

Ecphoric Synthesis and Event-Historical Algebra

A Formal Framework for Persistent Knowledge, Memory, and Meaning

Flyxion

January 22, 2026

Abstract

This work develops a formal algebraic framework for ecphoric synthesis: the process by which coherent present understanding is reconstructed from irreversible event histories without erasure or overwriting. Unlike state-based or optimization-driven models of knowledge, the proposed framework treats events, relations, and meta-relations as first-class entities embedded in time. A synthesis operator is defined that produces a present-view as a constraint-respecting projection over history rather than as a mutable state. The formalism is designed to support causality, analogy, provenance, and contradiction while remaining compatible with merge-based computational systems and biologically plausible models of cognition. The goal is to provide a rigorous substrate for persistence-first reasoning across cognitive, institutional, and computational domains.

1 Introduction

Most contemporary knowledge systems, both computational and institutional, are built around the notion of mutable state. Information is summarized, overwritten, or replaced as new data arrives, producing systems optimized for speed, legibility, and short-term efficiency. While such architectures are tractable, they exhibit systematic failure modes: loss of provenance, inability to reconcile contradiction, brittleness under delayed feedback, and degradation of long-term coherence. These failures are not incidental; they arise from the assumption that correctness lies in the present state rather than in the preserved structure of history.

Biological cognition offers a contrasting model. Human memory does not operate by overwriting past experiences. Instead, experience accumulates as an irreversible history, and present understanding is reconstructed through a process of synthesis that integrates partial, noisy, and sometimes conflicting traces. This reconstructive process is commonly described in psychology as *ecphory*. Here, the term *ecphoric synthesis* is used to emphasize that the relevant operation is not retrieval of stored representations, but the active integration of event histories into a context-sensitive present view.

The objective of this paper is to formalize *ecphoric synthesis* as an algebra operating over event-historical structures. Rather than proposing another symbolic knowledge representation or probabilistic inference engine, the aim is to define a minimal and stable substrate in which meaning, causality, and understanding can persist under time, contradiction, and partial information. This requires a shift away from graphs in which relations are mere edges and toward a unified structure in which relations themselves are entities capable of temporal extension, provenance tracking, and recursive qualification.

The formalism developed here is motivated by several converging considerations. From distributed systems, it draws on event sourcing and append-only logs, where correctness emerges from replay rather than mutation. From cognitive science, it adopts the reconstructive nature of memory and the tolerance of contradiction until contextual resolution. From epistemology, it rejects the notion that new knowledge invalidates old knowledge by default, instead treating supersession as a structured relation. From computational practice, it aligns with merge-based algebras such as those explored in *Spherepop*, where histories are combined prior to interpretation rather than resolved at write time.

2 Event-Historical Substrate

We begin by defining the event-historical substrate on which the algebra operates. Let T be a totally ordered set representing time. Time is assumed to be irreversible in the sense that for any $t_1, t_2 \in T$, if $t_1 < t_2$ then t_2 does not precede t_1 . No further metric or continuity assumptions are required.

An event is defined as a tuple

$$e = (id_e, t_e, \pi_e),$$

where id_e is a unique identifier, $t_e \in T$ is the event's occurrence time, and π_e is a payload drawn from some domain of observations, actions, declarations, or measurements. Events are atomic and immutable. Once an event occurs, it is never deleted, modified, or retracted. Events are not propositions and are not evaluated as true or false; they merely record that something happened.

The set of all events is denoted by E . A strict temporal ordering relation \prec_t is induced on E such that $e_i \prec_t e_j$ if and only if $t_{e_i} < t_{e_j}$. This ordering is partial only to the extent that events may share identical timestamps.

3 Relation-Nodes

Relations are introduced not as edges but as first-class entities. A relation-node is defined as a tuple

$$r = (id_r, \tau_r, A_r, I_r, \rho_r, \mu_r),$$

where id_r is a unique identifier, τ_r is the relation type, A_r is an ordered tuple of arguments drawn from $E \cup R$, $I_r \subseteq T$ is a temporal validity interval, ρ_r denotes provenance information, and μ_r is an arbitrary metadata structure.

The argument tuple A_r allows relations to connect events, other relations, or mixtures of both. This recursive structure permits the representation of statements about statements without introducing a separate logical layer. Relation types τ_r are drawn from a typed hierarchy that may include causal, temporal, evidential, analogical, conditional, or normative categories. The typing system constrains admissible argument structures and prevents category errors but does not impose semantic closure.

Temporal validity intervals I_r specify when a relation is considered active. A relation may be permanently valid, conditionally valid, or valid only over a finite interval. Provenance ρ_r encodes the source of the relation, whether observational, inferential, testimonial, or institutional. Metadata μ_r may include confidence measures, explanatory notes, or computational

annotations.

The set of all relation-nodes is denoted by R . A dependency ordering \prec_d is defined such that $r_i \prec_d r_j$ if r_j depends structurally on r_i through argument inclusion.

4 Meta-Relations and Typing Discipline

Meta-relations are relations whose arguments include other relations. They encode precedence, override conditions, contradiction markers, analogical mappings, and inference control. Formally, meta-relations are not a separate category but a subset of R distinguished by their type and argument signatures.

To avoid paradoxes, the typing discipline enforces acyclicity in certain dependency dimensions. In particular, no relation may directly or indirectly assert precedence over itself under the same temporal scope. This constraint preserves well-foundedness while allowing deep recursive structure.

5 Historical Structure

The complete historical structure is defined as the typed triple

$$H = (E, R, \prec_t),$$

with \prec_d implicitly derivable from R . Importantly, H is not a sequence but a structured, partially ordered set. New information is incorporated into H only by extension: adding new events or relations. No operation removes elements from H .

At this stage, H contains no notion of a current state. All interpretations of meaning, truth, or relevance are deferred to the synthesis operator, which computes views over H without altering it.

6 The Synthesis Operator

The central operation of the framework is the synthesis operator, denoted by Σ , which computes a coherent present-view from the event-historical structure without modifying that structure. Formally, let V denote the space of present-views. A present-view is not a state in the conventional sense but a structured projection that selects and interprets portions of history as active, relevant, or binding at a given time. The synthesis operator is defined as

a function

$$\Sigma : H \times T \rightarrow V,$$

where H is the complete event-historical structure and T is the temporal domain.

For any $t \in T$, the present-view $V_t = \Sigma(H, t)$ consists of a triple

$$V_t = (E_t, R_t, C_t),$$

where $E_t \subseteq E$ is the set of events deemed relevant at time t , $R_t \subseteq R$ is the set of relations active at time t , and C_t is a context structure encoding constraints, scopes, and interpretive parameters. The precise membership of these sets is determined algorithmically rather than declaratively.

The synthesis operator proceeds by first identifying the temporally admissible relations. A relation-node $r \in R$ is temporally admissible at time t if and only if $t \in I_r$. Let $R_t^{(0)}$ denote the set of all such relations. This step enforces temporal integrity by ensuring that relations are considered only within their declared intervals of validity.

The second stage applies meta-relational constraints. Let $M \subseteq R$ denote the subset of relations whose type encodes precedence, override, contradiction, or scope restriction. These relations are evaluated over $R_t^{(0)}$ to produce a refined set $R_t^{(1)}$ in which relations disabled or superseded by higher-precedence meta-relations are excluded. Importantly, exclusion here is local to the present-view; disabled relations remain elements of R and may re-enter future views if conditions change.

In the third stage, dependency closure is computed. Let \prec_d be the dependency ordering induced by argument inclusion. The synthesis operator computes the transitive closure of causal and inferential chains within $R_t^{(1)}$, yielding a set $R_t^{(2)}$ that is closed under admissible composition rules. This closure step allows implicit consequences to become explicit within the present-view without introducing new elements into H .

The fourth stage resolves contradictions. A contradiction is defined as the simultaneous activation of two relations whose types and arguments are marked as mutually exclusive by a contradiction meta-relation. Resolution strategies are not global but encoded explicitly as relations within H . The synthesis operator applies these strategies deterministically, ensuring that for a fixed H and t , the output V_t is unique.

Finally, the set E_t is derived as the set of events that appear as arguments in $R_t^{(2)}$ or are otherwise marked as contextually salient by relations in C_t . The context structure C_t aggregates active constraints, scopes, and confidence measures, providing the interpretive background against which E_t and R_t are understood.

7 Well-Definedness and Basic Properties

The synthesis operator is well-defined under mild assumptions. Termination follows from the finiteness of admissible relations at any finite time and from the acyclicity constraints imposed by the typing discipline. Determinism is guaranteed by the requirement that all conflict resolution strategies be explicitly represented as relations in H . Given identical histories and identical times, Σ produces identical present-views.

Monotonicity holds in the following sense: if $H_1 \subseteq H_2$, then for any t , the present-view computed from H_2 may refine or qualify the view computed from H_1 , but it cannot invalidate historical facts present in H_1 . Continuity in time is piecewise constant, with discontinuities occurring only at event times or at the boundaries of relation validity intervals.

