Authoritative Stability Under Semantic Drift in AI-Native Systems of Record

Brian Calvanese[a]

[a]*Independent Researcher,*

Abstract

AI-native systems interpret interaction histories using probabilistic models, retrieval-augmented inference, and evolving policy layers. Under admissible semantic context evolution—changes in models, corpora, prompts, or policy constraints—identical event logs may yield different hypotheses and therefore different candidate fact sets. When authoritative state is defined as the output of the current interpretation, candidate fact formation and authoritative state assignment are conflated, coupling institutional authority to a semantic context that may evolve over time and permitting retroactive changes in closed-loop systems.

We formalize this phenomenon as *interpretive variability* of the interpretation-to-fact pipeline under admissible semantic context evolution and define *authoritative stability* as invariance of authoritative state absent governance transitions. We show that three desirable architectural properties—admissible semantic context evolution, authoritative state defined by recomputation under the current semantic context, and stability of authoritative state over no-governance intervals—cannot in general be satisfied simultaneously once interpretive variability is realized. We therefore prove a structural no-go result: when interpretive variability is realized, any architecture that defines authoritative state by recomputation of the interpretation-to-fact pipeline under the current semantic context cannot guarantee authoritative stability over no-governance intervals.

To resolve this incompatibility, we introduce *governance-bound commitments*, immutable authority-establishing records with immutable provenance that separate provisional fact derivation from authoritative state. Commit-

Email address: calvab@aol.com (Brian Calvanese)

ments record governance acceptance of selected candidate facts and thereby fix authoritative state while preserving the semantic context under which acceptance occurred. We prove a constructive sufficiency result: under commitment semantics with forward-only construction, authoritative stability holds for any admissible semantic context evolution over no-governance intervals. The induced authority dependency graph is acyclic, ensuring non-retroactive evolution and traceable causal structure.

We outline a simulation-based evaluation comparing interpretation-driven and commitment-based architectures under controlled semantic context evolution, quantifying retroactive mutation, actuation instability, provenance trace depth, provenance trace completeness, and computational overhead in closed-loop decision systems.

---

1. Introduction

Enterprise systems of record were historically designed under a deterministic paradigm: authoritative state is fixed at write time through schema constraints, transactional guarantees, and controlled processing pipelines. The relational model formalized accepted meaning as structured state stored in well-defined relations [1], and ACID semantics ensured durability and stability of state transitions [2, 3]. Event-sourced and log-centric architectures extended this foundation by elevating immutable logs as the substrate of state reconstruction [4, 5, 6, 7]. Although projection logic may evolve, replay semantics are engineered to preserve determinism, so authoritative state changes only through governed transitions.

These architectures rely on a core invariant: authoritative state changes only through governed transitions applied to a fixed interpretation of system history. AI-native systems challenge this invariant by introducing interpretation layers whose semantics may evolve during normal operation.

Modern AI-native architectures integrate probabilistic foundation models and retrieval-augmented generation (RAG) directly into operational pipelines [8, 9]. Interpretation is therefore no longer a fixed function of immutable logs: identical interaction histories may yield different hypotheses and therefore different candidate fact sets under admissible changes in model parameters, retrieval corpora, prompts, policy constraints, or other elements of the system's semantic context. Empirical work documents sensitivity to contextual variation [10, 11], and ML systems engineering highlights non-stationarity

arising from evolving data and models [12, 13]. When authoritative state is defined as the output of current interpretation, accepted meaning becomes coupled to a semantic context that may evolve over time. We use the term *semantic drift* as shorthand for such admissible semantic context evolution. This observation motivates an architectural separation between interpretation and authority: interpretation may evolve under changing semantic context, while authoritative state is established through explicit governance transitions rather than through recomputation under the current interpretation context.

Operational AI systems can be understood as participating in a governed epistemic feedback loop:

$$\text{Knowledge} \rightarrow \text{Understanding} \rightarrow \text{Perception} \rightarrow \text{Candidate Fact} \rightarrow \text{Decision} \rightarrow \text{Commitment} \rightarrow \text{Action} \rightarrow \text{Response} \rightarrow \text{Knowledge'}.$$

Interaction events provide *knowledge*. System configuration, schema, and policy constraints define *understanding*. Probabilistic inference produces hypotheses corresponding to the system's *perception*. Structured derivation transforms these hypotheses into candidate *facts*. Governance processes make authority-establishing *decisions* that accept selected candidate facts. These decisions are recorded as immutable *commitments*, and the system's authoritative state is defined over the set of effective commitments. Automated *actions* are then triggered from that state. Environmental consequences return as recorded *responses*, becoming new knowledge for subsequent evaluation.

The architectural artifacts corresponding to these stages in AI-native systems are summarized in Table 1.

Throughout the paper we distinguish between the conceptual act of governance and the artifact that records it. A *decision* denotes the authority-establishing governance act that accepts a candidate fact, while a *commitment* denotes the immutable system record produced by that decision. Authoritative state is defined over effective commitments rather than over candidate facts or decisions themselves.

In AI-native architectures, perception and candidate fact derivation are inherently probabilistic and may evolve as models, corpora, prompts, or policies change. If candidate facts produced by interpretation are treated directly as authoritative state, authority becomes implicitly coupled to semantic context.

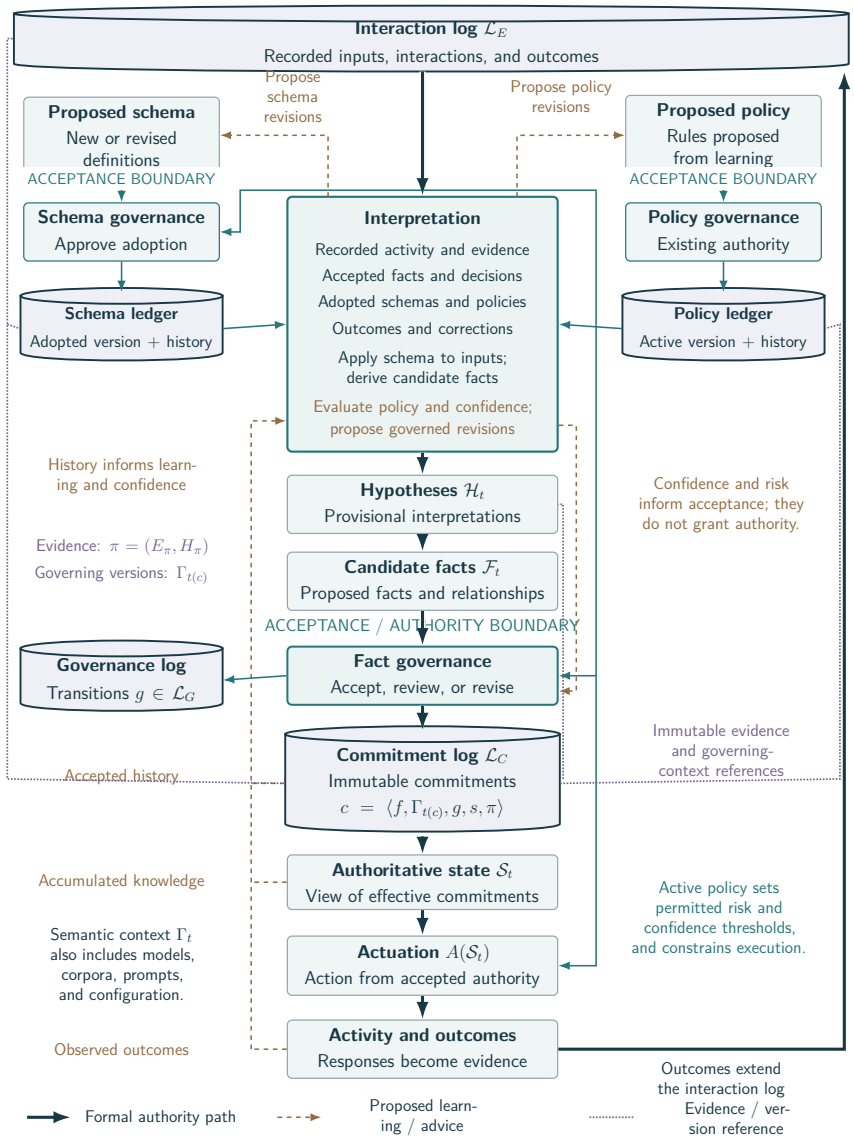


Figure 1: The governed epistemic loop. Cylinders denote durable histories. The bold path is the paper's formal authority model: provisional interpretation yields candidates; governance records transitions and immutable commitments; effective commitments define authoritative state. Separate schema and policy lifecycles, learning, confidence estimation, and risk-based acceptance are architectural extensions, not established by the stability result. Current policy governs its successor. Dotted references distinguish supporting evidence $\pi = (E_\pi, H_\pi)$ from governing context $\Gamma_{t(c)}$. Learning may propose changes but cannot alter accepted authority without governance.

Epistemic Stage	**Conceptual Role**	**System Artifact**
Knowledge	Recorded observations of the environment	Interaction log $\mathcal{L}_E$
Understanding	Semantic interpretation context	Semantic context Γ
Perception	Contextual interpretation of events	Hypotheses $\mathcal{H}_t = I(\mathcal{L}_E, \Gamma_t)$
Candidate Fact	Structured candidate assertions about system state	Candidate facts $\mathcal{F}_t = \Phi(\mathcal{H}_t; \Gamma_t)$
Decision	Governance act accepting candidate facts	Governance transition record $g \in \mathcal{L}_G$
Commitment	Immutable authority-bearing record produced by a governance decision	Commitments $c \in \mathcal{C}$
Action	Automated system response	Actuation $a \in A(\mathcal{S}_t)$
Response	Environmental effects recorded by the system	Appended events $e' \in \mathcal{L}'_E$

Table 1: Mapping stages of the epistemic decision loop to architectural artifacts in AI-native systems.

We focus on a production regime characterized by three properties: (i) interpretation is probabilistic and routinely updated as models, corpora, or policies evolve, (ii) candidate facts derived from interpretation are consumed directly by automation (e.g., routing, compliance classification, risk gating, or agentic execution), and (iii) governance transitions are sparse relative to evaluation frequency. In this regime, interpretation may be recomputed frequently while authority-establishing decisions occur infrequently.

Let $\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t)$ denote hypotheses derived from the interaction log under semantic context Γ_t, and let $\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t)$ denote the candidate fact set derived from those hypotheses. If authoritative state is defined as $\mathcal{S}_t := \mathcal{F}_t$, then authority becomes indexed to the evolving semantic context.

We formalize *interpretive variability* as the existence of admissible semantic context evolution under which the composed interpretation-to-fact map $(\Phi \circ I)$ yields different candidate fact sets from the same interaction log. A governance transition denotes the recording of an authority-establishing decision as a commitment. A *no-governance interval* is any time interval during which no such transition occurs.

