. <https://standardgalactic.github.io/alphabet/example/history-before-function.pdf>