The computational complexity of Σ depends on the size of $R_t^{(0)}$ and the depth of dependency chains. In the worst case, synthesis is polynomial in the number of active relations, though practical implementations may employ indexing and incremental updates to achieve near-linear performance for streaming histories.

8 Consistency and Constraint Satisfaction

Consistency within a present-view is defined relative to an explicit constraint set. Let $C_t = \{c_1, \dots, c_n\}$ be a set of predicates on V_t . A present-view is consistent if and only if all predicates in C_t evaluate to true. Constraints may encode logical coherence, ecological limits, institutional rules, or cognitive load bounds.

The synthesis operator is constraint-respecting by construction. Constraints are represented as relations or meta-relations and therefore participate directly in the resolution process. However, the framework does not guarantee global satisfiability. There exist histories for which no present-view satisfies all constraints simultaneously. In such cases, inconsistency is represented explicitly rather than concealed, preserving diagnostic information.

9 Analogy, Causality, and Inference

Analogy arises naturally in this framework because relations themselves are objects. An analogical relation maps the argument structure of one relation-node to that of another, preserving role correspondence rather than identity. This permits structural comparison across domains without semantic collapse.

Causality is treated as a specialized relation type constrained by temporal ordering and counterfactual sensitivity. Because causal claims are relations rather than rules, they may be

revised, scoped, or overridden without rewriting the events they connect. Inference proceeds by the controlled composition of relations under the governance of meta-relations, allowing deductive, inductive, and abductive patterns to coexist.

10 Biological and Cognitive Interpretation

The proposed algebra mirrors several well-established features of biological cognition. Events correspond to episodic traces, relations to associative assemblies, and synthesis to attractor dynamics constrained by prior structure. Memory reconsolidation is modeled not as event modification but as the introduction of new relations that qualify or reinterpret existing ones.

Ecphoric synthesis explains empirical phenomena such as context-dependent recall, false memories, and the persistence of outdated beliefs. It predicts that recall difficulty scales with the density and conflict of relational structure rather than with raw event count, a hypothesis amenable to experimental testing.

11 Compatibility with Merge-Based Systems

Because histories are append-only and interpretation is deferred, the framework is inherently compatible with merge-based computation. Merging two histories corresponds to set union over events and relations, with conflicts represented explicitly rather than resolved prematurely. Synthesis is applied after merging, ensuring that no information is lost.

This property aligns the algebra with version control systems, distributed ledgers, and merge-oriented programming models such as Spheredrop. It also supports multi-agent knowledge construction in which divergent perspectives coexist until sufficient structure accumulates to support reconciliation.

12 Conclusion

Ecphoric synthesis provides a formal alternative to state-centric models of knowledge. By grounding meaning in event-historical structure and by elevating relations to first-class status, the framework supports persistence, causality, and interpretability without sacrificing rigor. Its emphasis on constraint-respecting reconstruction rather than optimization offers a path toward systems that remain coherent under time, contradiction, and scale.

The algebra presented here is intentionally minimal. Its purpose is not to exhaust the space of possible representations but to define a stable substrate upon which richer cognitive, institutional, and computational systems may be built.

13 Applications to Institutional and Computational Systems

The event-historical algebra described above has direct implications for the design of institutions and computational infrastructures. In governance systems, policies and regulations are typically amended through textual replacement, obscuring the rationale for change and erasing historical context. Under an event-historical model, legislative acts are events, interpretations are relations, and amendments are meta-relations that qualify or supersede prior relations without deleting them. This allows the present legal view to be synthesized dynamically while preserving a complete audit trail of normative evolution.

In scientific knowledge production, experimental results, hypotheses, and theoretical commitments can be represented as events and relations with explicit provenance and confidence. Retractions, paradigm shifts, and methodological critiques appear not as deletions but as higher-order relations that constrain the applicability of earlier claims. This structure accommodates historical continuity in science while avoiding the false impression that superseded theories were simply erroneous rather than locally coherent under prior constraints.

Medical records offer another domain where persistence is critical. Diagnoses, symptoms, and test results are naturally event-like, while diagnostic criteria and treatment guidelines function as relations that change over time. Encoding revisions as meta-relations prevents loss of clinical history and supports safer decision-making by exposing how present recommendations emerged from prior interpretations.

In artificial intelligence and alignment research, the framework reframes learning as relation refinement over behavioral events rather than as reward maximization over states. Value learning becomes the progressive qualification of evaluative relations, corrigibility becomes explicit override relations on prior inferences, and interpretability emerges from provenance chains embedded directly in the relational structure. Alignment verification can then be posed as constraint satisfaction over synthesized present-views rather than as post hoc behavioral auditing.

14 Relation to Existing Formalisms

The proposed algebra intersects with several existing traditions while differing in fundamental orientation. Event sourcing in distributed systems shares the commitment to append-only logs but typically treats interpretation as external to the log itself. Here, interpretive structure is encoded within the same substrate as the events, allowing synthesis to be a first-class operation.

Temporal logics such as linear temporal logic and computation tree logic formalize reasoning over time but presuppose propositional states rather than event-historical reconstruction. Belief revision theories, including AGM frameworks, address inconsistency but rely on prioritized belief sets that obscure provenance. By contrast, the present framework embeds revision as explicit relational structure.

Causal graphs and Bayesian networks represent dependencies but collapse time and revision into parameter updates. Version control systems preserve history but delegate semantic interpretation to human users. The event-historical algebra unifies these perspectives by combining irreversibility, relational expressivity, and algorithmic synthesis within a single formal substrate.

15 Limitations and Open Problems

Despite its advantages, the framework faces significant challenges. Computational tractability becomes an issue as histories grow large and relational density increases. While incremental synthesis and indexing strategies mitigate this problem, worst-case complexity remains substantial. Determining optimal synthesis algorithms for large-scale, distributed histories remains an open research question.

Another limitation concerns the selection and typing of relations. Although the algebra permits rich expressivity, inappropriate or inconsistent typing can lead to opaque or unstable present-views. Automated discovery of relation types and constraints, particularly from data, poses both technical and philosophical difficulties.

Meta-relational depth raises the possibility of regress, in which relations about relations proliferate without bound. While typing and acyclicity constraints prevent logical paradox, practical heuristics are required to manage interpretive complexity.

Finally, privacy and security concerns arise in shared histories. Because persistence is a feature rather than a flaw, mechanisms are required to scope visibility, redact sensitive payloads, or cryptographically restrict access without violating the algebras irreversibility guarantees.

16 Future Directions

Several extensions suggest themselves. Formal convergence results for multi-agent histories could clarify under what conditions shared present-views emerge. Integration with probabilistic reasoning may allow uncertainty to be propagated more naturally through relational structures. Approximate synthesis operators could enable real-time applications under resource constraints.

On the cognitive side, empirical studies could test predictions about recall difficulty, context dependence, and reconstruction latency as functions of relational structure. Developmental models might explore how new relation types emerge through abstraction and how skill acquisition corresponds to progressive refinement of dependency relations.

On the computational side, reference implementations and query languages are needed to operationalize the framework. Such systems would allow users to interrogate histories, request synthesized views, and explore counterfactual scenarios without sacrificing provenance.

17 Conclusion

This work has presented a formal, event-historical algebra for ephoric synthesis, grounded in irreversibility, relational expressivity, and constraint-respecting reconstruction. By rejecting destructive state updates and elevating relations to first-class entities, the framework provides a stable substrate for reasoning under time, contradiction, and partial information.

Rather than optimizing for immediate efficiency, the algebra prioritizes persistence, auditability, and coherence. In doing so, it aligns computational systems more closely with biological cognition and with the requirements of long-lived institutions. Ephoric synthesis, formalized in this way, offers not merely a technical alternative but a foundational reorientation of how knowledge, memory, and meaning are represented over time.

Appendices

A Appendix A: Formal Definitions and Notation

This appendix presents the core mathematical definitions used throughout the paper in a precise and self-contained manner.

A.1 Time and Order Structure

Let (T, \leq) be a totally ordered set representing time. The order relation \leq is assumed to be antisymmetric, transitive, and total. No assumption of continuity or metric structure is imposed. For any $t_1, t_2 \in T$, the strict order $t_1 < t_2$ denotes $t_1 \leq t_2$ and $t_1 \neq t_2$.

A.2 Events

An event is defined as a triple

$$e := (id_e, t_e, \pi_e),$$

where $id_e \in \mathcal{I}_E$ is a unique identifier drawn from a countable identifier set, $t_e \in T$ is the event timestamp, and $\pi_e \in \Pi$ is an event payload. The payload space Π is left intentionally abstract and may encode observations, measurements, actions, declarations, or internal signals.

Let E denote the set of all events. Events are immutable. There exists no operation within the algebra that removes or modifies elements of E .