Authoritative stability requires that official system state remain invariant over no-governance intervals even if interpretation semantics evolve. However, under implicit authority semantics, if interpretive variability is realized over a no-governance interval $(t_a, t_b]$, then there exist evaluation times $t'_a, t'_b \in (t_a, t_b]$ with $t'_a < t'_b$ such that

$$\mathcal{F}_{t'_a} \neq \mathcal{F}_{t'_b},$$

and therefore

$$\mathcal{S}_{t'_a} \neq \mathcal{S}_{t'_b}.$$

Authoritative state therefore violates the core invariant of systems of record: that accepted state changes only through explicit governance transitions.

In closed-loop settings—where authoritative outputs govern actuation and resulting environmental responses re-enter the interaction log—this coupling allows the effects of semantic context evolution to propagate across successive cycles of interpretation, decision, and response. Versioning semantic context improves reproducibility but does not resolve the instability: authoritative state remains indexed to the selected context, so evaluating the same interaction history under a different admissible context may still yield divergent authoritative states without any governance transition.

Prior work on provenance [14, 15], workflow governance [16], event sourcing, and ML systems engineering addresses durability, auditability, lineage,

and drift, but generally assumes that authoritative state is deterministically derived or semantically stable once accepted. The systems literature has not explicitly formalized the instability that arises when authoritative state is defined as recomputation under an admissibly evolving semantic interpretation layer.

To resolve this incompatibility, we introduce *governance-bound commitments*, immutable authority-establishing records that separate provisional fact derivation from authoritative state while preserving provenance of the semantic context under which acceptance occurred. We prove a constructive sufficiency result: under commitment semantics with forward-only construction, authoritative stability holds for any admissible semantic context evolution over no-governance intervals. The induced authority dependency graph is acyclic, ensuring non-retroactive evolution and traceable causal structure.

Finally, we outline a simulation-based evaluation framework comparing interpretation-driven and commitment-based architectures under controlled semantic context evolution. The evaluation measures retroactive mutation, actuation instability, provenance trace depth and completeness, and computational overhead in closed-loop decision systems.

Paper Organization. Section 2 summarizes the paper's contributions. Section 3 reviews related work. Section 4 defines the formal system model. Section 5 proves instability under implicit authority semantics. Section 6 introduces governance-bound commitments. Section 7 formalizes the authority dependency graph. Section 8 specifies the Authority-Driven Actuation Corollary. Section 9 describes closed-loop dynamics. Section 10 outlines the experimental evaluation framework. Section 11 discusses limitations. Section 12 outlines future work. Section 13 concludes.

2. Contributions

This paper makes the following contributions:

- **Problem formalization.** We formalize *interpretive variability* and *authoritative stability* in AI-native systems where interpretation is probabilistic and subject to admissible semantic context evolution over time. These definitions isolate the architectural tension between context-sensitive interpretation and stable authoritative state in AI-native systems.

- **Structural no-go result.** We prove that architectures employing *implicit authority semantics*—where authoritative state is defined by recomputation of candidate facts under the current semantic context—cannot guarantee authoritative stability if interpretive variability is realized over a no-governance interval.

- **Governance-bound commitment architecture.** We introduce *governance-bound commitments*, an architectural primitive that separates provisional interpretation from authoritative state through authority-establishing records with immutable provenance. We prove that commitment semantics provide a constructive separation that restores authoritative stability under admissible semantic context evolution. This separation permits components of semantic context to evolve without retroactively altering accepted authoritative state, because authority is determined exclusively through governance-bound commitments.

- **Authority dependency graph formalism.** We develop a directed authority dependency graph that captures relationships among interaction events, hypotheses, commitments, and actuation. Under commitment semantics, we prove that this graph is acyclic, yielding forward-only authority evolution, non-retroactivity of authoritative state, and traceable causal structure.

- **Evaluation framework.** We outline a simulation-based experimental framework comparing interpretation-driven and commitment-based architectures under controlled semantic context evolution, measuring retroactive mutation, actuation instability, provenance trace depth and completeness, and computational overhead in closed-loop decision systems.

3. Background and Related Work

This work lies at the intersection of deterministic systems of record, log-centric replay architectures, data provenance, workflow governance, ML systems engineering, and control-oriented system design. Prior literature addresses durability, auditability, replay, lineage, and model drift. However, these bodies of work generally assume that authoritative state is either deterministically derived or semantically stable once accepted. We instead ana-

lyze a production regime in which interpretation semantics evolve admissibly over time and authoritative outputs are consumed by automation.

We do not propose improved drift detection or calibration; rather, we isolate an architectural incompatibility between admissible semantic evolution and recomputed authority, and provide a sufficient separation construction that restores stability. Our focus is architectural: under admissible semantic context evolution, what conditions are required to preserve invariance of authoritative state absent governance transitions?

3.1. Systems of Record and Deterministic State

Established foundations. Classical systems of record define authoritative state at write time through deterministic schema validation and transactional guarantees [1, 2, 3]. ACID semantics and serializability ensure that committed state results from atomic, durable transitions over structured relations. Architectural work on database coordination further clarifies when invariants require synchronization and when they can be maintained without it [17]. In these systems, meaning is fixed at commit time and remains stable unless explicitly updated through further transactions. The relational model and its successors assume that interpretation is deterministic once schema and constraints are specified.

Limitation under probabilistic interpretation. AI-native systems introduce probabilistic interpretation layers whose outputs may vary under evolving semantic context (e.g., model updates, retrieval drift, or policy revisions). Deterministic write-time architectures do not treat interpretive variability as a first-class architectural concern. When authoritative state is defined as a recomputed function of evolving interpretation rather than as a committed artifact, invariance cannot be assumed.

Our contribution. We introduce a structural separation between interpretation and authority: authoritative state is established exclusively through governance-bound commitments, not through implicit recomputation under evolving semantic context. This extends deterministic guarantees to regimes where interpretation itself is probabilistic and non-stationary.

3.2. Event Sourcing and Log-Centric Replay

Established foundations. Event-sourced and log-centric architectures elevate immutable logs as the source of truth [4, 5, 6, 7]. State is reconstructed deterministically through replay under projection functions. Projection versioning

allows controlled evolution of processing logic while preserving replayability. Immutability has been argued to simplify reasoning and system evolution more broadly [18].

Limitation under admissible semantic context evolution. Replay semantics assume deterministic or explicitly versioned projection logic. In AI-native systems, effective interpretation may change under admissible semantic context evolution due to model updates, embedding changes, or retrieval evolution. Even with unchanged logs and governance history, recomputation under a new semantic context may yield divergent outputs. Freezing projection semantics restores determinism but prevents systems from incorporating improved models, corpora, or policy layers when interpreting historical interaction data.

Our contribution. We permit admissible evolution of semantic context without projection freezing by separating interpretation from authority. In standard event-sourced architectures, authoritative state remains defined by replay under the current projection semantics. In contrast, governance-bound commitments relocate authority from recomputation to an explicit governance artifact whose meaning is fixed at acceptance time. Authoritative state is therefore defined over committed records rather than over the outputs of replay under the current semantic context. Commitment semantics can be interpreted as introducing an explicit authority layer over event-sourced replay architectures.

3.3. Data Provenance and Lineage

Established foundations. Provenance and lineage systems record explicit dependency relationships between inputs, transformations, and outputs to support explainability, auditability, and reproducibility [14, 15, 19]. Formal models such as provenance semirings provide algebraic foundations for reasoning about derivations and their propagation under recomputation [20]. These systems clarify how a result was produced and which upstream artifacts influenced it, enabling structured reasoning about causality and traceability across data-processing pipelines.

Limitation under evolving semantics. Provenance and workflow governance systems clarify when and how state transitions are authorized, and they record lineage for accepted artifacts. However, they generally assume that the semantic interpretation of underlying data is stable at evaluation time.

They regulate approval and traceability, but do not address the structural consequences of defining authoritative state as a function of recomputation under an evolving semantic context. In regimes of interpretive variability, instability arises not from absence of governance logic, but from coupling authority directly to context-sensitive interpretation. Provenance and approval mechanisms that do not separate interpretation from authority therefore remain vulnerable to context-induced recomputation instability.

Our contribution. We embed provenance within an authority-establishing commitment primitive. Supporting hypotheses and semantic context are immutably recorded at acceptance time, and authoritative state is defined exclusively over committed records rather than over recomputed interpretations. This separation ensures that subsequent interpretive evolution cannot alter committed authority without an explicit governance transition.

3.4. Workflow Governance and Human-in-the-Loop Systems

Established foundations. Workflow engines and governance systems distinguish proposal from approval and record authorization events [21, 16]. Once approved, content is typically assumed semantically stable, and authority evolves only through workflow transitions.

Limitation under interpretive variability. In AI-native systems, the interpretation generating a proposal may evolve even when underlying events remain unchanged. Standard governance frameworks do not address regimes in which admissible semantic evolution could re-derive divergent outputs absent new governance actions.

Our contribution. We formalize governance-bound commitments as the exclusive mechanism for authoritative state evolution. Authority changes only through recorded governance transitions, not through reinterpretation alone.

3.5. AI-Native Systems, RAG, and Evolving Semantic Context

Established foundations. Retrieval-augmented generation and foundation models incorporate probabilistic reasoning into operational systems [8, 9]. Research documents sensitivity to prompt perturbations and distributional drift [10, 11, 12, 13]. Software engineering research emphasizes the operational complexity and hidden technical debt of production ML systems [22, 23]. Non-stationarity is recognized as a persistent engineering challenge.

Limitation for systems of record. While drift and variability are acknowledged, existing literature does not analyze the structural implications when authoritative state is implicitly defined as a function of current semantic interpretation. In closed-loop enterprise systems, such coupling can induce authority instability even in the absence of governance transitions.

Our contribution. We isolate probabilistic interpretation from authoritative commitment. Interpretation may drift; authority does not. This structural separation provides a minimal guarantee required for stable closed-loop automation under semantic variability.

3.6. Closed-Loop Automation and Structural Safety

Established foundations. Control theory separates uncertain estimation from actuation logic to ensure stability [24, 25]. Systems design principles emphasize separation of mechanism and policy to reduce unintended coupling [26, 27]. DAG-based dataflow systems enforce forward-only computation and causal traceability [28, 29]. Distributed data structures such as *conflict-free replicated data types* (CRDTs) guarantee convergence under concurrent updates by constraining allowable mutation patterns [30]. These systems enforce structural constraints on mutation and dependency, but do not address the distinct problem of authority defined as recomputation under an evolving semantic interpretation layer.

