

Cognitive Memoisation (CM-2) Protocol

metadata (Normative)

Title:	Cognitive Memoisation (CM-2) Protocol
Former Title:	Cognitive Memoisation (CM-2) for Governing Knowledge in Human-AI Collaboration
Curator:	Ralph B. Holland
Affiliation:	Arising Technology Systems Pty Ltd
Contact:	ralph.b.holland [at] gmail.com
Publication Date:	2026-01-06T03:56Z
DOI:	2026-04-21T03:24Z 1.8.1 - https://zenodo.org/records/19681210 2026-01-26T10:19Z 1.7.3 - https://zenodo.org/records/18380070 2026-01-21T15:03Z 1.7.1 - https://zenodo.org/records/18322056
Version:	1.8.1 2026-04-21T03:24Z 1.8.1 - DOI update 2026-03-08T21:18Z 1.8.0 - renamed from "Cognitive Memoisation (CM-2) for Governing Knowledge in Human-AI Collaboration" 2026-02-13T12:20Z 1.7.4 - fixed typo near EO_BODY_END from 'my' be to 'may' be. 2026-01-26T10:19Z 1.7.3 - # EO_BODY_END (instead of 1. EO_BODY_END - mediawiki markup problem) 2026-01-24T02:30Z 1.7.2 - curated removed timezone metadata 2026-01-21T15:03Z 1.7.1 - released for DOI 2026-01-21T15:03Z 1.7.0 - included # EO_BODY_END sentinel
Updates:	2026-01-15T01:46Z 1.6.0 – Included (CM-2) in the title to avoid confusion with CM-1, updated scope 2026-01-13T09:49Z 1.5.1 – Included Scope statement 2026-01-13T03:03Z 1.5.0 – Title Change, Included Additional EO invariants 2026-01-11T13:43Z 1.4.2 – refactor EO/EA naming 2026-01-11T18:40Z 1.4.1 – normative protocol 2026-01-09T11:49Z 1.4.0 – Stop Press - terminology ratification for EO 2026-01-08T14:09Z 1.3.0 – amended vendor portability invariants (Appendix A) 2026-06-06T14:38Z 1.2.0 – included Appendix A 2026-06-06T09:43Z 1.1.0 – replaced Figure 1
Binding:	normative

The metadata table is CM-defined and constitutes the authoritative provenance record for this artefact.

All fields in that table (including artefact, author, version, date and reason) MUST be treated as normative metadata.

The assisting system MUST NOT infer, normalise, reinterpret, duplicate, or rewrite these fields. If any field is missing, unclear, or later superseded, the change MUST be made explicitly by the human and recorded via version update, not inferred.

Curator Provenance and Licensing Notice

This document name supersedes the former title "Holland R. B. (2026-01-06T03:56Z) Cognitive Memoisation (CM-2) for Governing Knowledge in Human-AI Collaboration".

As curator and author, I apply the **Apache License, Version 2.0** (<https://www.apache.org/licenses/LICENSE-2.0.txt%7C>), at publication to permit reuse and implementation while preventing enclosure or patent capture. This licensing action does not revise, reinterpret, or supersede any normative content herein.

Authority remains explicitly human; no implementation, system, or platform may assert epistemic authority by virtue of this license.

Cognitive Memoisation (CM-2) Protocol

Scope

Cognitive Memoisation (CM-2) is a **public domain** protocol specification under Apache 2.0. This document defines CM-2 as a vendor-neutral protocol for the externalisation, transport, rehydration, and governance of epistemic objects in stateless AI systems. The capitalised terms **MUST**, **MUST NOT**, **SHALL**, and **MAY** are used in their normative sense. Any system claiming CM-2 compliance is required to implement the protocol and invariants defined herein.

Scope Invariant (Normative — CM-2)

S1. Binding Layer

In CM-2, **all Appendices explicitly marked “Normative” constitute the binding protocol**. They define the complete and exclusive set of conformance requirements for any CM-2 implementation.

S2. Non-Binding Material

All other sections of CM-2 — including introductions, metaphors, examples, figures, and explanatory text — are **informative only** and **MUST NOT** be used to assert or deny conformance.