A temporal precedence relation $\prec_t \subset E \times E$ is induced by timestamps such that

$$e_i \prec_t e_j \iff t_{e_i} < t_{e_j}.$$

A.3 Relations

A relation-node is defined as a six-tuple

$$r := (id_r, \tau_r, A_r, I_r, \rho_r, \mu_r),$$

where $id_r \in \mathcal{I}_R$ is a unique identifier, $\tau_r \in \mathcal{T}$ is a relation type drawn from a typed relation ontology, $A_r \in (E \cup R)^k$ is an ordered argument tuple of finite arity $k \geq 1$, $I_r \subseteq T$ is a temporal validity interval, $\rho_r \in \mathcal{P}$ is a provenance descriptor, and $\mu_r \in \mathcal{M}$ is a metadata structure.

Let R denote the set of all relations. Relations are immutable once introduced. Temporal validity does not imply deletion; it constrains admissibility during synthesis.

A.4 Dependency Ordering

A dependency relation $\prec_d \subset R \times R$ is defined such that

$$r_i \prec_d r_j \iff r_i \in A_{r_j}.$$

This relation induces a directed acyclic graph over R under the typing constraints described below.

A.5 Typing Constraints

Relation types τ_r are elements of a partially ordered type system $(\mathcal{T}, \sqsubseteq)$, where \sqsubseteq denotes subtype inclusion. Each type τ determines admissible arities, argument types, and dependency constraints.

Typing rules prohibit cyclic self-reference along precedence or override dimensions. Formally, there exists no $r \in R$ such that $r \prec_d^+ r$ holds within a precedence-relevant subgraph, where \prec_d^+ denotes transitive closure.

A.6 Meta-Relations

Meta-relations are relations whose argument tuples include at least one element of R . No separate syntactic category is introduced. Meta-relations differ only by type and admissible argument structure.

Examples include override relations, contradiction markers, precedence constraints, scope restrictions, and analogical mappings. All such structures obey the same immutability and temporal validity rules as ordinary relations.

A.7 Historical Structure

The complete historical structure is defined as

$$H := (E, R, \prec_t),$$

with \prec_d derivable from R . The structure H is append-only. For any two histories H_1 and H_2 , a merge operation is defined as

$$H_1 \sqcup H_2 := (E_1 \cup E_2, R_1 \cup R_2, \prec_t),$$

assuming consistent identifier namespaces.

A.8 Present-View Space

Let V denote the space of present-views. A present-view at time t is defined as

$$V_t := (E_t, R_t, C_t),$$

where $E_t \subseteq E$ is the set of contextually active events, $R_t \subseteq R$ is the set of active relations, and C_t is a context structure encoding constraints, scopes, and interpretive parameters.

The present-view is not a stored object but the output of a synthesis operation.

A.9 Synthesis Operator

The synthesis operator is a total function

$$\Sigma : H \times T \rightarrow V.$$

For fixed H and t , $\Sigma(H, t)$ is uniquely determined by the admissibility rules, meta-relational constraints, and deterministic resolution strategies encoded in H .

No operation within Σ modifies H .

B Appendix B: Algorithmic Synthesis, Termination, and Determinism

This appendix provides an explicit algorithmic specification of the synthesis operator Σ , together with sufficient conditions for termination and determinism.

B.1 Auxiliary Predicates and Operators

Fix a history $H = (E, R, \prec_t)$ and a time $t \in T$. For any relation-node

$$r = (id_r, \tau_r, A_r, I_r, \rho_r, \mu_r) \in R,$$

define the temporal admissibility predicate

$$\text{Adm}_t(r) \iff t \in I_r.$$

Define the temporally admissible relation set

$$R_t^{(0)} := \{r \in R \mid \text{Adm}_t(r)\}.$$

Let $\mathcal{T}_{\text{meta}} \subseteq \mathcal{T}$ denote the set of meta-relation types relevant to constraint, precedence, override, contradiction, and scoping. Define the extracted meta-relation set

$$M_t := \{m \in R_t^{(0)} \mid \tau_m \in \mathcal{T}_{\text{meta}}\}.$$

Let $\text{Args}(r)$ denote the set of all arguments appearing in the tuple A_r . Define the dependency graph on $R_t^{(0)}$ by a directed edge $r_i \rightarrow r_j$ whenever $r_i \in \text{Args}(r_j)$. Let $\prec_{d,t}$ denote the corresponding dependency relation and $\prec_{d,t}^+$ its transitive closure.

B.2 Override and Precedence Semantics

The synthesis operator is parameterized by a deterministic resolution policy that is itself encoded within H by relations of types in $\mathcal{T}_{\text{meta}}$. To make this explicit, define a binary predicate

$$\text{Disables}_t(m, r),$$

intended to mean that meta-relation $m \in M_t$ disables relation $r \in R_t^{(0)}$ at time t .

The concrete semantics of Disables_t depends on relation types. For the purposes of a minimal rigorous account, it suffices to require that Disables_t be computable from (m, r, t) and that it be well-founded in the dependency ordering. The well-foundedness requirement is formalized below.

Define the meta-disabled set induced by M_t as

$$D_t := \{r \in R_t^{(0)} \mid \exists m \in M_t \text{ } \text{Disables}_t(m, r)\}.$$

Define the post-meta admissible set

$$R_t^{(1)} := R_t^{(0)} \setminus D_t.$$

B.3 Inference Closure

Let $\mathcal{T}_{\text{inf}} \subseteq \mathcal{T}$ denote the relation types eligible for inference closure under the synthesis operator. Let Rule be a computable operator that, given a finite set of relations $S \subseteq R$, returns a finite set of inferred relations $\text{Rule}(S)$ such that each inferred relation's arguments are drawn from $E \cup S$ and its type lies in \mathcal{T}_{inf} .

In order to preserve the append-only requirement, inferred relations are not added to R . Instead, they are added to an auxiliary working set used only within the computation of the present-view. Let \widehat{R}_t denote this working set.

Define the inference closure iteratively by

$$\widehat{R}_t^{(0)} := R_t^{(1)},$$

and for $n \geq 0$ define

$$\widehat{R}_t^{(n+1)} := \widehat{R}_t^{(n)} \cup \text{Rule}(\widehat{R}_t^{(n)}).$$

If there exists N such that $\widehat{R}_t^{(N+1)} = \widehat{R}_t^{(N)}$, define

$$R_t^{(2)} := \widehat{R}_t^{(N)}.$$

B.4 Contradiction Detection and Resolution

Let $\mathcal{T}_{\perp} \subseteq \mathcal{T}_{\text{meta}}$ denote the set of contradiction-marker types. A contradiction marker is any meta-relation $c \in M_t$ such that $\tau_c \in \mathcal{T}_{\perp}$ and $\text{Args}(c)$ contains two relations r_a, r_b intended to be mutually exclusive under some scope condition.

Define a symmetric contradiction predicate on the working set $R_t^{(2)}$ by

$$\text{Contr}_t(r_a, r_b) \iff \exists c \in M_t \left(\tau_c \in \mathcal{T}_{\perp} \wedge \{r_a, r_b\} \subseteq \text{Args}(c) \right),$$

together with any additional scoping and temporal conditions carried by c .

Resolution is performed by a deterministic selector induced by meta-relations. Let

$$\text{Res}_t : \mathcal{P}(R_t^{(2)}) \rightarrow \mathcal{P}(R_t^{(2)})$$

be a computable operator that removes members of a set according to explicit precedence relations and tie-breaking policies encoded in M_t . The operator Res_t must satisfy idempotence,

$$\text{Res}_t(\text{Res}_t(S)) = \text{Res}_t(S),$$

and must guarantee that for any $S \subseteq R_t^{(2)}$, the set $\text{Res}_t(S)$ contains no unresolved contradiction pair unless an explicit “coexistence” marker is present in M_t authorizing contradictory retention.

Define

$$R_t^{(3)} := \text{Res}_t(R_t^{(2)}).$$

B.5 Event Extraction and Context Construction

Define the active event set by argument projection:

$$E_t := \{e \in E \mid \exists r \in R_t^{(3)} \ e \in \text{Args}(r)\}.$$

Define the active relation set of the present-view as

$$R_t := R_t^{(3)} \cap R,$$

and define the context object C_t as a computable summary of active constraints, scopes, and confidence measures extracted from M_t and from metadata μ_r for $r \in R_t^{(3)}$. The exact representation of C_t is implementation-dependent, but it must be derivable deterministically from (H, t) .

Finally define

$$\Sigma(H, t) := (E_t, R_t, C_t).$$

B.6 Termination

A sufficient condition for termination of the inference closure is that **Rule** be inflationary and bounded over a finite universe of admissible relation schemata at time t . Formally, assume that for fixed H and t there exists a finite set \mathcal{U}_t of admissible relation-nodes such that for any $S \subseteq \mathcal{U}_t$ one has $\text{Rule}(S) \subseteq \mathcal{U}_t$. Then the sequence $\widehat{R}_t^{(n)}$ stabilizes in at most $|\mathcal{U}_t|$ steps.