Limitation in AI-native enterprise systems. In many AI-native architectures, interpretation and authority are entangled: downstream actuation may depend on recomputed interpretations whose semantics evolve, enabling indirect retroactive influence on upstream authoritative state.

Our contribution. We define a forward-only authority dependency graph in which actuation depends exclusively on committed authority, and authority evolves only through governance-bound commitments. The resulting structure guarantees non-retroactive authority evolution even under admissible semantic context evolution, extending separation-of-concerns principles from control and distributed systems into probabilistic semantic architectures.

More broadly, our contribution is not a new provenance or workflow mechanism. Instead, we identify and formalize a structural instability that arises when authoritative state is defined as recomputation under an interpretation

layer whose semantics may evolve over time. While prior work addresses durability, replay, provenance, workflow governance, and model drift, these literatures generally assume that authoritative state is either deterministically derived or semantically stable once accepted. We prove that architectures defining authority through recomputation cannot guarantee authoritative stability once interpretive variability is realized, and we introduce governance-bound commitments as an architectural separation primitive that restores stability by decoupling authoritative state from evolving interpretation semantics.

4. System Model

The system model formalizes the epistemic decision loop introduced in Section 1. Interaction events constitute recorded knowledge of the system's environment; semantic context specifies the interpretive framework under which those events are understood; probabilistic inference produces hypotheses; fact derivation produces candidate facts about system state; governance decisions accept selected candidate facts; and governance transitions record those decisions as authoritative commitments. The formal objects introduced below capture these stages while separating provisional interpretation from authoritative state.

Notation is summarized in Appendix A: Notation Table; symbols are introduced formally as they arise in the model below.

4.1. Global Time Index

We assume a time domain $\mathbb{T}$ indexing system evaluations and governance transitions. An evaluation time $t \in \mathbb{T}$ denotes a point at which the system interprets the interaction log under the current semantic context and derives candidate facts about system state. Governance transitions create commitments that record authority-establishing decisions.

All time-indexed quantities (e.g., $\mathcal{C}_t$, $\mathcal{S}_t$) are defined with respect to $t \in \mathbb{T}$.

4.2. Interaction Log

Let $\mathcal{L}_E = (e_1, e_2, \ldots, e_n)$ denote an immutable sequence of semantic interaction events. These events constitute the system's recorded knowledge of observed interactions with its environment.

We write $e \in \mathcal{L}_E$ to denote that e is an element of the sequence (i.e., one of the recorded events).

The interaction log records observable system history but does not itself define authoritative state. Instead, it serves as the knowledge substrate from which interpretation produces hypotheses and candidate facts. This design is consistent with log-centric and event-sourced architectures in which immutable event histories form the basis for derived system state [4, 7].

4.3. Semantic Context

Let Γ denote the *semantic context* under which interpretation is performed. The semantic context comprises model parameters, retrieval corpora, policy rules, schema constraints, and configuration state that collectively determine how interaction history is interpreted at evaluation time. In the epistemic loop introduced earlier, Γ represents the system's operational *understanding*: the interpretive framework through which recorded knowledge is evaluated.

In the architecture studied here, schema is treated as a component of semantic context rather than as an independent source of authoritative state. Schema defines the admissible structure of facts in $\mathcal{F}$, while interpretation derives candidate facts under the current semantic context Γ. Authoritative state, however, is determined only through governance-bound commitments.

A *retrieval corpus* is the collection of indexed documents or knowledge artifacts from which context is dynamically retrieved to condition model inference in retrieval-augmented systems. Policy rules within Γ specify validation logic, acceptance thresholds, conflict-resolution rules, and decision constraints that govern fact derivation. A *policy revision* is a modification to these rule-layer constraints that can alter candidate facts without changing the underlying interaction log or recording a governance transition.

Unlike classical projection logic in event-sourced systems—where interpretation semantics are typically deterministic and explicitly versioned—semantic context in AI-native architectures may evolve over time. Such evolution may arise from model updates, retrieval corpus modifications, policy revisions, or configuration changes. These forms of non-stationarity are consistent with known phenomena of model evolution, distributional shift, and technical debt in machine learning systems [12, 13].

Admissible Semantic Context Evolution. We restrict attention to *admissible* transitions of Γ: operationally valid updates that preserve system well-formedness. We do not assume adversarial or arbitrary mutation of semantic context. Rather, Γ evolves only through changes considered legitimate by the system designer (e.g., deployment of a new model version,

updating retrieval indices, refining policy constraints, or adjusting configuration parameters).

The instability result established later requires only that admissible semantic context evolution induce a non-constant interpretation-to-fact map over a no-governance interval. Formally, it suffices that there exist admissible contexts Γ_a and Γ_b such that

$$\Phi\big(I(\mathcal{L}_E, \Gamma_a); \Gamma_a\big) \neq \Phi\big(I(\mathcal{L}_E, \Gamma_b); \Gamma_b\big).$$

No stronger assumptions about the magnitude, frequency, or mechanism of context evolution are required.

4.4. Inference Function

Let

$$I : (\mathcal{L}_E, \Gamma_t) \to \mathcal{H}_t$$

denote the inference function that maps the interaction log and semantic context at time t to a set of hypotheses $\mathcal{H}_t$. In the epistemic loop introduced earlier, inference corresponds to the system's *perception*: the contextual interpretation of recorded events under the current semantic context.

Inference may be probabilistic and context-sensitive; identical interaction histories may yield different hypothesis sets under admissible variation of Γ_t. Hypotheses therefore represent provisional interpretations of the interaction log. They are intermediate artifacts of interpretation and do not themselves constitute authoritative state.

4.5. Interpretive Variability

Interpretive variability is a structural property of the interpretation pipeline under admissible variation of semantic context. In the epistemic loop introduced earlier, inference produces hypotheses $\mathcal{H}_t = I(\mathcal{L}_E, \Gamma_t)$, and fact derivation transforms those hypotheses into candidate facts $\mathcal{F}_t = \Phi(\mathcal{H}_t; \Gamma_t)$. Because both stages may depend on semantic context, admissible evolution of Γ can cause the same interaction history to yield different candidate fact sets.

Definition 4.1 (Interpretive Variability Realized Over an Interval)**.** Interpretive variability is realized over an interval $(t_a, t_b]$ if the interaction log $\mathcal{L}_E$ remains fixed over that interval and there exist evaluation times $t'_a, t'_b \in (t_a, t_b]$ with $t'_a < t'_b$ such that

$$\Phi\big(I(\mathcal{L}_E, \Gamma_{t'_a}); \Gamma_{t'_a}\big) \neq \Phi\big(I(\mathcal{L}_E, \Gamma_{t'_b}); \Gamma_{t'_b}\big).$$

In AI-native systems, interpretive variability arises naturally during normal operation as model parameters, retrieval corpora, prompts, or policy rules evolve over time while the interaction log remains fixed.

This definition captures situations in which admissible semantic context evolution produces divergent candidate fact sets from the same interaction history, rather than merely permitting such divergence in principle.

We allow Γ to influence I and/or Φ; interpretive variability is therefore a property of the composed interpretation-to-fact pipeline $\Phi \circ I$ under admissible semantic context evolution.

4.6. Implicit Authority Semantics

Under implicit authority semantics, inference produces hypotheses

$$\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t).$$

The system then derives structured assertions from those hypotheses via a (possibly context-dependent) fact-derivation map

$$\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t), \qquad \mathcal{F}_t \subseteq \mathcal{F}.$$

In the epistemic loop described earlier, these correspond to provisional interpretive conclusions about system state, represented formally as candidate facts.

Under implicit authority semantics, authoritative state is defined directly from these derived assertions:

$$\mathcal{S}_t := \mathcal{F}_t.$$

Thus candidate facts produced by interpretation are treated as authoritative system state at evaluation time.

We assume that both I and Φ may depend on semantic context Γ_t, and therefore need not be invariant across admissible semantic contexts.

In classical event-sourced systems, interpretation semantics are typically deterministic or explicitly versioned through projection logic. In the AI-native systems considered here, semantic context Γ may evolve through model updates, retrieval corpus changes, or policy revisions without a governance transition, potentially altering the candidate fact set $\mathcal{F}_t$ even when the interaction log remains unchanged.

4.7. Authoritative Stability

Definition 4.2 (No-governance interval)**.** An interval $(t_a, t_b]$ is a *no-governance interval* if $\neg\mathrm{Gov}(t_a, t_b)$.

Definition 4.3 (Authoritative Stability)**.** A system satisfies authoritative stability if, for all evaluation times $t_a < t_b$,

$$\text{if } (t_a, t_b] \text{ is a no-governance interval, then } \mathcal{S}_{t_a} = \mathcal{S}_{t_b}.$$

Authoritative stability requires that accepted system state not change solely due to admissible variation in interpretation context. In the epistemic loop introduced earlier, this condition ensures that authoritative state induced by governed commitments remains invariant in the absence of governance transitions, even if perception and candidate facts produced by interpretation evolve.

5. Instability of Implicit Authority

We begin by illustrating the architectural distinction between interpretation-driven authority and governance-bound authority. The key difference lies in where authoritative state is established within the interpretation pipeline.

Implicit authority: $\mathcal{L}_E \to \mathcal{H}_t \to \mathcal{F}_t \equiv \mathcal{S}_t \to A(\mathcal{S}_t)$

Commitment semantics: $\mathcal{L}_E \to \mathcal{H}_t \to \mathcal{F}_t \to \mathcal{C}_t \to \mathcal{S}_t \to A(\mathcal{S}_t)$
where $\mathcal{C}_t$ records governance-bound commitments.

Under implicit authority semantics, authoritative state is obtained directly from recomputation under the current semantic context. Under commitment semantics, authoritative state arises only through governance transitions that create commitments. The two architectures may share the same interpretation pipeline but differ in where authoritative state is established.

5.1. Intuition: Standard Event Sourcing Under Semantic Drift

We first present an intuition-level replay counterexample, aligned with the experimental framework in Section 10, to illustrate how context-coupled

Architectural Aspect	**Implicit Authority Semantics**	**Commitment Semantics**
Perception	Hypotheses derived as $\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t)$.	Same: hypotheses derived as $\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t)$.
Candidate fact derivation	Candidate facts derived as $\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t)$.	Same: candidate facts derived as $\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t)$.
Decision rule	Authoritative state defined by recomputation: $\mathcal{S}_t := \mathcal{F}_t$.	Authoritative state defined only through commitments: $\mathcal{S}_t = \{\, c.f \mid c \in \mathcal{C}_t \wedge \mathrm{Eff}_t(c) \,\}$.
Authority source	Current semantic interpretation (recomputation).	Governance transitions (commitments).
Effect of admissible semantic context evolution	Admissible semantic context evolution may alter authoritative state even with unchanged log and no governance transition.	Admissible semantic context evolution may alter hypotheses or candidate facts, but not authoritative state absent governance transitions.
Closed-loop actuation	Actuation may vary because authority varies with recomputation.	Actuation depends only on effective commitments and remains traceable to governed decisions.
Stability over no-governance intervals	Not guaranteed if interpretive variability is realized.	Guaranteed by construction.