S3. Precedence

If any conflict exists between the main body and a Normative Appendix, **the Normative Appendix prevails**.

S4. Portability Guarantee

Any claim of CM-2 compliance that does not fully implement the Normative Appendices is **invalid**, regardless of partial or informal alignment with the paper.

Abstract

Large language models and other probabilistic systems are increasingly used as collaborators in knowledge-intensive work. While these systems can generate fluent and contextually rich outputs, they are inherently stateless, non-authoritative, and prone to subtle forms of drift, rediscovery, and silent reinterpretation. These failure modes are not primarily errors of intelligence or capability; they are failures of knowledge governance.

This paper introduces Cognitive Memoisation (CM) and presents the first public articulation of CM-2, a governance-oriented framework for human–AI collaboration. Using the metaphor of collaborative shipbuilding, the paper illustrates why complex work fails when assumptions, context, and prior agreements are not explicitly externalised and governed. CM-2 formalises the separation between continuity and authority by defining memoised objects for grounded material, contextual attachments for governing facts, bounded Universes of Discourse for parallel reasoning, and an explicit human-initiated promotion boundary for durable knowledge.

Rather than proposing new model architectures or vendor-dependent memory mechanisms, CM-2 focuses on disciplined externalisation, idempotent identity, and explicit human authority as the foundations for reliable knowledge work. The paper positions CM-1 as an immediately applicable practice, while CM-2 establishes the normative structure required for interoperable tooling and future implementations.

This contribution establishes the scope, invariants, and practical utility of Cognitive Memoisation while deliberately leaving full formalisation and vendor alignment to subsequent work, allowing the method to be practised openly ahead of implementation support.

CM-2 is a normative protocol that defines conformance requirements for the externalisation, serialisation, interchange, and promotion of knowledge artefacts under stateless cognitive interaction.

1 Introduction

This specification is grounded in the Cognitive Memoisation framework introduced in “Progress Without Memory: Cognitive Memoisation as a Knowledge-Engineering Pattern for Stateless LLM Interaction”, which establishes the epistemic and governance assumptions used throughout.

Modern knowledge work increasingly involves collaboration between humans and probabilistic systems such as large language models. These systems are capable of generating fluent text, structured reasoning, and domain-specific material at scale, yet they remain fundamentally stateless, non-authoritative, and context-bounded. As a result, many failures observed in practice are not dramatic errors or obvious hallucinations, but quieter breakdowns: rediscovery of prior work, silent reinterpretation of agreements, loss of contextual boundaries, and gradual semantic drift.

These failures are often misdiagnosed as shortcomings of intelligence, memory, or model capability. In practice, they are failures of knowledge governance. Humans assume continuity where none exists, machines produce plausible output without authority, and neither side has a clear mechanism for deciding what has been agreed, what remains tentative, and what must persist beyond the current interaction.

To illustrate the problem, consider the metaphor of building a ship. Shipbuilding is not primarily an exercise in intelligence or creativity; it is an exercise in coordination under constraint. Multiple teams work in parallel, each within its own local context, yet their work must eventually align. Assumptions that remain implicit, drawings that are not externalised, or changes that are not governed quickly become expensive failures. No amount of individual competence compensates for poor control of what has been decided, what remains provisional, and what has been committed.

Human–AI collaboration exhibits the same structural risks. Conversations feel continuous even when they are not. Generated material feels authoritative even when it is not. Context appears shared even when it is only locally inferred. Without explicit mechanisms for externalisation, scoping, and promotion, collaborators are forced to rely on memory, inference, or trust where governance is required.

This paper introduces Cognitive Memoisation (CM) as a practice and governance framework for addressing these problems. CM does not attempt to improve model intelligence, provide artificial memory, or enforce global ontologies. Instead, it focuses on disciplined externalisation and explicit authority. Knowledge is treated as something that must be deliberately captured, bounded, and promoted, rather than implicitly accumulated.