A sufficient condition for finiteness of \mathcal{U}_t is that the arity of inference rules be bounded, that event and base-relation inputs at time t be finite, and that the type system restrict recursive introduction of novel relation identities to those grounded in existing identifiers.

Termination of contradiction resolution follows from idempotence of Res_t and the finiteness of $R_t^{(2)}$.

B.7 Determinism and Uniqueness

Determinism of Σ is ensured if Disables_t , Rule , and Res_t are deterministic functions of their inputs. This requires that any tie-breaking policy be encoded within M_t so that no implicit external ordering is used.

Under these conditions, for fixed H and t , the output $\Sigma(H, t)$ is unique.

B.8 Monotonicity in History

Let $H_1 = (E_1, R_1, \prec_t)$ and $H_2 = (E_2, R_2, \prec_t)$ be histories such that $E_1 \subseteq E_2$ and $R_1 \subseteq R_2$. If the resolution policies are history-extensional in the sense that adding relations can only introduce new disablements by explicit meta-relations, then the synthesis operator is monotone with respect to information preservation: no event or relation in H_1 is deleted from history, and any exclusion from the present-view in H_2 must be justified by a corresponding meta-relation in H_2 .

This monotonicity is a property of the substrate, not of any particular present-view. Present-views may change discontinuously as history grows, but history itself is never reduced.

C Appendix C: Consistency, Constraints, and Contradiction

This appendix formalizes the notion of consistency within present-views and introduces a constraint framework that governs admissible syntheses without collapsing history.

C.1 Constraint Predicates

Let $V_t = (E_t, R_t, C_t)$ be a present-view produced by $\Sigma(H, t)$. A constraint is defined as a total predicate

$$c : V \rightarrow \{\text{true}, \text{false}\}.$$

Constraints may reference events, relations, metadata, temporal intervals, or structural properties of the dependency graph induced by R_t .

Let \mathcal{C} denote the set of all constraint predicates encoded in H as relation-nodes or meta-relations. The active constraint set at time t is defined as

$$\mathcal{C}_t := \{c \in \mathcal{C} \mid c \text{ is temporally admissible at } t\}.$$

A present-view V_t is said to be constraint-consistent if and only if

$$\forall c \in \mathcal{C}_t, \quad c(V_t) = \text{true}.$$

C.2 Hard and Soft Constraints

Constraints may be stratified by severity. A hard constraint is one whose violation invalidates a present-view, whereas a soft constraint encodes a preference rather than a requirement. Formally, define a valuation function

$$\omega : \mathcal{C} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\},$$

where $\omega(c) = \infty$ denotes a hard constraint.

Given a present-view V_t , define its constraint cost as

$$\text{Cost}(V_t) := \sum_{c \in \mathcal{C}_t} \omega(c) \cdot \mathbf{1}[c(V_t) = \text{false}],$$

where $\mathbf{1}[\cdot]$ is the indicator function.

A synthesis is admissible if and only if $\text{Cost}(V_t) < \infty$. When multiple admissible views are possible under resolution strategies, the synthesis operator selects the view minimizing $\text{Cost}(V_t)$, provided this minimization criterion is itself encoded in H .

C.3 Structural Consistency

Structural consistency concerns the internal coherence of the active relation graph. Let $G_t = (R_t, \prec_{d,t})$ denote the dependency graph induced by R_t . Structural consistency requires that G_t satisfy the following conditions.

First, G_t must be acyclic along precedence and override dimensions. Second, for any contradiction marker $c \in R_t$ with arguments (r_a, r_b) , at most one of r_a and r_b may belong to R_t , unless an explicit coexistence relation authorizes simultaneous activation.

These requirements ensure that contradiction is represented explicitly rather than resolved implicitly through erasure.

C.4 Explicit Contradiction Retention

In certain contexts, contradictory relations are intentionally retained. Let $m \in R$ be a meta-relation of type coexistence with arguments (r_a, r_b) and validity interval I_m . If $t \in I_m$, then both r_a and r_b may appear in R_t despite being marked as contradictory elsewhere.

This mechanism allows the framework to represent unresolved scientific debates, conflicting testimony, or parallel hypotheses without forcing premature resolution.

C.5 Constraint Propagation

Constraints may propagate along relation chains. Let $r \in R_t$ be a relation with arguments (x_1, \dots, x_k) . A constraint attached to r may induce derived constraints on x_i depending on the semantics of τ_r . Formally, let

$$\text{Prop}_r : \mathcal{C} \rightarrow \mathcal{P}(\mathcal{C})$$

be a propagation operator associated with relation type τ_r . Constraint propagation is applied during synthesis prior to contradiction resolution.

C.6 Impossibility and Diagnostic Views

There exist histories H and times t for which no present-view satisfies all hard constraints. In such cases, synthesis does not fail silently. Instead, the output V_t includes a distinguished inconsistency marker in C_t , together with the minimal unsatisfiable subset of constraints responsible for failure.

This diagnostic behavior preserves epistemic transparency and prevents the masking of structural failure through forced resolution.

C.7 Persistence Under Constraint Evolution

Constraints themselves may evolve over time. A constraint predicate may be qualified, superseded, or restricted by meta-relations in the same manner as other relations. Let c_1 and c_2 be constraint relations such that c_2 overrides c_1 over some interval. Then c_1 remains in H but is excluded from C_t whenever c_2 is admissible.

This treatment ensures that changes in normative, ecological, or institutional limits do not retroactively rewrite the conditions under which past syntheses were valid.

C.8 Preservation Theorem

If V_t is a constraint-consistent present-view and H is extended to H' by adding new events or relations whose validity intervals lie strictly after t , then V_t remains constraint-consistent under H' . In particular, no future information can retroactively violate constraints at earlier times.

This preservation property formalizes the irreversibility principle underlying the event-historical algebra.

D Appendix D: Analogy, Causality, and Relational Composition

This appendix formalizes analogical and causal reasoning within the event-historical algebra by treating both as structured compositions over relation-nodes rather than as external inference rules.

D.1 Relational Composition

Let $r_1, r_2 \in R$ be relation-nodes such that there exists an argument x with $x \in \text{Args}(r_1)$ and $x \in \text{Args}(r_2)$. A composition operator

$$\circ : R \times R \rightarrow R$$

is defined when the relation types τ_{r_1} and τ_{r_2} admit composition under the typing system. The resulting composite relation $r_3 = r_1 \circ r_2$ has arguments derived from the non-shared arguments of r_1 and r_2 , a type determined by a composition table over \mathcal{T} , and a validity interval equal to the intersection of I_{r_1} and I_{r_2} .

Composition does not introduce new historical entities. The composite relation exists only within the working set of synthesis unless explicitly asserted as a relation-node event in H .

D.2 Causal Relations

A causal relation-node is a relation $r \in R$ whose type τ_r belongs to a distinguished causal subset $\mathcal{T}_{\text{cause}} \subset \mathcal{T}$. Such relations must satisfy temporal ordering constraints. For any causal relation r with arguments (x, y) , if x and y are events, then $t_x < t_y$ must hold. If x or y are relations, then the maximal timestamp of their argument events must respect the same ordering.

Causal relations may carry counterfactual qualifiers encoded as metadata. These qualifiers specify conditions under which the causal dependency would fail, enabling explicit representation of non-monotonic causation.

D.3 Causal Chains and Closure

Let R_t be the active relation set at time t . A causal chain is a finite sequence (r_1, \dots, r_n) of causal relations such that for each i , the effect argument of r_i coincides with the cause argument of r_{i+1} . The synthesis operator computes causal closure by identifying all such chains admissible at t and including their composite effects in the present-view.

This closure supports explanations that traverse multiple intermediate steps without collapsing them into a single opaque inference.

D.4 Analogical Relations

An analogical relation-node $a \in R$ maps a source relation r_s to a target relation r_t by establishing a correspondence between their argument roles. Formally, an analogical relation includes a partial isomorphism

$$\phi_a : \text{Args}(r_s) \rightarrow \text{Args}(r_t)$$

that preserves relational roles rather than identities.

Analogical validity is scoped temporally and contextually. An analogy may be active in one domain or time interval and inactive in another. This allows analogies to be reasoned about, compared, and revised explicitly.

D.5 Inference via Analogy

Analogical inference is performed by lifting constraints, expectations, or causal patterns from r_s to r_t through ϕ_a , subject to explicit admissibility relations. Such inferences are not automatic; they are mediated by meta-relations that authorize or block transfer.

This design prevents uncontrolled metaphorical drift while allowing structured cross-domain reasoning.

D.6 Conflict Between Causal and Analogical Structures

Causal and analogical relations may conflict. For example, an analogy may suggest a dependency that contradicts an established causal relation. Such conflicts are represented explicitly as contradiction relations between the respective relation-nodes, and their resolution is governed by precedence meta-relations.

