

Gallery Before Feed: Why Engagement-Optimized Interfaces Destroy Coherence and How Merge-Aware Systems Restore Meaning

Flyxion

January 10, 2026

Abstract

Social media platforms increasingly present themselves as neutral infrastructures for expression, memory, and social connection. Yet their dominant architectural form—the engagement-optimized feed—systematically undermines these functions. Users who attempt to treat timelines as galleries, archives, portfolios, or coherent identity surfaces encounter persistent interface intrusions, forced duplication, semantic fragmentation, and non-consensual content insertion. These failures are commonly attributed to poor interface design, inadequate moderation, or misaligned recommendation algorithms. This paper argues instead that they are necessary consequences of feed-based architectures that collapse semantic objects, events, projections, and interaction states into a single irreversible stream.

We develop a gallery-first, merge-aware alternative grounded in event-sourced semantics, explicit object identity, and lossless consolidation. In contrast to feed architectures—where posts function as ephemeral, feed-ordered atoms—gallery-first systems treat images, texts, and other artifacts as persistent semantic objects referenced by events over time. By separating objects from events, treating feeds as projections rather than sources of truth, and formalizing merge as a first-class operation, the proposed architecture permits repetition without fragmentation, resurfacing without duplication, and accumulation without semantic drift. Meaning strengthens through revisitiation rather than dissolving through recontextualization.

The paper situates this architectural critique within three complementary theoretical frameworks. First, conceptual blending theory (Fauconnier and Turner 2002; Turner 2014) explains how humans generate meaning through the integration of multiple mental spaces. We show that engagement-driven platforms inadvertently instantiate unconstrained blending as infrastructure, turning feeds into perpetual blend spaces in which identity cannot stabilize. Second, remix culture (Lessig 2008; Manovich 2001) clarifies how contemporary platforms privilege propagation and transformation at the expense of archival integrity. We argue that while remix logic is productive as a creative paradigm, it fails catastrophically as a memory substrate. Third, cognitive theories of analogy and essence recognition (Hofstadter and Sander 2013) illuminate why instance-bound representations prevent meaning from consolidating: when systems fail to acknowledge that multiple surface encounters refer to the same underlying essence, repetition produces noise rather than understanding.

Building on these foundations, the paper develops a formal apparatus for merge-aware systems. We define semantic objects, events, projections, reaction aggregation, and merge operations, and we specify explicit interface states—gallery, interaction, and annotation—with enforceable transition rules. Interaction affordances are treated as latent rather than persistent, appearing only after deliberate user action. Silence is reclassified as a valid and stable state rather than an error condition requiring remediation. These design commitments are shown to be architectural rather than stylistic: they determine whether systems conserve meaning or inevitably destroy it.

A formal appendix proves that semantic drift is unavoidable under minimal feed-based assumptions, independent of moderation quality or recommendation accuracy, while semantic accumulation is guaranteed under merge-aware constraints. The result is a structural theorem rather than an empirical claim: systems that bind reactions to instances rather than to semantic objects cannot preserve meaning over time.

The paper concludes that meaning cannot be optimized into existence through engagement metrics. It must be conserved by design. Gallery-first, merge-aware systems therefore represent not a refinement of existing social platforms, but a fundamentally different architectural regime—one aligned with human sense-making, durable identity, and the long-term preservation of meaning.

1 Introduction

Social media platforms increasingly present themselves as neutral infrastructures for expression, memory, and connection. They promise tools for sharing images, maintaining personal histories, and cultivating identity over time. Yet the everyday experience of using these systems tells a different story. What appears to the user as a personal timeline or gallery is, in practice, a surface optimized for engagement extraction, monetization, and behavioral modulation. The resulting tension is not merely aesthetic. It is architectural.

This paper originates in a concrete and recurrent frustration: the attempt to use a personal social media timeline as a coherent visual archive or spaced repetition surface, only to be repeatedly disrupted by advertising affordances, forced duplication, algorithmic intrusions, and non-removable content classes. These disruptions are often described in colloquial terms as “clutter,” “noise,” or “bad UI.” Such descriptions are accurate but incomplete. The deeper problem is that the system actively resists being used as a gallery, an archive, or a stable identity surface at all.

The central claim of this work is that this resistance is not accidental, nor is it remediable through preference toggles, content filters, or improved moderation. Instead, it reveals a fundamental incompatibility between engagement-optimized feed architectures and the requirements of coherent personal memory systems. Feed-based platforms are not merely poorly suited to archival use; they are structurally incapable of supporting it.

At the heart of this incompatibility lies a collapse of distinctions. In feed architectures, semantic objects, publication events, interaction signals, and presentation surfaces are fused into a single primitive: the post as feed-ordered instance. Identity becomes chronological. Meaning becomes reactive. Repetition becomes duplication rather than reinforcement. Under these conditions, accumulation is impossible. Drift is inevitable.

This paper argues that the familiar failure modes of contemporary platforms—semantic fragmentation, duplicated content, incoherent identity surfaces, compulsive interaction prompts, and non-consensual content insertion—are not independent pathologies. They are different manifestations of the same underlying architectural choice: treating engagement as the primary organizing principle of meaning.

In response, the paper develops an alternative paradigm: a gallery-first, merge-aware, event-sourced social substrate in which meaning accretes rather than fragments. Rather than treating posts as ephemeral feed entries, this architecture treats images, texts, and other artifacts as persistent semantic objects referenced by events over time. Repetition strengthens identity rather than splitting it. Interaction becomes deliberate rather than compulsory. Silence becomes a valid state rather than an error condition.

The argument unfolds in three parts.

The first part analyzes the structural failures of feed-based architectures. Through close examination of interface elements such as monetization prompts, duplicate posts, fragmented identity surfaces, and algorithmic intrusions, it shows how engagement optimization systematically undermines coherence. These failures are framed not as design mistakes but as logical consequences of instance-bound, feed-ordered systems.

The second part situates this critique within a broader theoretical context. Drawing on conceptual blending theory (Fauconnier and Turner 2002; Turner 2014), remix culture (Lessig 2008; Manovich 2001), and cognitive theories of analogy and essence recognition (Hofstadter and Sander 2013), it demonstrates that engagement-driven platforms instantiate unconstrained blending and remix as infrastructure. While these processes are generative in cognition and culture, they become destructive when misapplied as memory substrates. Meaning generation is confused with meaning preservation, and the latter collapses.

The third part presents the gallery-first alternative in detail. It introduces a formal apparatus of semantic objects, events, projections, and merge operations; develops design corollaries for interface behavior; specifies explicit UI state taxonomies and transition rules; and proves, under minimal assumptions, that semantic drift is inevitable in feed-based systems but avoidable in merge-aware ones. The result is not a stylistic proposal, but a different architectural regime with distinct guarantees.

Throughout, the paper maintains a deliberately narrow focus. It does not propose to replace real-time communication, ephemeral conversation, or news dissemination. Instead, it isolates a specific but critical use case that contemporary platforms systematically fail to support: the use of digital systems as durable, coherent repositories of personal meaning.

The frustration that motivates this work is therefore not incidental. It is diagnostic. It signals a deeper misalignment between human sense-making and engagement-driven infrastructure. Addressing this misalignment requires more than better interfaces or smarter algorithms. It requires a rethinking of the primitives on which social systems are built.

Meaning cannot be recovered by tuning a feed. It must be conserved by design.

2 From Feed to Memex: Search, Retrieval, and the Failure of Timeline Memory

The critique developed in this paper is not limited to visual presentation or interaction affordances. It also concerns retrieval. A system that cannot be searched coherently cannot function as memory, regardless of how expressive or engaging its surface may be. This observation motivates an additional claim: engagement-optimized timelines fail not only as galleries, but as mnemonic devices.

Users routinely attempt to use timelines as informal memory systems. They return to past posts to recall when something occurred, to resurface an image they remember making, or to re-enter a line of thought they once articulated. This behavior is not accidental. It reflects a natural tendency to treat personal archives as externalized cognition, in the sense described by distributed and situated cognition research (Hutchins 1995).

