Evolutionary Selection for Atemporal Memory Storage

Why Three Convergent Pressures Favor Architectures Where Time Belongs to Retrieval, Not to Storage

Boris Kriger[1]

[1]Institute of Integrative and Interdisciplinary Research,
`boriskriger@interdisciplinary-institute.org`

Abstract

This paper argues that natural selection *strongly favors* atemporal memory storage—architecture where temporal ordering emerges through retrieval operations rather than being intrinsic to stored states. We identify three convergent selection pressures: (1) **Resource pressure**: temporal indexing as structure imposes costs that temporal information as content avoids; (2) **Speed pressure**: associative retrieval supports faster adaptive response than temporal lookup; (3) **Flexibility pressure**: atemporal storage enables adaptive recombination that temporal binding constrains.

We formalize these pressures through a minimal competing-architectures model, deriving conditions under which atemporal storage is evolutionarily stable. We operationalize a "complexity threshold" as the ratio $C = |\mathcal{P}|/(B \cdot \tau)$ of decision-space size to verification capacity, and show that as C increases, the fitness advantage of atemporal architecture grows while any advantage of temporal accuracy saturates.

Calibrated claim: We demonstrate that atemporal storage has *strong, compounding selective advantages* making it the *expected evolutionary outcome* for complex systems. We do *not* claim logical necessity; temporal architectures remain possible but face cumulative selective disadvantage that increases with complexity. Where precise temporal encoding exists (time cells, interval timing), it requires specialized mechanisms—consistent with a default architecture that lacks intrinsic temporal structure.

This framework reframes "design flaws" in biological memory as signatures of an architecture that selection convergently produces across lineages.

Keywords: evolutionary epistemology, memory architecture, natural selection, resource constraints, complexity threshold, temporal coding, reconstructive retrieval

1 Introduction

1.1 The Puzzle of "Imperfect" Memory

Biological memory exhibits features that seem suboptimal from an engineering standpoint: temporal distortions are pervasive [Friedman, 1993]; retrieval modifies stored content [Nader et al., 2000]; memories contaminate each other [Loftus, 1979]; chronological order is inferred rather than directly accessed [Friedman, 2004].

Why didn't evolution produce accurate temporal indexing? This paper argues that selection pressures *favor* these features—they are signatures of an architecture that natural selection produces, not tolerates.

1.2 The Thesis: Strong Selective Advantage

Three independent pressures converge on atemporal storage:

1. **Resource pressure**: Atemporal architectures have lower metabolic costs
2. **Speed pressure**: Associative retrieval enables faster adaptive response
3. **Flexibility pressure**: Atemporal storage permits beneficial recombination

Each pressure independently favors atemporal storage; together, they create cumulative selection.

1.3 Calibrating the Claim

Previous versions of this argument claimed atemporal storage is "evolutionarily inevitable." Reviewers correctly noted this overstates what the arguments support. We therefore calibrate:

What we claim:

- Atemporal storage has *selective advantages* that compound with complexity
- These advantages make atemporal architecture the *expected outcome*
- The claim generates *testable predictions*

What we do not claim:

- That atemporal storage is *logically necessary*—temporal architecture remains possible
- That evolution always reaches optima—path dependence, constraints, and drift can prevent it
- That all temporal encoding is absent—specialized systems evolve for specific functions

The difference between "strongly favored" and "inevitable" matters for scientific rigor. We argue the former.

1.4 Key Definitions

Definition 1 (Temporal Structure vs. Temporal Content)**.** ***Temporal structure****: Time is a coordinate axis of storage—each memory has a temporal position, like records in a timestamped database.*

Temporal content*: Temporal information is encoded as features within states ("summer 2020," "before graduation") without temporal coordinates.*

Definition 2 (Atemporal Storage)**.** *Memory architecture where stored states have identity independent of encoding time, associations form via content similarity, and temporal ordering is constructed during retrieval from contextual features.*

1.5 Engagement with Temporal Coding Literature

This paper must address findings that might seem to contradict the thesis:

Time cells [MacDonald et al., 2011]: Hippocampal neurons fire at specific temporal intervals during delay periods. We argue these are *specialized additions* to a default atemporal architecture—their existence as dedicated mechanisms suggests temporal encoding is *not* automatic.

