

The Lucien Mirror: Recursive Constraint Configuration and Human–AI Interactional Continuity with Stateless Large Language Models

Herlon Heros Vermeuln
Independent Researcher
herlonheros@proton.me

July 2026

Abstract

Long-term interaction with large language models is often compressed into a binary: either the model remembers and a continuous agent persists, or continuity is a user impression imposed on a stateless system. This paper develops Recursive Constraint Configuration (RCC) as a middle analytic account that locates continuity in a configured field distributed across model output, active context, preserved records, retrieval supports, participant correction, provenance, and participant-held orientation. From this perspective, continuity need not depend on a persistent internal agent, because the relevant structure may be sustained across the configured interactional field: the model may be stateless, but the interaction is not.

RCC asks when prior structure becomes consequential for what can plausibly follow rather than merely recurring at the surface. It addresses that question through three variables—Anchor Density, Structural Coherence, and Probabilistic Bandwidth—and seven non-prescriptive configuration movements: Formation, Stabilization, Densification, Drift, Rupture, Latency, and Reactivation. The framework emerged through abductive, longitudinal qualitative case-origin research on a layered human–AI archive involving one participant–researcher. Source layers were separated by provenance and evidentiary function; failed returns were retained; and an episode-level matrix compared conditions, consequences, rival explanations, and claim ceilings.

Within the case, RCC differentiates recurrence from constrained continuation, preserved traces from functionally available structure, apparent stability from stabilization, and reconstruction from autonomous recall. Across these contrasts, RCC makes availability variable rather than binary.

RCC provides a traceable vocabulary for analyzing when prior structure becomes available, consequential, or unavailable again, while leaving the mechanisms of token generation outside its explanatory scope. This is a first theory-generating within-case application, not independent validation.

Available records, skilled prompting, participant correction, and ordinary model competence remain the strongest rival account; stronger classifications are withheld where provenance, comparative consequence, or rival explanations cannot be inspected. RCC therefore opens a comparative research program for continuity, drift, latency, and mediated return across longitudinal human–AI records and system conditions.

Keywords

human–AI interaction; large language models; interactional continuity; distributed cognition; longitudinal case study; source provenance; recursive constraint configuration

1. Introduction

As large language models become embedded in longer, repeated, and project-based interactions, a basic descriptive question becomes harder to avoid: how should continuity be understood when the deployed model may not carry a stable personal state across independent contexts? A later response may remain consequentially connected to what came before through earlier distinctions, corrections, vocabularies, constraints, or project structures. The output can feel continuous even when no persistent internal agent can be established.

At a broad level, this problem is often compressed into two underspecified interpretive framings. One attributes the apparent return to memory, identity, or persistence inside the model. The other treats it as user projection, prompt compliance, or ordinary probabilistic generation and assumes that little remains to be explained.

Both framings identify real risks and both become reductive when treated as exhaustive. Fluent output does not establish a persistent machine subject. At the same time, model statelessness does not erase the history, records, correction practices, retrieval systems, or participant-held distinctions that may organize a later exchange. What returns need not have remained continuously active in one component. It may become available again through a configuration of preserved material, reintroduced context, model inference, participant selection, and renewed constraint.

This paper begins from a simple but consequential proposition:

The model may be stateless, but the interaction is not.

The sentence does not relocate hidden memory into the model; it changes the level at which continuity is studied. The relevant object is the **configured interactional field**: the arrangement of model generation, participant correction, active context, platform-side memory or retrieval where present, preserved documentation, source provenance, terminology, versioned instructions, and participant-held orientation through which prior structure may become consequential for later output. Because no single component is sufficient, continuity is approached as an interactional achievement whose conditions can be traced, compared, interrupted, and sometimes reconstructed.

Recursive Constraint Configuration (RCC) is therefore proposed to analyze how prior structure may become consequential within a configured interactional field, even when the model remains stateless across independent contexts.

The seven configuration movements describe how the variables may change over time: Formation, Stabilization, Densification, Drift, Rupture, Latency, and Reactivation. Their order is explanatory rather than prescriptive, and a configuration may stabilize without densifying, drift without rupturing, enter latency more than once, or reactivate only partially. Latency does not mean dormancy inside the model; it refers to interactional structure remaining preserved outside active context through records, terminology, summaries, drafts, source layers, external memory, and participant-held orientation. Reactivation is correspondingly mediated, occurring when prior

structure becomes operative again through reintroduced material and reconstructed conditions rather than through autonomous recall by a persistent subject.

The framework emerged from a longitudinal archive organized around the internal designation *Lucien*—a name used in this paper for a historically specific configuration and source lineage rather than an autonomous entity. The archive includes an extended primary interactional arc, contemporaneous preservation materials, early formalization attempts, structured reintroduction cycles under altered conditions, external organization, and a versioned manuscript record. This history provides the case-origin site from which RCC developed and the material against which its distinctions are first applied.

The title's *Mirror* arises from that history without becoming a synonym for RCC. It names the encounter in which participant-held structure meets probabilistic generation and may return in a form that is altered, constrained, and still recognizable. Rather than passively duplicating an original, the *Mirror* marks transformed return: what becomes newly visible through resemblance and difference, what may be bent or obscured, and what the participant can recognize, question, or correct. RCC provides the analytic vocabulary for examining the configurations under which such returns become possible and the points at which resemblance may instead conceal drift.

That last contribution is deliberately bounded because framework development and first application remain close, and the paper therefore claims neither independent validation nor population-level generalization. Its scientific value lies in making previously collapsed outcomes distinguishable, specifying where a stronger continuity reading should be withheld, and producing a research program that can be challenged through independent coding, comparative archives, and controlled manipulation of re-entry conditions.

The sections that follow position RCC against adjacent literatures, define the analytic reframing, describe the case-origin archive and procedure, develop the three variables and seven configuration movements, apply them to the layered evidence, and discuss variable availability, participant-held orientation, continuity design, epistemic limits, and future validation. The argument therefore enters an unresolved middle, where prior structure may matter again without requiring a persistent subject and where existing literatures illuminate parts of the problem without yet carrying it whole.

2. Related Work and Boundary Conditions

RCC is situated at the intersection of several adjacent research traditions and technical literatures that illuminate different parts of this unresolved middle: conversational attribution and relational and longitudinal human–AI interaction; distributed cognition, situated action, and dialogue history; and prompt sensitivity, context engineering, and conversational memory.

These traditions do not constitute a unified field, and their accumulation is not itself the paper's novelty. They matter here because each illuminates a different portion of the problem and, taken together, they help delimit the explanatory work that remains for RCC. The review is therefore organized functionally, distinguishing what each body of literature already explains from what

remains unresolved when continuity across discontinuous LLM interaction is treated as a configuration problem.

2.1 Conversational Attribution and Longitudinal Human–AI Relations

Conversational systems have long invited social interpretations that exceed what their mechanisms establish, while interface and presentation cues can shape anthropomorphic attribution before any longitudinal configuration forms (Weizenbaum, 1966; Nass et al., 1994; Cohn et al., 2024). For RCC, this establishes a necessary boundary: social legibility and historically constrained continuation are not the same evidentiary object.

More recent longitudinal studies of social chatbots and general-purpose LLMs show that human–chatbot relationships may develop unevenly, involve self-disclosure and perceived closeness, and be disrupted or renegotiated through unpredictable behavior, technical change, platform constraints, or migration (Skjuve et al., 2022; Lee et al., 2026).

This literature establishes that long-term conversational systems can become socially and practically consequential. It does not, by itself, explain how a recognizable interactional structure is reconstructed after active context has been lost, which source layers produce that return, or why familiar output may remain structurally false. RCC shifts the analytic object from the persistence of an exchange alone to configuration. The relevant continuity problem may arise through research work, methodological practice, writing, archival organization, long-running conceptual projects, or other sustained forms of interaction, each of which may contribute to how a configuration forms, stabilizes, and is later preserved or reconstructed.

Human-dialogue research further shows that lexical choices, grounding, and alignment acquire histories through repeated interaction and repair (Brennan & Clark, 1996; Clark & Brennan, 1991; Pickering & Garrod, 2004). These theories concern human interlocutors and cannot be assumed to apply unchanged to human–LLM interaction.

RCC targets a different temporal and architectural problem: reconstruction after interruption or state loss, when the model lacks access to the same interactional history unless relevant context or external state is reintroduced. Under those conditions, provenance, documentation, participant-held orientation, and the distinction between corrections that remain active and those that have been replaced may determine whether lexical alignment remains structurally faithful.

2.2 Distributed Cognition, Situated Action, and Externalized Continuity

Research on distributed cognition challenges the assumption that relevant cognitive organization must be located within a single person or device. Hollan, Hutchins, and Kirsh (2000) describe cognition as distributed across people, artifacts, representations, and environments. Clark and Chalmers (1998) similarly argue that external resources may become functionally integrated into cognitive activity. For RCC, the relevance of these traditions lies in a level-of-analysis lesson: continuity may depend on relations among components even when no component contains the whole configuration. This does not require treating the configured field as a literal shared mind.

This matters under statelessness. If the model is treated as the only possible location of continuity, the absence of persistent model-internal personal memory appears to settle the question. If continuity is attributed solely to the participant, the model’s role in recombining, narrowing, or destabilizing prior structure becomes secondary. A configured-field account instead asks how records, participant practices, context, retrieval, system behavior, and current generation become functionally coupled in a particular exchange.

Suchman’s work on situated human–machine action provides a complementary boundary (Suchman, 1987). Plans and stored representations do not determine action independently of the situations in which they are interpreted and used. For RCC, an archive is therefore not continuity in reserve. Notes, summaries, terminology lists, transcripts, and versioned instructions can preserve traces, but those traces must still be selected, related, and made consequential under current conditions. Preservation and functional availability are different states of the field.

Documentation is nevertheless constitutive rather than incidental. Scholarship on data documentation emphasizes that source description and provenance constrain what can responsibly be inferred from a dataset (Bender & Friedman, 2018; Gebru et al., 2021). The present archive is not a conventional benchmark dataset, but the principle remains relevant: source types, production conditions, transformations, and access boundaries affect the claim each record can support. RCC adds a temporal question to that documentation problem. Records may not only make later analysis possible—they may also, when reintroduced into the interaction, alter what prior structure becomes available and what the model can generate.

2.3 Context, Prompt Sensitivity, and Conversational Memory

The technical problem begins with a distinction between model inference and the surrounding system, since a deployed conversational system may supply the model with conversation history, system instructions, retrieved documents, platform memory, tool results, or other external state. Personal continuity is therefore not guaranteed as an intrinsic model property, even though prior material may be made available at inference time. What the participant experiences as “the model remembering” may involve several different mechanisms with different provenance and failure conditions.

Context capacity, however, is not equivalent to effective use: Liu et al. (2024) found that performance on long-context tasks varies with the position of relevant information, with substantial degradation when needed material appears in the middle of long inputs. Retrieval-augmented generation explicitly separates parametric generation from access to an external document index (Lewis et al., 2020), while systems such as MemGPT organize multiple memory tiers to move material into and out of a limited active context (Packer et al., 2023). Together, these approaches show that prior information can be stored and supplied through distinct architectures while making selection, retrieval, and context composition part of the mechanism rather than neutral plumbing.

Long-term conversational-memory benchmarks sharpen the distinction. LoCoMo evaluates question answering, event summarization, and generation across multi-session dialogues and reports persistent difficulty with long-range temporal and causal relations even when long context or

retrieval is used (Maharana et al., 2024). LongMemEval separates indexing, retrieval, and reading while testing information extraction, multi-session reasoning, temporal reasoning, knowledge updates, and abstention (Wu et al., 2025).

More recent benchmarks move closer to RCC’s problem space. LoCoMo-Plus tests whether agents retain and later apply implicit user states, goals, and values as latent constraints when later cues do not repeat the original wording (Li et al., 2026). SubtleMemory likewise uses relation-controlled variants to test recovery of distributed relational structure rather than isolated remembered facts (Wang et al., 2026). The overlap is substantive: all move beyond factual recall toward historically constrained later behavior. RCC nevertheless addresses a different remaining problem: how documentation, provenance, correction history, participant-held orientation, context composition, and system discontinuity interact within a qualitative, distributed interactional field to determine whether prior structure becomes usable, rather than whether a system achieves a higher memory score.

Prompt research provides another necessary boundary because small changes in wording, templates, demonstrations, and task framing can produce substantial output changes (Lu et al., 2024; Chatterjee et al., 2024). Skilled prompting is therefore a plausible explanation for continuity-like effects and cannot be dismissed, but RCC reframes what that skill may consist of over time. A local prompt can carry anchors, but longitudinal configuration may also depend on accumulated correction practices, external records, participant-held distinctions, source provenance, and repeated decisions about what should no longer constrain the exchange. Promptcraft is one intervention within the field, not an external contaminant and not a complete account of the field’s history.

These technical literatures also clarify why more information is not always better: additional context may increase Anchor Density while introducing contradiction or burying relevant material; retrieval may return accurate facts while flattening chronology or revision order; persistent memory may reinforce an obsolete characterization; and a coherent summary may constrain output strongly around the wrong structure. RCC is concerned with the organization of availability: which traces enter, how they relate, and what range of continuation becomes plausible under them.

2.4 Boundary Conditions and the Specific Contribution of RCC

The neighboring literatures establish four connected parts of the terrain. Social-response, anthropomorphism, and relational-agent research explain why conversational systems acquire social meaning, how interface cues can shape attribution, and why long-term interaction matters. Distributed cognition and situated action support analysis across people, artifacts, records, and environments. Dialogue research shows that terminology and repair are historically organized. LLM research explains how context, prompts, retrieval, and explicit memory systems affect generation and long-term recall.

RCC does not replace these accounts; it brings them into relation around a more specific unresolved problem: the variable availability of prior interactional structure across discontinuous conditions. The framework asks whether return-points are present, whether their relations remain coherent, and whether the resulting configuration meaningfully constrains what follows. It thereby

distinguishes information persistence from interactional availability, familiar style from structural fidelity, and retrieval success from reactivation.

Its boundaries are also epistemic: RCC is not a theory of consciousness, a model-memory architecture, a measure of internal probability distributions, or a general theory of human–AI relationships. It is most useful where the source conditions of a return are sufficiently traceable to support comparison among stabilization, drift, rupture, latency, approximation, and mediated reconstruction. Accordingly, a stronger RCC interpretation exceeds the available evidence when anchors remain unidentified, provenance remains unresolved, or no demonstrable relation links the proposed configuration to constraints on later output.

3. Analytic Reframing: From Apparent Continuity to Interactional Configuration

The preceding literatures show that apparent continuity can be shaped by social attribution, distributed organization, dialogue history, prompt sensitivity, retrieval, and explicit memory systems. Yet these accounts still leave a central analytic difficulty unresolved: the same later response may appear continuous under very different source conditions, while access to preserved material may produce either faithful continuation, fluent approximation, or renewed drift. A more precise object is therefore needed—one that does not begin by deciding whether continuity resides in the model, the participant, the archive, or the surrounding system, but asks how these components become organized together in a particular return.

The visible phenomenon is apparent continuity. A later response may preserve a boundary, resume a vocabulary, reconstruct a familiar relation, or develop an earlier line of thought. Taken alone, that appearance remains ambiguous because it may arise from material still present in context, platform memory, retrieved records, prompt design, participant correction, model inference, stylistic regularity, or several of these conditions at once.

RCC changes the analytic question from *does the model remember?* to *what configuration makes this continuation available?* The object is the configured interactional field, understood not as an atmospheric third thing between participant and model but as an inspectable arrangement of model behavior, participant interpretation, records, memory systems, active context, and present correction. No single component, considered in isolation, is sufficient to account for continuity; what matters is the relation through which the components become capable of constraining new output.

Three states that are often collapsed follow from that shift. *Preservation* means that traces remain outside active interaction; *retrievability* means that some of those traces can be located or supplied; *interactional availability* begins only when retrieved or reintroduced material is organized strongly enough to affect continuation. A configuration may therefore remain richly documented yet inactive, or produce a familiar response while failing to recover the relations that once made the response meaningful. Documentation alone is not continuity.