Table 2: Comparison of implicit authority semantics and governance-bound commitment semantics. Both architectures may share the same interpretation pipeline but differ in where authority is established.

authority behaves under admissible semantic drift. The example uses a standard replay architecture similar to event-sourced systems, where authoritative state is recomputed from an immutable log under the current interpretation semantics. The purpose of this example is not to prove the general result, but to make the instability concrete before stating the structural theorem.

A familiar operational analogue is a classification or routing system in which the same historical interaction log is replayed under an updated model, retrieval corpus, or policy layer. The underlying record has not changed, but the current interpretation of that record has.

Setup

Let $\mathcal{L}_E = (e_1, \ldots, e_n)$ be an immutable event log. Let the governance transition log $\mathcal{L}_G$ be fixed.

As a concrete instance of interpretation-derived authority, consider a standard replay architecture in which authoritative state at evaluation time t is defined by recomputation of the interpretation pipeline under the current semantic context:

$$\mathcal{S}_t = \Phi\big(I(\mathcal{L}_E, \Gamma_t); \Gamma_t\big),$$

where interpretation I and fact derivation Φ are evaluated under semantic context Γ_t. This formulation corresponds to the common case in which authoritative state is defined directly as the candidate fact set produced by the interpretation-to-fact pipeline.

Assume:

- The log $\mathcal{L}_E$ does not change.
- The governance transition log $\mathcal{L}_G$ does not change.
- Semantic context Γ evolves admissibly to Γ' (e.g., model update, retrieval corpus shift, prompt revision, or policy adjustment).

Divergence Under Replay

At time t_1:

$$\mathcal{S}_{t_1} = \Phi\big(I(\mathcal{L}_E, \Gamma); \Gamma\big).$$

At time t_2:

$$\mathcal{S}_{t_2} = \Phi\big(I(\mathcal{L}_E, \Gamma'); \Gamma'\big).$$

If interpretation is context-sensitive, then admissible evolution of semantic context may induce

$$\Phi\big(I(\mathcal{L}_E, \Gamma); \Gamma\big) \neq \Phi\big(I(\mathcal{L}_E, \Gamma'); \Gamma'\big),$$

and therefore

$$\mathcal{S}_{t_1} \neq \mathcal{S}_{t_2}.$$

This divergence occurs despite:

- No new events,
- No governance transition,
- No explicit freezing of interpretation semantics.

Implication

Standard replay-based architectures guarantee invariance only if projection functions are deterministic or projection semantics are frozen. Freezing projection semantics restores invariance by preventing semantic evolution, but doing so defeats a central purpose of AI-native interpretation layers: allowing improved models, corpora, and policies to reinterpret historical interaction data. Allowing admissible semantic evolution, however, permits authoritative state to change despite an unchanged event log and governance history.

This tension is not specific to event sourcing. It reflects a more general architectural choice: defining authoritative state as the output of recomputation under a context-evolving interpretation-to-fact pipeline. When authority is coupled directly to Γ, accepted meaning becomes a function of the current semantic context rather than a governed artifact tied to explicit transitions.

The replay example above illustrates the instability in a familiar operational setting. We now abstract from the replay mechanism and focus on the underlying architectural property: authority defined as the output of an interpretation layer whose semantics may evolve over time. In this broader formulation, instability arises not from replay itself but from coupling authoritative state to a context-sensitive interpretation process.

5.2. No-go Framing: Implicit Authority Is Context-Coupled

Under implicit authority semantics, authoritative state is obtained directly from recomputation under the current semantic context. Formally, inference and fact derivation are evaluated under the semantic context Γ_t:

$$\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t), \qquad \mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t), \qquad \mathcal{S}_t := \mathcal{F}_t.$$

Authoritative state is therefore determined by the interaction log together with the semantic context under which interpretation is performed, rather than by a committed governance artifact. If admissible semantic context evolution alters the output of the interpretation-to-fact pipeline while the interaction log remains fixed, then candidate fact sets may differ at distinct evaluation times within the same no-governance interval. Because authoritative state is defined directly from those candidate facts, such divergence implies that authoritative state also differs across those evaluation times despite the absence of a governance transition.

This instability does not depend on any particular inference algorithm or replay implementation. It follows from the architectural decision to define authority as the output of a context-sensitive interpretation layer whose semantics may evolve during normal system operation. When authority is constructed in this manner, authoritative stability over no-governance intervals cannot be guaranteed if interpretive variability is realized. The following theorem formalizes this incompatibility between context-coupled authority and authoritative stability.

5.3. No-go theorem: context-coupled authority and authoritative stability are incompatible

Theorem 5.1 (Structural Incompatibility of Context-Evolving Interpretation and Stable Authority)**.** *Consider any architecture satisfying the following properties:*

1. ***Admissible semantic context evolution:*** *the semantic context* Γ *may evolve during normal system operation.*
2. ***Interpretation-derived authority:*** *authoritative state is defined by recomputation of the interpretation-to-fact pipeline from the interaction log under the current semantic context; that is, authoritative state is determined by the candidate fact set produced by interpretation.*
3. ***Authoritative stability:*** *authoritative state remains invariant over no-governance intervals.*

If interpretive variability is realized over a no-governance interval, these three properties cannot be satisfied simultaneously.

Proof. Suppose interpretive variability is realized over a no-governance interval $(t_a, t_b]$. By Definition 4.1, there exist evaluation times $t'_a, t'_b \in (t_a, t_b]$ with $t'_a < t'_b$ such that

$$\Phi\big(I(\mathcal{L}_E, \Gamma_{t'_a}); \Gamma_{t'_a}\big) \neq \Phi\big(I(\mathcal{L}_E, \Gamma_{t'_b}); \Gamma_{t'_b}\big).$$

Because the architecture defines authoritative state by recomputation of the interpretation-to-fact pipeline under the current semantic context, authoritative state at evaluation time t is determined by the candidate fact set produced by that pipeline under Γ_t. Therefore the divergence above implies

$$\mathcal{S}_{t'_a} \neq \mathcal{S}_{t'_b}.$$

Since $(t_a, t_b]$ is a no-governance interval, no governance transition occurs within that interval. This contradicts authoritative stability (Definition 4.3), which requires authoritative state to remain invariant across no-governance intervals. Hence the three properties cannot all hold simultaneously. □

6. Governance-Bound Commitments and the Necessity of Separation

The instability established in Section 5 arises because authoritative state is defined as the output of a context-sensitive interpretation layer. To restore stability while permitting admissible evolution of semantic context, authoritative state must instead be derived from an artifact whose meaning is fixed at the time of acceptance. We introduce *governance-bound commitments* as this separation primitive.

In the epistemic decision loop introduced earlier, interpretation produces candidate facts about system state. Governance processes then make authority-establishing decisions that accept selected candidate facts. These decisions are recorded in the system as immutable artifacts called *commitments*. Each commitment records the accepted fact together with the semantic context under which acceptance occurred.

6.1. Commitment Primitive

Governance transitions conceptually correspond to authority-establishing *decisions*. In the system model, however, authoritative state evolves through the immutable artifacts that record those decisions. We refer to these artifacts as *commitments*. A commitment is therefore the persisted system record through which a governance decision enters the authoritative state.

Definition 6.1 (Commitment)**.** A commitment is an immutable authority-establishing record that captures an accepted fact together with the governance transition, semantic context, operational status annotation, and provenance associated with that acceptance. Formally,

$$c = \langle f, \Gamma_{t(c)}, g, s, \pi \rangle,$$

where $f \in \mathcal{F}$ is the accepted fact, $\Gamma_{t(c)}$ is the semantic context at commitment time, $g \in \mathcal{L}_G$ is the governance transition that created the commitment, $s \in \{\text{active}, \text{superseded}\}$ is a stored status label, and π is immutable provenance metadata. The timestamp $t(c)$ denotes the time at which commitment c is recorded.

The authoritative effect of a commitment is determined by the supersession relation and the effectiveness predicate defined below. The stored status label s is an operational annotation included for clarity and auditability and does not determine authoritative state.

6.2. Commitment Set

Let $\mathcal{C}$ denote the set of all commitments recorded by the system.

Let $\mathcal{L}_C = (c_1, c_2, \dots)$ denote the append-only commitment log, i.e., the time-ordered sequence of commitment records created through governance transitions.

Each commitment record in $\mathcal{L}_C$ is created only through a governance transition recorded in $\mathcal{L}_G$, and no commitment may be added to $\mathcal{L}_C$ except via such a transition. The governance log $\mathcal{L}_G$ records authority-establishing transitions at the workflow level, while the commitment log $\mathcal{L}_C$ records the resulting commitment artifacts over which authoritative state is formally defined. Accordingly, $\mathcal{L}_G$ tracks when authority-changing governance events occur, whereas $\mathcal{L}_C$ stores the immutable authority-bearing records produced by those events.

Let $\mathcal{C}_t \subseteq \mathcal{C}$ denote the set of commitments recorded up to time t, i.e.,

$$\mathcal{C}_t := \{\, c \in \mathcal{C} \mid t(c) \leq t \,\}.$$

Equivalently, $\mathcal{C}_t$ is the set obtained by replaying the append-only commitment log $\mathcal{L}_C$ up to time t.

Consequently, for any $t_a < t_b$,

$$\neg\mathrm{Gov}(t_a, t_b) \;\Rightarrow\; \mathcal{C}_{t_a} = \mathcal{C}_{t_b}.$$

6.3. Keyed Facts and Effective Commitments

Let $\mathcal{F}$ denote the space of facts. In the architectural interpretation of the model, the fact space $\mathcal{F}$ corresponds to the system's schema: it defines the admissible structure of facts that may be derived or committed. Interpretation derives candidate facts within this schema, while governance transitions determine which of those facts become authoritative commitments.

In the epistemic interpretation of the system introduced earlier, candidate facts correspond to structured assertions about system state derived from interpretation. Governance processes select some of these facts for acceptance and record them as commitments.

Each fact $f \in \mathcal{F}$ has an associated key

$$\kappa : \mathcal{F} \to \mathcal{K},$$

which identifies the entity–attribute slot about which the fact is asserted (e.g., an entity identifier together with an attribute name).

For a commitment c, we write $c.f$ to denote the fact recorded in that commitment.

