

Selective Retention and the Geometry of Continuation

Flyxion

July 2026

Two earlier pieces approached the same unresolved question from progressively narrower angles and, in retrospect, from the wrong direction each time. The first treated it as a complaint about programming languages: state-primitive design in C++ and Swift accumulates historical complexity, and a history-primitive calculus like Spherepop was offered as a counterexample showing that much of that complexity is contingent rather than essential. The second treated it as a compression problem: Spherepop's Collapse operator was formalized as a rate-distortion optimization, trading storage cost against the expected cost of failing some future query, with ordinary mutable state and unbounded history retention falling out as limiting cases of a single Lagrangian. Both pieces were correct as far as they went, and both stayed inside a frame that made the question look smaller than it is. The question was never really about programming languages, and it was never really about compression in the information-theoretic sense either. It is a question that already has a name in the parts of this research program that have nothing to do with software: which distinctions must survive a transformation in order for a system's future continuations to remain admissible. Spherepop's Collapse, Git's commit graph, a consolidated memory trace, and a published scientific record are five different substrates answering the same question, and treating any one of them as the primary case obscures rather than clarifies what the question actually is.

1 Continuation as the Unit of Value

A history has no value in itself. Neither does a compressed representation of one. What has value is a system's continuation capacity: the space of admissible future trajectories still reachable from where it currently stands, evaluated with respect to whatever tasks, decisions, or reconstructions that system will actually be called on to perform. A retention policy, whatever substrate it operates on, is good exactly to the degree that it preserves the distinctions those future continuations will depend on, and bad exactly to the degree that it discards them, regardless of how much or how little it happens to store in absolute terms. This reframing matters because the earlier framing in terms of state versus history invited a category error: it made retention sound like a question about which representation is more fundamental, ontologically, when it is actually a question about which distinctions are functionally load-bearing for a specific population of futures. A state-primitive system that happens to retain exactly the distinctions its continuations require is not doing anything wrong by being state-primitive. A history-primitive system that retains distinctions no future continuation will ever need is not doing anything virtuous by being history-primitive; it is merely paying a storage cost for a form of safety it does not actually possess, since the distinctions it failed to identify as relevant are exactly as likely to be lost in an unhelpful compressed summary somewhere downstream as they would have been in a state-primitive system that never kept them at all.

2 From Duality to Triad

The natural correction is to stop asking whether history or state is the more fundamental substrate and to ask instead what mediates between them. History is not valuable because it is history; it is valuable because it is the raw material from which continuation-relevant distinctions can, in principle, still be recovered. Compression is not harmful because it discards information; it is harmful exactly when, and only when, it discards distinctions that some future continuation will actually require. The right picture is therefore not a duality between two competing primitives but a triad: history feeds into a compression step, and the compression step is answerable not to some abstract storage budget but to the continuations it must remain capable of supporting.

$$\text{History} \longrightarrow \text{Compression} \longrightarrow \text{Continuation}$$

Under this picture, the earlier essays' entire debate over whether computation is fundamentally the evolution of state or the accumulation of history was a dispute about which end of the triad to name first, not a dispute about what actually determines whether a system's design is any good. What determines that is the middle arrow, and specifically whether the compression it performs is calibrated to the continuations on the right.

3 The Continuation Deficit

This can be made precise using machinery this research program has already developed for a different purpose: the admissibility distortion of a projection, defined as the probability that a projection collapses together two states a given task needs to keep distinct, and the sharper notion of an ontological deficit, in which a distinction is not merely collapsed for particular inputs but is inexpressible in the target representation altogether, invisible from inside the system that has lost it. Both transfer directly to the temporal case.

Definition 1 (Continuation deficit). *Let H be a history, let $C = \varphi(H)$ be a compressed representation of it under some retention policy φ , and let π be a distribution over the population of future continuations a system will actually be asked to support, where each continuation τ determines a set $\text{Adm}_\tau \subseteq H \times H$ of history-pairs that must remain distinguishable for τ to remain admissible. The continuation deficit of φ is*

$$\delta(H, C) = \mathbb{E}_{\tau \sim \pi} \left[\Pr_{(h, h') \in \text{Adm}_\tau} [\varphi(h) = \varphi(h')] \right],$$

the expected, continuation-weighted probability that φ has destroyed a distinction some future continuation actually needed.

Perfect retention gives $\delta = 0$ trivially, since no distinction is ever collapsed. A retention policy calibrated correctly to π can also give $\delta \approx 0$ while discarding almost everything, provided what it discards is continuation-irrelevant — this is the case explored formally in the previous essay's worked example, where a running balance alone suffices whenever π places all its weight on queries about the current total. Catastrophic forgetting corresponds to large δ : a compression that happens to destroy exactly the distinctions a high-probability, high-stakes continuation depends on.