Feed-based architectures frustrate this use case at every level. Search is treated as an auxiliary feature rather than as a primary mode of navigation. Results are filtered through engagement heuristics, relevance scores optimized for popularity rather than recall, and temporal truncation. Identical or near-identical posts are returned as separate items, even when the user is clearly attempting to retrieve an underlying artifact rather than a specific instance.

As a result, searching one's own timeline often feels like rummaging through debris rather than consulting memory. The system cannot answer simple questions such as: "Show me that photograph I keep reposting," or "Return the object I was thinking about when I wrote these three related posts." The user remembers an essence. The system only indexes surfaces.

This is precisely the failure mode identified by theories of analogy and essence recognition. Humans recall by recognizing sameness across variation. Feed architectures, by binding search indices to instance-level posts, actively interfere with this process. They force the user to reconstruct equivalence manually, increasing cognitive load and eroding trust in the system as a reliable external memory.

A gallery-first, merge-aware system enables a different relationship between search and memory. Because semantic objects possess stable identity, search operates over objects rather than over ephemeral feed entries. Queries return objects with consolidated histories rather than disjoint instances. Repeated resurfacing strengthens retrievability instead of obscuring it.

In such a system, the timeline ceases to be the primary retrieval interface. It becomes one projection among many. Search becomes co-equal with browsing, and in many cases superior to it. A user may navigate by time, by object, by annotation, by thematic clustering, or by revisit frequency, without duplicating content or fragmenting meaning.

This architecture aligns naturally with the concept of a memex, as originally articulated by Vannevar Bush, not as a device for storing more information, but as a system for associative retrieval. In a merge-aware system, associations are explicit events linking stable objects. Search can therefore traverse meaningful paths rather than scanning flattened histories.

Crucially, memoization emerges as a system property rather than a user burden. Objects that are frequently revisited, annotated, or resurfaced naturally rise in prominence, not because they are

engaging in the moment, but because they have proven durable across time. This is a fundamentally different optimization target from engagement. It rewards persistence rather than reaction.

The design implications follow directly. Search must be promoted from a utility feature to a first-class navigation mode. Query results should privilege object identity, event history, and semantic continuity over recency or popularity. Timelines should be searchable as memory traces, not merely as chronological logs.

Without these changes, timelines cannot function as memexes or memoization boards, no matter how carefully they are curated. With them, the distinction between gallery, archive, and memory system begins to dissolve. The timeline becomes not a stream to be consumed, but a surface through which thought can be revisited, recomposed, and sustained.

This reframing completes the diagnostic picture. Engagement-optimized feeds fail visually, interactively, and mnemonically. Gallery-first, merge-aware systems succeed in all three domains for the same reason: they treat meaning as something that accumulates through recognition, retrieval, and return, rather than something that must be continually re-extracted from novelty.

3 The Timeline as Monetization Surface

The failure of feed-based systems as memory substrates is inseparable from their economic role. Timelines are not merely poorly suited to archival or mnemonic use; they are actively structured to prevent it. Their primary function is not to preserve meaning, but to expose monetizable surfaces.

A defining feature of contemporary platforms is the persistent embedding of monetization controls directly into user-created content. Buttons labeled with imperatives such as “Boost Post,” “Promote,” or “See Ads and Insights” are rendered with visual priority, often in high-contrast colors that dominate the surrounding composition. These elements are not optional overlays. They are inseparable from the artifact itself.

This design choice reclassifies personal expression as latent inventory. A photograph, rather than being treated as an object to be viewed, remembered, or revisited, is framed as an underperforming asset awaiting optimization. The interface does not ask whether the user wishes to sell or promote. It assumes that monetization is the natural continuation of expression.

From an architectural perspective, this introduces a categorical confusion between artifact and instrument. A gallery presents objects. A dashboard presents controls. By collapsing these into a single surface, the system ensures that artifacts are never encountered without evaluative framing. The user cannot see their own work without simultaneously being shown how it could have performed better.

The cognitive consequence of this conflation is subtle but corrosive. The artifact ceases to function as a stable reference point. Instead, it becomes a mutable performance snapshot, implicitly judged by metrics and surrounded by prompts to intervene. This undermines both memory and authorship. The creator is repositioned as a manager of engagement rather than a steward of meaning.

This design is not accidental. Engagement-driven platforms require continual stimulation to maintain signal flow. A gallery that allows artifacts to rest, to be revisited quietly, or to remain unchanged

over time is economically inert. It does not invite clicks. It does not generate experiments. It does not expose new surfaces for optimization.

As a result, the system must actively resist gallery-like use. Persistent prompts such as “Write a note,” visible reaction counters, and ever-present promotion affordances ensure that the timeline remains a site of production rather than conservation. Silence becomes suspicious. Completion becomes illegible.

The contrast with archived views is instructive. When posts are archived, monetization and engagement affordances are often suppressed. Images appear without reaction counts or comment prompts. The visual field stabilizes. The artifact becomes legible again as an object rather than a performance trace. The relief users report in these views is not subjective preference; it is the restoration of semantic integrity.

That this alternative projection already exists reveals the depth of the problem. The platform is capable of presenting content as a gallery. It simply chooses not to, because doing so would interfere with its primary economic function.

The same logic explains the persistent degradation of search. A search system optimized for memory would surface stable objects, consolidated histories, and frequently revisited artifacts. A search system optimized for monetization surfaces instead prioritizes recent, engaging, or promotable instances. Retrieval becomes another opportunity for exposure rather than recall.

In this sense, the timeline’s failure as a memex is not a side effect of monetization. It is a prerequisite for it. A system that allowed users to reliably retrieve and revisit their own artifacts would reduce dependence on the feed as a discovery mechanism. It would shift value from exposure to preservation.

Gallery-first systems invert this relationship. Monetization, where it exists, must attach to objects rather than to moments. Promotion becomes optional rather than assumed. Metrics become properties of consolidated histories rather than incentives to fragment them. Most importantly, the artifact is allowed to exist without solicitation.

This reframing clarifies the stakes of the architectural choice. The question is not whether platforms should monetize. It is whether monetization is permitted to define the ontology of content itself. Feed-based systems answer in the affirmative. Gallery-first systems refuse.

By insisting that artifacts precede metrics, and that memory precedes optimization, gallery-first architectures restore a distinction that engagement-driven systems must erase: the difference between something that is made and something that is sold.

4 Feed-Ordered Atoms and the Incoherence of Duplication

The structural failures described thus far converge on a single architectural decision: the treatment of posts as feed-ordered atoms. In engagement-driven systems, the post is the smallest unit of meaning, interaction, monetization, and retrieval. It is immutable, instance-bound, and irreducible. Once published, it can only be reacted to, promoted, or buried by newer instances.

This design choice has far-reaching consequences. Because each post is treated as a unique atomic entity, the system lacks any mechanism for recognizing equivalence across instances. Two posts

containing identical images are not understood as references to the same artifact. They are understood as distinct competitive units, each vying for attention in the feed.

From the user's perspective, this is incoherent. The user experiences the image as one thing encountered multiple times. The system experiences it as multiple unrelated things that happen to look similar. This mismatch forces the user to perform semantic reconciliation manually, while the system actively resists it.

The absence of mergeability is particularly striking when contrasted with versioned systems such as distributed source control. In systems like Git, files are not treated as disposable snapshots. They are tracked as evolving entities with histories, diffs, and merges. Multiple branches may diverge, but reconciliation is not only possible; it is foundational. Identity persists across time, even as content changes.

Feed architectures adopt the opposite stance. Divergence is trivial. Convergence is impossible. Once two posts exist, there is no operation that can unify them without erasing one entirely. Deletion is the only form of reconciliation offered, and it is destructive. History is lost. Reactions vanish. The system enforces a false dichotomy between duplication and erasure.

This limitation is not merely inconvenient. It destroys the possibility of cumulative meaning. Reactions, comments, and annotations that would naturally accrete around a stable artifact are instead dispersed across instances. Each instance accumulates a partial, context-specific interpretation. No global view is ever permitted to form.