The contribution of this paper is the first public articulation of CM-2, the normative layer of Cognitive Memoisation. CM-2 defines a small set of invariants that separate continuity from authority and practice from commitment. It introduces memoised objects for grounded material, attachment objects for governing facts, bounded Universes of Discourse for parallel reasoning, and an explicit human-initiated promotion boundary for durable knowledge. These elements are intentionally conservative and designed to be usable without specialised tooling.

At the same time, CM-1, the practice layer of Cognitive Memoisation, can be adopted immediately. CM-1 consists of habits and disciplines that reduce rediscovery and drift through explicit capture and stand-down, independent of vendor support or system integration. CM-2 formalises these practices sufficiently to enable interoperability, future tooling, and shared understanding, while deliberately avoiding premature formalisation.

The remainder of this paper elaborates these ideas in a pragmatic order. Section 2 examines common failure modes in human–AI knowledge work. Section 3 introduces the core CM concepts and their intended roles. Section 4 illustrates how these concepts operate together through practical scenarios. Section 5 discusses scope, non-goals, and limitations. The paper concludes by outlining how CM may be practised today and extended incrementally as tools and implementations mature.

Cognitive Memoisation treats date/time as an explicit, governed property of externalised cognition rather than an implicit feature of conversation or system state.

Note: that Epistemic Object (EO) and Epistemic Attribute (EA) and their Projection are defined in the normative Appendix A.

2 Failure Modes in Human–AI Knowledge Work

Many of the most damaging failures in human–AI collaboration are subtle rather than catastrophic. Outputs appear fluent, reasoning seems coherent, and conversations feel continuous. Yet over time, work degrades: prior conclusions are rediscovered, assumptions quietly change, and context boundaries blur. These failures are often tolerated until they accumulate into confusion or rework.

This section describes common failure modes observed in practice. These failures are not specific to any particular model or vendor, nor are they primarily technical defects. They arise from structural mismatches between how humans assume knowledge persists and how probabilistic systems actually operate.

2.1 Rediscovery and the Illusion of Continuity

Large language models do not retain memory across interactions unless explicitly supported by external mechanisms. Despite this, humans naturally assume continuity, especially when conversational tone and vocabulary are consistent. As a result, work that has already been completed is revisited, re-derived, or re-negotiated without awareness that it has occurred before.

This phenomenon, informally described as Groundhog Day, is not merely inefficient. It erodes confidence in prior decisions and encourages informal workarounds such as repetition, excessive prompting, or reliance on personal memory rather than shared artefacts.

2.2 Silent Semantic Drift

Even when material is revisited rather than rediscovered, its meaning may shift subtly over time. A term introduced informally may later be treated as defined. A speculative idea may harden into an assumption. A local explanation may be reused outside its original context.

Because these shifts occur incrementally and without explicit checkpoints, they are rarely detected until disagreement or error surfaces. The system appears cooperative and responsive throughout, masking the fact that meaning has changed.

2.3 Conflation of Fluency with Authority

Generated outputs are often well-structured, confident in tone, and responsive to instruction. This fluency can be mistaken for authority. Humans may treat generated material as agreed, settled, or endorsed simply because it appears complete.

In reality, probabilistic systems do not assert authority. They produce plausible continuations conditioned on input. Without explicit markers distinguishing tentative material from committed knowledge, collaborators are left to infer status based on style rather than governance.

2.4 Context Boundary Collapse

Complex work frequently involves parallel lines of reasoning, alternative hypotheses, or different audiences. Humans manage this by maintaining mental boundaries regarding what applies where and under what assumptions. In conversational systems, these boundaries are easily crossed.

Material introduced for illustration may later be reused as fact. Constraints scoped to a particular discussion may be applied globally. Without explicit containment, contexts bleed into one another, producing confusion rather than synthesis.

2.5 Implicit Promotion and Accidental Canon

Perhaps the most dangerous failure mode occurs when material becomes treated as durable knowledge without any explicit decision to do so. Notes become references. Draft explanations become guidance. Generated summaries become policy.

This implicit promotion bypasses human authority. It replaces deliberate commitment with gradual assumption. Once this occurs, correcting errors becomes socially and cognitively difficult, even if the underlying material was never intended to be authoritative.