Definition 2 (Ontological deficit, temporal case). *A retention policy φ has an ontological deficit with respect to a continuation τ if no state of the compressed representation's own vocabulary corresponds to the distinction Adm_τ requires — the distinction is not merely collapsed for some inputs but structurally inexpressible in C , regardless of which inputs are considered.*

The distinction between a continuation deficit and an ontological deficit is not a technicality. A nonzero $\delta(H, C)$ under an assumed π is an ordinary, priced risk: the designer of φ knew the trade-off and accepted some probability of failure in exchange for a smaller representation. An ontological deficit is different in kind, because it is invisible from within the system that has it: a compressed representation with an ontological deficit cannot, using its own resources, detect that anything has been lost, since the vocabulary needed even to notice the missing distinction was itself discarded along with the distinction. This is the sharper, more accurate name for what the earlier essay gestured at when it observed that a retention policy's real exposure is to a future query that falls entirely outside the support of π : that exposure is not merely a probability the model failed to price. It is very often an ontological deficit the model could not have priced, because the compressed representation it produced has no internal way of representing the shape of its own blind spot.

4 Case Study: Version Control

Git is worth returning to not because it is a programming tool but because it is the cleanest available instance of a system whose designers, without any of this vocabulary, correctly identified a narrow and specific continuation population and built a retention policy calibrated tightly to it. The continuations Git actually needs to support are few and well understood: reverting to a prior state, identifying which change introduced a regression, reconciling divergent lines of development, and attributing a change to its author and moment. For exactly that continuation population, Git retains full fidelity — every committed snapshot, every parent relationship, every author and timestamp — while discarding, correctly and without apology, everything that happened between commits: the keystroke-level editing history within a single uncommitted change is gone the moment the commit is made, because no continuation in Git's actual population depends on it. The continuation deficit of Git's retention policy is close to zero for the population it was designed around and undefined, rather than merely nonzero, for continuations outside that population, such as reconstructing exactly how a developer arrived at a particular line of code keystroke by keystroke — a real ontological deficit, not a priced risk, since Git's object model has no representation in which that distinction could even be expressed.

5 Case Study: Event Sourcing, Audit, and Legal Archive

Event-sourced architectures, database write-ahead logs, and legal or regulatory archives sit further along the same axis because their continuation population includes rare but severe events that Git's does not: a compliance audit, a legal dispute, a forensic reconstruction of exactly what a system believed and when. The retention policies these systems adopt — keep full transaction detail, keep signed and timestamped records, keep chain-of-custody metadata that no ordinary application would bother with — are not more virtuous than Git's leaner policy. They are calibrated to a different π , one in which the cost of an ontological deficit at audit time is not merely inconvenient but can be existential for the institution involved, which changes where the optimal trade-off between storage and continuation deficit sits without changing the shape of the underlying problem at all.

6 Case Study: Biological Memory

Memory consolidation supplies an instance of the same structure in a substrate that was never designed by anyone. Newly formed memories are initially hippocampus-dependent, rich in incidental contextual detail, and vulnerable to disruption; over time, systems consolidation transforms many of them into a more schematic, cortex-dependent form that is more stable but has demonstrably lost some of the original episodic specificity, a transformation whose signature in patients with selective hippocampal damage is precisely a pattern of retained schematic memory alongside lost episodic detail [2]. What determines whether a given consolidated memory still supports a given future continuation is not how much detail survived in some absolute sense but whether a retrieval cue relevant to that continuation can still reactivate a usable trace, a process classical memory theory calls *ecphory*: retrieval as the joint product of a cue and a stored engram, neither of which alone determines what is recovered [1]. A consolidated memory with a large continuation deficit relative to some cue population is functionally indistinguishable, for that population, from a memory that was never formed at all, regardless of how much of the original trace survives in principle — which is the same lesson Git’s discarded keystrokes and a compliance archive’s retained signatures teach in a completely different substrate: what matters is never the raw quantity of retained material but whether the surviving structure remains addressable by the cues, or queries, or continuations that will actually arrive.

7 Case Study: Scientific Records and Replication

A published scientific result is a retention policy applied to an underlying history of observation, and its continuation population includes, centrally, independent replication: a future researcher attempting to reproduce the result from what the publication chose to retain. The reproducibility difficulties documented across empirical psychology over the last decade are, on this reading, evidence of a systematic continuation deficit in what the field’s publication norms treated as sufficient to retain — methods sections compressed past the point where an independent team could reliably reconstruct the original procedure, effect sizes reported without the raw data or analytic choices that would let a later continuation distinguish a real effect from a researcher-degree-of-freedom artifact [5]. This is not a claim that any individual scientist acted in bad faith; it is a claim that a field-wide retention policy, evolved under incentives that priced brevity and novelty far more highly than continuation capacity, produced exactly the large- δ compressions this framework would predict, and that the subsequent push toward preregistration, open data, and open methods is legible as a field correcting its π after the fact, in exactly the sense the earlier essay’s closing objection anticipated but did not name.