The result is semantic dilution. The more an object is resurfaced, the less coherent its meaning becomes. What should be reinforcement becomes fragmentation. What should be memory becomes noise.

This failure mode mirrors known limitations in systems that lack conflict resolution primitives. In distributed systems, the absence of merge operations leads to divergence and eventual inconsistency. In social systems, the absence of merge leads to interpretive drift and identity fragmentation. The parallel is exact, even if the domain differs.

The feed-ordered atom model also explains why search fails so predictably. Because each post is indexed independently, search results return surfaces rather than referents. The user remembers an artifact. The system retrieves a pile of instances. No canonical result exists because the system does not believe canon exists.

This design further incentivizes duplication. Since visibility and interaction attach to instances rather than to objects, reposting becomes the only way to reassert presence. The system rewards fragmentation. It punishes consolidation. Over time, users internalize this logic and adapt their behavior accordingly, even when it conflicts with their own intentions.

Merge-aware systems reject this model entirely. They treat posts not as atoms but as events referencing persistent objects. Divergence is permitted at the level of events. Convergence is always possible at the level of objects. Identity survives repetition. History accumulates.

This distinction clarifies why feed architectures cannot be repaired incrementally. As long as posts are treated as irreducible atoms, no amount of interface refinement or algorithmic tuning can restore coherence. The problem is ontological, not procedural.

To support memory, search, and identity, systems must abandon the feed-ordered atom as their foundational unit. They must replace it with something that can be revisited, merged, and recognized as the same thing over time.

Until that shift occurs, duplication will remain incoherent by construction, and meaning will continue to dissipate with every repost.

5 Identity Surfaces and the Cover Photo Paradox

The incoherence introduced by feed-ordered atoms becomes especially visible when a single artifact is forced to inhabit multiple identity surfaces simultaneously. A canonical example is the use of a photograph as both a cover image and as one or more timeline posts. Although these surfaces are intuitively understood by users to refer to the same object, the platform treats them as unrelated entities.

Each surface maintains its own engagement counters, contextual framing, and lifecycle. A cover image persists as a quasi-static identity marker, while timeline posts appear, scroll away, and are periodically resurfaced through reposting. There is no structural relationship between them. No shared identity is acknowledged. No history is unified.

From the system's perspective, this separation is coherent. Each surface is an independent instance optimized for different engagement roles. From the user's perspective, it is deeply disorienting. The same image appears multiple times in close proximity, sometimes even adjacent in the timeline, each instance demanding evaluation, reaction, or promotion. The image ceases to function as an identity anchor and instead becomes visual clutter.

This phenomenon reveals a deeper failure: the absence of a formal notion of identity surfaces as bindings to semantic objects. In a gallery-first system, a cover image would not be a copy of an artifact, nor a separate post. It would be a role assignment: a persistent binding between an identity surface and an underlying object. Changing the cover would alter the binding, not duplicate the object. History would remain intact and discoverable.

Feed-based systems cannot represent this relationship because they lack object identity. Surfaces can only point to instances, not to referents. As a result, identity surfaces fragment rather than stabilize identity.

The cognitive cost of this fragmentation is significant. Users rely on repeated exposure to stable representations to form and maintain identity narratives. When the same artifact appears as multiple unlinked instances, the mind is forced to reconcile them manually. This reconciliation is not supported by the interface and is actively undermined by competing engagement signals.

The paradox intensifies when users attempt to curate their timelines deliberately. Removing one instance risks losing reactions or comments attached to it. Keeping multiple instances preserves history at the cost of coherence. The system presents a false choice between memory and order.

This failure generalizes beyond cover photos. Profile pictures, pinned posts, featured content, and highlights all exhibit similar behavior. Each is implemented as a specialized instance rather than as a view onto a shared object. The platform multiplies surfaces while refusing to unify what lies beneath

them.

In a merge-aware architecture, identity surfaces become first-class constructs. They are projections over objects, not containers of content. A single object may be bound to multiple surfaces without duplication. Engagement accumulates across surfaces. Search retrieves the object, not the surface through which it was last encountered.

This distinction is crucial for treating timelines as memoization boards or memex-like structures. Identity surfaces are among the strongest mnemonic anchors available. When they fragment, memory degrades. When they stabilize, recall improves.

The cover photo paradox therefore serves as a diagnostic lens. It shows, in a single everyday interaction, why feed-based systems cannot support coherent identity over time. They mistake surface for substance and instance for essence.

Gallery-first systems reverse this mistake. By treating surfaces as bindings and objects as referents, they allow identity to persist even as presentation changes. What is remembered is not where something appeared, but what it was.

This completes the critique of feed-based identity at the level of surfaces. The next step is to explain why this failure is not merely cultural or economic, but rooted in deeper theoretical misunderstandings about how meaning is generated and preserved.

6 From Structural Failure to Theory: Why These Problems Persist

The preceding sections have documented a set of failures that recur across engagement-optimized platforms: incoherent duplication, fragmented identity surfaces, degraded search and retrieval, forced interaction, and the systematic erosion of timelines as memory substrates. These failures present themselves at the level of interface, but they are not interface bugs. They are symptoms of deeper theoretical commitments.

At this point, it is tempting to attribute the problem to incentives alone. Monetization pressures, advertising markets, and growth imperatives undoubtedly shape platform behavior. Yet economic explanation, while necessary, is not sufficient. Even platforms that experiment with alternative business models reproduce many of the same failure modes. The persistence of these patterns suggests that something more fundamental is at work.

Specifically, feed-based architectures embody a particular theory of meaning—one that is rarely articulated but continuously enacted. In this theory, meaning is assumed to be generated by exposure, reinforced by reaction, and exhausted by novelty. Content is treated as stimulus. Context is treated as disposable. Identity is treated as an effect of circulation rather than as a stable reference.

This implicit theory explains why timelines are optimized for flow rather than recall, why repetition is punished rather than consolidated, and why search is subordinated to recommendation. It also explains why attempts to retrofit coherence through moderation tools, preference settings, or improved ranking algorithms consistently fail. These interventions operate at the wrong level. They attempt to regulate behavior without revising the underlying semantic model.

To understand why feed architectures behave as they do—and why gallery-first alternatives be-

have differently—it is therefore necessary to step outside the platform discourse and engage with broader theories of meaning formation, propagation, and recognition. The remainder of this paper undertakes that task.

Part II situates the architectural critique within three complementary theoretical frameworks. Conceptual blending theory explains how humans generate new meaning by integrating multiple mental spaces. Remix culture explains how contemporary media systems privilege transformation and propagation over preservation. The theory of analogy and essence recognition explains how meaning stabilizes through the recognition of sameness across variation.

Taken together, these frameworks clarify why engagement-driven platforms drift inexorably toward incoherence. They also clarify why merge-aware, gallery-first systems succeed where feeds fail. The difference is not merely technical. It is theoretical.

By making these commitments explicit, the paper shifts from diagnosis to explanation. The failures of feed architectures are not accidental outcomes of poor design choices. They are the predictable results of instantiating unconstrained blending and instance-bound representation as infrastructure.

The sections that follow therefore do not introduce new problems. They explain why the problems already identified could not have been otherwise under the prevailing model—and why a different model produces a fundamentally different outcome.

7 Conceptual Blending as Infrastructure: When a Cognitive Process Becomes a System Failure

Conceptual blending theory, developed primarily by Fauconnier and Turner, describes how humans generate meaning by integrating multiple mental spaces into a blended space that exhibits emergent structure (Fauconnier and Turner 2002; Turner 2014). Blends are not simple unions. They selectively project structure from inputs, compress relationships, and generate inferences that were not explicitly present in any single source.

At the level of cognition, this process is extraordinarily powerful. It enables metaphor, abstraction, humor, and creative insight. Humans routinely perform blends such as understanding time as money, arguments as war, or software bugs as physical defects. In each case, meaning arises not from literal correspondence but from structured alignment across domains.

The problem addressed in this paper is not that conceptual blending exists, but that engagement-optimized platforms unknowingly instantiate it as infrastructure.