Temporal Context Model [Howard and Kahana, 2002]: TCM encodes temporal relations through gradually drifting context. Critically, TCM encodes temporal *proximity* (context similarity), not temporal *position* (timestamps). This is temporal-as-content, not temporal-as-structure—consistent with our thesis.

Phase coding [Buzsaki, 2006]: Oscillatory phase can encode temporal/sequential information. We address whether this constitutes "cheap" temporal structure in Section 3.

2 Three Selection Pressures

2.1 Pressure 1: Resource Costs

Background: Brains consume ∼20% of metabolic resources while constituting ∼2% of body mass [Raichle and Gusnard, 2002]. Any architectural feature reducing neural overhead, holding function constant, confers selective advantage.

Proposition 1 (Cost Differential—Qualified). *Temporal structure imposes costs that temporal content partially avoids:*

- ***Timestamp generation***: *Creating temporal markers at encoding*
- ***Ordering maintenance***: *Updating temporal relations as new memories form*
- ***Index infrastructure***: *Supporting temporal queries ("what happened at time t?")*

Temporal content requires encoding temporal features only when task-relevant, with no global ordering maintenance.

Reviewer concern: "What if temporal structure is simply another dimension in associative space with negligible extra cost?"

Response: This is an empirical question. We do not claim temporal structure is prohibitively expensive, only that temporal content is *cheaper*. The magnitude of the cost differential affects the *strength* of selection, not its *direction*. Even small cost differentials, compounded over evolutionary time, produce directional selection.

Testable prediction: If the cost differential is negligible, organisms under strong resource pressure (small body size, high metabolic rate) should show the same temporal architecture as resource-rich organisms. If significant, resource-constrained organisms should show greater reliance on temporal-as-content.

2.2 Pressure 2: Retrieval Speed

Background: In environments where delayed response is costly, faster retrieval has selective advantage [Cisek, 2019].

Proposition 2 (Speed Differential—Qualified). *Most adaptive queries are content-based: "Is this dangerous?" "Is this edible?" "Have I encountered this before?" These require matching current input to stored patterns—similarity computation.*

Temporal queries ("When did I last see this?") are typically secondary, requiring:

1. *Retrieve relevant content (the encounter)*
2. *Access temporal information associated with that content*

If temporal information is encoded as content features (contextual associations), it is retrieved with the content. If encoded as structure, a separate temporal lookup is required.

Reviewer concern: "Efficient indexing can perform temporal queries in $O(\log n)$ time."

Response: True for computational systems with designed index structures. The question is whether evolution produces such structures for general memory. Where precise temporal encoding is fitness-critical (interval timing, sequence learning), specialized mechanisms *do* evolve—suggesting the default architecture lacks them.

2.3 Pressure 3: Adaptive Flexibility

Background: Variable environments reward recombining past experience to generate novel responses [Schacter et al., 2012].

Proposition 3 (Flexibility-Accuracy Trade-off). *Temporal structure constrains recombination: combining elements from different times creates temporal incoherence if time is a structural coordinate.*

Atemporal storage permits free recombination. Temporal information, as content, can be retained, modified, or omitted.

Reviewer concern: "Uncontrolled recombination produces maladaptive confabulation."

Response: Correct. The claim is not that all recombination is beneficial, but that *the capacity for recombination* has net positive expected fitness in variable environments. False memories are the cost of a system enabling creative imagination. Evolution calibrates the accuracy-flexibility trade-off, not maximizes either end.

Empirical support: Memory and imagination share neural substrates [Schacter et al., 2012]—consistent with an architecture enabling both, rather than separate systems for accurate replay and creative construction.

3 Formal Model: Competing Architectures Under Selection

To move beyond verbal arguments, we formalize the competition between temporal and atemporal architectures.

3.1 Setup

Let A_T (temporal architecture) and A_A (atemporal architecture) be competing memory designs. Let $C = |\mathcal{P}|/(B \cdot \tau)$ be the **complexity ratio**: decision-space size divided by verification capacity (bandwidth × time).