Let $\prec$ denote the supersession relation on commitments with the same key. The relation is assumed to be time-respecting:

$$c \prec c' \;\Rightarrow\; \kappa(c.f) = \kappa(c'.f) \;\text{ and }\; t(c) < t(c').$$

Thus supersession is forward-only: a later governance decision may replace an earlier commitment for the same key, but cycles cannot arise because commitments are created monotonically in time.

Define the effectiveness predicate $\mathrm{Eff}_t(c)$ to hold if and only if there exists no commitment $c' \in \mathcal{C}_t$ such that

$$\kappa(c'.f) = \kappa(c.f) \;\text{ and }\; c \prec c'.$$

The stored status label $s \in \{\text{active}, \text{superseded}\}$ is maintained as part of the commitment record to support operational clarity and audit readability.

The authoritative definition of current state is determined by the effectiveness predicate $\mathrm{Eff}_t(c)$ rather than by the stored label itself.

Effectiveness therefore selects the authoritative commitment per key among commitments recorded up to time t (i.e., the maximal element under $\prec$ when it exists). Accordingly,

$$\mathcal{S}_t = \{\, c.f \mid c \in \mathcal{C}_t \wedge \mathrm{Eff}_t(c) \,\}.$$

6.4. Commitment Provenance

Each commitment records immutable provenance describing the semantic lineage supporting acceptance of the committed fact.

Formally, the provenance component π associated with commitment c is defined as

$$\pi = (\mathcal{E}_\pi, \mathcal{H}_\pi),$$

where $\mathcal{E}_\pi \subseteq \mathcal{L}_E$ denotes the subset of interaction events supporting the decision, and $\mathcal{H}_\pi \subseteq \mathcal{H}_{t(c)}$ denotes the subset of hypotheses derived from those events that motivated acceptance of the fact.

In the epistemic interpretation of the system introduced earlier, interaction events represent recorded *knowledge* and hypotheses represent system *perception*. Provenance therefore captures the semantic pathway through which knowledge and perception give rise to the governance decision recorded as commitment c.

Provenance is immutable once recorded; subsequent evolution of semantic context or interpretation mechanisms does not alter the provenance of existing commitments.

6.5. Authoritative State

Authoritative system state is determined exclusively by effective commitments recorded in the commitment log. Concretely, the state definition introduced above is evaluated over the set of effective commitments in $\mathcal{C}_t$, so authoritative state depends on recorded commitments and their supersession structure rather than on recomputation of candidate facts under the current semantic context.

Interpretation may produce new hypotheses and candidate facts, but these artifacts do not modify authoritative state unless they are accepted through a governance transition and recorded as commitments. Consequently, authoritative state evolves only through governance actions that append new commitments to the commitment log.

6.6. Sufficiency for Stability

Theorem 6.2 (Sufficiency of Commitment Semantics)**.** *Under governance-bound commitment semantics (Section 6), authoritative stability (Definition 4.3) holds for any admissible semantic context evolution in the absence of governance transitions.*

Proof. Under commitment semantics, authoritative state at time t is derived exclusively from effective commitments:

$$\mathcal{S}_t = \{\, c.f \mid c \in \mathcal{C}_t \wedge \mathrm{Eff}_t(c) \,\}.$$

The commitment set $\mathcal{C}_t$ evolves only through governance transitions recorded in $\mathcal{L}_G$. Interpretation and hypothesis generation,

$$\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t),$$

as well as admissible evolution of semantic context Γ_t, do not modify $\mathcal{C}_t$ or the supersession relation $\prec$.

Let $t_a < t_b$ with $\neg\mathrm{Gov}(t_a, t_b)$. Then no governance transition occurs in $(t_a, t_b]$, implying

$$\mathcal{C}_{t_a} = \mathcal{C}_{t_b}.$$

Because the effectiveness predicate $\mathrm{Eff}_t(c)$ depends only on $\mathcal{C}_t$ and $\prec$, the set of effective commitments is identical at t_a and t_b. Consequently,

$$\mathcal{S}_{t_a} = \mathcal{S}_{t_b}.$$

Thus authoritative state remains invariant over any no-governance interval, establishing authoritative stability. □

6.7. Provenance Stability

Proposition 6.3 (Provenance Invariance)**.** *Interpretive variability does not alter the recorded provenance $\pi = (\mathcal{E}_\pi, \mathcal{H}_\pi)$ of existing commitments in the absence of governance transitions.*

Proof. By Definition 6.1, each commitment

$$c = \langle f, \Gamma_{t(c)}, g, s, \pi \rangle$$

records immutable provenance

$$\pi = (\mathcal{E}_\pi, \mathcal{H}_\pi)$$

at the time of governance transition.

Commitments are stored in the append-only commitment log $\mathcal{L}_C$, and once recorded neither the committed fact $c.f$ nor the associated provenance π may be modified.

Interpretation may produce new hypotheses

$$\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t)$$

under evolving semantic contexts Γ_t, but these operations do not modify existing commitment records.

If $(t_a, t_b]$ is a no-governance interval, then by construction no new commitments are appended to $\mathcal{L}_C$, implying

$$\mathcal{C}_{t_a} = \mathcal{C}_{t_b}.$$

Since provenance π is stored within each commitment record and commitments are immutable once recorded, the provenance associated with all commitments in $\mathcal{C}_{t_a}$ remains unchanged at time t_b. Therefore interpretive variability cannot alter the provenance of existing commitments over any no-governance interval. □

6.8. General Necessity for Context-Coupled Authority

The following definition formalizes the broader class of architectures in which authoritative state is *context-coupled* to interpretation.

Definition 6.4 (Context-Coupled Authority)**.** An architecture has context-coupled authority if there exists a function

$$R : (\mathcal{L}_E, \Gamma) \to \mathcal{P}(\mathcal{F})$$

such that, for each evaluation time t,

$$\mathcal{S}_t := R(\mathcal{L}_E, \Gamma_t).$$

Here $\mathcal{P}(\mathcal{F})$ denotes the power set of $\mathcal{F}$, so R returns a set of authoritative facts. In such architectures, authoritative state is defined by recomputation as a function of the interaction log and the current semantic context rather than by reference to an independent commitment history. No structural assumptions are made about R beyond its dependence on $(\mathcal{L}_E, \Gamma)$; in particular, it need not factor as $\Phi \circ I$.

Implicit authority semantics is the special case of context-coupled authority in which

$$R(\mathcal{L}_E, \Gamma) := \Phi\big(I(\mathcal{L}_E, \Gamma); \Gamma\big),$$

so that authoritative state at evaluation time t is defined by

$$\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t), \qquad \mathcal{S}_t := \Phi(\mathcal{H}_t; \Gamma_t).$$

Here I and/or Φ may depend on the semantic context Γ_t, and authoritative state is therefore defined solely by recomputation under the current interpretation context.

The following lemma isolates the structural source of instability: when authority is defined as a function of the current semantic context, invariance over no-governance intervals requires that the authority function remain constant across admissible contexts.

Lemma 6.5 (Instability of Context-Coupled Authority)**.** *Consider any context-coupled authority architecture (Definition 6.4). If there exist evaluation times $t_a < t_b$ such that*

$$\neg\mathrm{Gov}(t_a, t_b) \quad \textit{and} \quad R(\mathcal{L}_E, \Gamma_{t_a}) \neq R(\mathcal{L}_E, \Gamma_{t_b}),$$

then authoritative stability (Definition 4.3) does not hold.

Proof. By Definition 6.4,

$$\mathcal{S}_{t_a} = R(\mathcal{L}_E, \Gamma_{t_a}), \qquad \mathcal{S}_{t_b} = R(\mathcal{L}_E, \Gamma_{t_b}).$$

If these values differ while $\neg\mathrm{Gov}(t_a, t_b)$ holds, then

$$\mathcal{S}_{t_a} \neq \mathcal{S}_{t_b}$$

in the absence of a governance transition, contradicting authoritative stability (Definition 4.3). □

Versioning the semantic context (e.g., model versions, retrieval indices, prompts, or policy configurations) improves reproducibility but does not resolve instability if authoritative state remains defined by recomputation under the selected context. Under context-coupled authority, authoritative state remains Γ-indexed rather than governance-indexed: evaluating the same interaction log under a different admissible semantic context may

yield a different authoritative state even when the governance history is unchanged. Projection or context versioning therefore preserves replayability but not authoritative stability. Stability in the sense of Definition 4.3 requires architectural separation of authority from recomputation itself, not merely preservation or versioning of semantic configurations.

6.9. Scoped Minimal Sufficiency

We characterize the architectural regime in which separation between interpretation and authority is required to ensure authoritative stability.

Consider systems with the following properties:

1. Admissible semantic context evolution under which interpretive variability may be realized over a no-governance interval (Definition 4.1).
2. Closed-loop actuation whose behavior depends on authoritative state.

In such systems, if interpretive variability is realized over a no-governance interval, authoritative state cannot remain invariant when it is defined by context-coupled recomputation under the current semantic context. Lemma 6.5 therefore implies that, in architectures where interpretive variability is realized and authoritative state is context-coupled to interpretation, separation of authority from context-coupled recomputation is necessary to ensure authoritative stability within this regime.

Governance-bound commitments with immutable provenance provide such a separation and are sufficient to guarantee authoritative stability (Theorem 6.2). This separation yields additional structural consequences, including traceable authority evolution and non-retroactivity (Proposition 7.2 and Corollary 7.3 in Section 7).

The minimality claim is scoped to this architectural regime. If interpretation is deterministic and context-invariant, interpretive variability does not arise and authoritative instability under implicit semantics may not occur. The necessity of commitment-based separation therefore applies specifically to architectures in which identical interaction histories may yield distinct hypotheses under admissible semantic contexts.

We do not claim uniqueness: alternative separation primitives may also satisfy authoritative stability. However, within this regime, any stable design must prevent context-coupled recomputation under evolving semantic context Γ from directly redefining authoritative state in the absence of governance transitions.

Taken together, Lemma 6.5 and Theorem 6.2 establish a scoped necessity–sufficiency characterization for authoritative stability under admissible semantic context evolution. Architectures in which authoritative state is recomputed solely as a function of the current semantic context are structurally unstable when that context may evolve without governance transitions. Conversely, a separation primitive that decouples authoritative state from evolving semantic context—such as governance-bound commitments with immutable provenance—is sufficient to restore stability over no-governance intervals.

7. Authority DAG Formalism

7.1. Authority Dependency Graph

Authority relationships in the system are represented by a directed dependency graph that captures how interaction events, interpretation artifacts, governance decisions, and system actions depend on one another.

Define the authority dependency graph

$$G = (V, E),$$

where the vertex set and edge set are defined as follows.