7.1 Mental Spaces and Feed Contexts

In blending theory, mental spaces are temporary, task-relative structures constructed for local reasoning. They are lightweight, disposable, and context-dependent. A blend is successful precisely because it is not required to persist. Once insight is achieved, the mental scaffolding can dissolve.

Feed architectures recreate this pattern at the level of content presentation. Each post appears embedded in a local context defined by adjacent posts, current trends, recent interactions, and algo-

rithmic juxtapositions. Meaning emerges relationally from the surrounding feed environment rather than from the post alone.

This design choice implicitly treats each feed viewport as a mental space. Scrolling becomes the act of moving through a sequence of transient spaces. Meaning is constructed moment by moment through juxtaposition.

The critical error lies in assuming that what works for cognition will work for memory.

Mental spaces are not archives. They are not required to support retrieval, accumulation, or identity persistence. Feed systems, by contrast, are used as long-term personal repositories whether or not they are designed to be. When ephemeral blend spaces are mistaken for storage substrates, instability becomes inevitable.

7.2 Emergent Meaning Without Referential Anchors

In conceptual blending, the loss of explicit referential identity is acceptable because the mind can always reconstruct or abandon a blend as needed. Infrastructure does not have this luxury. When platforms remove stable referential anchors, they remove the conditions under which meaning can consolidate.

Feed systems exacerbate this problem by binding reactions, comments, and metrics to blended instances rather than to underlying referents. Each post becomes meaningful only within its immediate blended context. When that context dissolves, meaning dissipates with it.

Repeated exposure does not reinforce understanding because the system does not recognize repetition. The same image encountered in different feed contexts is treated as a new blend input each time. No generic space is preserved across encounters. No compression occurs.

In blending terms, the system continuously creates new blends while refusing to maintain a shared generic space. The result is combinatorial proliferation without convergence.

7.3 A Worked Example: Blended Feed Versus Merge-Aware Representation

Consider a single photograph posted three times over the course of a year: once with a caption about place, once with a hashtag for discovery, and once without text as a visual reminder. In a feed-based system, these appear as three unrelated posts.

Each post accumulates reactions in a different social and temporal context. Each is interpreted differently by viewers. The system records three distinct meaning traces. No operation exists to align them.

From a blending perspective, each appearance creates a separate blended space. The image is integrated with different contextual cues each time. But because the system does not preserve the generic structure—the recognition that these encounters refer to the same object—no stable concept emerges.

In a merge-aware system, the photograph is a single semantic object. Each posting is an event referencing that object. Contextual variation is preserved as event metadata, but reactions aggregate at the object level. The generic space is explicit and persistent.

Here, blending occurs where it belongs: in interpretation, not in storage. The system supports multiple readings without fragmenting identity.

7.4 Why Feeds Cannot Support Accumulation

Conceptual blending is generative precisely because it tolerates ambiguity, redundancy, and loss. Infrastructure must do the opposite. It must constrain ambiguity, reconcile redundancy, and preserve lineage.

Feed architectures fail because they refuse these constraints. They maximize local blending opportunities at the expense of global coherence. Every scroll is an invitation to re-blend. Nothing is allowed to settle.

This explains why feed-based systems inevitably drift even under ideal conditions. Drift is not a moderation failure. It is the emergent result of unbounded blending applied to a domain that requires consolidation.

7.5 Bounded Blending as an Architectural Principle

Gallery-first, merge-aware systems can be understood as implementing bounded blending. They permit contextual interpretation and recontextualization, but only at the level of events. The underlying object remains stable.

This corresponds directly to the separation between mental spaces and long-term memory in human cognition. We generate meaning locally, but we store it relationally. We do not overwrite concepts every time we think about them.

By making this separation explicit, merge-aware systems align infrastructure with cognition rather than parodying it. They allow blending to occur where it is useful while preventing it from destroying what must persist.

7.6 Summary

Conceptual blending explains how meaning is created. It does not explain how meaning is preserved. Feed-based platforms collapse these two functions and thereby fail at both.

Gallery-first systems restore the distinction. They treat blending as a cognitive operation and identity as an infrastructural one. In doing so, they resolve a category error that engagement-optimized architectures have normalized.

This insight prepares the ground for the next section, which examines remix culture as the cultural analogue of unbounded blending and shows why its logic, while generative, is equally unsuitable as a memory substrate.

8 Remix Culture and the Collapse of Referential Stability

Where conceptual blending explains how meaning is generated at the cognitive level, remix culture explains how meaning is propagated at the cultural level. The two are deeply aligned. Both privi-

lege recombination, reinterpretation, and contextual reframing over preservation of original structure. Both treat transformation as a primary value. Both assume abundance rather than scarcity.

Remix culture emerged alongside digital media's capacity for frictionless copying, editing, and redistribution. In music, video, and visual art, the ability to sample, splice, quote, and recontextualize content produced new forms of creativity and new aesthetic norms. Authorship became distributed. Lineage became playful rather than juridical. Canon became provisional.

As a mode of cultural production, this shift was generative rather than destructive. Remix culture allowed communities to explore meaning collaboratively, to respond to one another rapidly, and to treat media as conversational material rather than as finished monuments (Lessig 2008; Manovich 2001). Transformation became a way of thinking.

The difficulty arises when remix logic is mistaken for an archival substrate.

8.1 Propagation Versus Preservation

Remix culture optimizes for propagation. Artifacts are expected to move, mutate, and proliferate. Persistence is optional. Attribution may be partial, ironic, or deliberately obscured. What matters is not that an object remains intact, but that it continues to circulate.

Memory systems, by contrast, require consolidation. They depend on stable referents, durable identities, and retrievable histories. A memory that cannot be reliably revisited is not a memory at all. It is a stimulus.

Feed-based platforms implicitly adopt remix culture as their infrastructural default. Every post is treated as a remixable, disposable unit. Even original content is immediately absorbed into a flow of reinterpretation, quotation, and algorithmic reframing. The system does not distinguish between transformation as expression and transformation as erosion.

This conflation produces a predictable outcome. Meaning spreads, but it does not settle.

8.2 Platform Case Studies: Reblogs and Quote-Tweets

The logic of remix as infrastructure is especially visible in systems that explicitly encourage resharing with commentary. Reblogging systems, such as those historically associated with ~~4chan~~, allow content to propagate through chains of transformation. Each reblog adds context, tags, or commentary, but the original object remains loosely traceable through embedded references.

This model excels at conversational creativity. It allows communities to build layered interpretations and to participate in shared meaning-making. However, it struggles with consolidation. Long reblog chains become unwieldy. Original context is gradually obscured. Search retrieves fragments rather than wholes. Identity diffuses.

Quote-based resharing, as popularized on platforms like ~~Twitter~~, intensifies this effect. Quoted content is reproduced as an image or snippet, detached from its original object identity. Engagement accrues to the quote instance rather than to the quoted artifact. Over time, the same content exists in dozens or hundreds of partially overlapping forms, none of which can be reconciled.

In both cases, remix logic succeeds at expression but fails at memory. The platforms reward transformation without providing a substrate for reunification.

8.3 Why Remix Logic Fails as Infrastructure

The core problem is that remix culture assumes that meaning is produced through divergence. Infrastructure must assume the opposite. It must assume that meaning stabilizes through convergence.

In remix systems, there is no obligation to preserve equivalence across transformations. Two artifacts that “feel” related may or may not be structurally linked. Recognition is left to human inference. The system itself does not enforce sameness.

This assumption is tolerable in cultural production because humans are skilled at reconstructing context when they care to. It is intolerable in large-scale systems where volume overwhelms attention. As scale increases, manual reconciliation becomes impossible. Meaning fragments faster than it can be reassembled.

Feed-based platforms amplify this failure by coupling remix logic to engagement optimization. Transformations that provoke reaction are rewarded, regardless of whether they preserve referential integrity. Over time, the system selects for mutations that maximize visibility rather than coherence.

8.4 Intentional Remix Versus Accidental Dissolution

A gallery-first, merge-aware system does not reject remix. It reclassifies it.