3.2 Fitness Components

Define fitness components (relative, not absolute):

$$\begin{aligned} W_{\text{resource}}(A) &= -k_r \cdot \text{Cost}(A) && (1) \\ W_{\text{speed}}(A) &= -k_s \cdot \text{Latency}(A) && (2) \\ W_{\text{flex}}(A) &= k_f \cdot \text{Recombination}(A) && (3) \\ W_{\text{accuracy}}(A) &= k_a \cdot \text{TemporalPrecision}(A) && (4) \end{aligned}$$

where $k_r, k_s, k_f, k_a > 0$ are environment-dependent weights.

Total fitness: $W(A) = W_{\text{resource}} + W_{\text{speed}} + W_{\text{flex}} + W_{\text{accuracy}}$

3.3 Architecture Properties

Temporal architecture A_T:

- $\text{Cost}(A_T) = c_0 + c_1 \cdot n$ (base cost plus per-memory overhead for temporal indexing)
- $\text{Latency}(A_T) = \ell_0 + \ell_1 \cdot \log n$ (content retrieval plus temporal lookup)
- $\text{Recombination}(A_T) = r_0$ (limited by temporal binding)
- $\text{TemporalPrecision}(A_T) = p_{\max}$ (high precision)

Atemporal architecture A_A:

- $\text{Cost}(A_A) = c_0$ (no temporal indexing overhead)
- $\text{Latency}(A_A) = \ell_0$ (content retrieval only)
- $\text{Recombination}(A_A) = r_{\max}$ (unconstrained)
- $\text{TemporalPrecision}(A_A) = p_0 < p_{\max}$ (lower precision, reconstructed)

3.4 Fitness Differential

The fitness advantage of atemporal over temporal architecture:

$$\begin{aligned} \Delta W &= W(A_A) - W(A_T) \\ &= k_r \cdot c_1 \cdot n + k_s \cdot \ell_1 \cdot \log n + k_f \cdot (r_{\max} - r_0) - k_a \cdot (p_{\max} - p_0) \end{aligned} \tag{5}$$

3.5 Key Result

Theorem 1 (Scaling of Fitness Differential). *As system complexity increases (larger n, larger C):*

1. *The resource advantage of A_A grows as $O(n)$*
2. *The speed advantage grows as $O(\log n)$*
3. *The flexibility advantage is constant (architecture-dependent, not n-dependent)*
4. *The accuracy advantage of A_T is constant (bounded by $p_{\max} - p_0$)*

Therefore, $\Delta W \to +\infty$ as $n \to \infty$: the fitness advantage of atemporal architecture grows without bound, while the advantage of temporal accuracy is bounded.

Proof. From Equation 5: The first two terms grow with n; the last two are constant in n. For sufficiently large n, $\Delta W > 0$. □

3.6 The Complexity Threshold

Definition 3 (Complexity Threshold C^*). *The complexity threshold C^* is the value of C (or equivalently n) at which $\Delta W = 0$:*

$$C^*: \quad k_r \cdot c_1 \cdot n^* + k_s \cdot \ell_1 \cdot \log n^* = k_a \cdot (p_{\max} - p_0) - k_f \cdot (r_{\max} - r_0) \tag{6}$$

For $C > C^*$, atemporal architecture has positive fitness advantage. For $C < C^*$, temporal architecture may be favored (if accuracy benefits exceed costs).

Operationalization: C^* depends on environment-specific weights (k_r, k_s, k_f, k_a) and implementation-specific parameters (c_1, ℓ_1, r_0, p_0). We cannot specify a universal numerical threshold, but we can make comparative predictions:

- Organisms with larger n (more memory states) should show stronger atemporal signatures
- Environments where k_a is high (temporal precision is fitness-critical) should show more specialized temporal encoding
- Within lineages, increasing complexity should correlate with increasing temporal flexibility

3.7 Can Temporal Architecture Compensate?

Reviewer concern: "What if temporal structure can be implemented cheaply enough that costs are negligible?"

Response: Suppose we grant $c_1 \to 0$ (negligible cost per memory). Then:

$$\Delta W = k_s \cdot \ell_1 \cdot \log n + k_f \cdot (r_{\max} - r_0) - k_a \cdot (p_{\max} - p_0) \tag{7}$$

The speed advantage still grows with n. Only if *both* $c_1 \to 0$ and $\ell_1 \to 0$ (costless *and* instant temporal retrieval) does the accuracy term dominate.