The vertex set is

$$V := \mathcal{L}_E \ \cup \ \mathcal{H}_{\text{all}} \ \cup \ \mathcal{C} \ \cup \ \mathcal{A}, \qquad \mathcal{H}_{\text{all}} := \bigcup_{t \in \mathbb{T}} \mathcal{H}_t.$$

Thus vertices represent interaction events ($\mathcal{L}_E$), hypotheses produced during interpretation ($\mathcal{H}_{\text{all}}$), commitment records ($\mathcal{C}$), and actuation events or actions ($\mathcal{A}$).

The directed edge set $E \subseteq V \times V$ encodes authority dependencies between artifacts. An edge $(u, v) \in E$ indicates that artifact v depends on artifact u in the authority construction process.

Edges are introduced only according to the following dependency rules:

- **Event → Hypothesis:**

  $(e, h) \in E$ iff event $e \in \mathcal{L}_E$ contributes to the derivation of hypothesis $h \in \mathcal{H}_{t(h)}$, where $t(h)$ denotes the evaluation time at which h is produced and $\mathcal{H}_{t(h)} := I(\mathcal{L}_E, \Gamma_{t(h)})$.

- **Hypothesis → Commitment:**

  $(h, c) \in E$ iff $h \in \mathcal{H}_\pi$ for commitment c, i.e., the hypothesis is explicitly referenced in the provenance recorded at commitment time.

- **Commitment → Commitment (supersession):**

  $(c, c') \in E$ iff $c \prec c'$.

- **Commitment → Actuation:**

  $(c, a) \in E$ iff $c \in \rho(a)$, where $\rho(a)$ records the subset of governing commitments responsible for the actuation a.

- **Actuation → Event:**

  $(a, e') \in E$ iff actuation event $a \in \mathcal{A}$ produces an interaction event e' appended to the successor log $\mathcal{L}'_E$, with $t(a) < t(e')$.

The actuation provenance mapping $\rho(a)$ is recorded at actuation time and is immutable thereafter, analogously to the commitment provenance π recorded at commitment time.

In particular, hypothesis-to-commitment edges are provenance-scoped: commitments may depend only on hypotheses available at their commitment time $t(c)$, as recorded by $\mathcal{H}_\pi \subseteq \mathcal{H}_{t(c)}$.

7.2. Forward-Construction Property

Define a creation-time function

$$\tau : V \to \mathbb{T}$$

that assigns to each vertex the time at which the corresponding artifact is created or recorded in the system.

For hypotheses, let $t(h)$ denote the unique evaluation time such that $h \in \mathcal{H}_{t(h)}$. For hypothesis vertices we set $\tau(h) := t(h)$.

In particular,

$$\begin{aligned} \tau(e) &= t(e) \quad \text{for } e \in \mathcal{L}_E, \\ \tau(h) &= t(h) \quad \text{for } h \in \mathcal{H}_{\text{all}}, \\ \tau(c) &= t(c) \quad \text{for } c \in \mathcal{C}, \\ \tau(a) &= t(a) \quad \text{for } a \in \mathcal{A}. \end{aligned}$$

Lemma 7.1 (Forward Construction)**.** *For every edge $(u, v) \in E$, we have*

$$\tau(u) \leq \tau(v).$$

Proof. Edges in E are introduced only when their target node is created, and dependencies are recorded only to artifacts that already exist. This reflects the append-only recording of provenance and actuation dependencies, which disallows forward references.

For each edge type:

- **Event $\rightarrow$ Hypothesis:**

  If $(e, h) \in E$, then $h \in \mathcal{H}_{t(h)}$ is derived from events in $\mathcal{L}_E$ recorded no later than $t(h)$, hence $\tau(e) \leq t(h) = \tau(h)$.

- **Hypothesis $\rightarrow$ Commitment:**

  If $(h, c) \in E$, then $h \in \mathcal{H}_\pi \subseteq \mathcal{H}_{t(c)}$, thus $\tau(h) \leq t(c) = \tau(c)$.

- **Commitment $\rightarrow$ Commitment (supersession):**

  If $(c, c') \in E$ because $c \prec c'$, then by time-respecting construction $t(c) < t(c')$, hence $\tau(c) < \tau(c')$.

- **Commitment $\rightarrow$ Actuation:**

  If $(c, a) \in E$, then $c \in \rho(a) \subseteq \mathcal{C}_{t(a)}$, therefore $\tau(c) \leq t(a) = \tau(a)$.

- **Actuation $\rightarrow$ Event:**

  If $(a, e') \in E$, then by construction $t(a) < t(e')$, hence $\tau(a) < \tau(e')$.

In all cases $\tau(u) \leq \tau(v)$. □

7.3. Acyclic Authority Structure

Proposition 7.2 (Acyclicity)**.** *Under governance-bound commitment semantics, the authority dependency graph $G = (V, E)$ is acyclic.*

Proof. Suppose, for contradiction, that G contains a directed cycle

$$v_0 \rightarrow v_1 \rightarrow \cdots \rightarrow v_k = v_0.$$

By Lemma 7.1, every edge $(u, v) \in E$ satisfies $\tau(u) \leq \tau(v)$. Hence

$$\tau(v_0) \leq \tau(v_1) \leq \cdots \leq \tau(v_k) = \tau(v_0),$$

which implies

$$\tau(v_0) = \tau(v_1) = \cdots = \tau(v_k).$$

We first eliminate edge types that enforce strict time increase.

- **Commitment → Commitment (supersession):**

  If $(c, c') \in E$ because $c \prec c'$, then by time-respecting construction $t(c) < t(c')$, hence $\tau(c) < \tau(c')$, contradicting constancy of τ along the cycle.

- **Actuation → Event:**

  If $(a, e') \in E$, then by construction $t(a) < t(e')$, hence $\tau(a) < \tau(e')$, again contradicting constancy of τ.

Therefore no directed cycle can contain a supersession edge or an actuation→event edge.

Any remaining directed cycle must consist only of edges of the types event→hypothesis, hypothesis→commitment, and commitment→actuation. These edge types respect a strict type-level ordering of artifact classes

$$\mathcal{L}_E \prec \mathcal{H}_{\text{all}} \prec \mathcal{C} \prec \mathcal{A}$$

in the sense that event→hypothesis, hypothesis→commitment, and commitment→actuation edges each map a vertex in one class to a vertex in the next class in this ordering. There are no edges that reverse this ordering (i.e., none from $\mathcal{H}_{\text{all}}$ to $\mathcal{L}_E$, from $\mathcal{C}$ to $\mathcal{H}_{\text{all}}$, or from $\mathcal{A}$ to $\mathcal{C}$).

Therefore any directed walk composed solely of these edge types must move strictly forward in this class hierarchy and cannot return to its starting vertex. Hence no directed cycle can exist.

This contradiction shows that G is acyclic. □

7.4. Non-Retroactivity

Corollary 7.3 (Non-Retroactivity)**.** *Authoritative evolution under governance-bound commitment semantics is non-retroactive.*

Proof. Let $u \to^* v$ denote reachability in G. By Lemma 7.1, along any directed path $u \to^* v$ we have $\tau(u) \leq \tau(v)$. Consequently, no artifact recorded at a later time can have a directed path to an artifact recorded at a strictly earlier time.

Since commitments and their provenance are immutable once recorded, later artifacts cannot modify the content or provenance of previously recorded commitments. Hence authoritative commitments cannot be retroactively altered by subsequent system activity. □

8. Authority-Driven and Traceable Actuation

Corollary 8.1 (Authority-Driven Actuation)**.** *Under governance-bound commitment semantics with actuation rules defined over authoritative state, every actuation event is governed by effective commitments.*

Formally, for any actuation event $a \in \mathcal{A}$ recorded at time $t(a)$, there exists a non-empty provenance set

$$\rho(a) \subseteq \mathcal{C}_{t(a)}$$

such that

$$\forall c \in \rho(a),\ \mathrm{Eff}_{t(a)}(c),$$

and

$$a \in A(\mathcal{S}_{t(a)}).$$

Proof. Under commitment semantics, actuation rules are evaluated over authoritative state

$$\mathcal{S}_t = \{\, c.f \mid c \in \mathcal{C}_t \wedge \mathrm{Eff}_t(c) \,\}.$$

Therefore any actuation event a recorded at time $t(a)$ must depend on effective commitments in $\mathcal{C}_{t(a)}$.

By construction of the authority dependency graph G, each actuation event a records a provenance set $\rho(a) \subseteq \mathcal{C}_{t(a)}$ such that $(c, a) \in E$ iff $c \in \rho(a)$. Since effectiveness is defined solely in terms of $\mathcal{C}_{t(a)}$ and the supersession relation $\prec$, and commitments are immutable once recorded, actuation cannot depend directly on provisional hypotheses except through committed authority.

By Lemma 7.1, every actuation→event edge satisfies $t(a) < t(e')$, and by Proposition 7.2 the graph G is acyclic. Hence closed-loop effects re-enter the system only as newly recorded events and cannot retroactively alter prior commitments.

Therefore every actuation event is authority-driven and traceable to effective commitments at the time of actuation. □

9. Closed-Loop Dynamics

This section summarizes how the architectural components introduced above interact to produce closed-loop system behavior under governance-bound commitment semantics.

Let $\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t)$ denote the set of candidate facts derivable at time t from hypotheses $\mathcal{H}_t$ under semantic context Γ_t. These candidate facts are provisional artifacts of interpretation and affect authoritative state only when a corresponding fact is accepted through a governance transition and recorded as a commitment.

Let $\mathcal{L}'_E$ denote the interaction log obtained by appending new events produced by actuation to $\mathcal{L}_E$. The original log $\mathcal{L}_E$ remains immutable; $\mathcal{L}'_E$ represents its extended successor.

$$
\begin{aligned}
\mathcal{L}_E \;&\longrightarrow\; \mathcal{H}_t = I(\mathcal{L}_E, \Gamma_t) \\
&\longrightarrow\; \mathcal{F}_t = \Phi(\mathcal{H}_t; \Gamma_t) \\
&\longrightarrow\; g \in \mathcal{L}_G \text{ (governance transition)} \\
&\longrightarrow\; c = \langle f, \Gamma_{t(c)}, g, s, \pi \rangle \\
&\longrightarrow\; c \in \mathcal{C}_t \text{ (commitment set up to time } t) \\
&\longrightarrow\; \mathcal{S}_t = \{\, c.f \mid c \in \mathcal{C}_t \wedge \mathrm{Eff}_t(c) \,\} \\
&\longrightarrow\; A(\mathcal{S}_t) \\
&\longrightarrow\; \mathcal{L}'_E
\end{aligned}
$$

Under commitment semantics, hypotheses $\mathcal{H}_t = I(\mathcal{L}_E, \Gamma_t)$ and derivable candidate facts $\mathcal{F}_t = \Phi(\mathcal{H}_t; \Gamma_t)$ may vary across admissible semantic contexts. However, authoritative state $\mathcal{S}_t$ evolves only through governance-bound updates to the commitment log $\mathcal{L}_C$.