In such a system, remix is an explicit event referencing a stable object. Transformation is recorded as lineage rather than as replacement. A derivative work does not erase or compete with its source. It points to it.

This distinction allows remix to remain creative without becoming destructive. It supports intentional divergence while preventing accidental dissolution. Users can explore variation without losing the ability to return to a shared referent.

This model mirrors best practices in other domains. In version-controlled software, branching allows experimentation, but merging preserves coherence. In scholarly citation, interpretation proliferates, but canonical references remain identifiable. In museums, reinterpretation occurs through curation, not through destruction of artifacts.

Feed architectures ignore these lessons. They treat all transformation as equal and all instances as disposable.

8.5 Remix, Memory, and the Illusion of Choice

It is often argued that users prefer remix-driven environments and that ephemerality is a feature rather than a bug. This objection misunderstands the proposal.

Gallery-first systems do not force permanence. They enable it. They allow users to choose when something should be ephemeral and when it should persist. Feed-based systems make the opposite choice by default. Everything is ephemeral, whether the user intends it or not.

The result is not freedom, but loss of control.

8.6 Summary

Remix culture explains why feed-based platforms feel vibrant, expressive, and alive. It also explains why they fail as archives, galleries, or memory systems.

By adopting remix logic as infrastructure rather than as a mode of expression, engagement-driven platforms guarantee semantic drift. Meaning spreads, but it cannot settle. Identity circulates, but it cannot stabilize.

Gallery-first, merge-aware systems resolve this tension by separating propagation from preservation. They allow remix to flourish where it belongs, while ensuring that what matters can still be found, recognized, and remembered.

This prepares the ground for the next section, which turns from cultural dynamics to cognitive foundations and examines how analogy and essence recognition explain why instance-bound systems systematically undermine understanding.

9 Analogy, Essence Recognition, and Why Instance-Bound Systems Undermine Understanding

The deepest explanation for why feed-based architectures fail as memory systems lies not in economics or interface design, but in cognition. Humans do not remember by retrieving exact instances. They remember by recognizing essences across variation. This insight is developed at length by Hofstadter and Emmanuel Sander, who argue that analogy is not a special mental operation but the core mechanism of thought itself (Hofstadter and Sander 2013).

In their account, meaning arises when the mind recognizes that two superficially different situations are “the same in the right way.” Surface features vary; underlying structure persists. To understand something is to compress multiple encounters into a single conceptual referent.

This model of cognition has direct implications for how digital systems should represent content if they are to support understanding rather than merely stimulation.

9.1 Surfaces Versus Essences

Hofstadter and Sander distinguish between surfaces—the particular sensory or contextual manifestations of an idea—and essences—the abstract structure that unifies those manifestations. A child learns the concept of “dog” not by memorizing one dog, but by encountering many different dogs and recognizing their shared essence. Variation is not noise; it is the raw material of abstraction.

Crucially, this process requires that the mind be able to recognize sameness across difference. When every encounter is treated as unrelated, learning stalls. When equivalence is recognized, understanding accelerates.

Feed-based systems systematically block this process. By treating each post as a unique, irreducible instance, they refuse to acknowledge essence. Two identical images are not “the same thing seen again.” They are separate objects competing for attention. The system encodes difference where the mind seeks sameness.

As a result, repetition does not strengthen memory. It fragments it.

9.2 Instance Binding as a Cognitive Anti-Pattern

In feed architectures, reactions, comments, and search indices bind to instances rather than to referents. This forces the user into an unnatural cognitive posture. Instead of recognizing an object and recalling its accumulated history, the user must remember where and when a particular instance appeared.

This is not how human memory works. We do not recall facts by reconstructing the precise circumstances of first exposure. We recall them by recognizing them again.

Instance-bound systems therefore impose a constant tax on cognition. Each encounter requires re-identification. Each resurfacing demands re-evaluation. Over time, this produces fatigue, frustration, and disengagement—not because the content lacks value, but because the system refuses to support recognition.

The user remembers an essence. The system returns surfaces.

9.3 Analogy as the Basis of Search and Memoization

Search, when properly understood, is not a lookup operation. It is an analogical act. The user does not ask, “Show me post 1437 from March 2022.” They ask, implicitly, “Show me that thing I’m thinking about again.”

In a merge-aware system, this request is satisfiable because semantic objects correspond to essences. Search operates over referents rather than instances. Queries retrieve objects whose histories compress multiple encounters into a single result.

This enables memoization. Each time an object is revisited, its conceptual weight increases. It becomes easier to find, easier to recognize, and easier to think with. Memory is reinforced through use.

Feed-based systems cannot support memoization because they do not acknowledge equivalence. Searching one’s own timeline returns a scatter of instances, each partially relevant, none canonical. The system fails to perform the analogical compression that cognition expects.

This explains why users often resort to external tools—folders, screenshots, notebooks, or personal archives—to store what platforms refuse to stabilize. They are reconstructing essence manually because the system does not.

9.4 Why Engagement Optimization Conflicts with Analogy

Engagement-driven systems are optimized for surface differentiation. Novelty, contrast, and surprise increase reaction likelihood. Recognizing sameness reduces stimulation. From an engagement perspective, essence recognition is a liability.

This creates a structural conflict. Cognition seeks compression. Optimization seeks variation. The more successfully a platform extracts engagement, the more aggressively it must prevent content from settling into stable referents.

The result is a system that actively undermines understanding. Not by accident, but by necessity.

9.5 Merge as Formalized Analogy

Merge-aware systems resolve this conflict by formalizing analogy as an infrastructural operation. When two events reference the same semantic object, the system asserts that they are “about the same thing.” This assertion is preserved, not inferred anew each time.

Merge is therefore not merely a technical convenience. It is a cognitive alignment. It encodes, at the system level, the same operation the mind performs when it recognizes essence across variation.

Reaction aggregation, consolidated search results, and identity stabilization all follow from this single commitment.

9.6 Implications for Learning and Memory

The implications extend beyond social media. Systems that support analogy and memoization naturally support learning. Repeated exposure strengthens concepts. Variation enriches understanding without fragmenting identity. Recall improves with use.

Instance-bound systems do the opposite. They punish repetition, reward novelty, and force users to relearn what they already know. Over time, this produces a sense of overload rather than mastery.

The frustration users experience when attempting to use timelines as memory boards is therefore not incidental. It is the felt experience of a system working against cognition.

9.7 Summary

Hofstadter and Sander’s theory makes clear that meaning stabilizes through the recognition of sameness across difference. Feed-based architectures refuse to recognize sameness. They therefore refuse to support meaning.

Gallery-first, merge-aware systems restore this recognition explicitly. They align infrastructure with cognition by treating semantic objects as essences and events as surfaces. Search becomes analogical. Memory becomes cumulative. Repetition becomes reinforcement.

This alignment does not simulate human thought. It merely stops interfering with it.

With this cognitive foundation in place, the paper can now return to design and formalism, showing how interface commitments and transition rules either preserve or sabotage the conditions for understanding.

10 Design Corollary Revisited: Interfaces That Respect Analogy and Memory

The cognitive analysis of analogy and essence recognition yields a precise design corollary. If meaning stabilizes through repeated recognition of the same underlying referent, then interfaces must

minimize surface noise that interferes with recognition. Any element that competes for attention at the moment of encounter weakens the formation of essence.

From this perspective, many familiar interface features reveal themselves as actively hostile to understanding. Persistent like counters, comment prompts, share buttons, promotion affordances, and composer bubbles do not merely clutter the visual field. They inject evaluative and performative context at the precise moment when recognition should occur.

The artifact is never encountered alone. It is always framed as an opportunity for action.

This framing disrupts analogy. Instead of recognizing “this is that image again,” the user is prompted to ask, “What should I do with this?” The system forces a shift from recognition to response, from memory to performance.

The effect is cumulative. Over time, users cease to experience their own timelines as places of return. They become sites of obligation.

10.1 Latent Affordances as a Cognitive Necessity

The proposal to make interaction affordances latent rather than persistent is therefore not an aesthetic preference. It is a cognitive requirement. Recognition precedes action. Essence precedes evaluation.