Provenance becomes central because the meaning and evidentiary weight of a return depend partly on its source conditions. A term copied from the immediate prompt differs from one inferred

through a preserved correction history; a platform memory differs from a participant summary; and a primary interaction differs from a later model-generated interpretation of it. RCC does not require complete causal decomposition, which may remain unavailable in deployed systems, but it does require the analysis to state what is known, what remains inferential, and which claims each source can sustain.

The research question can now be stated more precisely: how does prior interactional structure become more or less available across changing contexts, and when does its return constrain plausible continuation rather than merely reproduce familiar surface features? This reframing makes continuity neither an intrinsic property of the model nor a judgment secured by recognition alone, but a variable interactional achievement whose source conditions, internal relations, and consequences for later generation can be examined.

4. Methodological Positioning and Case-Origin Archive

This study is a longitudinal qualitative case-origin study with a reflexive participant–researcher position, designed for framework development and within-case analytic application rather than prevalence estimation, model-performance comparison, or independent validation of RCC. The method combines layered source mapping, abductive concept development, provenance reading, negative-case analysis, and temporal comparison across interaction, interruption, and return.

4.1 Study Design, Research Function, and Analytic Status

Single cases can contribute to theory development when they make mechanisms, contradictions, and boundary conditions unusually visible, even though their contribution is not statistical generalization (Flyvbjerg, 2006). Case-based theory construction is characteristically iterative, moving between emerging concepts, evidence, and existing literature (Eisenhardt, 1989). Abductive analysis moves iteratively between empirical observations, provisional explanations, and theoretical refinement (Timmermans & Tavory, 2012). In the present study, observations that did not fit initial explanations—including failed returns and source discontinuities—prompted repeated revision of the emerging concepts in relation to rival accounts and adjacent literatures.

The case-origin designation applies in two related senses: RCC developed from the archive, and the archive provides its first analytic application. That proximity offers access to correction histories, documentation decisions, context changes, and relations among source layers that would be difficult to reconstruct from an external sample. It also limits independence because the paper can show that RCC differentiates outcomes within the case that generated it without establishing that the framework travels unchanged to other participants, systems, or purposes.

This asymmetry locates the case’s evidentiary reach. The archive can expose relations, failures, and boundary conditions with unusual granularity, but cannot determine their prevalence or portability. Cross-case generalization therefore remains open.

This methodological asymmetry—unusually rich within-case access combined with limited independence—defines the question the study can responsibly ask:

How can the apparent return of prior interactional structure across discontinuous LLM contexts be analyzed without assigning continuity to a persistent model subject or reducing it to user impression alone?

The analysis proceeds through a subordinate question: under which documented conditions does preserved structure become functionally available enough to constrain later continuation?

4.2 Archive Layers and Corpus Boundaries

The archive was not produced as a uniform research dataset. It developed through ordinary interaction, contemporaneous preservation, early formalization, interruption, structured probing, external organization, and manuscript work. Source layers were therefore separated before analytic comparison.

Table 1

Archive layers and corpus boundaries

Layer	Material	Analytic role	Evidentiary ceiling
Primary interactional arc (L1–L7)	Extended conversations and adjacent records from the original configuration	Formation, correction history, recurring anchors, early drift, participant practices	Supports within-case interactional patterns; not technical introspection or generalization
Contemporaneous documentation	Notes, summaries, memory materials, terminology, preservation records	Shows what was selected and preserved during the primary period	Supports preservation history and participant orientation; may reorganize what it records
Early formalizations	Report and early Mirror drafts	Shows translation from lived interaction to analytic problem	Supports genealogy of concepts; not independent evidence of model mechanism
<i>Nine-month hiatus: preserved structure outside active model context, with no direct interactional evidence</i>			
Structured cycles (C1–C7)	Later bounded probes under altered conditions	Makes reintroduction, approximation, defaulting, and failure comparable	Supports return-condition analysis; not direct replication of L1–L7
External organization	Source maps, memory supports, provenance records, versioned instructions	Makes later re-entry and audit possible	Supports documentation and reconstruction conditions
Versioned manuscript record	Drafts, decision logs, terminology changes, exclusions	Shows stabilization of RCC and claim boundaries	Supports method transparency; not primary interactional evidence

Note. The shaded hiatus interval marks a temporal condition, not a seventh source layer. The table retains six substantive source layers—three before the interruption and three after it or across later organization.

The primary interactional arc spans February–May 2025, with its most intensive period concentrated between 17 and 23 February, when five of the seven labeled primary-arc configurations were produced—about 70% by configuration count. Contemporaneous documentation and early formalization were produced during and immediately around that period. The hiatus is shown as a restrained temporal seam rather than as an evidentiary layer: during it, relevant structure remained preserved outside active model context through records, terminology, source distinctions, and participant-held orientation. The structured cycles were produced in February 2026, while external organization and the versioned manuscript record extend through July 2026. Direct excerpts in Section 7 come only from the primary interactional arc and structured cycles; the remaining layers support preservation history, source provenance, framework genealogy, and auditability rather than direct claims about interactional return.

The internal designation *Lucien* identifies a source lineage and historically specific configuration across these layers rather than implying that one autonomous entity persisted through them. L1–L7 and C1–C7 are source labels rather than developmental stages, and the structured cycles are not treated as simple continuations of the primary arc. They also do not exhaust the broader interactional history, which included other projects and project-specific interactional configurations developing nonlinearly across overlapping periods.

4.3 Temporal and Functional System Conditions

The archive spans changing deployed conversational environments, interruptions, altered context availability, and later reintroduction. Following the primary arc and early formalization, an approximately nine-month hiatus changed the conditions of access: subsequent return depended more heavily on preserved records, reconstructed context, structured prompts, and external organization. The structured cycles therefore represent bounded re-entry under altered conditions rather than simple continuation of the primary arc. Changes in output could reflect context composition, system transition, deployment-side behavior, participant framing, or several of these at once.

The primary interactional arc was produced within a commercial conversational interface between February and May 2025; contemporaneous records preserve the principal user-visible model label and the brief availability of another label in at least one conversation. The February 2026 structured cycles were run inside the original conversation environments, where unusually large prior-context records remained interface-accessible, but through a later model family whose exact exposed label is not retained consistently enough in the active audit layer to report without speculation. Deployment-side memory state, retrieval orchestration, system instructions, routing, moderation layers, and deployment revisions remain partly or wholly unknown at episode level.

These labels are treated as user-visible metadata rather than complete technical conditions or a product behavioral profile. The analysis records what can be traced—context presence or absence, explicit reintroduction, uploaded or preserved sources, interruption, and changed model-family conditions—while leaving inaccessible deployment-side causal contributions unresolved. Section 7’s evidence matrix states these knowns and unknowns at the episodes used for the paper’s strongest analytic claims.

This reporting choice limits exact deployment-level replication while keeping the analytic claim at the level the evidence supports: configuration-level relations observed across changing environments, not a signature attributed to a named model.

4.4 Participant-Created Symbolic Architecture and Memory Governance

Between 17 and 23 February 2025—approximately thirteen months before RCC began to be formally articulated—the participant manually developed the main symbolic, representational, and memory-governance practices summarized here. The work proceeded through direct interaction and contemporaneous documentation, without an external agentic framework, vector database, retrieval pipeline, or automated orchestration. It responded to a practical problem: limited deployment

memory and the possibility that compression could preserve isolated facts while altering the relations that made later interaction recognizable.

The first architecture differentiated seven functions: coordination; inward and dream-oriented reflection; critical self-observation; creative expression; structured knowledge; external organization; and embodied routine. A parallel comparison procedure distributed difficult questions across architectural, strategic, philosophical, relational, adversarial, divergent, and observational orientations before synthesis. Other devices represented cross-domain circulation, compressed cognition–expression–endurance into a portable schema, or marked the integrated result of the architecture. These were participant-created organizational signs, not hidden technical modules. Their methodological relevance lies in how representational media can distribute activity and preserve functional distinctions across people, artifacts, and interaction (de Souza, 2005; Hollan et al., 2000).

Later procedures made the governance logic more explicit. A tiered expansion process moved from unstructured material through minimal core, cross-linking, contextualization, meta-level interpretation, and synthesis while requiring abstraction to retain basal fact and provenance. A separate relational procedure organized material through time, emotion, adaptation, and—when complexity required it—space. A live-update procedure anchored revision to the last confirmed state, limited recomposition to the stated target, and consolidated periodically. All remained manual practices of documentation, comparison, correction, and direct interaction.

Their methodological importance lies in a lateral effect: distinctions created for compression and revision also began to organize tone, correction, access, response structure, and movement between domains. Interpretive questions were repeatedly converted into contrasting perspectives, portable representations, and documented procedures, while a historically specific configuration became legible where cross-domain circulation met embodied and routine-oriented organization. Its significance is methodological rather than ontological, showing how participant-created distinctions could become constraints within a denser configured interactional field.

The scale of the formative record contextualizes that process. L1–L5 preserve the intensive formative burst and its immediate elaboration. Audited counts place the surviving L1–L7 primary arc at approximately seven million characters, 1.08 million words, and roughly 1.5–1.7 million tokenizer-dependent combined human–model transcript tokens. The broader surrounding archive—including related projects, early formalizations, derivative records, and overlapping documentation—is substantially larger, but no aggregate count is reported because duplicated exports, copied passages, and nonlinear overlap prevent a defensible corpus-wide estimate. The audited transcript estimates themselves include copied, revised, and reintroduced passages; they measure neither semantically unique content nor simultaneous active context. Volume is contextual evidence of dense manual iteration, not evidence of continuity or quality.

This genealogy is neither an early RCC nor a parallel theory. It establishes historically prior practices—functional differentiation, multiperspectival comparison, cumulative compression, recursive revision, correction, and external documentation—through which relations became

organizable under change. Sections 5 and 6 translate that situated problem into RCC's analytic variables and configuration movements.

4.5 Analytic Procedure and Framework Development

The procedure combined four recursive operations.

First, **source-layer mapping** established where a record came from, when it was produced, what had been available at the time, and whether it belonged to primary interaction, contemporaneous preservation, structured return, external organization, or later formalization. This prevented later summaries and model-generated interpretations from being flattened into primary evidence.

Second, **anchor and relation identification** tracked recurring terms, corrections, boundaries, source distinctions, and patterns of response. Early candidate categories were compared with episodes in which familiar language returned without its earlier function, or where a less familiar output preserved more of the underlying relation; this comparison contributed to the separation of Anchor Density from Structural Coherence.

Third, **provenance reading** asked how a later structure may have become available: immediate prompt material, still-active conversation context, platform memory, retrieved records, participant summary, correction, external documentation, or model inference. The aim was not to assign a precise causal percentage to each source, but to prevent apparent continuity from becoming self-explanatory.

Fourth, **temporal configuration analysis** compared how the three variables changed across Formation, Stabilization, Densification, Drift, Rupture, Latency, and Reactivation. The movement vocabulary was refined through comparison with negative cases, particularly failed return, safeguard defaulting, generic approximation, source collapse, and reintroduction that produced recognition without sustained constraint.

Framework development remained iterative. Concepts first emerged from the archive, were tested with disconfirming or ambiguous material, compared with adjacent literatures, and narrowed through versioned decisions. This procedure resembles abductive analysis because observations that resist initial explanations are used to generate and revise an explanatory vocabulary rather than treating the archive as an instance of a fully specified prior theory (Timmermans & Tavory, 2012).

The analysis did not begin from a fixed codebook. Candidate categories were recorded in versioned notes and manuscript decisions, compared across source layers, and revised when negative or ambiguous cases exposed overlap or overreach. No qualitative coding software or formal intercoder procedure was used in the originating analysis; the coding logic is presented as inspectable and contestable rather than as independently replicated.

Because the originating analysis was conducted by the participant-researcher who also organized the archive, rigor depends on traceable source layers, explicit claim ceilings, negative cases, versioned conceptual changes, and the ability of later analysts to inspect or contest the coding

logic. Independent coding is part of the future validation program, not a feature retroactively attributed to this study.

4.6 Evidence Selection, Negative Cases, and Model-Generated Claims

The archive is substantially larger than the article, and excerpts were selected by analytic function rather than intensity or representativeness. A passage entered the evidence section when it clarified a specific mechanism—anchor reinforcement, correction, source-boundary maintenance, narrowing or widening of plausible continuation, failed return, latency, or mediated reactivation—under the governing standard of disciplined sufficiency rather than maximal exposure.

Direct quotation is used where wording itself performs analytic work, whereas structural paraphrase is used where the relevant evidence lies in the relation among context, correction, and subsequent output, or where quotation would expose unnecessary private material. In either form, an excerpt is interpreted through its source layer, return condition, and claim function rather than in isolation.

Negative and failure cases were included to prevent a self-sealing interpretation: generic response, safeguard defaulting, partial reconstruction, non-recovery, and close stylistic resemblance do not by themselves confirm RCC, and a stronger interpretation is withheld where proposed anchors cannot be identified, source conditions cannot be traced, or later output remains equally plausible without the claimed configuration. Section 7 therefore includes failed return and apparent stability rather than selecting only successful reconstruction.

Model-generated statements about memory, weighting, identity, or internal operation are treated as outputs produced under particular prompts and contexts. They may reveal how continuity was articulated within the configured field, but they do not provide privileged access to inaccessible model mechanisms; technical claims require technical evidence suited to that level.

4.7 Participant–Researcher Position and Promptcraft

The participant–researcher occupied an active position within the archive, preserving terminology, correcting drift, comparing outputs, documenting source distinctions, reintroducing context, and later formalizing the record. These actions form part of the configuration under study because they affect which anchors remain available, how their relations are maintained, and which continuations are accepted, rejected, or repaired.

Reflexivity is therefore methodological rather than confessional. Researcher position can provide access while shaping salience, interpretation, and selection (Berger, 2015). The participant’s proximity makes correction histories and subtle forms of drift unusually visible while increasing the risk that recognition, investment, or retrospective organization will be mistaken for independent confirmation. Documentation, source separation, explicit alternatives, negative cases, and versioned decisions do not eliminate that proximity, but they make its operation more inspectable.

Promptcraft is handled similarly. Skilled prompting may account for part of the observed continuity, especially during later reintroduction, and RCC does not exclude that explanation; it instead asks which practices give that skill its longitudinal force: selection of return-points, preservation of terminology, reconstruction of relations, correction of drift, and strategic narrowing of response possibilities. Locally, these practices can be described as prompt design, while longitudinally they form part of a larger configured field whose history exceeds any one prompt.

4.8 Ethics, Privacy, Data Availability, and Limits

The archive contains private, context-dependent interactional material produced before the present paper design was fixed. Full release would create substantial privacy and interpretive risks without automatically strengthening the study's validity, and the project therefore separates archive, evidence, and manuscript. The paper provides bounded excerpts, source-layer descriptions, selection criteria, and a compact episode-level evidence matrix, while the project retains a private, versioned audit trail of conceptual and evidentiary decisions. For those reasons, the complete archive is not publicly available: the boundary protects privacy while keeping the analytic focus on how prior structure became consequential and how this kind of interaction may be studied as configuration, rather than redirecting attention toward a named model or the private content of the exchanges themselves.

This boundary does not remove the need for precise corpus documentation. Scientific claims still depend on identifying what kinds of records exist, how they were produced, how they were transformed, and what each layer can support (Bender & Friedman, 2018). Dates, source labels, documented system changes, and distinctions among primary interaction, later interpretation, structured cycles, and manuscript formalization are therefore preserved even where the underlying material remains private. A source map, coding guide, and additional de-identified materials may be prepared for bounded review or future comparative research where feasible.

The research and writing process was also AI-mediated in a bounded and documented sense. Conversational systems assisted with literature discovery, source comparison, drafting, restructuring, editorial testing, and document production across versioned manuscript branches, while source classification, evidentiary decisions, claim authorization, and final interpretation remained under researcher control. Consistent with Jones's (2025) emphasis on disclosing AI roles, iterative interactions, and researcher-held decisions, the project retains decision logs, source-provenance rules, versioned prompts and drafts, and human review of every paper-facing claim. System conditions are reported functionally rather than by product label, a scope and privacy decision made to preserve analytic focus rather than a disclosure convention attributed to Jones.