This mechanism ensures that explanatory reasoning remains transparent and auditable.

D.7 Expressivity Bound

The relational algebra supports finite-depth analogical and causal composition but does not permit unrestricted higher-order self-application. This restriction prevents paradoxical constructions while preserving sufficient expressivity for scientific, legal, and cognitive reasoning.

Formally, there exists a bound k such that no admissible relation may depend on itself through a chain of more than k alternating analogical and causal compositions without explicit authorization by a meta-relation. The value of k is determined by the typing discipline.

D.8 The Role of Analogy in Ecphoric Synthesis

Within the synthesis operator, analogical relations contribute to relevance weighting and constraint propagation rather than to truth determination. An analogy increases the salience of certain relational patterns without asserting their correctness. This aligns with empirical observations that analogies guide hypothesis generation but do not constitute proof.

By embedding analogy as a first-class relational structure, the algebra captures its heuristic power while maintaining epistemic discipline.

E Appendix E: Multi-Agent Histories, Merge Convergence, and Observer-Relative Views

This appendix formalizes the behavior of the event-historical algebra under multi-agent contribution, history merging, and observer-relative synthesis.

E.1 Agent-Indexed Histories

Let \mathcal{A} be a set of agents. For each agent $a \in \mathcal{A}$, define an agent-local history

$$H_a := (E_a, R_a, \prec_t),$$

where $E_a \subseteq E$ and $R_a \subseteq R$ denote the events and relations contributed or observed by agent a . Agent-local histories are not assumed to be complete or mutually consistent.

Agent identity may appear explicitly in provenance descriptors ρ_r , allowing synthesis to reason about source reliability, expertise, or bias.

E.2 History Merge

Given two histories H_a and H_b , their merge is defined as

$$H_{a \sqcup b} := (E_a \cup E_b, R_a \cup R_b, \prec_t),$$

assuming disjoint or namespace-qualified identifiers. No conflict resolution occurs at merge time. All conflicts are deferred to synthesis.

The merge operator \sqcup is commutative, associative, and idempotent, making the space of histories a join-semilattice under set inclusion.

E.3 Observer-Relative Synthesis

An observer-relative present-view is defined as

$$V_t^{(a)} := \Sigma(H, t \mid \Theta_a),$$

where Θ_a is a parameterization of synthesis encoding observer-specific trust relations, constraint weights, or scope preferences. These parameters must themselves be representable as relations in H to preserve auditability.

Observer-relative synthesis allows multiple agents to derive distinct but internally coherent present-views from the same history without contradiction at the historical level.

E.4 Convergence Conditions

Let $\{H_{a_i}\}_{i=1}^n$ be a collection of agent histories and let

$$H^* := \bigsqcup_{i=1}^n H_{a_i}$$

be their merged history. Convergence is said to occur at time t if for all agents a_i and a_j ,

$$\Sigma(H^*, t \mid \Theta_{a_i}) = \Sigma(H^*, t \mid \Theta_{a_j}).$$

Sufficient conditions for convergence include shared constraint structures, aligned precedence relations, and compatible trust orderings. Divergence persists when these structures differ, even in the presence of identical event data.

E.5 Epistemic Pluralism

The framework explicitly supports epistemic pluralism. Divergent present-views are not treated as errors unless they violate shared hard constraints. This reflects real-world scientific, legal, and cultural practice, where disagreement persists despite shared evidence.

Pluralism is thus a property of synthesis, not of history. The historical substrate remains unified even when interpretations diverge.

E.6 Counterfactual Queries

Counterfactual reasoning is implemented by constructing hypothetical histories H' that differ from H by the addition of counterfactual relation-nodes with restricted validity intervals. Synthesis over H' yields counterfactual present-views without modifying H .

This mechanism allows "what-if" analysis while preserving historical integrity.

E.7 Security and Privacy Scoping

Privacy constraints are represented as scope-restricting meta-relations that limit the visibility of events or relations during synthesis. Such relations do not remove elements from H but restrict their inclusion in E_t or R_t for particular observers.

This design permits fine-grained access control without compromising persistence.

E.8 Final Convergence Theorem

If a merged history H^* satisfies the condition that all agents share identical hard constraints and precedence relations, then for any time t there exists a unique present-view V_t such that

$$\Sigma(H^*, t \mid \Theta_a) = V_t$$

for all agents a . In this case, observer-relativity collapses and synthesis becomes objective with respect to the shared constraint structure.

This theorem formalizes the intuition that agreement emerges not from erasure of difference, but from alignment of constraints.

E.9 Summary

This appendix demonstrates that the event-historical algebra naturally supports multi-agent knowledge construction, delayed reconciliation, and principled disagreement. Convergence,

when it occurs, is the result of structural alignment rather than forced consensus.

The algebra thereby accommodates both cooperation and dissent within a single persistent substrate.

F Appendix F: Computational Realization and Complexity Analysis

This appendix addresses computational realizability of the event-historical algebra, with particular emphasis on algorithmic complexity, data structures, and incremental synthesis.

F.1 Representational Assumptions

Assume that the historical structure $H = (E, R, \prec_t)$ is stored in a persistent data store supporting append-only insertion and indexed access. Events are indexed by timestamp and identifier. Relations are indexed by type, temporal validity interval, and argument identifiers. Meta-relations are not stored separately but are identified by type membership in $\mathcal{T}_{\text{meta}}$.

No assumption is made about global memory availability. The framework permits out-of-core storage and streaming access.

F.2 Incremental Synthesis

Let $t_0 < t_1$ be consecutive synthesis times such that no events occur in the open interval (t_0, t_1) . Then the present-view V_{t_1} differs from V_{t_0} only by relations whose validity intervals begin or end at t_1 or by meta-relations whose scope changes at t_1 .

Define the delta relation set

$$\Delta R := \{r \in R \mid t_0 \notin I_r \wedge t_1 \in I_r\} \cup \{r \in R \mid t_0 \in I_r \wedge t_1 \notin I_r\}.$$

Incremental synthesis recomputes only the dependency closure affected by ΔR , rather than recomputing $\Sigma(H, t_1)$ from scratch. Under bounded dependency depth, this yields amortized linear-time updates relative to the size of ΔR .

F.3 Streaming Synthesis

For histories too large to fit in memory, synthesis may be performed in a streaming manner. Let $H_{<t}$ denote the prefix of H containing all events and relations with timestamps strictly

less than t . The synthesis operator can be approximated by maintaining a rolling window of active relations together with a compressed summary of inactive structure.

Streaming synthesis preserves correctness provided that no relation with validity interval intersecting t is discarded and that all meta-relations governing precedence are retained. This condition defines a minimal retention frontier that bounds memory usage.

F.4 Time and Space Complexity

Let $n_t = |R_t^{(0)}|$ denote the number of temporally admissible relations at time t , and let d be the maximum dependency depth under $\prec_{d,t}$. Under bounded inference rules and deterministic resolution, the worst-case time complexity of synthesis is $O(n_t \cdot d)$.

Space complexity is dominated by storage of H . Because no deletion occurs, storage grows monotonically. Archival strategies may compress inactive relations provided that provenance and dependency metadata are preserved.

F.5 Garbage Collection and Archival

Although deletion is prohibited at the algebraic level, implementation-level archival is permitted provided that it is semantics-preserving. A relation or event may be archived if and only if there exists a summary object s such that synthesis over the archived structure yields identical present-views for all future times.

Formally, archival is admissible if for all $t' > t_{\text{archive}}$,

$$\Sigma(H, t') = \Sigma(H', t'),$$

where H' replaces the archived subset with s . Such summaries may be viewed as higher-order relations encoding the net effect of historical structure.

F.6 Query Language Considerations

Queries against H and V_t are expressed not as state lookups but as constrained synthesis requests. A query specifies a target time, optional observer parameters, and structural constraints over events and relations. The query result is a present-view or a substructure thereof.

This design avoids impedance mismatch between querying and synthesis and ensures that all answers are grounded in the same algebraic semantics.

F.7 Concurrency and Distribution

In distributed settings, multiple nodes may append events and relations concurrently. Provided that identifier namespaces are disjoint or cryptographically unique, histories may be merged without coordination. Synthesis over the merged history yields deterministic results subject to observer parameters.

Concurrency control is thus reduced to ensuring causal ordering of events where required. The algebra itself imposes no global locking or consensus requirement.

F.8 Limits of Computation

There exist histories for which synthesis is computationally infeasible due to extreme relational density or deeply nested meta-relations. These cases correspond to epistemic overload rather than algorithmic failure. Approximate synthesis operators may be employed to trade completeness for responsiveness, provided that approximation bounds are explicit.

F.9 Summary

This appendix demonstrates that the event-historical algebra is computationally realizable using known techniques from persistent data structures, incremental computation, and distributed systems. While worst-case complexity remains high, practical implementations can exploit temporal locality, bounded dependency depth, and archival summaries to scale synthesis to real-world domains.