In a gallery-first system, the default encounter with an object is quiet. The image, text, or artifact is allowed to present itself without solicitation. Only after the user performs an explicit gesture—selecting, focusing, or entering interaction mode—do affordances appear.

This sequencing mirrors cognition. We recognize first. We decide what to do second.

Persistent affordances reverse this order. They demand decision before recognition. The result is shallow engagement and poor recall.

10.2 The Composer Bubble as a Category Error

The ubiquitous “Write a note” or “What’s on your mind?” prompt embedded in personal timelines exemplifies this error in its purest form. Its constant presence asserts that the timeline is not a record or gallery, but an unfinished utterance.

This assumption is incompatible with any notion of completion, curation, or rest. A finished gallery does not ask the artist to speak again. It waits.

By refusing to allow silence, the interface denies the possibility of memory. Every empty space is treated as a failure state. The system does not distinguish between absence and sufficiency.

In a merge-aware system, creation and curation are separate modes. The absence of a composer is not a lack. It is a declaration that the current projection is complete.

10.3 Why Archived Views Feel Better

The accidental success of archived views deserves renewed emphasis. When engagement controls disappear, users report immediate relief. The content becomes legible. Navigation becomes calmer. Recall improves.

Nothing about the artifact has changed. Only the projection has.

This confirms the central thesis of the paper: coherence is not produced by content quality, but by architectural restraint. When the system steps back, meaning steps forward.

10.4 Design Implication

Interfaces must be evaluated not by how much interaction they elicit, but by how well they support recognition across time. An interface that maximizes immediate response while undermining recall is anti-cognitive, regardless of its commercial success.

The design corollary is therefore simple and non-negotiable: systems that aim to support memory, identity, or understanding must privilege gallery states over interaction states, silence over prompts, and recognition over response.

This corollary motivates a more formal treatment of interface state, to which we now turn.

11 Interface States as Semantic Commitments

The distinction between gallery, interaction, and annotation states introduced earlier can now be sharpened. These states are not merely modes of presentation. They are semantic commitments about what kind of cognitive operation the system is inviting.

Each state privileges a different relationship between surface and essence.

11.1 Gallery State: Recognition Without Obligation

The gallery state exists to support analogical recognition. Its function is to allow the user to encounter an object as itself, without pressure to act.

In this state, the system must do as little as possible. No metrics, no prompts, no calls to action. Navigation is allowed; solicitation is not. The user may scroll, browse, or search without producing new semantic events.

This state supports memoization. Repeated encounters strengthen familiarity. Objects settle into memory.

11.2 Interaction State: Deliberate Engagement

The interaction state exists to support intentional response. It is entered explicitly. Its appearance signals a shift in purpose.

Here, affordances become visible. Reactions, comments, and lightweight responses are enabled. Importantly, these interactions bind to objects, not to instances. They extend history without fragmenting it.

Because entry into this state is deliberate, interaction regains meaning. A reaction is no longer an accidental byproduct of scrolling. It is a choice.

11.3 Annotation State: Meaningful Extension

The annotation state is reserved for durable contributions. It supports the addition of explanation, interpretation, linkage, or revision.

This state produces semantic structure. It modifies the object–event graph. It is therefore distinct from casual interaction and must be visually and functionally separated.

Annotation is not conversation. It is authorship.

11.4 State Separation as Drift Prevention

Feed-based systems collapse all three states into one. Viewing, reacting, and modifying occur simultaneously and continuously. This collapse guarantees drift. Every encounter becomes a blend. Nothing stabilizes.

State separation prevents this by construction. It ensures that recognition can occur without interference, that interaction is intentional, and that modification is explicit.

This is not a limitation. It is a safeguard.

11.5 Summary

Interface states are not cosmetic layers. They encode a theory of cognition. Systems that fail to distinguish recognition from action force users into perpetual response and thereby destroy memory.

Gallery-first systems respect the order in which understanding forms. They allow essence to be recognized before surface is acted upon. In doing so, they align interface behavior with the cognitive foundations of meaning itself.

12 Formal Transition Rules Revisited: Guaranteeing Recognition, Memory, and Search

The interface state taxonomy introduced earlier becomes meaningful only when coupled to explicit transition rules that constrain how and when semantic events may be generated. Without such constraints, state separation collapses back into feed behavior under pressure. This section therefore reframes the transition rules as semantic guarantees rather than interaction guidelines.

The purpose of these rules is threefold. First, they preserve the conditions for analogical recognition by protecting gallery encounters from interference. Second, they enable memoization by ensuring that repeated encounters strengthen existing referents rather than generating new ones. Third, they ensure that search operates over stable objects rather than over transient instances.

12.1 States as Semantic Filters

Recall the state set $\mathcal{S} = \{G, I, A\}$, corresponding to gallery, interaction, and annotation. Each state defines not only visible affordances but also the admissibility of semantic events.

This distinction is crucial. A system may appear calm while still generating semantic noise in the background. To prevent this, state must gate event production itself.

In the gallery state G , the system must be observational. No events targeting semantic objects may be produced. Navigation, scrolling, and search queries are permitted, but they leave no trace in the object–event graph. This preserves the purity of recognition. Encounter does not imply commitment.

In the interaction state I , lightweight events are permitted, but only those that do not alter object identity or payload. Reactions, acknowledgments, and conversational responses extend history without modifying referents. These events bind to objects and therefore contribute to memoization rather than fragmentation.

In the annotation state A , durable semantic modification is permitted. Events generated here explicitly extend the object–event graph by adding structure that must persist across projections. Because these events alter meaning, entry into this state must be explicit and reversible.

12.2 Transition Rules as Commitments

Define the transition function

$$\delta : \mathcal{S} \times \mathcal{U} \rightarrow \mathcal{S}$$

where \mathcal{U} denotes intent-bearing user actions.

The key constraint is that passive actions do not trigger state transitions. Scrolling, dwell time, or viewport exposure are insufficient. Only deliberate gestures—selection, invocation, or mode switching—may alter state.

This requirement eliminates accidental engagement by construction. It prevents the system from interpreting attention as action, a foundational error of engagement-optimized platforms.

12.3 Analogy Preservation Under Transition

Let $o \in \mathcal{O}$ be a semantic object encountered multiple times across projections. The transition rules must preserve the invariant that all encounters with o reinforce the same referent.

Formally, for any sequence of states $\{s_1, s_2, \dots\}$ and projections $\{\pi_{s_1}, \pi_{s_2}, \dots\}$, the object identity function must satisfy

$$\text{id}_o(\pi_{s_i}) = \text{id}_o(\pi_{s_j}) \quad \forall i, j.$$

State transitions may alter affordances and visibility, but they must not introduce new identities or shadow copies. This invariant is what allows the system to perform the analogical compression that cognition expects.

12.4 Memoization Guarantee

Memoization requires that repeated encounters with the same object increase its retrievability rather than dispersing it across instances.

Let $f(o)$ be a retrieval weight function derived from event history, such as visit count, annotation density, or revisit frequency. Transition rules must ensure that all admissible events contribute monotonically to $f(o)$.

Because gallery encounters produce no events, they do not inflate $f(o)$ artificially. Because interaction and annotation events bind to objects, they strengthen retrieval without duplication.

Thus, the system satisfies the memoization condition:

$$\forall e \in \mathcal{E}, e \text{ targets } o \Rightarrow f(o) \text{ increases.}$$

No event ever decreases retrievability or splits it across competing referents.

12.5 Search Correctness Under State Transitions

Search queries are projections over $\mathcal{G} = (\mathcal{O}, \mathcal{E})$. Correctness requires that search results be invariant under interface state.

Let Q be a search query and $\text{Search}(Q, \pi_s)$ the result set under projection π_s . Then for all $s \in \mathcal{S}$,

$$\text{Search}(Q, \pi_s) = \text{Search}(Q, \mathcal{G}),$$

up to presentational ordering.

This ensures that search retrieves objects, not surfaces. State affects how results are displayed, not what exists.