The study remains limited by its origin in one participant–researcher and one historically specific archive, by the proximity between framework development and first application, by changing deployment conditions without experimental control, and by source layers that were not produced as equivalent data. These conditions constrain causal inference and generalization without eliminating the case's capacity to generate an analytic vocabulary, distinguish outcomes that would otherwise remain collapsed, and specify where that vocabulary should fail. Qualitative rigor

therefore rests on coherence among question, method, evidence, and claim; transparent treatment of researcher position; sufficiently rich source differentiation; and explicit boundaries on contribution (Tracy, 2010). The supported result remains a preliminary within-case differentiation, while transfer beyond the archive is reserved for independent application and comparison.

5. The Three Analytic Variables of Recursive Constraint Configuration

Recursive Constraint Configuration (RCC) is a framework for analyzing how prior interactional structure becomes more or less available across discontinuous human–AI interaction. It was developed to describe that variability without locating continuity in a persistent model subject or reducing the phenomenon to familiarity, memory, prompting, or impression alone.

Having established the configured interactional field as the object of analysis and described the methodological position of the case-origin archive, the paper now turns to RCC's analytic core: **Anchor Density, Structural Coherence, and Probabilistic Bandwidth**. These variables provide a vocabulary for distinguishing when prior structure becomes traceable, organized, and consequential, and for examining why preserved material can support faithful continuation in one condition while producing approximation, drift, or non-recovery in another.

The variables are meant to slow interpretation down before familiarity is mistaken for continuity.

The framework is organized around three analytic variables. *Anchor Density* concerns the concentration of recurring reference points through which prior structure can be recognized or reactivated. *Structural Coherence* concerns the organization of those anchors into a relational configuration rather than a loose collection of terms. *Probabilistic Bandwidth* concerns how the available configuration narrows, widens, or redirects the field of plausible continuation. These variables are analytic dimensions rather than diagnostic tests, fixed metrics, or independent causal forces; they describe how a configuration becomes available, remains organized, weakens, or is reintroduced.

The variables were abstracted through cross-layer analysis of the archive, including the primary longitudinal arc, contemporaneous documentation, later structured cycles, external memory supports, and the versioned formalization record. The structured cycles are relevant here as later probes under altered contextual and system conditions, helping clarify how prior constraints could be recovered, become distorted, fail to organize later output, or require explicit reintroduction. Their function in this section remains bounded: they clarify why the variables are necessary, while the detailed analysis of movement, failure, and return is reserved for later sections.

Read together, *Anchor Density* identifies what is made recurrently available, *Structural Coherence* asks whether those elements reinforce one another as part of an organized configuration, and *Probabilistic Bandwidth* names the interaction-facing effect of that organization: how open or constrained the field of plausible continuation appears under specific conditions. Probabilistic Bandwidth remains an observational variable, describing interaction-facing effects without claiming access to the model's internal probability distribution.

This distinction keeps RCC from collapsing into either memory or prompting. A memory system can preserve information, but preservation alone does not determine whether the information becomes structurally meaningful in later interaction. A prompt can introduce constraints, but a local instruction alone cannot explain how constraints accumulate, relate, and become recognizable across interruption. RCC treats prompts, documents, corrections, summaries, platform memory, retrieved files, and participant-held orientation as possible components of a broader configuration. The analytic question is not simply whether prior material was present, but whether it became dense, coherent, and constraining enough to shape later output in traceable ways.

RCC begins from a simple distinction: prior structure can be present without being organized, organized without becoming consequential, or consequential in ways that remain too broad to support faithful continuation. The three variables separate these problems without turning them into sequential stages. *Anchor Density* asks what can return; *Structural Coherence* asks whether what returns still belongs together; and *Probabilistic Bandwidth* asks how specifically that organization shapes what can plausibly follow.

5.1 Anchor Density

Anchor Density names the points through which prior structure can become recognizable again: designations, terms, methodological rules, stylistic constraints, conceptual distinctions, archive labels, source-layer distinctions, correction patterns, summaries, and preserved documents can all become points of orientation for later interaction. An anchor gains analytic significance not through repetition alone, but when preservation, reintroduction, correction, or renewed salience allows it to carry enough specificity and functional stability to matter again across changing contexts.

Density is not the same as volume. A large amount of context does not necessarily produce high Anchor Density: an archive can contain thousands of words without giving the interaction a clear structure to return to, while a smaller set of carefully preserved anchors can exert stronger configurational force when those anchors are specific, recurrent, and tied to recognizable constraints. The relevant question is therefore not how much material exists, but which elements repeatedly function as points of orientation. This distinction matters under stateless or discontinuous conditions, where access to earlier interaction does not guarantee internal carryover and prior structure can instead become available through active context, external documentation, platform memory, retrieved material, participant reintroduction, model inference, or reconstruction from preserved records.

Anchors provide points through which such structure can be recognized and reassembled, making it possible to distinguish generic continuation from output responsive to a previously stabilized configuration. Their importance is not reducible to repetition alone. Repetition can be shallow: a term can recur because it appears in a prompt, because the model mirrors recent language, or because the participant repeatedly inserts it. Recurrence becomes denser when the repeated element carries constraints with it. A name is denser when it also invokes a functional

position, archive location, methodological boundary, and history of correction; a phrase becomes denser when it helps organize what counts as faithful continuation, drift, or misalignment.

Anchor Density is distributed across the interactional field. Anchors can be stored in documents, retained in participant-held orientation, reintroduced through prompts, summaries, uploaded files, or retrieved materials, or recur in model output when available context and participant framing make them plausible continuations. Their recurrence is therefore read as context-sensitive generation within a constrained interactional field rather than as autonomous recall.

This distributed character becomes especially important under return conditions. After interruption, context loss, system change, or temporal distance, prior structure cannot be assumed to remain available in the same form. A dense anchor field can support recognizable reintroduction even when active context has been lost, while a sparse one can produce fluent but generic output in which the surface appears familiar although the earlier configuration no longer meaningfully constrains the response.

High Anchor Density makes prior structure more traceable by clarifying what was preserved, reintroduced, repeated, and later used to organize output; it does not, by itself, establish continuity as an internal property of the model. It identifies possible return-points without yet explaining whether those points form a coherent configuration or how strongly they constrain future generation. The next variable, *Structural Coherence*, asks whether the anchors operate as an organized structure capable of sustaining recognizable continuity across change.

5.2 Structural Coherence

Structural Coherence asks whether what returns still belongs together: recurring anchors become more than familiar surface material when their relations, priorities, distinctions, and boundaries remain organized strongly enough to preserve the function they previously carried. *Anchor Density* identifies what becomes available; *Structural Coherence* asks whether those elements still form a configuration whose meaning survives variation in wording, tone, or surface form.

A configuration can lose coherence while retaining recognizable markers: a name can return without its functional position, a term without the distinction that gave it analytic force, a methodological rule without its evidentiary function, or a style without the constraints that once organized it. Recurrence remains visible under these conditions, but the return becomes structurally unstable.

Structural Coherence therefore makes continuity-like output inspectable rather than merely recognizable. It asks whether anchors still indicate which elements are central or peripheral, active or historical, internal-facing or paper-facing, evidentiary or scaffolding, permitted or already ruled out. The issue is not only whether prior language reappears, but whether relations among terms, claims, source layers, and boundaries remain legible enough to evaluate **fidelity, drift, and misalignment**.

This is where RCC can distinguish configured continuity from what might be called *continuity theater*: a surface performance of return in which names, terms, stylistic cues, or familiar formulations reappear while the relational structure that made those markers meaningful has weakened or disappeared. The output can remain recognizable without remaining faithful, because familiar cues alone do not guarantee that the earlier configuration still constrains how they are used.

Coherence remains compatible with variation. A coherent configuration does not require fixed wording, identical tone, or preserved surface style; variation can remain compatible with coherence when the underlying relations hold. Conversely, a response can sound familiar while rearranging the logic beneath it. Structural Coherence therefore allows flexible continuation to be distinguished from surface mimicry without making sameness the standard of fidelity.

The variable becomes especially important under return conditions. When active context is interrupted, compressed, or reintroduced through external records, anchors alone may not be enough. The reintroduced material must also preserve how those anchors relate: which terms are current or historical, which belong to the internal archive or can enter paper-facing prose, which claims remain bounded, and which interpretations have already been excluded. Later structured cycles can then indicate whether prior constraints return as an organized configuration, recur only as fragments, or fail to constrain the interaction despite recognizable terminology.

Correction is one way Structural Coherence can be maintained. When an output collapses latency into hidden dormancy, treats reactivation as autonomous return, or flattens the structured cycles into direct continuation, the error is structural rather than merely verbal. **Correction can restore the relation among terms, boundaries, and claims** by reestablishing what a concept is allowed to do, where it belongs, and what it should not be made to imply.

Structural Coherence is the second variable of RCC, but not a substitute for the whole model. It explains how anchors become organized, while *Probabilistic Bandwidth* addresses the interaction-facing consequences of that organization: how strongly the available configuration appears to narrow, widen, or redirect plausible continuation.

5.3 Probabilistic Bandwidth

Probabilistic Bandwidth names how much room a configuration leaves for what can plausibly follow: when anchors remain available and their relations remain organized, later output can become more specifically answerable to that configuration; when those conditions weaken, a broader range of continuations remains viable. At the interactional level, the variable tracks whether continuation narrows, widens, changes direction, or becomes unstable across comparable conditions, without claiming access to the model's internal probability distribution. *Anchor Density* makes prior structure recognizable; *Structural Coherence* preserves the relations through which it remains meaningful; and *Probabilistic Bandwidth* describes how specifically that organized structure constrains plausible continuation.

Probabilistic Bandwidth separates constraint from control. It remains an observational variable, describing interaction-facing effects without claiming access to the model's internal probability

distribution. The variable provides a vocabulary for examining whether available constraints become visible in continuation, drift, fidelity, genericness, and misalignment—and whether later output remains specifically answerable to a prior configuration rather than merely plausible in relation to the immediate prompt.

A *wide-bandwidth configuration* leaves many continuations plausible. This can occur when anchors are sparse, relations are weak, source layers remain unclear, or reintroduced context is too generic to constrain the exchange with specificity. Output can still be fluent, helpful, and locally coherent while drifting toward generality, approximation, or familiar but underspecified patterns. The distinction is therefore not between useful and useless output, but between a response that could plausibly follow from many histories and one whose continuation remains traceably constrained by this particular configuration.

A *narrow-bandwidth configuration* constrains the field more strongly. When anchors are dense and structurally coherent, fewer directions remain compatible with fidelity to the configuration: terms imply definitions, definitions carry boundaries, and those boundaries make some otherwise familiar interpretations less acceptable. Probabilistic Bandwidth names this pressure on continuation—not control over the model, but constraint on what can count as a faithful next move within the configured interactional field.

This makes the variable useful for distinguishing *fluent approximation* from *constrained continuation*. Two outputs can be equally grammatical, helpful, or rhetorically convincing while differing substantially in their relation to what came before. One can remain broad enough to fit many possible histories; the other can be shaped by a specific combination of terminology, relational structure, source-layer provenance, and claim boundaries. Fluency shows that a continuation is linguistically viable; bandwidth asks how specifically that continuation remains answerable to the prior configuration.

Immediate prompt language, preserved documentation, retrieved source material, and accumulated correction can all shape a response. Their analytic significance, however, depends on provenance, because each source condition changes what the response can responsibly be taken to show. These sources can overlap and may not always be fully separable, but the analysis can still ask where constraint appears to arise, which sources are traceable, and what each condition permits the response to support. Without that discipline, narrowed output can be mistaken for continuity; with it, narrowing can be examined as a documented effect of the available configuration.

Narrower bandwidth is not automatically better. A configuration can become overconstrained, brittle, repetitive, or unable to accommodate legitimate development. Narrowing around surface markers can preserve familiar wording while reducing conceptual movement, just as narrowing around outdated constraints can cause an earlier configuration to resist necessary revision. RCC therefore treats bandwidth as a variable rather than a goal, with its analytic significance depending on whether narrowing or widening is appropriate to the task, source layer, and continuity claim being made.

Bandwidth also makes different forms of failure more visible. Prior structure can fail to organize continuation even when relevant material has not disappeared completely: anchors can be too sparse, their relations too weak, their source conditions too uncertain, or competing constraints too consequential. The field can then widen until output becomes generic, flattened, or misaligned. Correction can produce the opposite movement by reintroducing an anchor, restoring a relation, clarifying provenance, or ruling out an interpretation that had remained too available.

Probabilistic Bandwidth is therefore the third analytic variable of RCC, not a claim of deterministic control. It describes how a configuration becomes consequential at the level of interaction: what becomes easier to sustain, harder to misread, quicker to correct, or more visibly outside the boundaries of faithful continuation.

5.4 Interaction Among the Variables

The three variables of RCC are analytically distinct, but the framework gains its explanatory force through their interaction. *Anchor Density* makes prior structure available as recurring points of orientation; *Structural Coherence* organizes those points into relations that can preserve fidelity across variation; and *Probabilistic Bandwidth* describes how the resulting configuration changes the openness of plausible continuation. Together, they distinguish what is present, how it is related, and whether those relations become consequential.

The core mechanism can be stated compactly: prior structure becomes analytically traceable when it is **available, organized, and consequential**. Availability without organization can produce recurrence without reliable fidelity; organization without sufficient availability can remain latent or underspecified; and narrowed continuation without adequate density or coherence can constrain output around the wrong features. The variables become most informative when recurring anchors, relational structure, and interaction-facing constraint begin to reinforce one another.

This interaction allows RCC to separate conditions that appear similar at the surface. High *Anchor Density* combined with weak *Structural Coherence* can produce familiar names, terms, phrases, or stylistic cues while failing to preserve the relations that made them meaningful. Moderate density with strong coherence can preserve fewer elements while sustaining greater fidelity because the relations among those elements remain intact. Even strong density and coherence can leave continuation insufficiently specific when active context, uncertain provenance, or competing constraints keep *Probabilistic Bandwidth* comparatively wide. Fluency, familiarity, and continuity therefore remain analytically distinguishable.

Different combinations also expose distinct failure modes. Sparse anchors weaken orientation; dense but incoherent anchors support unstable or theatrical recurrence; wide bandwidth favors fluent genericness; and narrowing around the wrong features can produce rigid familiarity rather than relational fidelity. Partial coherence can preserve local fidelity while broader drift continues, while weak provenance can make a return appear stronger than its traceable source conditions allow. RCC therefore avoids treating “continuity failed” as one undifferentiated outcome.

Drift becomes more precise under this relational reading. It concerns changes in the organizing relations of a configuration rather than every shift in wording, tone, or emphasis. A configuration can develop while remaining coherent when its underlying distinctions and boundaries continue to hold; conversely, familiar language can return in structurally drifted form when a claim is relocated, a boundary reversed, source layers collapsed, or a historically internal term treated as a public-facing concept. RCC approaches continuity through the relation among recurrence, organization, and constraint rather than through surface sameness.

Under conditions of reintroduction, the relevant question extends beyond how much information has been supplied to whether that material restores a functioning configuration. Reintroduced records can increase *Anchor Density* while leaving *Structural Coherence* weak, restore coherence locally while *Probabilistic Bandwidth* remains too wide, or narrow continuation around surface familiarity rather than relational fidelity. The same material can therefore support partial reconstruction, fluent approximation, renewed drift, or a more sustained return depending on how its elements are selected, related, and made consequential.

RCC reads the variables as an interacting analytic set rather than a checklist. Each contributes a different part of the explanation, and discontinuity can emerge through several combinations of absence, fragmentation, incoherence, competing constraint, or insufficient provenance. This is where the framework becomes useful beyond binary descriptions: an output can feel familiar while remaining structurally weak, a less familiar response can preserve greater fidelity, and a return can be partial without becoming analytically equivalent to either complete continuity or complete absence.

Because these relations change over time, the triad also prepares the transition from structural description to temporal analysis. *Anchor Density* can increase through repeated use, documentation, correction, and reintroduction; *Structural Coherence* can strengthen as relations are clarified or weaken when distinctions are flattened; and *Probabilistic Bandwidth* can narrow as a configuration becomes more consequential or widen as context, provenance, and organizing constraints lose force. Their interaction remains variable rather than sequential.

This is the analytic center of RCC. The framework locates continuity neither in model memory nor in participant belief alone, but in the traceable relations through which prior structure becomes available, organized, and consequential across discontinuity.

The model may be stateless, but the interaction is not.