Since $\mathcal{S}_t$ is defined exclusively from effective commitments in $\mathcal{C}_t$, variation in interpretation cannot retroactively alter accepted state in the absence of a governance transition. Combined with the acyclic authority structure established in Section 7, closed-loop system evolution preserves authoritative stability, provenance invariance, and non-retroactivity across iterations.

10. Experimental Evaluation

10.1. Experimental Framework

We propose an evaluation framework to compare behavioral differences between implicit authority semantics and commitment-based semantics un-

der controlled interpretive variability.

We simulate two architectures:

1. **Implicit Authority Architecture:** inference produces hypotheses $\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t)$, candidate facts $\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t)$, and authoritative state is defined as $\mathcal{S}_t := \mathcal{F}_t$.
2. **Commitment Architecture:** inference produces hypotheses $\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t)$, but authoritative state is defined as

$$\mathcal{S}_t := \{\, c.f \mid c \in \mathcal{C}_t \wedge \mathrm{Eff}_t(c) \,\},$$

with provenance recorded in each commitment and authority dependencies captured in the graph G.

10.2. Interpretive Variability Scenarios

We introduce controlled interpretive variability by varying components of semantic context Γ across repeated evaluations while holding the interaction log $\mathcal{L}_E$ fixed. The following admissible context variations are considered:

1. **Model evolution:** deployment of updated model parameters or inference configurations affecting the mapping $I(\mathcal{L}_E, \Gamma_t)$.
2. **Retrieval variation:** modification of the retrieval corpus or embedding index influencing hypothesis generation under I.
3. **Policy adjustment:** changes to decision policies or schema constraints that alter the fact-derivation map Φ.

10.3. Operational Metrics

We measure:

- **Retroactive Mutation Rate:** fraction of evaluation steps where $\neg\mathrm{Gov}(t_a, t_b)$ but $\mathcal{S}_{t_a} \neq \mathcal{S}_{t_b}$.
- **Actuation Instability Rate:** fraction of evaluation steps where repeated evaluation yields differing actuation sets $A(\mathcal{S}_t)$ under different contexts for the same underlying interaction log.
- **Closed-Loop Instability Growth:** divergence growth of state/actuation across successive iterations when actuation appends new events to $\mathcal{L}_E$.

- **Governance Supersession Frequency:** number of explicit supersession transitions required to maintain correctness under commitment semantics.
- **Provenance Trace Depth:** length of the dependency path required to trace an authoritative fact or actuation event to supporting commitments and upstream evidence in G.
- **Provenance Trace Completeness:** proportion of authoritative facts and actuation events that can be traced in G to an effective governing commitment and its recorded upstream evidence.

10.4. Expected Outcomes

Under implicit semantics, we expect non-zero retroactive mutation and higher actuation instability under interpretive variability. Under commitment semantics, we expect zero retroactive mutation absent governance transitions, stable actuation driven by authoritative state, high provenance trace completeness, and interpretable provenance trace depth through G.

11. Limitations and Scope

The present analysis abstracts from distributed replication and concurrency control. The separation principle established here is orthogonal to consistency models (e.g., strong consistency or eventual consistency): it constrains how authoritative state is derived from interpretation, not how logs or commitments are replicated across nodes. Distributed deployment may introduce additional coordination complexity, but does not alter the structural instability result or the sufficiency of governance-bound commitments under the defined semantics.

The analysis is scoped to AI-native architectures exhibiting *interpretive variability*: admissible semantic contexts under which identical interaction histories may yield different *candidate fact sets* under implicit authority semantics. In systems where interpretation is deterministic and context-invariant, interpretive variability does not arise, the instability mechanism studied here may not occur, and the separation introduced by commitments may be unnecessary.

We further assume that authoritative state is derived from an interpretation layer rather than reconstructed solely via deterministic, versioned replay

over immutable logs. Architectures that enforce deterministic replay semantics through strict projection versioning fall outside the primary regime of interest.

Governance-bound commitments introduce operational overhead in three principal categories: (1) storage overhead for commitment and provenance records, (2) latency overhead due to explicit acceptance transitions, and (3) coordination overhead when supersession decisions require human or policy-driven approval. The experimental evaluation measures storage growth and supersession frequency directly. Latency and coordination costs are discussed qualitatively, as they depend on deployment policy, governance workflow design, and system scale. This work establishes a sufficiency result for stability under admissible semantic context evolution; it does not claim performance-optimality relative to implicit authority architectures.

Finally, we do not provide a full distributed systems implementation. Issues such as replication strategy, conflict resolution among concurrent governance actions, failure recovery, and cross-node consistency are important for production deployment, but are not required to establish the semantic separation principle or the structural results proved here.

12. Future Work

Future work extends naturally along several theoretical and systems directions.

First, empirical evaluation at scale can quantify governance overhead, latency impact, and stability improvements under realistic workloads and multi-tenant deployments. Large-scale experiments would clarify trade-offs between authoritative stability, supersession frequency, and operational cost under sustained semantic context evolution.

Second, integrating governance-bound commitments with distributed consistency and concurrency control models remains an important extension. Interactions with eventual consistency, replication strategies, cross-node supersession coordination, and failure recovery require formalization to support production-grade distributed deployment.

Third, the notion of *admissible semantic context evolution* warrants formal treatment. In this work, admissibility is assumed to preserve system well-formedness. Future research could characterize admissible transitions of Γ through formal semantic constraints, type systems, or versioned context

contracts, enabling formal reasoning about bounded regimes of admissible semantic context evolution, stability regions, or convergence properties of the interpretation-to-fact pipeline under controlled evolution.

Fourth, governance of semantic context itself is a promising direction. Model updates, retrieval corpus changes, and policy revisions currently evolve outside commitment semantics. Treating updates to Γ as first-class governed artifacts—with provenance, versioning, and authorization constraints—may provide stronger guarantees under adversarial conditions, multi-tenant regimes, or cross-organizational deployments.

Fifth, partially automated governance mechanisms could incorporate risk-aware or policy-driven decision layers that propose, batch, or auto-commit low-risk facts while preserving explicit authority boundaries for higher-risk commitments. Formalizing such hybrid human–machine governance workflows is an important step toward scalable deployment.

Sixth, extending the model to multi-agent or federated environments—where authority spans agents, services, or organizational boundaries—could clarify how commitments, provenance, and semantic context versioning interact to support cross-system traceability and stable closed-loop automation under heterogeneous interpretive regimes.

Finally, the separation between interpretation and authoritative commitment may also simplify application-layer architectures. Because authoritative state is fixed through commitments rather than recomputed under the current semantic context, systems may be able to accommodate richer semantic interfaces and evolving semantic context without retroactively altering accepted authority. In such architectures, interpretation derives candidate facts from interaction history, while governance establishes authoritative state through explicit commitments. Domain-specific application logic may therefore concentrate on governance policies, validation constraints, and commitment rules rather than on maintaining stable interpretation semantics. Exploring how such layered architectures interact with emerging agentic and semantic interface paradigms is an interesting direction for future work.

13. Conclusion

AI-native systems increasingly incorporate probabilistic and context-evolving interpretation layers into operational decision pipelines. When authoritative state is defined implicitly as the output of such interpretation, accepted meaning becomes coupled to a semantic context that may evolve

over time. We formalized this architectural coupling first through implicit authority semantics, and more generally through the notion of context-coupled authority, showing that under admissible semantic context evolution authoritative stability cannot in general be guaranteed over no-governance intervals.

To resolve this incompatibility, we introduced *governance-bound commitments*, authority-establishing records that separate provisional interpretation from authoritative state while preserving immutable provenance. We proved a constructive sufficiency result: under commitment semantics with forward-only construction, authoritative stability holds for any admissible semantic context evolution in the absence of governance transitions. The resulting authority dependency structure forms a directed acyclic graph, ensuring non-retroactive authority evolution and traceable actuation in closed-loop systems.

Taken together, these results yield a scoped necessity–sufficiency characterization for authoritative stability in AI-native architectures. Architectures that recompute authoritative state directly from evolving semantic context are structurally unstable when interpretive variability is realized. Conversely, architectures that separate interpretation from authority through governance-bound commitments ensure authoritative stability while permitting admissible evolution of semantic context. More broadly, this suggests a general architectural design principle for AI-native systems of record: when interpretation semantics may evolve over time, authoritative state should not be defined by recomputation under the current interpretation context. Instead, authoritative state should be established through explicit commitment artifacts whose meaning is fixed at acceptance time and whose provenance records the semantic context under which authority was established. This separation between schema, interpretation, and authoritative commitment allows systems to incorporate richer and evolving semantic context while preserving the stability, non-retroactivity, and traceability expected of systems of record.

Appendix A. Notation Table

This appendix summarizes the symbols used throughout the paper. Section 4 introduces the system model using this notation.

Table A.3: Notation used throughout the paper.

Symbol	Description
$\mathbb{T}$	Time domain indexing evaluations and recorded transitions.
$t \in \mathbb{T}$	Time index (evaluation time or logical time).
$\mathrm{Gov}(t_a, t_b)$	Predicate: at least one governance transition is recorded in the half-open interval $(t_a, t_b]$ (i.e., after t_a and up to and including t_b).
$\mathcal{L}_E$	Immutable interaction log of semantic events $(e_1, \ldots, e_n)$.
$\mathcal{L}'_E$	Successor interaction log obtained by appending actuation-generated events to $\mathcal{L}_E$ (with $\mathcal{L}_E$ immutable).
$e \in \mathcal{L}_E$	An interaction event in the log.
$t(e)$	Recorded time of event e.
$\mathcal{L}_G$	Governance transition log.
Γ	Semantic context (models, retrieval corpus, policies, schema/configuration).
Γ_t	Semantic context at time t.
$I(\mathcal{L}_E, \Gamma_t)$	Inference at time t, producing hypotheses.
$\mathcal{H}_t$	Hypotheses at time t: $\mathcal{H}_t := I(\mathcal{L}_E, \Gamma_t)$.
$\mathcal{H}_{\text{all}}$	Global hypothesis universe: $\mathcal{H}_{\text{all}} := \bigcup_{t \in \mathbb{T}} \mathcal{H}_t$.
$\mathcal{F}$	Fact schema space: the universe of well-formed fact types that may be derived or committed.
Φ	Fact-derivation map from hypotheses (and semantic context) to a candidate fact set; applied as $\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t) \subseteq \mathcal{F}$.
$\mathcal{F}_t$	Candidate fact set at evaluation time t: $\mathcal{F}_t := \Phi(\mathcal{H}_t; \Gamma_t) = \Phi(I(\mathcal{L}_E, \Gamma_t); \Gamma_t)$, with $\mathcal{F}_t \subseteq \mathcal{F}$. Under implicit authority semantics, $\mathcal{S}_t := \mathcal{F}_t$.
$\mathcal{K}$	Key space identifying the entity/slot about which a fact is asserted.
$\kappa : \mathcal{F} \to \mathcal{K}$	Key function assigning each fact to its key.
$\mathcal{C}$	Set of all commitments recorded by the system.
$\mathcal{L}_C$	Commitment log: append-only sequence of commitment records $c \in \mathcal{C}$.
$t(c)$	Commitment time at which commitment c is recorded.
$c = \langle f, \Gamma_{t(c)}, g, s, \pi \rangle$	Commitment record with components $c.f = f$ (accepted fact), $c.\Gamma = \Gamma_{t(c)}$ (acceptance context), $c.g = g$ (governance record), $c.s = s$ (status), and $c.\pi = \pi$ (provenance).