8 Case Study: Spherepop, Demoted

Spherepop’s Collapse operator, formalized in the previous essay as a rate–distortion-optimal compression of an event history, is simply the continuation deficit above stated in a different but fully compatible vocabulary, with the earlier essay’s distortion term $D(\varphi)$ and this essay’s $\delta(H, C)$ measuring the same quantity under different names, related by treating the earlier essay’s query population as this essay’s continuation population and its per-query cost function as an instance of Adm_7 -violation [6]. What has changed is not the mathematics but the essay’s own sense of where the interesting content lives. Spherepop is not the origin of this problem and does not deserve to be treated as its natural home; it is one clean, formally tractable instance of a much older and more

general question, useful precisely because its four primitives make the continuation deficit easy to state precisely, not because computation is somehow the right domain in which to first understand it.

9 Why State Versus History Was the Wrong Frame

Once retention is understood as a question about which distinctions a compression protects rather than about which substrate a system is built on, the state-versus-history debate that organized the first essay in this sequence dissolves rather than resolves. A state-primitive language is not conceptually impoverished for compressing aggressively; it is impoverished only in the specific cases where its compression discards a distinction some actual continuation needed, which is a claim about a mismatch between φ and π , not a claim about state as such. A history-primitive system is not conceptually superior for retaining everything; retaining continuation-irrelevant distinctions is waste, not safety, since a distinction nobody has identified as relevant is exactly as likely to go unexamined in an undifferentiated historical record as it would have been if it were never recorded, the record's mere existence doing no work unless something eventually queries it correctly. The earlier essay's fifth objection asked why history-native architectures keep appearing at the systems level while languages remain state-centric, and offered a layering explanation that was correct as far as it went. The sharper answer available now is that the ontology was never the load-bearing variable. What varies between a loop counter and a compliance ledger is not whether the underlying computation is "really" state or "really" history but the shape of π each one actually faces, and a good design in either idiom is one whose φ tracks that shape, while a bad design in either idiom is one whose φ does not.

10 What Remains Open

This reframing does not resolve the problem the earlier essay left open; it relocates it to where it was always going to have to be solved. π , the distribution over future continuations a retention policy is calibrated against, is not observable in advance in any of the five case studies above — not for Git's designers, not for a compliance officer setting a retention window, not for the nervous system consolidating a memory, not for a field setting publication norms, and not for an implementation of SpheroPop's Collapse choosing what to compress. Estimating π , revising it as evidence accumulates that it was wrong, and detecting an ontological deficit from inside a system that by definition cannot see what it is missing, are not problems this essay or its predecessors have solved. They are, if anything, closer to the actual center of gravity this research program has been circling under other names all along — what must be preserved for repair to remain possible, which compressions leave the admissible manifold and which do not, what a system can and cannot recover once a given distinction has been let go — and the honest conclusion is that this essay has not introduced a new problem so much as it has finally stopped mistaking a programming-language dispute for the more general one underneath it.

11 Postscript: A Fourth Shift

Read as a sequence, the three pieces this argument has passed through do not just accumulate; each one exposes an assumption the previous one was quietly resting on. The first took history's value as given and asked only when it should be collapsed. The second took the future query population

π as given and asked only how to compress against it optimally. This essay has taken continuation capacity itself as the thing worth protecting, and treated π as an input rather than a further question — which is exactly the assumption this postscript wants to flag rather than resolve. π is doing a great deal of work throughout this essay, and it does not come from nowhere. A continuation is not relevant because it is a fact about the future; it is relevant because some system, at some point, values it enough to have built a retention policy around protecting it. Git’s designers valued revertibility and blame-attribution over keystroke-level provenance. A compliance regime values auditability over storage economy. A nervous system values whatever consolidation was shaped by selection pressure to value, which is not obviously the same thing an individual would choose on reflection if asked. Asking why continuation τ_1 is weighted more heavily than continuation τ_2 in some system’s π is no longer a question this essay’s vocabulary can answer, because it is not a question about memory, compression, or retention at all. It is a question about preference. The chain implied by that observation runs the opposite direction from how this trilogy was written: preferences determine which futures matter, which futures determine which distinctions matter, which distinctions determine what must be retained, and only then does retention determine what may be safely compressed. That ordering reconnects this entire line of argument to the parts of this research program concerned directly with preference fields and admissibility, where the question of why one reachable future is weighted over another was already the primary object of study rather than an assumed input. Whether π itself is estimable, learnable, or must instead be treated as a further primitive on the same footing as continuation is not a question this essay can answer, but it is now, at least, the visible next assumption to interrogate, and continuation deficit and ontological deficit both look like special cases of a broader question about preference over reachable futures once the chain is followed that far back.

References

- [1] Endel Tulving. *Elements of Episodic Memory*. Oxford University Press, 1983.
- [2] Larry R. Squire and Pablo Alvarez. Retrograde Amnesia and Memory Consolidation: A Neurobiological Perspective. *Current Opinion in Neurobiology*, 5(2):169–177, 1995.
- [3] Scott Chacon and Ben Straub. *Pro Git*, 2nd edition. Apress, 2014.
- [4] Martin Fowler. Event Sourcing. *martinfowler.com*, 2005.
- [5] Open Science Collaboration. Estimating the Reproducibility of Psychological Science. *Science*, 349(6251):aac4716, 2015.
- [6] Claude E. Shannon. A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(3):379–423, 1948.