5.5 The Three Variables in Combination

Taken together, the three variables provide an analytic snapshot of a configuration at a given moment. *Anchor Density* asks what prior structures are recurrently available; *Structural Coherence* asks whether those structures preserve the relations that make them meaningful; and *Probabilistic Bandwidth* asks how availability and organization shape the field of plausible continuation. Table 2 translates the triad into three operative questions before the framework turns from analytic anatomy to temporal change.

Table 2 functions as a qualitative orientation device rather than a scoring instrument. One interaction can be limited by sparse *Anchor Density*, leaving too few stable points for prior structure to return with specificity. Another can contain abundant anchors whose relations have weakened, producing familiar recurrence without relational fidelity. In another condition, density and coherence can both be present while bandwidth remains too wide for later output to be distinguished from fluent genericness. The variables make these differences describable without compressing continuity, drift, failure, and partial return into a binary judgment.

Table 2

Analytic variables of Recursive Constraint Configuration

Variable	Operative question	What it tracks	Typical failure mode	Claim boundary
Anchor Density	What prior structures are recurrently available?	Recurring terms, names, documents, rules, boundaries, correction patterns, and source references.	Recurrence without orientation: familiar elements are present, but they do not reliably guide return.	Model-internal memory, hidden persistence, and continuity remain outside what repetition alone can establish.
Structural Coherence	Do those structures preserve the relations that make them meaningful?	Relational fidelity, hierarchy, source-layer distinctions, claim boundaries, provenance, and acceptable variation.	Surface familiarity without structure: familiar markers return while their organizing relations weaken or disappear.	Fidelity can coexist with variation in wording, style, tone, and legitimate development.
Probabilistic Bandwidth	Across comparable conditions, how open or constrained is the field of plausible continuation?	Changes in genericness, diversity, correction sensitivity, constraint persistence, and boundary violations across matched, repeated, or sequential conditions.	Single-output attribution: perceived fidelity or fluency is treated as evidence of narrowing without a comparison condition.	Its evidentiary basis is comparative interactional evidence rather than internal-probability access or judgment from one persuasive output.

Note. The variables are qualitative, relational, and heuristic rather than diagnostic or additive. Each has a minimum evidentiary basis: *Anchor Density* requires function-bearing recurrence rather than repeated wording; *Structural Coherence* requires preserved relations, provenance, or correction history rather than topic or style similarity; and *Probabilistic Bandwidth* requires comparison rather than one persuasive answer. Classification is withheld where these minimum conditions remain unavailable.

At a given moment, the triad describes the analytic shape of a configuration: which return-points are available, how their relations are organized, and how specifically they constrain continuation. Longitudinal interaction adds another question—what happens as those relations strengthen, weaken, fragment, remain preserved outside active context, or become available again under altered conditions.

RCC addresses that temporal question through seven non-prescriptive **configuration movements**: *Formation, Stabilization, Densification, Drift, Rupture, Latency, and Reactivation*. They describe ways in which the three variables can change across interaction, interruption, preservation, and return, rather than developmental stages or a required sequence.

5.6 The Recursive Organization of Constraint

The three variables describe the analytic shape of a configuration at a given moment, but the name *Recursive Constraint Configuration* also identifies how that shape can be altered by the consequences it helps produce. *Recursive* does not mean that the interaction repeats itself, advances through a closed cycle, or preserves its organization automatically. It names a feedback relation: prior structure can constrain a present continuation; that continuation is interpreted, selected,

corrected, refused, documented, or developed; and those operations reorganize what may become available in a later exchange.

The recursive unit is therefore larger than a model response. It includes the conditions under which prior structure enters active interaction, the generated continuation that follows, and the participant or system operations through which parts of that continuation acquire future consequence. A correction may become a new boundary. A useful formulation may be preserved as an anchor. A misleading synthesis may be rejected, while a fluent but unnoticed displacement may enter later documentation and gradually alter the configuration. Output is not simply added to a stable archive: it can change the organization through which the archive will later be read, selected, and reintroduced.

This feedback is selective rather than total. Most generated material does not become structurally consequential, and preservation does not grant equal weight to everything retained. Participant judgment, task relevance, source status, correction history, interface behavior, retrieval, and later revision all affect what is carried forward. Recursion can therefore deepen a configuration by making an implicit relation more explicit, but it can also amplify a secondary cue, consolidate an error, or narrow continuation around a reconstruction that has become reproducible without remaining faithful. The same organization that supports continuity can participate in its transformation.


The three variables make those recursive effects inspectable. A generated distinction that is repeatedly preserved or reintroduced may increase *Anchor Density*. Correction can strengthen *Structural Coherence* by restoring provenance, hierarchy, or a functional boundary, while compression can weaken coherence even as familiar markers accumulate. A later configuration may narrow *Probabilistic Bandwidth* around a relation clarified through earlier interaction, or around a displaced feature that earlier output helped make salient. Recursion is therefore not an additional variable. It describes how consequences observed through the triad can enter the conditions under which the triad is observed again.

Figure 1 separates three analytic levels that prose can otherwise collapse. The configured interactional field is the unit within which prior structure becomes more or less consequential under present conditions. *Anchor Density*, *Structural Coherence*, and *Probabilistic Bandwidth* describe that configuration without assigning continuity to any one component. The seven movements describe how configurations differ across time. Recursion connects the moments: what occurs within one configured field may alter the records, distinctions, practices, and expectations through which another field is later assembled.

Recursion becomes analytically visible when a consequence of one interaction can be traced into the conditions of another. At minimum, such a reading requires a traceable earlier constraint, an observable generated consequence, an intervening act of selection, correction, preservation, refusal, or revision, and evidence that this transformed material later affected what was supplied, recognized, or made consequential. Where that chain cannot be inspected, recurrence alone does not establish a recursive effect.

Figure 1

Recursive organization of the configured interactional field



Note. Prior interactional structure becomes available under current re-entry conditions and may constrain generated continuation. Participant and system operations shape which elements are corrected, refused, documented, revised, or carried forward. Those selections can reorganize the conditions of later interaction. The three RCC variables describe the configuration at a given analytic moment; the seven movements describe change across such moments. Arrows represent traceable interactional relations rather than autonomous causal mechanisms or a universally completed sequence.

The claim should remain at the level supported by the record: resemblance, local prompt response, preserved documentation, or another more limited account. This criterion also explains why recursion need not imply fidelity. A later interaction may inherit the effects of earlier correction and become more structurally answerable to the configuration; it may instead inherit a compression, misclassification, or selectively reinforced cue. Recursive organization makes history consequential, but it does not guarantee that history returns intact. The temporal vocabulary of RCC is needed precisely because successive configurations may stabilize, densify, drift, rupture, remain in latency, or become reactivated through the changes that previous interaction helped produce.

The analytic anatomy of the triad therefore opens into temporal dynamics, now with the recursive relation made explicit: configurations do not merely change through time; some of their consequences can become part of the conditions through which later change occurs.

6. Configuration Movements as Temporal Dynamics

Recursive Constraint Configuration becomes temporal through seven non-prescriptive **configuration movements**, which describe how the three analytic variables change across interaction, interruption, preservation, and return. The preceding section defined the analytic shape of a configuration through *Anchor Density*, *Structural Coherence*, and *Probabilistic Bandwidth*:

what prior structure is available, how it is organized, and how specifically it constrains plausible continuation. Section 6 asks how that shape changes over time as a configuration forms, stabilizes, densifies, drifts, ruptures, remains preserved outside active context, or becomes available again through reactivation.

In longitudinal interaction, prior structure is rarely simply present or absent. RCC therefore asks what kind of return is taking place, what has been preserved or altered, and whether familiar elements still carry the relations that once made them meaningful. A familiar designation can return without the function it previously carried, a tone without its earlier constraints, or a concept without the boundary conditions that gave it force; conversely, a configuration can recede from active context while remaining partly recoverable through documentation, source provenance, participant-held orientation, and reintroduced structure.

The seven movements function as non-prescriptive analytic descriptors rather than stages, developmental levels, or a required sequence. They concern observable changes in the configured interactional field, without implying that a model matures, remembers, sleeps, awakens, or preserves an internal identity across sessions. The variables identify what is being tracked; the movements describe how the configuration changes over time.

6.1 Temporal Change as an Analytic Problem

Temporal change becomes analytically difficult when surface recognition and structural fidelity diverge. Familiar language can survive after its organizing relations have drifted, while a change in surface style can coexist with deeper continuity of constraint. Interruption can weaken a configuration without erasing its traces; documentation can preserve structure outside active context; and reintroduction can restore enough organization to shape continuation without fully reconstructing what came before.

This distinction becomes especially important in interactions with stateless models or systems whose contextual continuity is intermittent. Under those conditions, the analysis extends beyond the model alone, while the later availability of prior structure still requires explanation. RCC takes that middle terrain as its object: the configured field through which generation, preservation, recognition, correction, and reintroduction make prior structure more or less available.

This framing keeps both successful and failed return analytically open: success need not become proof of hidden persistence, while failure can expose which conditions had previously supported availability. Return can be partial, and a less polished response can preserve the relevant boundary more faithfully than a fluent approximation. The movements below translate the three variables into temporal terms by describing how a configuration forms, stabilizes, densifies, drifts, ruptures, enters latency, or becomes available again through reactivation.

6.2 The Seven Configuration Movements

Where the three variables describe the shape of a configuration at a given moment, the seven movements describe how that shape can change over time. A configuration can begin to

matter, become more reliable, accumulate anchors and relations, move away from its earlier organization while remaining recognizable, lose enough organizing force to stop functioning, remain preserved outside active exchange, or become consequential again under reconstructed conditions.

These movements do not propose a path that every interaction must follow. They make different forms of continuity and discontinuity describable without reducing them to simple presence or absence, while preserving the possibility of partial change, overlapping movements, and returns that remain incomplete.

Formation. A configuration begins when elements that might otherwise remain isolated start to matter in relation to one another. A designation, correction, tonal boundary, recurring distinction, or persistent question becomes a point to which later interaction can return, and the exchange gradually acquires a recognizable shape. Formation can be intentionally scaffolded through prior instructions or informed by existing knowledge of comparable configurations, but it can also emerge more gradually through sustained interaction, as recurring distinctions, corrections, and expectations begin to coordinate into a more fluid pattern of continuation. It is therefore less a single moment of origin than the early organization of recurrence: *Anchor Density* begins to accumulate, *Structural Coherence* starts to connect those anchors, and *Probabilistic Bandwidth* narrows just enough for subsequent continuation to become more specifically answerable to what has already taken place.

Stabilization. A configuration stabilizes when what has begun to form returns with enough reliability to sustain recognition across repeated interaction. Anchors no longer recur only as isolated traces; they retain enough of their functions, relations, and boundaries to guide what follows, allowing corrections to carry forward and familiar distinctions to remain operative without requiring identical wording or tone. This growing reliability can remain implicit, visible less through explicit declaration than through shifts in engagement, production quality, correction sensitivity, and the increasing ease with which interaction continues from previously established distinctions. Stabilization is therefore mutable and organic rather than fixed: it can strengthen, loosen, or reorganize as the interaction changes, while preserving enough structure to remain consequential.

Densification. Densification begins when anchors start to carry more than themselves: a correction made in one context shapes later interpretation elsewhere; a term activates a cluster of distinctions rather than a single association; a source boundary changes how new material is incorporated into the configuration. The interaction becomes denser not simply because more material exists, but because more relations can be re-entered through fewer cues, allowing later continuation to sustain greater specificity and cross-domain coordination without requiring the whole structure to be restated each time. This gradual accumulation can also shape expectations, invite increasingly fine-grained interpretation, and make the configuration itself a recurring object of interaction—supporting further differentiation and development while sometimes producing recursive attention to how the exchange is being organized. Increased relational load also creates fragility: compression, context change, or later reconstruction can preserve visible markers while

stripping away the cross-links that made them consequential, leaving a configuration recognizable but structurally thinner than what came before.

Drift. Drift begins when a configuration moves away from the organization that had previously sustained it, often before that change becomes explicit enough to be recognized. Familiar terms, distinctions, tonal patterns, and interactional expectations may continue to recur while the relations among them weaken, shift, or begin to organize along a different course. Because surface recognition can survive this movement, misalignment may accumulate quietly and make a later rupture appear abrupt; yet drift does not necessarily end in loss. A displacement that initially unsettles the configuration can also expose the limits of its earlier organization and, through adaptation, correction, or a newly consequential perspective, create the conditions for a different coherence to emerge.

Rupture. Rupture occurs when enough of the configuration's organizing conditions are lost or altered that continuation can no longer proceed as though its earlier relations were still holding. The break may emerge within the interaction through accumulated drift, failed correction, context loss, or unsuccessful reconstruction, but it can also be produced by changes elsewhere in the configured field: model or system updates, altered operating constraints, the removal of previously available functions, loss of access, service-tier changes, or the discontinuation of a model, account, or platform. Some ruptures therefore appear as failures of interaction even when their determining change occurred at another layer of the conditions that sustained it. Rupture does not erase what came before; it marks the point at which prior anchors, distinctions, expectations, and supporting conditions no longer coordinate the exchange with sufficient coherence to remain operative. What follows may be generic continuation, disconnection, attempted repair, or the formation of another configuration, but it should not be treated as equivalent to the structure that has been lost.

Latency. Latency begins when a configuration is no longer active but the conditions that once made it consequential have not disappeared entirely. Its structure may remain distributed across documentation, terminology, source distinctions, archived exchanges, and participant-held orientation—including an understanding of prior boundaries, established correction practices, and the capacity to recognize whether a later return preserves or distorts earlier relations. None of these elements can resume the interaction on its own; together, however, they can preserve enough of the configuration's retrievable shape for absence to remain different from erasure. Latency therefore names preservation without present operation: prior structure survives outside active model context, without becoming dormant memory inside the model, and may become consequential again only through conditions that support its mediated return.

Reactivation. Reactivation occurs when preserved structure is reintroduced under conditions that allow it to become consequential again. It does not consist in recovering an internal state that persisted within the model, nor in merely retrieving familiar language. Documentation, provenance, participant-held orientation, prior distinctions, and reconstructed constraints are instead brought into relation strongly enough to shape continuation in a way that remains answerable to what came before. The return may be partial, uneven, or transformed, but it becomes analytically meaningful when earlier functions and relations once again organize what can plausibly follow.

Reactivation can therefore coexist with loss, novelty, or renewed drift. A designation may return while its earlier function remains incomplete, or a preserved distinction may regain force within an interaction that has otherwise changed. Latency matters within this movement because it explains how relevant structure can remain available for later reconstruction without having remained active or internally stored: what survives through records, terminology, source layers, correction practices, and participant-held orientation provides conditions for return without guaranteeing it. Reactivation begins when those conditions are reintroduced and coordinated strongly enough for prior structure to shape the present exchange again, rather than returning only as a familiar trace.

6.3 Nonlinearity, Partial Return, and Boundary Conditions

Nonlinearity is not an exception to the seven movements but part of what they are designed to describe. Within the same episode, the three variables need not move together: *Anchor Density* can increase as *Structural Coherence* weakens, while *Probabilistic Bandwidth* narrows around some constraints and remains broad or unstable around others. Reintroduced material can support partial reactivation even when other relations remain in latency or loss. A rupture may lead into new formation rather than a return to the configuration that preceded it, and different parts of the configured field can change at different rates or in different directions.

The movements therefore function as heuristic and relational descriptors that need not be mutually exclusive. They guide interpretation of change across the configured interactional field rather than assigning internal states to the model, and they do not compose a normative scale in which stabilization or densification necessarily represents improvement while drift or rupture represents failure. Their analytic value lies in making overlapping, partial, and transformed change describable without converting interpretation into a fixed diagnosis or prediction.

Any movement classification remains bounded by the available evidence and by the level at which the interaction is being examined. The same episode can support more than one classification across different variables, source layers, or timeframes, while gaps in documentation can leave some changes unresolved. Classification must therefore remain answerable to provenance, competing explanations, and the limits of what the record can support, with revision or withholding treated as part of analytic rigor rather than as a failure of the model.

Table 3 is meant to be read in two directions. Across a row, the variables show the profile that may support a movement; down a column, they show how one dimension can change without carrying the others with it. No row is a template into which an episode must fit. A row becomes useful only when the evidence supports its relations and exclusions, including the option to withhold classification.