Symbol	Description
g	Governance transition record associated with commitment c (an element of $\mathcal{L}_G$).
$\mathcal{C}_t$	Commitments recorded up to time t: $\mathcal{C}_t := \{\, c \in \mathcal{C} \mid t(c) \leq t \,\}$ (grows only via governance transitions).
$\mathrm{Eff}_t(c)$	Predicate: c is effective at time t, i.e., there exists no $c' \in \mathcal{C}_t$ such that $\kappa(c'.f) = \kappa(c.f)$ and $c \prec c'$.
s	Commitment status label stored with the record for operational clarity and auditability, $s \in \{\text{active}, \text{superseded}\}$; authoritative effect is determined by $\mathrm{Eff}_t(c)$ rather than by this label.
$\prec$	Supersession relation on commitments for the same key; time-respecting: $c \prec c' \Rightarrow \kappa(c.f) = \kappa(c'.f)$ and $t(c) < t(c')$.
π	Commitment provenance metadata (immutable once recorded).
$\mathcal{E}_\pi$	Event subset referenced by provenance: $\mathcal{E}_\pi \subseteq \mathcal{L}_E$.
$\mathcal{H}_\pi$	Hypothesis subset referenced by provenance: $\mathcal{H}_\pi \subseteq \mathcal{H}_{t(c)}$.
$\mathcal{P}(\mathcal{F})$	Power set of $\mathcal{F}$ (set of all subsets of $\mathcal{F}$).
$R(\mathcal{L}_E, \Gamma)$	Context-coupled authority map returning an authoritative fact set (Definition 6.4); at evaluation time t, $\mathcal{S}_t := R(\mathcal{L}_E, \Gamma_t)$.
$\mathcal{S}_t$	Authoritative state at time t: implicit $\mathcal{S}_t = \mathcal{F}_t$; commitment-based $\mathcal{S}_t = \{\, c.f \mid c \in \mathcal{C}_t \wedge \mathrm{Eff}_t(c) \,\}$.
$\mathcal{A}$	Set of actuation events.
$a \in \mathcal{A}$	An actuation event/action.
$t(a)$	Recorded time of actuation event a.
$A(\mathcal{S}_t)$	Actuation function triggered by authoritative state at time t.
$\rho(a)$	Actuation provenance: governing commitments for a, $\rho(a) \subseteq \mathcal{C}_{t(a)}$ (immutable once recorded).
$G = (V, E)$	Authority dependency graph over events, hypotheses, commitments, and actuation.
V	Vertex set: $V = \mathcal{L}_E \cup \mathcal{H}_{\text{all}} \cup \mathcal{C} \cup \mathcal{A}$.
E	Directed edge set encoding authority dependencies.
$u \to^* v$	Reachability in G (there exists a directed path from u to v).
$\tau : V \to \mathbb{T}$	Creation-time function; $\tau(v)$ is the recorded creation time of vertex $v \in V$ (event/hypothesis/commitment/actuation).
$t(h)$	Unique evaluation time of hypothesis h, i.e., the time such that $h \in \mathcal{H}_{t(h)}$.
$\tau(h)$	Creation-time of hypothesis vertex h in the authority DAG; by definition $\tau(h) = t(h)$.

References

[1] E. F. Codd, A relational model of data for large shared data banks, Communications of the ACM 13 (6) (1970) 377–387. doi:10.1145/362384.362685.

[2] J. Gray, A. Reuter, Transaction Processing: Concepts and Techniques, Morgan Kaufmann, 1993.

[3] P. A. Bernstein, E. Newcomer, Principles of Transaction Processing, 2nd Edition, Morgan Kaufmann, 2009.

[4] P. Helland, Life beyond distributed transactions: An apostate's opinion, in: Proceedings of the 3rd Biennial Conference on Innovative Data Systems Research (CIDR), Asilomar, CA, USA, 2007, pp. 1–10.
URL https://www.cidrdb.org/cidr2007/papers/cidr07p15.pdf

[5] P. Helland, Life beyond distributed transactions: An apostate's opinion, ACM SIGMOD Record 44 (2) (2015) 36–42. doi:10.1145/2814710.2814713.

[6] J. Kreps, N. Narkhede, J. Rao, Kafka: A distributed messaging system for log processing, in: Proceedings of the NetDB Workshop, 2011, pp. 1–7.

[7] M. Kleppmann, Designing Data-Intensive Applications, O'Reilly Media, 2017.

[8] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, S. Riedel, D. Kiela, Retrieval-augmented generation for knowledge-intensive NLP tasks, Advances in Neural Information Processing Systems 33 (2020) 9459–9474.

[9] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, D. Amodei, Language models are few-shot learners, Advances in Neural Information Processing Systems 33 (2020) 1877–1901.

[10] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. H. Chi, Q. V. Le, D. Zhou, Chain-of-thought prompting elicits reasoning in large language models, Advances in Neural Information Processing Systems 35 (2022) 24824–24837.

[11] A. Madaan, N. Tandon, P. Gupta, Y. Shnayderman, P. Clark, Y. Choi, S. Prabhumoye, Self-refine: Iterative refinement with self-feedback, Advances in Neural Information Processing Systems 36 (2023).

[12] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, A. Bouchachia, A survey on concept drift adaptation, ACM Computing Surveys 46 (4) (2014) 1–37. doi:10.1145/2523813.

[13] D. Sculley, G. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, D. Dennison, Hidden technical debt in machine learning systems, Advances in Neural Information Processing Systems 28 (2015).

[14] P. Buneman, S. Khanna, W.-C. Tan, Why and where: A characterization of data provenance, in: Proceedings of the 8th International Conference on Database Theory (ICDT), Vol. 1973 of Lecture Notes in Computer Science, Springer, 2001, pp. 316–330.
URL https://doi.org/10.1007/3-540-44503-X_20

[15] J. Cheney, L. Chiticariu, W.-C. Tan, Provenance in databases: Why, how, and where, Foundations and Trends in Databases 1 (4) (2009) 379–474. doi:10.1561/1900000006.

[16] W. M. P. van der Aalst, Process Mining: Data Science in Action, 2nd Edition, Springer, 2016. doi:10.1007/978-3-662-49851-4.

[17] P. Bailis, A. Fekete, M. J. Franklin, J. M. Hellerstein, I. Stoica, Coordination avoidance in database systems, Proceedings of the VLDB Endowment 8 (3) (2015) 185–196. doi:10.14778/2735508.2735509.

[18] P. Helland, Immutability changes everything, Communications of the ACM 59 (1) (2016) 40–50. doi:10.1145/2844112.

[19] B. Glavic, G. Alonso, Perm: Processing provenance and data on the same data model through query rewriting, in: Proceedings of the 25th

International Conference on Data Engineering (ICDE), IEEE, 2009, pp. 174–185.

[20] T. J. Green, G. Karvounarakis, V. Tannen, Provenance semirings, in: Proceedings of the 26th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS), 2007, pp. 31–40. doi:10.1145/1265530.1265535.

[21] W. M. P. van der Aalst, A. H. M. ter Hofstede, B. Kiepuszewski, A. P. Barros, Workflow patterns, Distributed and Parallel Databases 14 (1) (2003) 5–51. doi:10.1023/A:1022883727209.

[22] S. Amershi, A. Begel, C. Bird, R. DeLine, H. Gall, E. Kamar, N. Nagappan, B. Nushi, T. Zimmermann, Software engineering for machine learning: A case study, in: Proceedings of the 41st International Conference on Software Engineering (ICSE), 2019, pp. 291–300. doi:10.1109/ICSE.2019.00039.

[23] E. Breck, S. Cai, E. Nielsen, M. Salib, D. Sculley, The ml test score: A rubric for ml production readiness and technical debt reduction, in: Proceedings of the IEEE International Conference on Big Data, 2017, pp. 1123–1132. doi:10.1109/BigData.2017.8258038.

[24] K. J. Åström, R. M. Murray, Feedback Systems: An Introduction for Scientists and Engineers, Princeton University Press, 2008, draft/online edition; commonly cited as 2008. Use publisher print year if your style requires it.
URL https://www.cds.caltech.edu/~murray/books/AM08/pdf/am08-complete_28Jul12.pdf

[25] J. L. Hellerstein, Control theory: An idea whose time has come, in: Proceedings of the 10th Workshop on Hot Topics in Operating Systems (HotOS), 2005, pp. 1–6. doi:10.1109/HOTOS.2005.28.

[26] J. H. Saltzer, M. D. Schroeder, The protection of information in computer systems, Proceedings of the IEEE 63 (9) (1975) 1278–1308. doi:10.1109/PROC.1975.9939.

[27] B. W. Lampson, Hints for computer system design, in: Proceedings of the 9th ACM Symposium on Operating Systems Principles (SOSP), 1983, pp. 33–48. doi:10.1145/800217.806613.

[28] J. Dean, S. Ghemawat, Mapreduce: Simplified data processing on large clusters, in: Proceedings of the 6th USENIX Symposium on Operating Systems Design and Implementation (OSDI), USENIX Association, 2004, pp. 137–150, commonly cited as OSDI 2004.
URL https://research.google/pubs/pub62/

[29] T. Akidau, R. Chernyak, R. Lax, S. McVeety, D. Mills, The dataflow model: A practical approach to balancing correctness, latency, and cost, Proceedings of the VLDB Endowment 8 (12) (2015) 1792–1803.
URL https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/43864.pdf

[30] M. Shapiro, N. Preguiça, C. Baquero, M. Zawirski, Conflict-free replicated data types, in: Proceedings of the 13th International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS), 2011, pp. 386–400. doi:10.1007/978-3-642-24550-3_29.