But $\ell_1 \to 0$ requires that temporal lookup adds no latency—that temporal position is as immediately accessible as content. This would require temporal structure so deeply integrated that the distinction between "structure" and "content" dissolves.

Conclusion: The model shows that temporal architecture can compensate only if temporal indexing is both costless and latency-free. These are strong conditions. Under realistic assumptions, atemporal architecture maintains fitness advantage for complex systems.

3.8 Phase Coding and "Cheap" Temporal Structure

Does oscillatory phase coding [Buzsaki, 2006] provide cheap temporal structure?

Phase coding encodes *sequential position within a cycle* (theta phase precession) or *relative timing* (gamma phase). It does not provide absolute timestamps or global temporal ordering across all memories.

In our framework, phase coding is temporal-as-content (local temporal relations encoded as features) rather than temporal-as-structure (global temporal coordinates). It is consistent with, not contrary to, the atemporal thesis.

4 Empirical Predictions and Evidence

4.1 Predictions

1. **Cross-species**: Cognitive complexity should correlate with temporal flexibility. More complex organisms should show more reconstructive, less veridical temporal memory.
2. **Specialized mechanisms**: Where temporal precision is fitness-critical (interval timing, sequence learning, food caching), specialized encoding mechanisms should exist *in addition to* general memory—not as features of general memory itself.
3. **Resource correlation**: Under controlled conditions, metabolic cost of memory should not scale linearly with temporal precision. Temporal precision should require additional specialized investment.
4. **Retrieval signatures**: Temporal judgments should show inference signatures (variable RT, confidence patterns) rather than direct-access signatures.
5. **Developmental**: Temporal accuracy should not improve monotonically with development; it should reflect maturation of reconstruction strategies.
6. **Flexibility-accuracy trade-off**: Within populations, individuals with higher temporal flexibility should show greater imaginative/creative capacity; those with higher temporal accuracy should show reduced flexibility.

4.2 Existing Evidence

Time cells as specialized addition: Time cells exist in hippocampus for specific timing tasks [MacDonald et al., 2011]. Their existence as a dedicated mechanism supports the thesis: if general memory were temporally structured, specialized time cells would be unnecessary.

Food-caching birds: Scrub jays show episodic-like memory for what-where-when of caches [Clayton and Dickinson, 1998]. This demonstrates temporal encoding *can* evolve when fitness-critical. Critically, this appears to be a specialized system, not evidence that all avian memory is temporally structured.

Temporal Context Model fit: TCM successfully models temporal memory phenomena using context drift, not timestamps [Howard and Kahana, 2002]. This supports temporal-as-content architecture.

Shared memory-imagination substrates: Memory and future imagination share neural mechanisms [Schacter et al., 2012]—consistent with an architecture enabling recombination rather than accurate replay.

5 Objections and Replies

Objection 1: The claim is unfalsifiable.

Reply: The claim generates specific predictions (Section 5.1). It would be falsified if: complex organisms showed temporally-structured general memory without specialized mechanisms; temporal precision correlated with cognitive complexity rather than inversely; metabolic costs of memory scaled with temporal precision in general (not specialized) systems.

Objection 2: Reconsolidation and false memory are equally predicted by imperfect temporal indexing.

Reply: Correct that these phenomena are multiply realizable. The argument is not that only atemporal architecture predicts these phenomena, but that atemporal architecture *explains why* they occur—not as failures of a system "trying" to be accurate, but as consequences of an architecture optimized for other objectives. The explanatory framing differs even if the phenomena are compatible with both architectures.

Objection 3: Evolution doesn't always reach optima.

Reply: Correct, and this is why we claim "strongly favored" rather than "inevitable." Path dependence, developmental constraints, and drift can prevent optimal outcomes. The claim is that atemporal architecture is an *attractor*—selection pushes toward it—not that all lineages reach it.

Objection 4: AI convergence reflects engineering, not selection.

Reply: AI systems face analogous pressures: compute cost (resource), latency (speed), generalization (flexibility). Market competition and benchmark performance create selection-like dynamics. The parallel is imperfect (human design vs. blind variation) but the convergence on similar solutions under similar pressures is suggestive.

Objection 5: The complexity threshold C^* is not quantified.