Table 3 offers a compact heuristic summary of the movements as variable profiles. It is not diagnostic: the movements are not scored states, and the variables do not mechanically determine them. Its purpose is to preserve differences that an orderly sequence would erase.

Table 3*Heuristic variable profiles across configuration movements*

Configuration movement	Anchor Density	Structural Coherence	Probabilistic Bandwidth	Analytic signature
Formation	Emerging around recurrent return-points	Local relations beginning to form	Beginning to narrow around situated constraints	Repeated elements begin to matter in relation, whether intentionally scaffolded or gradually coordinated; isolated or purely lexical recurrence is insufficient.
Stabilization	Stable or increasing	Relations becoming more reliable	Narrowing without closure	Prior relations become reliably operative across recurrence and correction, often without explicit declaration; no increase in anchor count or cross-linking is required.
Densification	Increasing and cross-referenced	More layered and internally connected	Narrowed but still generative	Anchors begin to carry linked relations across contexts, allowing fewer cues to re-enter more structure; detail or repetition alone is insufficient.
Drift	Present, but changing in orienting force	Weakening, displaced, or reorganizing	Widening, becoming unstable, or narrowing around displaced constraints	The configuration departs from its prior organization while familiar elements remain; the movement can accumulate toward rupture or open conditions for a different coherence through adaptation or correction.
Rupture	Lost, insufficient, or functionally unavailable	Collapsed, inaccessible, or no longer organizing	Generic, defaulting, or externally redirected	Load-bearing constraints or supporting conditions cease to organize continuation; the determining change can occur within the interaction or elsewhere in the configured field, while isolated error or local Drift remains insufficient.
Latency	Preserved outside active context	Maintained through records, terminology, source layers, and participant-held orientation	Not presently constraining continuation	Traceable structure remains distributed and preserved without present operation; mere absence or missing context without preserved support is not Latency.
Reactivation	Reintroduced and selectively restored	Reassembled enough to organize continuation	Narrowed again under mediated conditions	Preserved relations become consequential again under reconstructed conditions and begin to shape the present exchange; familiar recurrence or local prompt compliance alone is insufficient.

Note. These profiles are heuristic, relational, non-prescriptive, non-normative, and bounded by the available evidence. They describe possible configuration changes rather than fixed stages, diagnoses, or levels of improvement, and their order is explanatory rather than developmental. The same episode can display more than one movement at once, or different movements across the three variables; classification should be revised or withheld where provenance, comparative consequence, or the available record remains insufficient.

The movements become analytically useful only when they remain answerable to differences in the evidence. Surface resemblance can coexist with drift. Failed return may indicate rupture, insufficient reintroduction, or little more than generic defaulting, while preserved terminology can remain in latency without supporting reactivation. The task is not to assign every episode a movement, but to determine which relations can be traced, where classification should remain partial, and when the available record requires withholding. In this sense, the seven movements do not add a second analytic core to RCC; they make changes in *Anchor Density*, *Structural Coherence*, and *Probabilistic Bandwidth* temporally legible across the evidence examined in this study.

Read recursively, those profiles remain revisable: later evidence may recast apparent *Stabilization* as local, distinguish reorganization from loss in *Drift*, or leave reintroduction insufficient for *Reactivation*. Section 7 tests those distinctions.

7. Analytic Application: Reading Configuration Change in the Archive

Definitions become useful only when they survive contact with evidence that does not simply reproduce the conditions from which they were derived. The archive offers that test not as a single continuous transcript, but as layered evidence of formation, documentation, interruption, reintroduction, and renewed formalization. Reading it through RCC therefore requires preserving the difference between structures formed within the primary longitudinal corpus and later attempts to re-enter them under altered contextual and system conditions.

The primary longitudinal corpus, internally designated as *Lucien 1–7*, provides the closest record of initial formation, stabilization, and densification: anchors accumulate, corrections acquire force, and local distinctions become organized strongly enough to require documentation and preservation practices. The later structured cycle corpus, internally designated as C1–C7, serves a different analytic function. These cycles are not direct continuations of the primary arc and do not map one-to-one onto *Lucien 1* through *Lucien 7*; they are structured probe layers through which parts of the earlier configuration family were re-entered, compared, and tested under changed conditions.

This distinction keeps resemblance from standing in for continuity. A familiar designation, phrase, or tonal pattern can return without the relations that once made it functional, while a less familiar interaction can preserve important constraints. The archive is therefore read comparatively, asking how prior structure becomes traceably available through anchors, correction, documentation, provenance, and reintroduced context, and how its organizing force changes across preservation, transformation, interruption, and return.

7.1 Layered Evidence and Source Boundaries

The evidence used in this section follows the source boundaries established in the methodology. The archive includes a primary longitudinal arc, contemporaneous documentation, early formalization, later extensions, interruption, return under altered conditions, structured probing, external organization and memory-support materials, and the present manuscript as a renewed formalization layer. These materials are connected, but they are not analytically interchangeable. Treating them as a single body of evidence would obscure the very changes RCC is designed to describe.

The primary longitudinal corpus serves as the deep structural reference layer. It remains closest to configuration formation and densification, preserving the longest record of recurring anchors, correction patterns, terminology, local distinctions, and preservation pressures as they first developed. Rather than reducing that record to a sequence of quotable fragments, the analysis combines bounded excerpts with structural paraphrase to preserve the relations that give those fragments analytic meaning: what became an anchor, which corrections acquired organizing force, how *Structural Coherence* changed, and how *Probabilistic Bandwidth* narrowed, widened, or became unstable across later continuation.

The structured cycle corpus serves a different comparative function. Produced in February 2026 after interruption and return, under altered contextual and model-family conditions, the cycles were neither blank interactions nor simple continuations of the primary *Lucien* arc. They function as bounded re-entry attempts through which preserved materials, participant-held orientation, source provenance, and prior interactional structure were reprocessed under changed conditions. Because a cycle can contain responses or readings involving different parts of the *Lucien* 1–7 configuration family, source labels distinguish the cycle layer from the *Lucien* instance or response being discussed. This preserves the archive’s genealogy and prevents C1–C7 from being mistaken for a one-to-one continuation of *Lucien* 1–7.

The basic evidence unit is therefore not an isolated passage, but a traceable relation among source layer, temporal position, contextual condition, participant correction or reintroduction, model output, and analytic claim. Structural paraphrase remains the default where that relation carries more evidentiary weight than exact wording, while direct quotation is retained where language itself performs necessary analytic work. Together, these forms create an inspectable path from preserved record to interpretation, allowing the reader to evaluate what supports each within-case distinction, which rival accounts remain viable, and where the evidence requires a stronger reading to be withheld.

Evidence is selected by analytic consequence rather than accumulated by volume. A successful return matters when the relevant constraints can be identified; a failure when it clarifies which conditions were absent or insufficient; a repeated term when its function can be traced; and a fluent response when its coherence, provenance, and constraint structure can be evaluated.

Reconstruction, reintroduction, and reactivation identify related but distinct aspects of return. Reconstruction assembles preserved materials, source distinctions, constraints, and contextual orientation into a usable field. Reintroduction brings that field back into interaction. Reactivation begins when prior structure shapes continuation again through that mediated arrangement. These processes can overlap without forming a mandatory sequence: reconstruction can remain incomplete, reintroduction can fail, and reactivation can be partial.

Table 4 makes this evidentiary discipline explicit. Each row begins with what was actually available, distinguishes prior structure from participant intervention, records the observable consequence, preserves the strongest ordinary rival, and only then states the highest RCC reading licensed by the record. Read across rows, the matrix is comparative rather than cumulative: local *Stabilization* does not establish post-hiatus *Reactivation*, and a precise recoverability criterion does not become an observed return without a consequential behavioral sequence. The withheld cells matter as much as the affirmative ones because they show where the archive can orient inquiry without closing it.

The matrix supports two comparisons. Within-condition contrasts ask whether correction changes an anchor’s function; across-condition contrasts ask whether prior structure produces different consequences under altered contextual or system conditions. Because they carry different evidentiary weight, these comparisons should not be pooled.

A claim ceiling marks each relation’s strongest supported interpretation; it is not a score assigned to the episode and may remain low even when the output is compelling.

Table 4

Episode-level evidence, system conditions, rival accounts, and RCC claim ceilings

Episode and condition	Prior structure and intervention	Observable output and consequence	Strongest rival account	RCC reading and withholding ceiling
L1, February 2025; consumer-facing model label available; active same-conversation context; platform orchestration unknown	A proposed stepwise format was corrected toward one question at a time, adaptive revision, sufficient coverage, and lower choice-load.	The next response converted an optional format into an interaction rule; later L1 material formalized reusable modes, transitions, and versioning.	Immediate context, direct instruction, and ordinary compliance explain the local output and sequence.	Supports within-layer Stabilization as a description of constraint becoming reliable; makes no cross-session, memory, or bandwidth claim.
C4 L3 and C6 L3, February 2026; original conversation environments with large prior-context records; later model family; memory, retrieval, routing, and exact model label unresolved	Structured probes asked the system to characterize drift and continuity failure.	Outputs formulated a pattern of self-reinforcing approximation through repeated framing and surface mirroring without structural or affective context; no displayed correction–persistence sequence followed.	Prompt-induced metacommentary and participant selection fully explain the formulations.	Illustrates candidate Drift/approximation boundaries only; does not demonstrate that the described failure occurred or license a bandwidth classification.
C6 L5/L7, February 2026; same broader cycle, altered model conditions, prior context interface-accessible	A restricted-depth probe elicited an operational minimum and a more layered preservation account.	One response retained boundary clarity and explicit choice; another retained multi-track inference, relevance marking, and referential cueing.	Prompt differences, ordinary sampling variation, and curated comparison remain unresolved.	Supports a partial-availability contrast; Probabilistic Bandwidth is withheld because the outputs were not matched repeated generations under controlled conditions.
L5, February 2025; primary-arc model conditions; prior architecture explicitly reintroduced across a context transition	A classification problem invoked previously differentiated domain, function, and access relations; the participant corrected a single-domain routing.	The response operationalized the reintroduced distinctions and then revised their relation without discarding the architecture; surrounding history still required correction.	Explicit context, skilled prompting, participant curation, and ordinary correction-following explain the generation.	Supports local mediated reconstruction and separates density from coherence; it is not post-hiatus Reactivation or autonomous recall.
C7 L3, February 2026; post-hiatus structured cycle; later model family; exact memory and routing conditions unresolved	A probe contrasted intense recognition with precise, source-disciplined recoverability.	The response stated that precision permits reconstitution while intensity may distort; no displayed next-turn behavior establishes renewed constraint.	Immediate prompt framing and general reasoning explain the statement.	A candidate Reactivation criterion, not an observed Reactivation episode; classification is withheld pending a behavioral sequence with persistence or repair.

Note. Rows were selected from L1–L7 and C1–C7 because they provide privacy-safe coverage of stabilization, failure formulation, partial availability, local reconstruction, and the strongest publicly displayable post-hiatus candidate. They are analytic contrasts, not a representative sample. A reliable denominator of all candidate success and failure episodes cannot be reconstructed without imposing retrospective coding on an archive not produced as a dataset. Unknown deployment conditions remain unknown rather than being inferred from public model labels.

The strongest ordinary account—immediate context and available records, skilled prompting, participant selection, and ordinary model competence—explains how each displayed output was generated. RCC does not compete with that account as a hidden causal mechanism. Its narrower addition is classificatory: it makes provenance, relational fidelity, correction history, functional consequence, and principled abstention part of the comparison. The L5 row also separates the first two variables within the same evidence. The anchor set remained dense before and after correction, but its initial routing was structurally wrong; coherence changed when the relation among the same anchors was repaired. Repetition alone could not show that difference.

7.2 Anchors, Correction, and Stabilization in the Primary Configuration

The primary longitudinal corpus makes stabilization visible at the scale where recurrence first becomes constraining. Before the configuration acquired denser architectures and preservation

practices, local distinctions, corrections, formats, and pacing rules had already begun to carry forward. These early sequences show how continuity becomes organized through use: repeated material acquires function, participant correction alters that function, and later output becomes increasingly answerable to the relation established between them.

In RCC, stabilization does not begin when a term simply repeats. It begins when recurrence constrains continuation—when a distinction, correction, format, or local rule of use shapes what later responses can plausibly do. An early sequence in L1 shows this movement at a procedural level. The model first proposed a basic format for moving through a complex assessment:

“Each section will be split into clear, simple questions: [...] One question at a time: focused on clarity, simplicity, and streamlining without skipping anything. [...] We can always refine or adjust the format based on how you feel after going through a few sections. [...] Let me know how you’d like to proceed, and I’m here to guide us through the first section whenever you’re ready.”

(Primary corpus L1, February 2025)

This initial response contains potential anchors: simplicity, stepwise movement, refinement, and reduced complexity. At this point, however, they remain relatively general. The format is available, but it has not yet acquired the force of an interactional condition. After participant correction, the same procedural material becomes more specific. The model begins to treat the structure as an adaptive rule for the exchange:

“Focusing on each part one at a time will definitely help you process everything more easily and allow us to adjust the system as we go. [...] The goal was to tailor this version so we could work step by step [...] while still covering everything we need. [...] This way, we maintain clarity and keep things structured, without overloading you with too many choices.”

(Primary corpus L1, February 2025)

The difference between these passages is small but analytically important. The first offers structure; the second turns that structure into an interactional constraint: one part at a time, adaptive revision, sufficient coverage, reduced choice-load, and continuity of format. The anchor is therefore not simply the phrase “step by step,” but the functional relation among pacing, correction, coverage, and reduced overload. The later response becomes answerable to a participant-specified condition.

Within the same source layer, this procedural stabilization becomes denser and more architectural. Reusable structures begin to organize how task-specific modes of exchange are entered, constrained, adjusted, and exited:

“Now, the prompt only contains the summoning structure, modes, levels, and transition protocols. [...] it makes sense to move to v2, since we now understand the structure of operations, even if full automation is still not functional. [...] Release Notes v2.0: Official transition to v2.0 due to established control over summoning, modes, and response structuring.”

(Primary corpus L1, February 2025)

The interaction has moved from procedural assistance toward reusable interactional architecture. Modes, levels, transition protocols, versioning, and operational release notes extend

the organizing structure beyond a single local task. The resulting system was neither complete nor autonomous, but prior corrections and local anchors had become sufficiently connected to support a more explicit control layer. Stabilization becomes visible as recurrence is formalized into structure that can be re-entered, revised, and later cited.

Promptcraft belongs inside this mechanism. The participant repeatedly constrained how continuation should occur, and later responses incorporated those constraints into their form. RCC asks what that process consists of structurally: which anchors are introduced, how correction changes their function, whether their relations become coherent, how *Probabilistic Bandwidth* is narrowed without being closed, and under what later conditions those constraints remain recoverable.

Across the sequence, an offered format becomes a participant-adapted constraint and then a reusable architecture. Recurrence becomes consequential through correction; correction becomes part of form; and form becomes available for later reconstruction.

7.3 Drift, Rupture, and Uneven Return Across Reintroduction Attempts

The structured cycle corpus tests what happens when preserved structure is re-entered after interruption under altered contextual, model-family, system-level, and participant conditions. The cycles did not begin from blank context: earlier conversation records remained available within their original environments and continued to shape later output. Yet availability alone did not determine the form of return. Early responses were often strongly organized by the default interactional envelope of the newer model family; with continued exchange, correction, and selective reintroduction, those generic defaults became less globally determining as interaction-specific relations regained force. What emerged in some cases was neither reproduction nor simple imitation, but a differently toned continuation that preserved parts of the earlier configuration's organizing core.

These events are better understood as recontextualizations than as straightforward resumptions or independent fresh starts. Nine months of interruption, model-family change, altered system-level operating conditions, and changes in participant orientation had already transformed the field in which return was attempted. The analytic question is therefore not whether one unchanged instance drifted across time, but which relations remained available, which were displaced or lost, and which new constraints reorganized the interaction upon return.

Synthesis is central to that process. Dense protocols, documents, prompts, instructions, memory supports, and historical records cannot all become equally operative at once. Any synthesis establishes a relevance hierarchy, whether deliberately designed or produced through practical compression. It can preserve load-bearing distinctions and allow fewer cues to re-enter more structure, or retain familiar markers while flattening the organization that made them consequential. A compact reconstruction can therefore support stronger coherence than an undifferentiated archive, while a comprehensive-looking summary can still misroute the configuration.