Feed-based systems violate this condition by allowing engagement context and recency to distort retrieval. Gallery-first systems enforce it by construction.

12.6 Why These Rules Cannot Be Relaxed

It may be tempting to weaken these constraints for convenience or growth. Doing so reintroduces the very failure modes this architecture is designed to eliminate.

Allowing gallery encounters to generate events pollutes memoization. Allowing passive attention to trigger interaction collapses state separation. Allowing search to operate over instances reintroduces fragmentation.

Each rule is therefore necessary. Together, they define a minimal semantic discipline under which meaning can accumulate.

12.7 Interpretation

These transition rules formalize restraint. They encode the principle that not everything that can be measured should be recorded, and not everything that can be acted upon should be solicited.

By enforcing deliberate transitions, the system aligns with cognition's natural order: recognition precedes action, and action precedes modification. Memory is protected not by limiting use, but by structuring it.

With these guarantees in place, the system satisfies the conditions for analogy, memoization, and coherent search simultaneously. What remains is to synthesize these results and clarify their implications.

13 Potential Objections and Responses

Any proposal that challenges feed-based architectures at a foundational level will encounter predictable objections. Many of these objections are reasonable, and addressing them directly helps clarify both the scope and intent of the gallery-first, merge-aware model. This section considers the most common critiques and explains why they do not undermine the core argument.

13.1 Objection: Users Want Chronological Feeds for News and Real-Time Events

A frequent objection is that users rely on chronological feeds for breaking news, live updates, and time-sensitive communication. According to this view, feed architectures are necessary because they support immediacy.

This objection conflates two distinct use cases: ephemeral communication and durable memory. The proposal advanced in this paper does not deny the legitimacy of real-time streams. It argues that streams should not be mistaken for archives.

Gallery-first systems do not eliminate feeds. They demote them from being the primary semantic substrate to being one projection among many. A real-time feed can exist as an ephemeral view optimized for immediacy, while the underlying semantic objects remain stable and retrievable. News can flow without overwriting memory.

In other words, the critique is not of chronology, but of chronology as ontology.

13.2 Objection: Merge Decisions Are Subjective and Ambiguous

Another objection concerns merge ambiguity. Determining whether two posts “should” be merged may appear subjective, particularly for expressive or personal content. Critics argue that automated merge risks collapsing meaningful distinctions, while manual merge introduces friction.

This objection misunderstands the role of merge. Merge is not required to be automatic or universal. It is required to be possible.

Merge-aware systems can support multiple merge regimes: automatic merging for identical payloads, suggested merging for near-identical content, and manual merging where intent is ambiguous. Crucially, refusing to merge must itself be an explicit decision rather than an architectural impossibility.

Feed-based systems already make irreversible decisions about equivalence by refusing to recognize it at all. Merge-aware systems simply return agency to the user.

13.3 Objection: Many Users Prefer Ephemerality and Drift

It is often argued that not all users want coherence, memory, or persistence. Some prefer ephemerality, fluid identity, and continual reinvention. Feed-based systems, on this view, are successful precisely because they support drift.

This objection confuses availability with enforcement. Gallery-first systems do not force persistence. They make persistence available.

Users who prefer ephemerality can choose projections that emphasize flow, novelty, and impermanence. What they cannot do in current systems is choose coherence even when they want it. Drift is mandatory. Memory is optional at best.

The proposal therefore increases expressive freedom rather than restricting it. It allows different temporalities to coexist without collapsing them into a single optimization regime.

13.4 Objection: Engagement Metrics Are Valuable for Creators

Creators often rely on engagement metrics to understand audience response, refine their work, or secure income. Critics may worry that suppressing visible metrics or deferring interaction affordances harms creators.

This objection again targets presentation rather than structure. Gallery-first systems do not eliminate metrics. They relocate them.

Metrics become properties of objects rather than of moments. They are accessible on demand rather than persistently visible. Creators can inspect performance without forcing every viewer into an evaluative frame.

This distinction preserves informational value while removing cognitive pressure. It treats metrics as tools, not as ambient signals that distort perception.

13.5 Objection: Feed Architectures Scale Better

A more technical objection holds that feed architectures scale more efficiently. Instance-level posts are easy to index, rank, and distribute. Object-level identity, merge operations, and event graphs appear more complex.

This objection mistakes engineering convenience for inevitability. Distributed systems routinely manage far more complex identity and merge constraints, from version control to replicated databases. The formal appendix demonstrates that merge-aware systems are not only tractable, but structurally simpler with respect to semantic guarantees.

More importantly, scalability without coherence is not success. Systems that scale engagement while destroying meaning merely scale entropy.

13.6 Objection: This Is a Niche Use Case

Finally, it may be argued that gallery-first, memex-style usage reflects a niche population—artists, researchers, or unusually reflective users—and should not dictate platform architecture.

This objection reverses causality. The rarity of coherent usage is a consequence of architectural hostility, not of lack of demand. When systems punish memory, users adapt by abandoning it.

The widespread emergence of external note-taking tools, personal archives, screenshot collections, and parallel knowledge systems suggests unmet demand rather than niche eccentricity. Users are building memory outside platforms because platforms refuse to support it.

13.7 Synthesis

None of these objections challenge the central claim of the paper. They instead clarify it. Feed-based systems enforce a single temporal logic optimized for engagement. Gallery-first systems permit multiple temporalities to coexist.

The question is therefore not whether feeds should exist. It is whether feeds should define what content is.

Once that distinction is made explicit, the architectural choice becomes clear. Systems can optimize for stimulation, or they can support meaning. They cannot do both by default.

Addressing objections does not weaken the argument. It sharpens it. Coherence is not an aesthetic preference, a productivity hack, or a nostalgic impulse. It is a structural property that must be enabled deliberately or it will be destroyed automatically.

The absence of such enablement is not neutral. It is a choice—and one whose consequences are now impossible to ignore.

14 Minimal Conditions for Semantic Drift and Accumulation

This section formalizes the central claim of the paper: that feed-based architectures necessarily exhibit unbounded semantic drift, while gallery-first, merge-aware systems necessarily permit semantic accumulation. The argument is constructive and relies only on minimal structural assumptions. It is not contingent on user behavior, moderation quality, or recommendation accuracy. The result is therefore architectural rather than empirical.

14.1 Preliminaries

Let \mathcal{C} denote the set of content instances presented to a user over time. Let \mathcal{E} denote the set of events, and let \mathcal{R} denote the set of reactions.

We distinguish between two architectural regimes.

In a *feed-based system*, content instances are primary, reactions bind to instances, and ordering is chronological.

In a *merge-aware system*, semantic objects are primary, events reference objects, and feeds are projections.

14.2 Assumptions for Feed-Based Systems

A feed-based system satisfies the following conditions.

First, each content instance $c \in \mathcal{C}$ has a unique identity even when two instances have identical payloads.

Second, reactions bind to instances rather than to payload equivalence classes. That is, if $c_1 \neq c_2$, then reactions to c_1 and c_2 are disjoint, regardless of content similarity.

Third, the system admits repeated introduction of semantically equivalent content instances over time.

Fourth, there exists no canonical merge operator μ such that semantically equivalent instances are unified into a single referent.

These assumptions are satisfied by all mainstream engagement-driven feed architectures.

14.3 Definition of Semantic Drift

Let $M(c)$ denote the meaning associated with content instance c , defined as the aggregate of reactions, annotations, and contextual associations accumulated over time.

Semantic drift occurs if, for a fixed payload p , the variance of $M(c)$ over instances c carrying p grows unbounded as the number of instances increases.

Formally, drift occurs if

$$\lim_{n \rightarrow \infty} \text{Var}(\{M(c_i) \mid c_i \text{ carries } p\}) = \infty.$$

14.4 Drift Theorem for Feed-Based Systems

Theorem 1. In any feed-based system satisfying the assumptions above, semantic drift is inevitable for any payload that is reintroduced infinitely often.

Proof. Let $\{c_1, c_2, \dots\}$ be an infinite sequence of content instances carrying identical payload p .