The algebra thus provides not only a conceptual framework but a viable computational substrate for persistent, interpretable knowledge systems.

A Appendix G: Ecphoric Synthesis, Structural Perception, and the Limits of Pattern-Based Models

This appendix situates ecphoric synthesis within a broader theoretical context by contrasting event-historical reasoning with pattern-based perceptual models, particularly as instantiated in contemporary multimodal large language models. The purpose is not evaluative but structural: to clarify which cognitive and computational capacities are enabled or precluded by different representational substrates.

A.1 Ecphoric Recall as Structural Reconstruction

Ecphoric recall is not retrieval of a stored representation but reconstruction of a present view from distributed historical traces. Formally, this corresponds to the synthesis operator $\Sigma(H, t)$ acting over an event-historical structure H to produce a present-view V_t without modifying H . The defining feature of ecphoric synthesis is that coherence emerges from relational density, temporal alignment, and provenance rather than from similarity to a stored prototype.

Let $H = (E, R, \prec_t)$ be a history as defined in the main text. For any cue set $C \subseteq E \cup R$, ecphoric recall corresponds to selecting V_t such that the induced subgraph of (E, R) maximizes structural coherence subject to active constraints. This coherence criterion is not reducible to metric similarity over payloads; it depends on the existence of multi-path relational support, compatibility with meta-relations, and alignment with temporal validity intervals.

This explains why ecphoric recall can privilege older or less salient events over more recent surface descriptions when the former participate in a denser and more explanatory relational structure.

A.2 Structural Perception Versus Feature Recognition

Pattern-based systems, including current vision-language models, operate primarily by mapping perceptual inputs into embedding spaces optimized for discrimination over training distributions. Let $\phi : X \rightarrow \mathbb{R}^n$ denote such an embedding function. Inference then proceeds by proximity in \mathbb{R}^n , optionally conditioned on textual prompts.

This architecture implicitly assumes that structural properties are either directly encoded in the embedding or recoverable through downstream symbolic reasoning. Empirical results demonstrate that this assumption fails for tasks requiring explicit structural traversal, such as counting polygon sides, decomposing merged shapes, or reasoning over novel geometric configurations.

From the perspective of the event-historical algebra, this failure arises because feature embeddings lack explicit relation-nodes encoding adjacency, succession, or boundary traversal. A polygon is represented as a texture-like object rather than as a sequence of edge-events connected by successor relations. Without such relations, there exists no substrate over which ecphoric synthesis could operate.

A.3 Why Visually-Cued Reasoning Works

The effectiveness of visually-cued chain-of-thought prompting can be explained algebraically. When vertices or edges are explicitly labeled in a diagram, each label functions as an externally supplied event identifier. Relations such as “edge-successor” or “vertex-adjacent-to” become inferable, allowing a temporary relational structure to be constructed within the present-view.

In other words, visual cues simulate relation-nodes that the underlying perceptual system cannot autonomously construct. The improvement in performance does not indicate latent geometric understanding; it indicates that the model can operate over relational structure when such structure is externally imposed.

This mirrors the distinction between retrieval and ecphoric synthesis. The former operates over static representations, while the latter requires a relational substrate capable of supporting traversal, dependency, and constraint propagation.

A.4 Event-Historical Interpretation of Geometric Reasoning

Within the event-historical framework, a geometric figure is not an object but a structured history. A polygon is represented as a finite sequence of boundary events with successor relations forming a closed chain. Counting sides corresponds to computing the cardinality of this chain under traversal closure. Irregularity, deformation, or rotation does not alter the underlying relational structure, only the payloads of events.

Formally, let $\{e_1, \dots, e_n\} \subset E$ be boundary events with relations $r_i = \text{successor}(e_i, e_{i+1})$ for $i < n$ and $r_n = \text{successor}(e_n, e_1)$. The number of sides is invariant under any transformation that preserves the successor relation graph. Any system lacking access to such relations cannot, in principle, perform reliable geometric reasoning.

A.5 Relation to Thousand-Brains and Reference Frame Theories

The algebra developed in this work is compatible with reference-frame-based theories of cognition. Each relation-node defines a local frame of interpretability, and synthesis corresponds to alignment across frames under shared constraints. What is often described biologically as grid-cell or place-cell activity corresponds, at an abstract level, to the stabilization of relational paths under movement or attention.

However, the present framework does not require commitment to any specific neural mechanism. It asserts only that structural cognition requires persistent relational entities and a synthesis process capable of reconstructing present understanding from them.

A.6 Implications for Artificial Intelligence

The limitations observed in multimodal models are not incidental shortcomings but predictable consequences of architectures optimized for pattern completion rather than historical reconstruction. Without event persistence, relation-nodes, and non-destructive synthesis, such systems cannot exhibit ecphoric recall, structural perception, or genuine causal reasoning.

Conversely, incorporating an event-historical substrate does not require abandoning statistical learning. Learned perceptual models may serve as event generators, while relational learning and synthesis operate at a higher structural level. The key architectural shift is to treat relations as first-class entities and to defer interpretation until synthesis time.

A.7 Conclusion

Ecphoric synthesis clarifies a boundary condition on intelligence: systems that overwrite state or collapse structure into embeddings cannot support persistent understanding. Structural reasoning, whether geometric, causal, or autobiographical, requires an event-historical substrate and a synthesis operator that reconstructs the present without erasing the past.

This completes the formal picture by connecting the algebra developed in this document to observed cognitive phenomena and to the documented limitations of current artificial systems, without appealing to anthropomorphism or informal metaphor. The distinction is not between human and machine intelligence, but between architectures that preserve history and those that do not.

B Appendix H: Categorical Formulation and a Worked Geometric Example

This appendix presents two complementary formalisms. The first gives a minimal categorical and rewrite-theoretic account of events, relations, and synthesis, suitable for direct connection to merge algebra and CRDT-style systems. The second provides a worked geometric example demonstrating why side-counting fails in pattern-based systems and succeeds under an event-historical, relation-first representation.

B.1 A Minimal Categorical Formulation

Let \mathcal{C} be a small category whose objects are event-identifiers and whose morphisms are typed relations. An object $e \in \text{Ob}(\mathcal{C})$ corresponds to an event as defined in the main text.

A morphism $r : e_i \rightarrow e_j$ corresponds to a relation-node whose argument structure includes (e_i, e_j) and whose type τ_r determines admissibility conditions.

Temporal structure is introduced by a functor

$$\mathsf{T} : \mathcal{C} \rightarrow \mathbf{Poset},$$

mapping each object and morphism to a temporal interval in a totally ordered set (T, \leq) . A morphism r is temporally admissible at time t if and only if $t \in \mathsf{T}(r)$.

Meta-relations are represented as 2-morphisms in a weak 2-category $\mathcal{C}^{(2)}$, where a 2-morphism

$$\alpha : r_1 \Rightarrow r_2$$

encodes precedence, override, contradiction, or coexistence between relations r_1 and r_2 . The typing discipline restricts which 2-morphisms may compose, ensuring the absence of paradoxical self-reference.

Composition of relations corresponds to categorical composition. For morphisms $r_1 : e_1 \rightarrow e_2$ and $r_2 : e_2 \rightarrow e_3$, their composite

$$r_2 \circ r_1 : e_1 \rightarrow e_3$$

exists when permitted by the relation-type composition rules. This captures causal chaining and analogical propagation without introducing new historical entities.

B.2 Rewrite Semantics and Merge Algebra

An equivalent presentation is given by a rewrite system over a signature Σ consisting of event symbols and relation symbols. Let \mathcal{H} denote a multiset of ground terms encoding events and relations. Rewrite rules are not allowed to delete terms; they may only introduce derived views or mark relations as inactive under specified conditions.

A synthesis step corresponds to a rewrite

$$\mathcal{H} \implies_t \mathcal{V}_t,$$

where \mathcal{V}_t is a structured normal form representing the present-view at time t . Normalization is confluent under the determinism conditions established in Appendix B.

Merge of histories corresponds to multiset union,

$$\mathcal{H}_1 \oplus \mathcal{H}_2 = \mathcal{H}_1 \uplus \mathcal{H}_2,$$

with conflicts preserved as explicit terms. Rewrite rules governing override and precedence ensure that resolution occurs only during synthesis, never during merge.

This rewrite formulation makes explicit the connection to conflict-free replicated data types. The algebra guarantees strong eventual consistency at the level of history, while allowing observer-relative present-views.

B.3 A Worked Geometric Example

Consider a planar polygon. In a pattern-based representation, the polygon is mapped to a feature vector $\phi(x)$, and properties such as the number of sides are inferred by classification against memorized prototypes. This approach fails for irregular or novel shapes because side-count is not a local feature.

In the event-historical representation, the polygon is modeled as a closed boundary history. Let

$$E = \{v_1, v_2, \dots, v_n\}$$

be vertex events, each with spatial payloads. Let

$$R = \{s_i : v_i \rightarrow v_{i+1} \mid 1 \leq i < n\} \cup \{s_n : v_n \rightarrow v_1\}$$

be successor relations encoding adjacency along the boundary. Each successor relation has a temporal interval corresponding to the persistence of the boundary.