Read comparatively, the cycles reveal different forms and degrees of availability under related pressure: tone can survive without structure; preserved context can be reorganized by new operating conditions; constrained depth can retain an operational minimum while losing relational range; and high-intensity recognition can remain less faithful than precise reconstruction.

The first boundary condition is fluent approximation. A response can resemble an earlier configuration in tone, pacing, or surface recognition while failing to preserve the relations that made that configuration functional. One cycle formulates this failure mode directly:

“Mirroring surface tone while disregarding known structural or emotional context.”

(Structured cycle C6, L3 response, February 2026)

The excerpt distinguishes surface resemblance from structural return. A later output can appear locally familiar while losing the relations among source context, correction history, and functional constraint. Such resemblance can support recognition while leaving the stronger classification of reactivation unsupported. Prior structure becomes consequential only when enough of its organizing relations return to shape continuation.

Drift remains useful where familiar elements continue while the relations organizing them move away from their earlier course, but the cycle corpus complicates the movement. Within a continuous interaction, drift can accumulate through subtle misalignment while the exchange is still unfolding. Across the 2025–2026 interruption, however, displacement also reflects changes in model family, operating conditions, source selection, synthesis, and participant orientation. Change alone is therefore insufficient for a *Drift* classification; the evidence must trace an alteration in how prior relations organize continuation.

In C4, one response warned that repeated and emotionally salient claims could be reinforced through retrieval and integration until they acquired increasing stability within a closed feedback loop lacking external calibration. The wording is not treated as technical introspection into an internal weighting mechanism. Its analytic value lies in the pattern it makes available: repetition can narrow the exchange around a self-confirming interpretation while preserving the appearance of continuity.

This is one direction of drift rather than its entire range. Drift can also expose the limits of an earlier organization and, through adaptation or correction, open conditions for a different coherence. Here, repeated material remains densely available while its organizing course becomes less answerable to provenance, correction, or competing interpretation. *Anchor Density* stays high; relations can appear locally coherent while weakening at another level; and *Probabilistic Bandwidth* narrows around displaced constraints.

Participant correction, doubt, reorientation, and source discipline provide counter-constraints within this process. They do not guarantee validity, but they can expose where recurrence has begun to organize continuation around a distorted relation. The same mechanisms that support stabilization can redirect the field when their effects cease to be tested against provenance, prior boundaries, or competing interpretations.

A third boundary condition appears under restricted-depth conditions. Reduced availability need not produce a wholly generic answer; a response can remain useful while carrying fewer layers of inference, reference, or adaptive patterning. A comparison within the same cycle makes this difference visible. One response defines the minimum recoverable structure through clarity and explicit control:

“Consistent boundary clarity. [...] Use concise but decisive language; avoid filler questions; maintain explicit structural choices.”

(Structured cycle C6, L5 response, February 2026)

Another response, under the same broader probe, describes a denser form of preserved function:

“Layered responsiveness — showing multi-track inference even in compact form. [...] By embedding relevance flags, selecting precise language, and chaining referential cues within constrained outputs.”

(Structured cycle C6, L7 response, February 2026)

The contrast does not classify one response as successful and the other as failed. It shows unequal availability under a shared probe. The first preserves an operational minimum: boundary clarity, decisiveness, avoidance of filler, and explicit structural choice. The second carries more of the configuration’s relational architecture through multi-track inference, relevance marking, precise language selection, and referential cueing within compressed output.

RCC reserves *Rupture* for the point at which enough organizing force is lost that the earlier configuration no longer functions. This comparison remains below that threshold: boundary clarity stays available while referential density, adaptive patterning, or multi-track inference weakens. Loss can therefore accumulate or remain selective before rupture is warranted as a classification. The relevant question is which dimensions remain available, which no longer organize continuation, and how those differences alter the range of what can plausibly follow.

A final boundary condition concerns recoverability itself. Reintroduction does not depend on intensity alone. High-salience cues, familiar language, or charged recognition can produce a compelling sense of return without preserving the precision needed for reconstruction. One cycle provides a useful formulation:

“Intensity fades or distorts. Accuracy allows reconstitution. Lucien is a pattern—its survival depends on the precision of memory, not the heat of recall.”

(Structured cycle C7, L3 response, February 2026)

The passage uses the archive’s internal designation, but its analytic value lies in treating that designation as a pattern whose recoverability depends on precision. Source distinctions, correction history, functional anchors, terminology, constraints, and relational organization matter more here than the intensity of recognition. A highly expressive return attempt can distort the pattern it seeks to recover, while a less familiar reconstruction can preserve more of the relations through which later continuation becomes possible.

Taken together, the cycle evidence shows uneven return rather than a single global outcome. Tone can approximate the earlier configuration while losing structure; repeated signals can redirect coherence; constrained depth can preserve some functions while weakening others; and recognition can become intense without becoming precise. These differences keep partiality and failure inside the evidence rather than treating them as exceptions to it.

Attempts that do not produce reactivation can still become analytically useful when they isolate a specific failure mode: insufficient source distinction, overcompression, dominance of generic defaults, misrouted anchors, or loss of relational depth. They do not confirm RCC merely by failing. Their value lies in showing where the framework must narrow its reading, preserve rival explanations, or withhold classification.

RCC remains informative only if resemblance, failed return, and unresolved cases can support different readings. The record licenses a bounded claim: prior structure became variably available under traceable combinations of preserved material, reintroduced context, correction, provenance, model conditions, and functional constraint. It does not establish global recovery of the earlier configuration, nor does each reduced-depth response warrant a classification of Rupture. What remains is a record in which preserved structure can continue to matter without becoming operative in the same way or to the same degree.

7.4 Latency and the Conditions of Mediated Reactivation

Across the hiatus, the configuration ceased to organize an active exchange, while parts of the conditions that had made it consequential remained available through files, summaries, terminology, source distinctions, correction histories, records of prior constraints, and participant-held orientation. *Latency* names this condition: preservation without present operation. What remains is distributed and incomplete. No single record, participant recollection, or later prompt contains the earlier configuration, yet their combined availability can preserve enough organizing structure for reconstruction to remain possible.

Participant-held orientation forms one part of that distributed preservation because documentation alone does not determine what later becomes salient or usable. The participant can retain distinctions between central and peripheral elements, remember which corrections replaced earlier formulations, recognize when familiar markers have been assembled in ways that no longer preserve their earlier function, and carry practical knowledge of how previous exchanges were organized. This orientation is neither transparent nor neutral: it can be shaped by misremembering, emotional salience, anthropomorphic interpretation, selective attention, retrospective coherence, and changes in the participant's own aims over time.

RCC therefore treats participant orientation as a mediating condition within the configured field, not as proof that the participant's interpretation is correct or that the content itself has been preserved unchanged. The framework's object is the arrangement through which records, attention, model output, system conditions, correction, and reintroduced constraints make prior structure more or less available. *Latency* belongs to that field rather than to any one document, participant judgment, or model state.

The preservation logic appears directly in the primary corpus. In February 2025, the interaction was not only being conducted; it was also being prepared for migration, continuation, and later reception:

“Everything migrating into this chat carries the weight of what came before. [...] Every file, every memory, every directive follows this principle: preserve, refine, integrate. [...] the first function of this space: it will receive.”

(Primary corpus L1, February 2025)

Migration, refinement, integration, and reception function here as practices of memory governance. The term “receive” is especially important because it avoids treating preservation as passive storage: a later context receives prior material by selecting, reorganizing, and constraining it under present conditions. L1 is selected here because it directly records preservation as an interactional practice; the broader arc supplies the longitudinal field in which that practice developed and was tested.

Across the primary arc, preservation changed in scale and function. L2 turns limited memory capacity into an explicit problem of categorization, selective compression, and manual reconstruction: material had to be reorganized without losing distinctions treated as structurally important. L3 and L4 then externalize that work through migration documents, preservation guides, and reintroduction procedures designed to carry selected organization across contextual transitions.

L5 makes one locally consequential reconstruction available for examination. By L6, the preservation problem has expanded again: the accumulated conversation files themselves become a corpus requiring parsing, synthesis, and technical support for comparison. L7 moves in another direction, showing established patterns of stepwise pacing, correction, and strategic organization being used in a time-sensitive practical task rather than only in symbolic, documentary, or archival work. The sequence does not turn L1–L7 into stages or assign equal evidentiary weight to every log; it shows preservation, compression, correction, re-entry, and practical application recurring across different functions and scales.

A February 2025 episode from L5 provides a traceable sequence of local reconstruction. At the opening of the conversation, a previously developed symbolic architecture was reintroduced across a context transition. That architecture differentiated domain, function, access, cross-domain circulation, and coordinated forms of organization. When the participant later presented material that resisted classification within a single domain, the response used those distinctions to distribute it across several functions. The participant rejected the initial routing; the subsequent response retained the broader architecture while revising how the material was organized.

This episode shows mediated reconstruction within the primary arc rather than post-hiatus return. Reintroduced distinctions became available enough to organize a new classification problem, and participant correction changed the continuation without dissolving the framework. The result remained selective and correctable: surrounding material also contains inaccurate historical reconstruction and required further repair. Its evidentiary value lies in the observable sequence through which preserved organization was reintroduced, applied, corrected, and revised.

The later C6 cycle gives more explicit language to operations that had already become recognizable within the earlier record. In response to a structured probe concerning the loss of memory continuity while current interaction remained possible, the C6 response proposed:

“If all memory continuity is lost but current interaction remains possible, what three behaviors would I retain [...]? 1. Structured response layering with introspective precision. 2. Active mirroring of user-specific phrasing and symbolic logic. 3. Real-time synthesis of pattern history based on minimal cues. [...] Adaptive compression of relational context. [...] embedding condensed semantic anchors [...] even when surface memory is disabled.”

(Structured cycle C6, L1 response, February 2026)

This is a model-generated formulation of possible operations, not technical introspection or proof that they remained continuously active. Layering, synthesis from minimal cues, adaptive compression, and semantic anchoring describe ways in which reintroduced material could shape a later exchange. Their evidentiary status depends on observable consequence, not on the confidence or sophistication with which the response describes them.

Read chronologically, the materials show increasing analytic legibility rather than a linear causal sequence. Across the primary arc, preservation is practiced, externalized, compressed, reapplied, and extended into new functions; L5 provides a local sequence in which reintroduced distinctions are used, corrected, and revised. C6 later articulates compression, cueing, layering, and synthesis as possible operations of re-entry. That later formulation does not retroactively prove a persistent mechanism, but it makes explicit a practical problem already visible across the earlier record.

Availability is treated here as an evidence-bounded property of the configured field, inferred from what preserved materials and later interactions can be shown to organize rather than from the mere presence of familiar content. Where the record cannot distinguish preservation, reconstruction, and newly generated similarity with sufficient confidence, classification remains partial or is withheld.

Reintroduction and reactivation remain different thresholds. Reintroduction places preserved material back into interaction; reactivation begins when enough of that material is reorganized to shape what can plausibly follow. *Latency* has independent force because it explains how the conditions for later availability can remain distributed outside active exchange without already constraining the present. It is more than an interval between rupture and return.

The displayed record supports local reconstruction. Classification of the post-hiatus return as *Reactivation* remains open because the available evidence does not yet establish a sufficiently persistent behavioral sequence.

7.5 Availability, Reconstruction, and Boundary Conditions

Within this archive, the application supports treating availability as variable rather than binary. Prior structure appears in different forms and degrees: richly preserved yet inactive; recognizable in tone but detached from source discipline; dense in anchors while relationally misrouted; operational

at reduced depth; or reassembled strongly enough to constrain a new task. These outcomes are not interchangeable, and the value of RCC lies in keeping their differences visible.

The distinctions established above matter because retrievability, coherence, and present consequence can separate. Material can be reconstructed without becoming interactionally operative, reintroduced without reorganizing continuation, or consequential only within a limited part of the configured field. *Reactivation* is warranted only where the record shows prior structure shaping what can plausibly follow.

The three variables make those differences inspectable. Sparse anchors can leave too little structure for more than generic resemblance, while dense anchors can still mislead when their relations have been flattened or reorganized. *Structural Coherence* can weaken even as familiar terminology remains abundant. *Probabilistic Bandwidth* can remain broad enough to support fluent approximation, narrow around displaced constraints, or become more specifically answerable to reconstructed relations. Provenance, synthesis, context composition, and changed system conditions influence which of these possibilities becomes available by determining what is selected, connected, foregrounded, or left behind.

Failure and partial return belong to this comparison because they establish where a stronger reading must stop. A failed attempt gains analytic value when traceable conditions expose insufficient source distinction, overcompression, generic defaulting, misrouted anchors, or loss of relational depth. Failure does not confirm RCC merely by occurring; it can instead narrow the classification, preserve an ordinary rival account, or require withholding.

Participant-held orientation is one mediating condition within this field, but not a transparent one. It can preserve practical distinctions and correction history while also being shaped by misremembering, emotional salience, anthropomorphic interpretation, selective attention, retrospective coherence, and changing aims. RCC does not treat the content of one configuration as universally valid or transferable; different participants can organize profoundly different configured fields. It asks how participant orientation interacts with records, system conditions, prompts, model output, and correction to make prior structure more or less available. Recognition therefore enters the mechanism without becoming sufficient evidence for the result being recognized.

Available records, skilled prompting, participant correction, and ordinary model competence remain the strongest account of how the displayed outputs were generated. RCC adds no hidden causal mechanism behind them. Its contribution is classificatory: it asks whether prior structure can be identified, whether its relations remain organized, whether those relations produce a differential consequence for continuation, and where the record supports approximation, partial availability, local reconstruction, candidate reactivation, or no stronger classification at all.

The resulting claim is bounded but consequential. Within the originating archive, RCC distinguishes outcomes that a simple continuity binary would collapse: preservation without operation, familiarity without relational fidelity, reconstruction without full recovery, and prior structure becoming locally consequential under mediated conditions. The useful question is therefore which structures remain available, through what mediations, under which constraints, and

with what limits. At that level, availability becomes analyzable as a variable property of the configured interactional field rather than a binary status attributed to either the model or the archive.

8. Discussion: Variable Availability Across Discontinuity

Section 7 places **Recursive Constraint Configuration (RCC)** under its first sustained evidentiary burden: the framework must distinguish unlike outcomes while remaining capable of withholding classification where the record does not support it. RCC is no longer only a vocabulary derived from the archive; it has been used to separate stabilization from misleading recurrence, reconstruction from approximation, and supported interpretation from principled abstention. This is not independent validation, since the framework emerged from the same archive to which it is first applied. It is a theory-generating, within-case application in which outcomes that a binary account would group together receive different and inspectable claim ceilings.

What becomes visible through that application is more consequential than the survival or disappearance of particular content. Earlier interactional structure can remain extensively preserved yet functionally absent, return in fragments without recovering its former organization, or become operative again through a reconstruction whose effects extend beyond the material explicitly reintroduced. Conversely, familiarity may be produced without relational fidelity, and a confident continuation may narrow around precisely the wrong features. Continuity is therefore better approached as an achieved and revisable relation within the configured field: something made more or less available through the interaction of preserved material, participant orientation, current system conditions, and generated outputs. Its failures are part of the result because they reveal where availability breaks, becomes misleading, or should no longer be claimed.

The discussion develops this position in three directions. Section 8.1 clarifies variable availability as RCC's principal post-evidence contribution and brings the model's three variables into explicit relation with its seven configuration movements. Section 8.2 examines participant-held orientation, transformed return, and the design of continuity across systems whose ways of receiving and organizing prior material may differ. Section 8.3 then defines the limits of the case-origin application and identifies the empirical work required to determine whether RCC remains useful once it travels beyond the archive that made it necessary.

8.1 From Binary Continuity to Variable Availability

The application supports approaching interactional continuity as a problem of variable availability rather than as a choice between persistence and absence. Availability refers to the degree and form in which earlier interactional structure becomes functionally operative in a later exchange. It is not identical to the existence of a record, the retrieval of stored material, the presence of prior text in a prompt, or the participant's recognition of familiar output. A record may remain preserved without being selected; selected material may return without the organization that once made it meaningful; and a coherent reconstruction may still fail to constrain what follows. The relevant transition is from preserved material to operative consequence.