By assumption, reactions bind to instances, not to payloads. Therefore, each c_i accumulates a distinct reaction history $M(c_i)$.

Because instances appear in different temporal, social, and contextual neighborhoods, the distributions of reactions differ across instances. No mechanism exists to reconcile or unify these histories.

Since the system admits repeated introduction without merge, the set $\{M(c_i)\}$ grows without bound in cardinality and diversity. Therefore, the variance of meaning associated with payload p increases monotonically with the number of instances.

Hence semantic drift is inevitable. □

14.5 Assumptions for Merge-Aware Systems

A merge-aware system satisfies the following conditions.

First, there exists a set of semantic objects \mathcal{O} such that all events reference objects rather than instances.

Second, there exists an equivalence relation \sim over payloads such that semantically identical payloads reference the same object.

Third, reactions bind to objects via events.

Fourth, there exists a merge operator μ that is associative, commutative, and idempotent over events referencing the same object.

14.6 Definition of Semantic Accumulation

Let $M(o)$ denote the meaning associated with object o , defined as the aggregation of all reactions and annotations across events referencing o .

Semantic accumulation occurs if repeated references to o increase $|M(o)|$ without increasing the number of distinct semantic referents.

14.7 Accumulation Theorem for Merge-Aware Systems

Theorem 2. In any merge-aware system satisfying the assumptions above, repeated reference to a semantic object results in accumulation rather than drift.

Proof. Let $o \in \mathcal{O}$ be a semantic object with payload p . Let $\{e_1, e_2, \dots\}$ be an infinite sequence of events referencing o .

By construction, all reactions generated by these events bind to o . The merge operator μ ensures that event histories are consolidated without duplication or loss.

Therefore, the meaning function $M(o)$ grows monotonically while the referent set remains singleton. No variance across semantic identities is introduced.

Hence repetition strengthens meaning rather than fragmenting it. □

14.8 Corollary: Incompatibility of Feed Optimization and Semantic Conservation

Corollary. Any system that optimizes for engagement via instance-level metrics while lacking a merge operator over semantic equivalence classes cannot conserve meaning over time.

Justification. Instance-level optimization incentivizes duplication, resurfacing, and recontextualization without consolidation. Without merge, these operations necessarily increase semantic variance.

Therefore, semantic conservation and engagement optimization are structurally incompatible unless the latter is subordinated to object-level identity.

14.9 Interpretation

These results are intentionally minimal. They do not depend on user behavior, moderation quality, or recommendation accuracy. They follow directly from structural choices about identity, binding, and merge.

Feed-based drift is not an accident. It is a theorem.

Merge-aware accumulation is not a feature. It is a consequence.

This completes the formal argument that coherence must be designed, not tuned.

15 Concluding Synthesis: Meaning as an Architectural Property

This paper began with an ordinary frustration: the inability to use a personal timeline as a coherent gallery, memory surface, or memoization board without constant interference from prompts, metrics, duplication, and unsolicited content. That frustration proved to be diagnostic. When examined carefully, it revealed not a collection of isolated design mistakes, but a single architectural incompatibility between engagement-optimized feed systems and the conditions required for durable meaning.

Part I demonstrated that feed-based architectures systematically fail as memory substrates. By treating posts as feed-ordered atoms, collapsing identity surfaces into instance-level representations, subordinating search to exposure, and embedding monetization affordances directly into artifacts, these systems prevent repetition from reinforcing understanding. What should accumulate instead fragments. What should stabilize instead drifts. Silence is treated as error, and completion is rendered illegible.

Part II explained why these failures are inevitable under the prevailing model. Conceptual blending theory clarified how engagement-driven feeds function as perpetual blend spaces, mistaking a cognitive process optimized for insight as an infrastructural substrate for storage. Remix culture showed how propagation and transformation, while generative as cultural practices, fail catastrophically as archival logic when adopted wholesale. The theory of analogy and essence recognition demonstrated why instance-bound systems directly undermine cognition by refusing to acknowledge sameness across variation. The user remembers essences. The system returns surfaces.

Part III established that these outcomes are not intrinsic to digital media. They follow from specific design commitments that can be revised. By separating semantic objects from events, treating feeds as projections rather than sources of truth, and making merge operations explicit, gallery-first systems allow repetition to strengthen meaning rather than dilute it. By rendering interaction affordances latent rather than persistent, and by enforcing explicit transitions between gallery, interaction, and annotation states, such systems align interface behavior with the natural order of cognition: recognition first, action second, modification last.

The formal apparatus made these claims precise. Transition rules were shown to be semantic invariants rather than interface conventions. Memoization emerged as a guaranteed property rather than a user workaround. Search correctness followed directly from object-level identity rather than from ranking heuristics. The formal appendix demonstrated that semantic drift in feed-based systems is not an empirical tendency but a theorem, while semantic accumulation in merge-aware systems is a structural consequence.

Taken together, these results support a strong conclusion. Meaning is not a byproduct of engagement. It is not produced by exposure, amplified by reaction, or sustained by novelty. Meaning is conserved—or destroyed—by architecture.

Systems optimized for engagement must continuously prevent content from settling into stable referents. They must fragment identity, reward duplication, and suppress silence. Systems designed for coherence must do the opposite. They must allow artifacts to rest, histories to accrete, and recognition to occur without obligation.

This distinction is not ideological. It is structural. No amount of moderation, personalization, or algorithmic refinement can make a feed-based system behave like a gallery, a memex, or a memory board. Those functions require different primitives.

The proposal advanced here does not demand that all communication be slow, permanent, or curated. It demands only that systems distinguish between propagation and preservation, between stimulation and understanding, between surfaces and essences. It demands that meaning be allowed to remain.

The frustration that motivated this paper is therefore not a personal complaint. It is an early warning signal of a deeper mismatch between human sense-making and engagement-driven infrastructure. Addressing that mismatch requires not better feeds, but fewer of them—and better galleries in their place. Meaning cannot be optimized into existence. It must be conserved by design.

References

- [1] V. Bush. As We May Think. *The Atlantic Monthly*, 176(1), 1945.
- [2] G. Fauconnier and M. Turner. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books, 2002.
- [3] M. Turner. *The Origin of Ideas: Blending, Creativity, and the Human Spark*. Oxford University Press, 2014.
- [4] D. R. Hofstadter and E. Sander. *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking*. Basic Books, 2013.
- [5] H. A. Simon. The Architecture of Complexity. *Proceedings of the American Philosophical Society*, 106(6), 1962.
- [6] E. Hutchins. *Cognition in the Wild*. MIT Press, 1995.
- [7] S. Brand. *How Buildings Learn: What Happens After They're Built*. Penguin Books, 2018. Originally published 1994.
- [8] L. Manovich. *The Language of New Media*. MIT Press, 2001.
- [9] L. Lessig. *Remix: Making Art and Commerce Thrive in the Hybrid Economy*. Penguin Press, 2008.
- [10] Y. Benkler. *The Wealth of Networks*. Yale University Press, 2006.
- [11] T. Gillespie. *Custodians of the Internet*. Yale University Press, 2018.
- [12] S. Zuboff. *The Age of Surveillance Capitalism*. PublicAffairs, 2019.
- [13] L. Floridi. *The Fourth Revolution: How the Infosphere Is Reshaping Human Reality*. Oxford University Press, 2014.
- [14] D. A. Norman. *The Design of Everyday Things*. Basic Books, Revised Edition, 2013.
- [15] E. R. Tufte. *The Visual Display of Quantitative Information*. Graphics Press, 2001.
- [16] R. Rogers. *Digital Methods*. MIT Press, 2013.
- [17] J. Packer and S. B. Crofts Wiley (eds.). *Communication Matters: Materialist Approaches to Media, Mobility and Networks*. Routledge, 2013.
- [18] L. Lamport. The Part-Time Parliament. *ACM Transactions on Computer Systems*, 16(2), 1998.
- [19] M. Fowler. Event Sourcing. martinfowler.com, 2017. Online essay.
- [20] M. Muller et al. Designing for Meaningful Social Interaction. *Proceedings of CSCW*, 2019.