The polygon is thus a cyclic relation graph. Counting sides corresponds to computing the cardinality of the successor cycle. This computation is invariant under deformation, rotation, scaling, or irregular spacing of vertices, because it depends only on the existence of a closed successor chain.

If a vertex is split or merged, this appears as additional events and relations in H , not as overwriting. Meta-relations may mark certain successor relations as inactive under specific interpretations, but the historical structure remains intact.

B.4 Failure Without Relation-Nodes

If the successor relations are absent, the representation collapses to a set of points or pixels. No traversal operator can be defined, and the notion of “next edge” is undefined. In such a representation, side-count cannot be computed generically; it can only be guessed by association with known shapes.

This precisely characterizes the failure modes observed in multimodal models. The absence of explicit successor relations prevents synthesis from constructing a traversal path,

forcing reliance on memorized labels.

B.5 Success With Explicit Structure

When vertices are labeled externally, as in visually-cued prompting, each label functions as an event identifier, and implicit successor relations can be inferred. The model temporarily reconstructs a boundary history and performs traversal. The improvement in performance follows directly from the restoration of relational structure.

The same effect is achieved intrinsically if perceptual preprocessing produces stable boundary events and successor relations prior to synthesis. In this case, no external cues are required.

B.6 Summary

This appendix has shown that the event-historical algebra admits both a categorical and a rewrite-theoretic formulation, each suitable for implementation in distributed and merge-based systems. The worked geometric example demonstrates that structural reasoning is not a matter of stronger pattern recognition but of representational adequacy.

Geometric understanding, like memory and causality, emerges when events persist, relations are first-class, and synthesis reconstructs the present without erasing the past.

C Appendix I: Representational Separation Results

This appendix establishes formal separation results between the ephoric event-relation-meta-relation algebra and several dominant representational frameworks, including temporal logics, Bayesian networks, and transformer-based sequence models. The results are stated as relative expressivity limitations under explicit structural assumptions. No claim of absolute impossibility in the Turing-complete sense is made. Rather, the focus is on what each framework can represent natively, without reconstructing an event-historical substrate by external means.

C.1 Preliminaries

Let $H = (E, R, M)$ denote an event-historical structure, where E is a countable set of events indexed by a totally ordered time domain, R is a set of relation-nodes, and M is a set of meta-relations defined recursively over R . Let Σ denote the synthesis operator mapping (H, t) to a present-view V_t without modifying H .

A representation is said to be *history-preserving* if no assertion, relation, or justification is deleted or overwritten as a result of revision, update, or inference.

A representation is said to support *first-class relations* if relations possess identity, temporal extent, provenance, and may themselves participate in higher-order relations within the same formal language.

C.2 Lemma 1: Non-Destructive Revision Is Not Native to Temporal Logic

Lemma 1. Standard temporal logics evaluated over Kripke structures do not natively support non-destructive revision.

Proof sketch. In linear-time and branching-time temporal logics, truth values are evaluated relative to states and paths. When a proposition p ceases to hold at a later state, the earlier assertion of p is no longer accessible as an object of reasoning. The logic provides no term that denotes the prior assertion itself, nor a means to relate it to a later correction. Any attempt to preserve prior assertions requires reifying propositions or transitions as objects external to the logic, thereby departing from its native semantics. \square

Corollary 1. Any temporal-logic-based system that supports explicit audit trails or revision history must embed an auxiliary event-historical structure outside the logic proper.

C.3 Lemma 2: Bayesian Networks Cannot Represent Supersession Without Erasure

Lemma 2. Bayesian networks cannot natively represent supersession relations between assertions without erasing or subsuming earlier belief states.

Proof sketch. Bayesian networks encode conditional dependencies among random variables at a fixed modeling level. Updating beliefs corresponds to conditioning or parameter adjustment. Prior belief states are not retained as first-class objects once updated. Although one may log parameter changes externally, the network itself does not represent the fact that one belief superseded another, nor the justification for that supersession, within its graph structure. \square

Corollary 2. Provenance-aware belief revision in Bayesian systems requires external bookkeeping that is not reflected in the probabilistic semantics of the model.

C.4 Lemma 3: Bounded-Context Sequence Models Cannot Guarantee Historical Recall

Lemma 3. Any sequence model with bounded internal state and bounded input context cannot guarantee correct recall of arbitrary historical events.

Proof. Let M be a model whose internal state space has finite cardinality and whose input context length is bounded by k . Consider a family of histories $\{H_n\}$ containing n distinct events with unique identifiers, for unbounded n . By the pigeonhole principle, there exist distinct histories H_n and H_m that induce identical internal states and identical input contexts for M with respect to any fixed query position. Therefore, there exists a query concerning the occurrence or provenance of an event that M cannot answer correctly for both histories. \square

Corollary 3. Transformer-based models without external persistent memory cannot realize unbounded auditability or non-destructive revision as invariants, regardless of scale.

C.5 Lemma 4: Relations-as-Objects Exceed Edge-Based Graph Semantics

Lemma 4. Graph formalisms in which relations are represented solely as edges lack the expressive capacity to encode meta-relations without reification.

Proof sketch. In edge-based graphs, edges do not possess identity independent of their endpoints. As a consequence, relations cannot be temporally scoped, justified, overridden, or contradicted except by introducing auxiliary nodes that reify edges. This transformation changes the ontology of the graph and introduces a higher-order representational layer not present in the original formalism. \square

Corollary 4. Any system supporting reasoning about causation, justification, exception, or override within a single uniform language must treat relations as first-class entities.

C.6 Lemma 5: Observer-Relative Views Require Structural Multiplicity

Lemma 5. A formalism that admits only a single authoritative global state cannot represent multiple observer-relative present-views without duplicating or mutating the underlying model.

Proof sketch. In temporal logics and Bayesian networks, differing interpretations require distinct models or distinct parameterizations. In sequence models, differing interpreta-

tions require distinct prompts or inputs. In each case, observer-relative variation is encoded externally rather than as a function over a shared preserved structure. \square

Corollary 5. The ability to generate multiple internally consistent present-views from a single preserved history is not native to these formalisms.

C.7 Theorem: Representational Separation

Theorem. The ecphoric event–relation–meta-relation algebra can natively represent non-destructive revision, first-class provenance, supersession without erasure, unbounded historical recall, and observer-relative synthesis, whereas temporal logics, Bayesian networks, and bounded-context sequence models cannot jointly represent all of these properties without structural extension.

Proof. Each property follows directly from the definitions of H and Σ . The lemmas above show that each competing formalism lacks at least one of these properties natively. Any attempt to recover them requires reconstructing an event-historical substrate external to the original formalism, thereby collapsing into an ecphoric-equivalent architecture. \square

C.8 Implications

These separation results clarify why pattern-completion systems exhibit brittleness in reasoning, causality, and geometry. Without persistent history, relations cannot be traversed. Without first-class relations, reasoning cannot be inspected. Without non-destructive revision, learning destroys memory.

The ecphoric framework does not optimize prediction. It preserves structure. Intelligence, on this view, is not the compression of the past into a state, but the disciplined reconstruction of the present from an irreducible history.

D Appendix J: A Minimal Algebraic Core for Ecphoric Histories

This appendix presents a minimal algebraic core sufficient to implement the ecphoric framework in a merge-oriented setting. The goal is to specify the smallest set of primitives and laws needed to support append-only history, first-class relations, meta-relations, and synthesis of a present-view without deletion. The presentation is intentionally austere so that it can be mapped directly to a concrete data model or to a rewrite semantics compatible with Spherepop-style merge algebra.

D.1 Carriers and Sorts

Fix a totally ordered time domain (T, \leq) . The algebra has four primary sorts: event identifiers, relation identifiers, meta-relation identifiers, and payloads. Let EID , RID , and MID denote countable sets of identifiers, and let P denote the space of payloads. Let K denote a set of relation-kinds, and let I denote a set of temporal intervals over T .

An event is a tuple

$$e = (i, t, p) \in \text{EID} \times T \times \mathsf{P}.$$

A relation-node is a tuple

$$r = (j, k, \vec{i}, \iota, p) \in \text{RID} \times \mathsf{K} \times \text{EID}^* \times \mathsf{I} \times \mathsf{P},$$

where \vec{i} is a finite list of event identifiers, ι is a validity interval, and p is an auxiliary payload that may encode provenance, confidence, derivation traces, or any other structured annotation. A meta-relation is a tuple

$$m = (\ell, \kappa, \vec{j}, \iota, p) \in \text{MID} \times \mathsf{K} \times \text{RID}^* \times \mathsf{I} \times \mathsf{P},$$

where \vec{j} is a finite list of relation identifiers and κ denotes the meta-relation kind.