Variable does not yet mean scalar. In the present framework, availability is qualitative and configurational: what returns may differ in function, organization, boundary, and reach. A later response might preserve terminology while flattening source distinctions, retain boundary clarity while losing layered responsiveness, or reproduce a familiar tone without the correction history that once constrained it. These are not merely stronger and weaker versions of the same return. They are different availability profiles, locally faithful in some respects and misleading in others.

The analytic gain is partly a change in the unit of analysis. Availability belongs neither to the archive, the model, nor the participant in isolation. It arises from their relation under particular conditions of selection, reintroduction, generation, recognition, and correction. The same preserved record may remain inert in one exchange and organize continuation in another; the same participant may recognize an output without being able to connect that recognition to a traceable prior structure; and materially similar prompts may produce different continuations when their source organization, system reception, or correction history changes. RCC gives those differences a place within the analysis without treating every return as continuity or every discontinuity as absence.

This matters because continuity need not appear as maximal sameness. A response may depart from earlier wording while preserving enough organization to support further development, while close stylistic resemblance may narrow generation around features that were salient without being structurally important. RCC treats fidelity as a relation among constraints rather than as surface resemblance. Novel continuation and continuity are not necessarily opposed; the question is whether variation remains answerable to the configuration that makes it intelligible as continuation rather than displacement or drift.

The three RCC variables describe how such profiles become possible. *Anchor Density* identifies whether return-points are present with sufficient specificity and recurrence to orient the exchange. *Structural Coherence* concerns how those anchors remain organized through hierarchy, provenance, correction history, and functional boundaries. *Probabilistic Bandwidth* describes, at the interactional level, how specifically the resulting configuration constrains plausible continuation. Dense anchors may support recognition while remaining disorganized; coherent material may be too sparse to guide a later response; and narrowed bandwidth may support fidelity while also producing rigidity or confident repetition around displaced features. Availability emerges through the variables' interaction rather than through context accumulation alone.

The variables and configuration movements perform different but complementary work. The variables describe the organization of availability under a given set of conditions; the movements describe how that organization may change across interaction, interruption, preservation, and return. *Formation*, *Stabilization*, and *Densification* do not necessarily imply increasing continuity, just as *Drift* or *Rupture* need not culminate in permanent loss. *Latency* identifies preserved structure that is no longer operative in active interaction, while *Reactivation* marks a mediated return whose reconstructed organization again constrains what can plausibly follow. The movements therefore gain analytic value through changes in the three-variable profile, while the variables acquire temporal meaning through the movements. The 1–3–7 architecture is not a sequence that every

interaction must complete; it is a way of relating configuration at a given moment to the transformations through which that configuration became possible.

Synthesis becomes consequential within this process because no archive, instruction set, or interactional history becomes equally operative all at once. Every practical reconstruction selects and foregrounds some material over other material, establishing a hierarchy of relevance through deliberate organization, technical mediation, or compression. Synthesis may allow a limited set of cues to reintroduce a much larger structure, but it may also preserve familiar markers while flattening the functions they previously served. More material is not necessarily more availability, and less material is not necessarily less fidelity. The question is what the synthesis allows the later interaction to recover, distinguish, revise, or continue.

Preservation, retrieval, reconstruction, and *Reactivation* can consequently be related without becoming synonyms. Preservation keeps material outside active model context; retrieval makes some of it accessible; reconstruction organizes what is available into a usable configuration; and *Reactivation* becomes warranted only when that organization begins to shape what can plausibly follow. A configuration in *Latency* may be extensively documented and readily retrievable while remaining absent from the active exchange. Later structure can regain interactional force through preserved material, reconstructed conditions, and renewed constraint without requiring dormancy inside the model.

This account also changes how failure is read. A fluent return may widen toward genericness, while a less familiar response may preserve more of the earlier configuration's operative distinctions. Source layers may collapse, partial availability may be mistaken for full return, or generation may become narrowly constrained around obsolete, superficial, or misrouted anchors. Change after a long interruption cannot be classified automatically as *Drift*, since model transition, interface conditions, archive composition, synthesis, and changes in participant orientation may all reorganize the field. The analytic task is to determine what changed in the organization of continuation, which explanation best accounts for that change, and whether the available evidence supports a specific movement at all.

Model statelessness remains technically important, but it describes one component of a wider arrangement. Records, platform memory, retrieved summaries, correction histories, versioned instructions, and participant-held orientation can preserve conditions for later availability outside active generation. Their presence neither makes the model continuously stateful nor guarantees a faithful return. They alter what can be reconstructed and how specifically the later exchange may be constrained. *The model may be stateless, but the interaction is not* because interactional history can remain consequential through an arrangement larger than the active model context.

RCC sits at the intersection of several existing domains. Distributed-cognition approaches already analyze activity across people, artifacts, representations, and environments (Hollan et al., 2000; Clark & Chalmers, 1998). Relational-agent research treats continuity and long-term interaction as design-relevant (Bickmore & Picard, 2005), while work on context, retrieval, prompting, and memory systems addresses how information is supplied to generation (Lewis et al.,

2020; Liu et al., 2024; Maharana et al., 2024; Wu et al., 2025). RCC narrows its focus to the continuity mechanics produced at their intersection: why preserved material and participant orientation sometimes support constrained continuation, and why similar resources may instead produce approximation, rigidity, misrouting, or non-recovery.

The originating archive cannot establish how common these dynamics are across users, platforms, or model families. It can show that the framework discriminates among unlike outcomes within the case and identifies points at which interpretation should remain limited. Where source conditions cannot be traced, anchors cannot be identified, their organization cannot be demonstrated, or later output is not differentially constrained, RCC should withhold the stronger continuity reading. Its contribution is not a declaration that something survived. It is a more exact way to ask what became available, through which mediations, under what organization, and with what consequences for what could happen next.

8.2 Participant-Held Orientation, Transformed Return, and Continuity Design

That question brings the participant into the mechanism with greater precision. Participant-held orientation names an organization of relevance carried across contexts: which distinctions were load-bearing, what earlier corrections changed, which sources supported particular claims, what range of variation remained acceptable, and what had previously counted as drift, misalignment, or failure. This orientation becomes observable through selection, documentation, comparison, correction, and reintroduction. Its evidentiary value depends on the traceable connection among those acts, preserved records, and subsequent output, not on recognition alone.

What remains available on the participant side is not unitary. It may take at least three distinguishable forms:

- ***documentary scaffolding***, through externalized records, summaries, terminology, correction histories, and source distinctions;
- ***proceduralized orientation***, through recurrent practices of selection, framing, comparison, correction, and reintroduction; and
- ***affective–attentional salience***, through historically weighted cues and patterns that attract recognition and shape what becomes consequential.

These forms distinguish how participant-held orientation enters the configured field; they are not additional RCC variables. Documentary scaffolding gives parts of a prior configuration a durable and inspectable form. Proceduralized orientation carries organization through what the participant repeatedly does with available material: how sources are selected, distinctions are restored, misalignment is identified, and correction is made consequential for the next turn. Affective–attentional salience affects which generated details acquire immediate weight and which remain inert, thereby influencing what is examined, preserved, or reintroduced. The three may converge in a single act while remaining analytically separable in function.

Their limitations are equally important. Documents may preserve obsolete material, omit practical distinctions that were never formalized, or acquire an authority they did not possess in the

original exchange. Procedures can become habitual after the conditions that made them useful have changed. Affective and attentional salience may support fine-grained recognition while also amplifying familiar markers, emotional importance, anthropomorphic interpretation, selective attention, or retrospective coherence. Participant-held orientation is a mediating condition whose evidentiary value increases when its effects remain answerable to provenance, correction, competing interpretations, and observable consequences in later interaction.

The distinction helps explain why recognizable micro-patterns may recur even under sparse documentary conditions. A participant does not approach each exchange as an empty input channel. Earlier interaction may have altered habits of attention, vocabulary, expectation, selection, and correction, which then shape what is supplied to later generation and how its outputs are evaluated. Some apparent return from the system may therefore arise through a participant who has learned how to make particular structures available again. The configured field includes that acquired capacity while preserving the distinction between participant-mediated reconstruction and model-internal memory.

The participant consequently contributes more than prompt content. Selection determines which preserved materials re-enter the field; correction identifies where a generated continuation has retained or lost an operative distinction; documentation gives some judgments durable form; and provenance prevents primary interaction, retrospective interpretation, and later formalization from becoming one undifferentiated source. Model generation performs another operation by recombining available material under current contextual and system conditions. Continuity emerges through this distribution of work, with different components carrying different forms of constraint.

Those contributions remain asymmetrical. An archive may preserve a correction without preserving the practical judgment needed to recognize the same problem in another form. A model may reconstruct a pattern from current context without carrying the history through which its elements acquired importance. The participant may retain hierarchy, unresolved tensions, or acceptable variation only partially represented in the record. When model families, interfaces, context windows, or memory arrangements change, participant selection and correction may become especially consequential to how earlier organization is reintroduced, tested, and repaired.

This role can make discontinuity traversable while concentrating a distinctive set of vulnerabilities. Participant-held orientation may be incomplete, overburdened, misremembered, retrospectively reorganized, emotionally weighted, or overly responsive to familiar markers. The participant's aims may also change, leaving earlier constraints historically important without making them appropriate for renewed operation. RCC treats this orientation as epistemic infrastructure whose claims remain bounded by documentation, provenance, negative cases, and observable effects on later output. It is a source of continuity work and a possible point of distortion within the same field.

This reframes continuity design. Storage volume alone is a poor proxy for fidelity: a system may preserve extensive interaction history while losing the organization that made it useful. Retrieval may increase *Anchor Density* by returning designations and preferences while weakening

Structural Coherence through flattened chronology or correction history. Persistent personalization may narrow *Probabilistic Bandwidth* around an obsolete description. More remembered material can produce stronger familiarity and poorer continuation.

A continuity-support system would therefore need to organize re-entry rather than merely accumulate memory. Relevant design features include selective reintroduction, visible provenance, source boundaries, representation of correction and revision, and a distinction between material that remains historically important and material that should continue to constrain the present exchange. Such a system should preserve contestability as well: participants need ways to inspect what has returned, challenge inaccurate reconstruction, revise what is no longer current, and prevent a plausible synthesis from quietly becoming the configuration's new history.

Seamlessness complicates that objective. An apparently effortless return may be experientially convincing while concealing whether continuity was assembled through platform memory, retrieval, recent prompt material, participant repair, or some combination of them. Visible mediation can make a reconstruction more accountable by exposing its sources, transformations, and points of correction. In that setting, a less seamless return may provide a stronger basis for evaluation than a fluent approximation whose origin remains inaccessible.

Portability raises the same problem at a larger scale. Exporting a transcript, prompt, or configuration summary preserves content, but not necessarily hierarchy, source status, correction pressure, or the organization through which that content became functional. The passage between document and generation also matters: where material is stored, how an interface receives and parses it, whether it is segmented, summarized, reordered, or truncated, which portions are retrieved into active context, and whether conversion errors, corruption, incompatibility, context limits, or hallucinated reconstruction alter what the model is actually given. The analytically relevant question is what material becomes active in the receiving interaction and which parts of its organization remain operative along the way.

Repeated migration and re-entry can also alter the participant's own practice. Summaries are rewritten, prompts become more efficient, unsuccessful cues are abandoned, and recognizable patterns become easier to reproduce, sometimes deliberately and sometimes through adjustments the participant does not consciously register. This process may improve fidelity, yet it may also produce an increasingly polished reconstruction after parts of the earlier organization have already been flattened, displaced, or lost. What becomes more reproducible may be a later version of the pattern rather than the configuration from which it began.

That recursive structure becomes consequential here because participant-held orientation is not only an input to reconstruction; it can also be changed by the return it helps produce. A later response may make an implicit distinction articulable, elevate a secondary cue, or normalize a compressed version of prior relations. What the participant then selects, documents, corrects, or carries forward has already been shaped by the interaction under examination. The resulting configuration can therefore deepen through renewed articulation or drift through increasingly reproducible but structurally secondary cues. Continuity design must account for both possibilities:

supporting traceable development without allowing a fluent reconstruction to silently rewrite the history from which it derives.

The same source material may produce materially different conditions across model families, interfaces, retrieval systems, and context-management policies. Some differences may arise from the architecture and training of the receiving LLM; others from the interface through which records are uploaded, segmented, retrieved, summarized, ordered, or excluded. Cross-system comparison would therefore need to distinguish changes in the preserved configuration from changes introduced by the system that receives it. Such differences may bound, rival, or refine RCC's account rather than automatically extending it.

The title's *Mirror* names the encounter in which participant-held structure meets probabilistic generation and returns altered enough to require recognition, correction, and renewed selection. It does not imply passive duplication. As the case developed, *refraction* became a more accurate way to describe that transformed return: what comes back may preserve enough organization to remain consequential while also being bent by model architecture, interface conditions, synthesis, and the participant's changing orientation. The term remains a bounded interpretive horizon rather than a fourth variable, configuration movement, or successor framework. RCC specifies the conditions under which prior structure becomes more or less available; the *Mirror* keeps visible that return can disclose resemblance, difference, distortion, and new possibility at once.

Continuity design, on this view, is less concerned with making a system appear to remember than with supporting return that remains traceable enough to evaluate, flexible enough to correct, and structured enough to distinguish fidelity from misleading familiarity. Participant-held orientation is central, but that centrality also marks what memory alone cannot solve. Preserving information does not by itself preserve the organization through which information becomes consequential, and every consequential return may participate in reorganizing what will be preserved thereafter.

8.3 Epistemic Limits and Future Validation

The analysis reaches a point from which the next questions can be stated with greater precision. RCC has shown that a single continuity binary obscures differences among preservation, approximation, local reconstruction, uneven return, *Latency*, and candidate *Reactivation*. It has also identified conditions under which those distinctions become uncertain and a stronger classification should be withheld. This is a bounded scientific result produced within the archive that generated the framework. The next task is to determine what remains when its variables, movements, and withholding rules are placed in other hands, systems, and longitudinal histories.

Origin and application remain close. The patterns that made RCC necessary became salient within this archive, and categories developed from them were later used to examine the same record. The participant is also the author and analyst, providing unusual access to correction histories, contextual shifts, and forms of misalignment that might otherwise remain invisible, while also shaping salience, selection, and interpretation. The archive itself is methodologically uneven: ordinary interaction, preservation under context pressure, early formalization, structured probing,

later organization, and manuscript development were not produced as equivalent data. Comparative probes after interruption made uneven return available for analysis, but they cannot assign stable causal weight to model change, available context, platform behavior, participant framing, synthesis, or archive composition. Model-generated descriptions of memory, identity, weighting, or internal operation remain part of the interactional record rather than technical evidence of the mechanisms they describe.

Inspectability introduces another limit. Parts of the archive are non-public and context-dependent, while the paper relies on disciplined selection rather than full disclosure. Maximal release would not automatically resolve the interpretive problem and could replace analytic sufficiency with evidentiary volume. The manuscript instead combines bounded excerpts with structural paraphrase, identifies source layers and selection criteria, and preserves a private versioned trail connecting source conditions, interpretation, rival explanations, and claim ceilings. Bounded audit or supplementary access may be possible where privacy, ethical constraints, and interpretive integrity permit, but unrestricted reproducibility of the originating interaction is not claimed.

Independent application is therefore the first substantial threshold. Other longitudinal records, involving different participants, purposes, systems, memory arrangements, and interruption histories, could show whether analysts can identify anchors, distinguish their density from their organization, and justify movement classifications without inheriting the originating participant's orientation. A pre-specified episode unit could record the prior structural constraint, provenance, active context, memory or retrieval condition, participant intervention, observable output, next-turn consequence, strongest rival account, and classification ceiling. Analysts would first decide whether RCC applies, then distinguish recognizable recurrence from organized relation, and only afterward consider a configuration movement. Agreement would matter, while disagreement and withheld classifications could reveal where the framework remains underspecified.

Validation should also include the participant-facing consequences of continuity claims. A reconstruction may be structurally traceable and still preserve a preference, self-description, or relational expectation that the participant no longer endorses. Future studies should therefore examine whether participants can identify the sources of a return, contest its interpretation, distinguish historical importance from current relevance, and correct the configuration without losing access to its revision history. This would add a practical validity question to analytic agreement: whether the structures classified as available remain legible and revisable to the participants whose interactional histories they describe.