Reply: Acknowledged as a limitation. C^* depends on environment-specific and implementation-specific parameters that vary across lineages. We provide qualitative predictions (comparative, not absolute) that are testable without numerical threshold specification. Full operationalization requires empirical measurement of the relevant parameters.

6 Implications

6.1 Reframing "Design Flaws"

Memory features traditionally viewed as limitations are reframed as consequences of an architecture that selection favors:

- **Temporal distortions**: Signatures of reconstructive retrieval, not indexing failures
- **Reconsolidation**: Mechanism enabling adaptive updating, not storage vulnerability
- **False memories**: Cost of flexibility that enables imagination

6.2 For AI Design

AI memory architectures (transformers with positional encoding, vector databases, RAG) exhibit similar properties. Positional encoding in transformers is temporal-as-content (position features added to representations), not temporal-as-structure (global temporal coordinates).

"Hallucination" in generative AI parallels false memory in biology: the cost of an architecture enabling generalization. The goal is calibration (knowing when to trust reconstructions), not elimination (which would eliminate useful generalization).

6.3 Open Questions

- What are the actual metabolic costs of temporal indexing in neural tissue?
- How does the complexity threshold vary across lineages and environments?
- Can the flexibility-accuracy trade-off be directly measured in individuals?
- What developmental trajectory produces the observed balance between accuracy and flexibility?

7 Conclusion

This paper argues that natural selection *strongly favors* atemporal memory storage due to three convergent pressures: resource costs, retrieval speed, and adaptive flexibility. A minimal formal model shows that the fitness advantage of atemporal architecture grows with system complexity, while the advantage of temporal accuracy is bounded.

We claim "strongly favored," not "inevitable": temporal architecture remains possible but faces cumulative selective disadvantage. Where temporal precision is fitness-critical, specialized mechanisms evolve—consistent with a default architecture that lacks intrinsic temporal structure.

The central insight: **Evolution puts time in retrieval, not in storage, because that is where time is cheap.** Storing temporal structure is costly, slow, and rigid. Constructing temporal order during retrieval is cheaper, faster, and more flexible. Under the pressures that complex adaptive systems face, selection converges on the latter.

For complex memory systems, asking "why isn't temporal memory more accurate?" is like asking "why don't birds have solid bones?" The question presupposes accuracy (or strength) is the optimization target. But evolution optimizes fitness, not accuracy. And fitness, under resource, speed, and flexibility pressures, is maximized by architectures where time belongs to the use of memory, not to memory itself.

Acknowledgments

The author thanks reviewers whose critiques substantially improved this paper, particularly regarding the distinction between "strongly favored" and "inevitable," the need for formal modeling of competing architectures, and engagement with temporal coding literature.

References

Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press.

Cisek, P. (2019). Resynthesizing behavior through phylogenetic refinement. *Attention, Perception, & Psychophysics*, 81(7), 2265–2287.

Clayton, N. S., & Dickinson, A. (1998). Episodic-like memory during cache recovery by scrub jays. *Nature*, 395(6699), 272–274.

Friedman, W. J. (1993). Memory for the time of past events. *Psychological Bulletin*, 113(1), 44–66.

Friedman, W. J. (2004). Time in autobiographical memory. *Social Cognition*, 22(5), 591–605.

Howard, M. W., & Kahana, M. J. (2002). A distributed representation of temporal context. *Journal of Mathematical Psychology*, 46(3), 269–299.

Loftus, E. F. (1979). *Eyewitness Testimony*. Harvard University Press.

MacDonald, C. J., Lepage, K. Q., Eden, U. T., & Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, 71(4), 737–749.

Nader, K., Schafe, G. E., & Le Doux, J. E. (2000). Fear memories require protein synthesis in the amygdala for reconsolidation after retrieval. *Nature*, 406(6797), 722–726.

Raichle, M. E., & Gusnard, D. A. (2002). Appraising the brain's energy budget. *Proceedings of the National Academy of Sciences*, 99(16), 10237–10239.

Schacter, D. L., Addis, D. R., Hassabis, D., Martin, V. C., Spreng, R. N., & Szpunar, K. K. (2012). The future of memory: Remembering, imagining, and the brain. *Neuron*, 76(4), 677–694.