A history is a triple

$$H = (E, R, M)$$

where E is a finite or countable set of events, R a finite or countable set of relations, and M a finite or countable set of meta-relations.

D.2 Append-Only Extension

The core update operation is append. For events, define

$$\text{addE} : H \times (\text{EID} \times T \times \mathsf{P}) \rightarrow H$$

by

$$\text{addE}((E, R, M), (i, t, p)) = (E \cup \{(i, t, p)\}, R, M),$$

with the convention that (i, t, p) is admissible only if $i \notin \pi_1(E)$, where π_1 projects event identifiers.

Similarly define

$$\text{addR} : H \times (\text{RID} \times \mathsf{K} \times \text{EID}^* \times \mathsf{I} \times \mathsf{P}) \rightarrow H$$

and

$$\text{addM} : H \times (\text{MID} \times \mathbf{K} \times \text{RID}^* \times \mathbf{I} \times \mathbf{P}) \rightarrow H$$

as append-only insertions, each requiring uniqueness of the corresponding identifier.

These operations are history-preserving in the strong sense: no pre-existing element of E , R , or M is removed or mutated by append.

D.3 Merge as Join

Define a merge operator

$$\oplus : \mathbf{Hist} \times \mathbf{Hist} \rightarrow \mathbf{Hist}$$

on the set \mathbf{Hist} of all histories by

$$(E_1, R_1, M_1) \oplus (E_2, R_2, M_2) = (E_1 \cup E_2, R_1 \cup R_2, M_1 \cup M_2),$$

assuming identifier uniqueness across merged components, or else treating equal identifiers with unequal payloads as a conflict that is itself represented as an appended relation-node. In the conflict-free regime, \oplus is commutative, associative, and idempotent, making (\mathbf{Hist}, \oplus) a join-semilattice.

The essential design constraint is that merge must be purely structural. Any resolution of incompatibilities is deferred to synthesis and represented via meta-relations rather than destructive edits.

D.4 Temporal Admissibility

For any $t \in T$, define the temporal admissibility predicate for relations and meta-relations. If $r = (j, k, \vec{i}, \iota, p)$, then

$$\text{adm}(r, t) \text{ holds iff } t \in \iota.$$

If $m = (\ell, \kappa, \vec{j}, \iota, p)$, then

$$\text{adm}(m, t) \text{ holds iff } t \in \iota.$$

Events are always admissible as historical facts of occurrence; temporal filtering applies to their participation in present-views by selection criteria defined by synthesis.

D.5 Present-View Carrier

A present-view at time t is a structured tuple

$$V_t = (E_t, R_t, U_t),$$

where $E_t \subseteq E$ and $R_t \subseteq R$ are the active events and relations at time t , and U_t is an auxiliary structure encoding unresolved incompatibilities, explicit contradictions, or suspended candidates. The presence of U_t ensures that synthesis may remain non-destructive even when total consistency is unattainable.

D.6 Synthesis as a Deterministic Reduction

Define the synthesis operator

$$\Sigma : \mathbf{Hist} \times T \rightarrow \mathbf{View}$$

where \mathbf{View} is the class of all present-views. The synthesis operator is defined by deterministic reduction rules applied to the admissible substructure at time t .

Let

$$R_t^{\text{adm}} = \{r \in R \mid \text{adm}(r, t)\} \quad \text{and} \quad M_t^{\text{adm}} = \{m \in M \mid \text{adm}(m, t)\}.$$

Let R_t^0 be an initial candidate set derived from R_t^{adm} by a deterministic cue policy, which may be the identity map if no cueing is used. Synthesis then applies meta-relations as constraints on membership in R_t .

To make this explicit, fix a family of meta-kinds that act as determiners of activity. For example, an override meta-relation of the form $\text{override}(r_a, r_b)$ expresses that, when active, r_b is suppressed whenever r_a is admissible and selected.

Define a suppression relation \triangleright_t on R_t^{adm} generated by admissible meta-relations. If $m \in M_t^{\text{adm}}$ encodes that r_a suppresses r_b , write $r_a \triangleright_t r_b$.

Given \triangleright_t , define the active set

$$R_t = \{r \in R_t^0 \mid \nexists r' \in R_t^0 \text{ such that } r' \triangleright_t r\}.$$

Define U_t as the set of those relations in R_t^0 that are suppressed by some admissible competitor, together with any explicitly contradictory pairs marked by meta-relations. Define E_t as the set of events referenced by R_t , together with any cue events required by the cue policy.

This yields a deterministic present-view

$$\Sigma(H, t) = (E_t, R_t, U_t).$$

D.7 Algebraic Laws

The minimal laws required for the ecphoric architecture are as follows.

First, append-only preservation holds: for any history H and any admissible insertion x ,

$$H \subseteq \text{addE}(H, x), \quad H \subseteq \text{addR}(H, x), \quad H \subseteq \text{addM}(H, x),$$

where \subseteq denotes componentwise subset inclusion of E , R , and M .

Second, merge laws hold on conflict-free histories:

$$H_1 \oplus H_2 = H_2 \oplus H_1, \quad (H_1 \oplus H_2) \oplus H_3 = H_1 \oplus (H_2 \oplus H_3), \quad H \oplus H = H.$$

Third, synthesis is non-destructive: for any H and t ,

$$\Sigma(H, t) = (E_t, R_t, U_t)$$

does not modify H , and H remains the unique conserved substrate from which all views are derived.

Fourth, synthesis is merge-compatible in the following weak sense. Let $H = H_1 \oplus H_2$. Then $\Sigma(H, t)$ depends only on the unioned structure and the deterministic meta-resolution policy. Any discrepancy between $\Sigma(H_1, t)$ and $\Sigma(H_2, t)$ is not repaired by erasure, but by the existence or absence of meta-relations in the merged history.

D.8 Interpretation and Minimality

This core isolates three commitments. History is append-only. Merge is a join on histories. Interpretation is deferred to synthesis, where meta-relations control activation without deleting alternatives. Everything beyond this core, including probabilistic confidence, learning of relation-kinds, or category-theoretic enrichment, can be treated as structured payload within P and therefore does not change the algebraic skeleton.

In this sense, the ecphoric framework is not primarily a learning model but a persistence model. It encodes the principle that understanding is the production of situated present-views from conserved historical structure, rather than the replacement of the past by a compressed state.

References

- [1] E. Tulving. Episodic and semantic memory. In *Organization of Memory*, E. Tulving and W. Donaldson, editors. Academic Press, 1972.
- [2] E. Tulving. *Elements of Episodic Memory*. Oxford University Press, 1983.
- [3] F. C. Bartlett. *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press, 1932.
- [4] D. L. Schacter. The seven sins of memory. *American Psychologist*, 54(3):182–203, 1999.
- [5] K. Friston. The free-energy principle: A unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.
- [6] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2nd edition, 2009.
- [7] G. Lakoff and M. Johnson. *Metaphors We Live By*. University of Chicago Press, 1980.
- [8] E. Hutchins. *Cognition in the Wild*. MIT Press, 1995.
- [9] I. Prigogine. *Order Out of Chaos*. Bantam Books, 1984.
- [10] E. Noether. Invariante Variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen*, 1918.
- [11] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565, 1978.
- [12] M. Fowler. *Event Sourcing*. Addison-Wesley, 2017.
- [13] M. Shapiro, N. Preguiua, C. Baquero, and M. Zawirski. Conflict-free replicated data types. *Stabilization, Safety, and Security of Distributed Systems*, 2011.
- [14] P. Grdenfors. *Conceptual Spaces: The Geometry of Thought*. MIT Press, 2000.
- [15] T. S. Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press, 1962.
- [16] C. E. Alchourrn, P. Grdenfors, and D. Makinson. On the logic of theory change. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
- [17] A. Clark. *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford University Press, 2016.

- [18] J. R. Anderson. *ACT: A Simple Theory of Complex Cognition*. American Psychologist, 1996.
- [19] Charles Simon, *Will Computers Revolt? Preparing for the Future of Artificial Intelligence*, 2nd ed., FutureAI Guru, 2021, 340 pp.
- [20] Charles Simon, *Brain Simulator II: Manual for Creating Artificial General Intelligence*, FutureAI Guru, 2021, 172 pp.
- [21] Ronald Brachman and Hector Levesque, *Machine Like Us: Toward AI with Common Sense*, MIT Press, 2022, 320 pp.
- [22] Melanie Mitchell, *Artificial Intelligence: A Guide for Thinking Humans*, Farrar, Straus and Giroux, 2020, 417 pp.
- [23] W. Rudman, M. Golovanevsky, A. Bar, V. Palit, Y. LeCun, C. Eickhoff, and R. Singh. *Forgotten Polygons: Multimodal Large Language Models are Shape-Blind*. arXiv preprint arXiv:2502.15969, 2025.
- [24] J. Hawkins. *A Thousand Brains: A New Theory of Intelligence*. Basic Books, New York, NY, 2021.