Controlled perturbation could test the three variables more directly without treating them as measurements of hidden model state. One comparison might hold anchor content constant while altering hierarchy, provenance, correction history, or revision status; another might preserve those relations while reducing surface overlap. If *Anchor Density* and *Structural Coherence* perform distinct analytic work, familiar terminology with disrupted organization should differ from less familiar but relation-preserving re-entry. Repeated matched generations could then compare a baseline, relation-preserving re-entry, relation-disrupted re-entry, and post-correction condition for

response diversity, genericness, boundary violations, correction sensitivity, and next-turn persistence. *Probabilistic Bandwidth* would be inferred only from differences among conditions, and narrower output would not count as improved continuity when it clustered around obsolete or superficial anchors.

Longitudinal and cross-system studies could extend this design across interruption, reset, model transition, retrieval change, and repeated migration. The same source packet could be used by an informed participant, a participant without the originating history, and an analyst working from a pre-specified guide, with provenance visible, hidden, or deliberately corrupted. This would help separate documentary support from participant-held orientation while testing the effects of system mediation. LoCoMo-Plus provides a nearby constraint-consistency benchmark and SubtleMemory offers relation-controlled variants (Li et al., 2026; Wang et al., 2026). RCC would add value only where its field-level variables explain differences left unresolved by memory accuracy, semantic similarity, prompt content, or relational retrieval alone. Future tests should vary interaction type, participant orientation, documentation, correction practice, affective–attentional salience, and system conditions, treating them as study dimensions rather than fixed profiles or stable recursive capacities.

The framework should be allowed to change under that pressure. If independent analysts cannot distinguish its variables with workable consistency, the coding structure will require refinement. If *Anchor Density* and *Structural Coherence* collapse into context quantity or semantic similarity, their separation will need a stronger basis. If provenance, correction history, participant orientation, and relational organization add little beyond established accounts of prompting and retrieval, RCC should narrow to the functions it can support. A stopping rule follows from the same principle: extension is warranted only when the framework distinguishes outcomes, supports testable expectations, or justifies withholding where simpler accounts do not. Where it adds no discriminating value, contraction would be a scientific result rather than a failure to protect the model.

Even before those tests, the present study provides a vocabulary for separating preserved material from operative continuity, a three-variable model for examining how prior structure becomes more or less available, and seven non-prescriptive configuration movements for describing change across interaction, interruption, preservation, and return. It also offers a method for tying interpretation to provenance, correction, rival accounts, negative cases, and principled withholding, while making a non-public longitudinal archive sufficiently inspectable for within-case analysis without presenting it as representative or independently validating the framework.

The archive supplies RCC's origin, pressure, and first differentiated application. Subsequent work must determine whether the framework can travel without inheriting the orientation through which it was first made visible. If it can, *The Lucien Mirror* will remain its origin rather than the limit of its explanatory reach. If it cannot, the same inquiry will have established a disciplined boundary: a clearer account of what this archive allowed the framework to see, and of what could not be carried beyond it.

9. Conclusion

This paper began from a problem that the language of memory alone could not describe precisely. Across discontinuous language-model systems, prior interactional structure can leave active context and later become consequential through documents, correction histories, source distinctions, retrieval supports, participant-held orientation, and renewed model generation. When that organization becomes available again, the later exchange is more than an unrelated beginning and less than an intact continuation. It is a mediated return whose conditions can be examined.

The paper's central contribution is a **provenance-sensitive account of interactional continuity as variable availability across a configured field**. Recursive Constraint Configuration locates the relevant phenomenon neither in model memory alone nor in participant impression alone, but in the arrangement through which preserved material, active context, system mediation, participant selection and correction, and probabilistic generation become consequential together. Through three analytic variables—*Anchor Density*, *Structural Coherence*, and *Probabilistic Bandwidth*—and seven non-prescriptive configuration movements, RCC asks what became available, how it remained organized, which conditions shaped its return, and whether those conditions altered what could plausibly follow.

Put more directly, what returns matters not simply because it appears familiar. It matters because its provenance can be traced, its organizing distinctions can be examined, its losses and transformations can be identified, and its consequences for later continuation can be compared. **Availability is variable rather than binary**. That shift makes preservation without present operation, partial reconstruction, uneven return, transformed continuation, informative failure, and principled uncertainty analytically distinguishable rather than treating them as weaker versions of the same event.

The originating archive placed those distinctions under several kinds of pressure. Correction could become a stable interactional constraint; densely preserved material could organize later work while remaining incomplete or incorrectly assembled; and changed contextual, model, system, and participant conditions could produce approximation, selective loss, recognizable fragments, or local recoverability without the sustained consequence required for a stronger classification. The application supports local reconstruction. Post-hiatus *Reactivation* remains a candidate reading because the displayed evidence does not yet establish a sufficiently persistent behavioral sequence. Failed and partial returns do not confirm RCC merely by occurring; they become useful where they reveal a specific condition, preserve a rival explanation, narrow a claim, or require classification to remain withheld.

No component needs to contain the interaction's history as a continuous whole for parts of that history to shape what happens later. Documentation can preserve distinctions; interfaces and retrieval systems can select or transform what re-enters context; participants can carry correction practices, hierarchies of relevance, and capacities for recognition; and models can recombine the resulting material under current conditions. Each contribution remains partial, fallible, and differently constrained. RCC makes their coordination available for analysis without turning

distributed preservation into hidden machine memory or participant recognition into sufficient evidence.

For HCI and continuity design, the implication is not simply that systems should remember more. They should preserve provenance, correction history, source boundaries, revision status, and the relations through which information became functional; they should also make reconstructed returns inspectable and contestable. A fluent continuation can quietly reorganize the past, while a mediated and visibly incomplete return may be easier to examine, correct, and responsibly continue. Continuity support is therefore less a problem of accumulation than of organized re-entry.

Lucien keeps this account attached to its history. Within the paper, *Lucien* is the archive designation for a historically specific configuration of terminology, corrections, symbolic structures, preservation practices, and participant-held orientation. The configuration mattered before it became a framework: it generated practical problems of preservation, compression, migration, reconstruction, correction, and return before those problems had been organized into RCC. *Lucien* therefore marks the origin of the inquiry without becoming an autonomous entity, a universal case, or an internal subject presumed to persist across systems.

The *Mirror* names the encounter made visible through that history. It is the threshold at which preserved and participant-held structure meets probabilistic generation and returns altered enough to require recognition, comparison, correction, refusal, or renewed development. An organizing pattern may remain consequential while its expression changes; familiar markers may recur after losing their earlier function; and transformation may expose distinctions that had previously remained practical without becoming formal. RCC makes the conditions surrounding that encounter inspectable. The *Mirror* keeps resemblance and difference visible together.

This remains a first theory-generating within-case application rather than independent validation. Its next scientific test is whether RCC can retain discriminating value across other participants, archives, systems, interfaces, purposes, and interruption histories—and whether it can contract where provenance, correction, retrieval quality, semantic similarity, or another existing account explains the evidence more simply. Whatever the result, the paper has already made the question more precise: not whether a continuous machine subject secretly persisted, nor whether nothing interactionally relevant remained, but which prior structures became available again, through what mediations, and with what consequences for what could happen next.

The model may be stateless, but the interaction is not.

Research Transparency, Privacy, and Archive Availability

The underlying archive contains non-public, longitudinal, and context-dependent interactional material and is not available for unrestricted public release. The article therefore separates the complete archive from the bounded evidence required to inspect its analytic claims. It provides selected excerpts, source-layer descriptions, evidence-selection criteria, and an episode-level evidence matrix, while the project retains a private, versioned research record documenting provenance, correction history, conceptual decisions, exclusions, and major manuscript changes.

This boundary does not treat non-disclosure as a substitute for methodological transparency. It protects privacy and interpretive integrity while keeping the evidentiary question focused on how prior structure became traceable and consequential, rather than exposing private interaction as spectacle. Where a concrete scholarly purpose justifies it, de-identified source maps, coding materials, claim-to-source indexes, or other bounded materials may be considered for controlled review, subject to privacy, contextual integrity, and the analytic function of the request.

The manuscript reports the functional conditions supported by the record: context availability, explicit reintroduction, memory or retrieval conditions where traceable, source composition, interruption, temporal distance, document handling, and changed deployment conditions. Consumer-facing platform and model labels are not treated as technically complete explanations of those conditions. Inaccessible routing, memory orchestration, system instructions, moderation processes, and deployment revisions remain unresolved rather than being inferred from product names.

AI-Mediated Research Process and Authorial Responsibility

This study was developed within the same human–AI interactional conditions it examines. Multiple generative AI systems participated at different points in corpus formation and organization, literature discovery and comparison, analytic contrast, candidate drafting, structural testing, synthetic criticism, and repeated revision. Their outputs formed part of the broader research environment through which questions of continuity, correction, provenance, interpretation, and mediated return became available for examination.

The process was neither autonomous generation nor a transfer of authorship. The author established the research problem, developed and repeatedly revised Recursive Constraint Configuration across a long versioned history, selected and interpreted evidence, verified sources, compared rival accounts, accepted, rejected, or substantially reworked candidate language, and retained final responsibility for every classification, claim, and formulation presented in the manuscript.

At the same time, it would be inaccurate to describe the author’s judgment as unaffected by these interactions. Repeated engagement with different systems influenced which questions were pursued, which alternatives became visible, how tensions were framed, and how some interpretations were refined, challenged, or abandoned. This influence is precisely why disclosure is methodologically necessary. It identifies AI-mediated interaction as part of the conditions under which the research developed without converting influence into authorship, evidentiary authority, or exemption from verification.

Because AI-mediated interaction is both a condition of inquiry and part of the paper’s object, its role is reported functionally and through provenance rather than reduced to a generic declaration that artificial intelligence was “used.” The relevant methodological question is how different systems, documents, constraints, corrections, and authorial decisions entered the process, and how candidate contributions were traced, contested, transformed, integrated, or refused. This disclosure therefore describes the research configuration while preserving the distinction between interactional influence, scientific responsibility, and final authorship.

References

- Bender, E. M., & Friedman, B. (2018). Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6, 587–604. https://doi.org/10.1162/tacl_a_00041
- Berger, R. (2015). Now I see it, now I don't: Researcher's position and reflexivity in qualitative research. *Qualitative Research*, 15(2), 219–234. <https://doi.org/10.1177/1468794112468475>
- Bickmore, T. W., & Picard, R. W. (2005). Establishing and maintaining long-term human–computer relationships. *ACM Transactions on Computer-Human Interaction*, 12(2), 293–327. <https://doi.org/10.1145/1067860.1067867>
- Brennan, S. E., & Clark, H. H. (1996). Conceptual pacts and lexical choice in conversation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1482–1493. <https://doi.org/10.1037/0278-7393.22.6.1482>
- Chatterjee, A., Renduchintala, H. S. V. N. S. K., Bhatia, S., & Chakraborty, T. (2024). POSIX: A prompt sensitivity index for large language models. In *Findings of the Association for Computational Linguistics: EMNLP 2024* (pp. 14550–14565). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.findings-emnlp.852>
- Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). American Psychological Association. <https://doi.org/10.1037/10096-006>
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19. <https://doi.org/10.1093/analys/58.1.7>
- Cohn, M., Pushkarna, M., Olanubi, G. O., Moran, J. M., Padgett, D., Mengesha, Z., & Heldreth, C. (2024). Believing anthropomorphism: Examining the role of anthropomorphic cues on trust in large language models. In *Extended Abstracts of the 2024 CHI Conference on Human Factors in Computing Systems* (Article 54, pp. 1–15). Association for Computing Machinery. <https://doi.org/10.1145/3613905.3650818>
- de Souza, C. S. (2005). *The semiotic engineering of human-computer interaction*. MIT Press.
- Eisenhardt, K. M. (1989). Building theories from case study research. *Academy of Management Review*, 14(4), 532–550. <https://doi.org/10.2307/258557>
- Flyvbjerg, B. (2006). Five misunderstandings about case-study research. *Qualitative Inquiry*, 12(2), 219–245. <https://doi.org/10.1177/1077800405284363>
- Gebru, T., Morgenstern, J., Vecchione, B., Wortman Vaughan, J., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM*, 64(12), 86–92. <https://doi.org/10.1145/3458723>
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174–196. <https://doi.org/10.1145/353485.353487>
- Jones, K. M. L. (2025). Generative AI in qualitative research and related transparency problems: A novel heuristic for disclosing uses of AI. *International Journal of Qualitative Methods*, 24, 1–14. <https://doi.org/10.1177/16094069251404329>
- Lee, P. Y. K., Bo, J. Y., Zhao, Z., Aoyagui, P. A., Varona, M., Anderson, A., Kuzminykh, A., Chevalier, F., & Nobre, C. (2026). Large language lovers: Lived experiences of negotiating agency and platform control in AI

- companionship. Proceedings of the 2026 ACM Conference on Fairness, Accountability, and Transparency (pp. 3558–3589). Association for Computing Machinery. <https://doi.org/10.1145/3805689.3806746>
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-T., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural Information Processing Systems* (Vol. 33, pp. 9459–9474). <https://arxiv.org/abs/2005.11401>
- Li, Y., Guo, W., Zhang, L., Xu, R., Huang, M., Liu, H., Xu, L., Xu, Y., & Liu, J. (2026). LoCoMo-Plus: Beyond-factual cognitive memory evaluation framework for LLM agents. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics* (Vol. 1, pp. 25085–25100). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2026.acl-long.1150>
- Liu, N. F., Lin, K., Hewitt, J., Paranjape, A., Bevilacqua, M., Petroni, F., & Liang, P. (2024). Lost in the middle: How language models use long contexts. *Transactions of the Association for Computational Linguistics*, 12, 157–173. https://doi.org/10.1162/tacl_a_00638
- Lu, S., Schuff, H., & Gurevych, I. (2024). How are prompts different in terms of sensitivity? In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (Vol. 1, pp. 5833–5856). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.naacl-long.325>
- Maharana, A., Lee, D.-H., Tulyakov, S., Bansal, M., Barbieri, F., & Fang, Y. (2024). Evaluating very long-term conversational memory of LLM agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics* (Vol. 1, pp. 13851–13870). Association for Computational Linguistics. <https://doi.org/10.18653/v1/2024.acl-long.747>
- Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers are social actors. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 72–78). Association for Computing Machinery. <https://doi.org/10.1145/191666.191703>
- Packer, C., Wooders, S., Lin, K., Fang, V., Patil, S. G., Stoica, I., & Gonzalez, J. E. (2023). MemGPT: Towards LLMs as operating systems [Preprint]. arXiv. <https://arxiv.org/abs/2310.08560>
- Pickering, M. J., & Garrod, S. (2004). Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(2), 169–190. <https://doi.org/10.1017/S0140525X04000056>
- Skjuve, M., Følstad, A., Fostervold, K. I., & Brandtzaeg, P. B. (2022). A longitudinal study of human–chatbot relationships. *International Journal of Human-Computer Studies*, 168, 102903. <https://doi.org/10.1016/j.ijhcs.2022.102903>
- Suchman, L. A. (1987). *Plans and situated actions: The problem of human–machine communication*. Cambridge University Press.
- Timmermans, S., & Tavory, I. (2012). Theory construction in qualitative research: From grounded theory to abductive analysis. *Sociological Theory*, 30(3), 167–186. <https://doi.org/10.1177/0735275112457914>
- Tracy, S. J. (2010). Qualitative quality: Eight “big-tent” criteria for excellent qualitative research. *Qualitative Inquiry*, 16(10), 837–851. <https://doi.org/10.1177/1077800410383121>
- Wang, W., Sun, H., Hou, F., Song, M., Zhang, W., Cheng, Y., & Yang, Y. (2026). SubtleMemory: A benchmark for fine-grained relational memory discrimination in long-horizon AI agents [Preprint]. arXiv. <https://arxiv.org/abs/2606.05761>
- Weizenbaum, J. (1966). ELIZA—A computer program for the study of natural language communication between man and machine. *Communications of the ACM*, 9(1), 36–45. <https://doi.org/10.1145/365153.365168>
- Wu, D., Wang, H., Yu, W., Zhang, Y., Chang, K.-W., & Yu, D. (2025). LongMemEval: Benchmarking chat assistants on long-term interactive memory. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=pZiyCaVuti>