2.6 Why These Failures Persist

These failure modes persist because neither humans nor machines are equipped, by default, with mechanisms for explicit knowledge governance. Humans rely on memory and shared understanding; machines rely on inference and pattern continuation. In the absence of deliberate externalisation and promotion boundaries, both sides compensate informally.

The result is not dramatic failure, but slow erosion of clarity.

3 Cognitive Memoisation

Cognitive Memoisation (CM) is a governance-oriented approach to human–AI knowledge work. It is concerned not with how systems generate content, but with how knowledge is externalised, scoped, and promoted in environments where reasoning is probabilistic and memory is not intrinsic.

CM begins from a simple observation: humans routinely assume continuity, authority, and shared understanding in collaborative work, while probabilistic systems provide none of these guarantees. Cognitive Memoisation addresses this mismatch by making continuity explicit, authority deliberate, and durability conditional.

3.1 Separation of Continuity and Authority

A central principle of CM is the separation between continuity and authority. Continuity refers to the ability to recognise and reuse material across interactions, preventing unnecessary rediscovery. Authority refers to the human decision to stand behind material as durable knowledge.

In CM, continuity may be supported automatically, but authority is never inferred. No artefact becomes authoritative without explicit human action. This separation allows work to proceed fluidly without silently accumulating commitments.

3.2 Epistemic Objects (EO)

CM introduces memoised Epistemic objects, referred to as EO, as the first-class units of continuity. An EO captures grounded material in the form it was introduced: prose, logic, constraints, rules, examples, or other structured or unstructured content. EO preserve locality and provenance, avoiding the fragmentation that occurs when material is decomposed into abstract conceptual fragments.

In CM, memoisation is used in its original technical sense: the automatic caching of results to avoid unnecessary recomputation. EO are therefore automatically memoised in a client-side cache to support continuity and prevent rediscovery across interactions. This memoisation is non-authoritative and may be ephemeral or persistent, but it is always subordinate to explicit human governance. Memoisation enables reuse without implying memory, agreement, or durability.

EO are memoised to prevent rediscovery. When externalised, their identities are idempotent: repeated recognition of the same EO resolves to the same object rather than creating duplicates. EO are fact-like in the structural sense, but they do not assert truth or authority.

EO are dynamically extensible. Their internal structure is not constrained by a rigid schema. Interoperability is achieved through serialisation conventions, stable formats, and plain-language parsing rather than fixed ontologies.

Externalised Objects are explicitly anchored in real clock time. Each EO carries an inviolate creation date/timestamp and governed update timestamps to ensure auditability, portability, and cross-vendor interoperability. These timestamps are part of the object's identity and are not inferred from session order, conversation flow, or system behaviour. The normative requirements governing date/time semantics are defined in Appendix A.

3.3 Epistemic Attachment (EA)

CM distinguishes EO from Epistemic Attachment/Attribute objects, referred to as EA. EA are non-authoritative artefacts that qualify, constrain, or contextualise EO. They express governing facts such as scope, provenance, lifecycle state, usage cautions, or interpretive notes.

EA exist to regulate how EO are interpreted and used without polluting the EO themselves. They are context-bound, disposable, and never promoted independently. If an attachment acquires enduring significance, it is promoted explicitly into an EO by human action.

As a semantic subtype of Epistemic Objects, Epistemic Attributes inherit the same time anchoring and portability requirements. (Note: Epistemic Attachment (EA) have subsequently been interchanged with Epistemic Attributes in the normative section - apologies this needs to be reverse propagated to this Prose section.)

3.4 Thought Bubbles as Universes of Discourse

Reasoning rarely occurs in a single global context. CM models this by introducing Thought Bubbles, implemented as a subtype of EO. A Thought Bubble defines a bounded Universe of Discourse within which EO are interpreted and related.

Thought Bubbles allow multiple lines of reasoning, audiences, or hypotheses to coexist without forced unification. EO may participate in multiple Thought Bubbles, but their interpretation remains local to each context. Thought Bubbles must not merge implicitly; aggregation is always explicit.

3.5 Persistence and Promotion

CM distinguishes between ephemeral continuity and durable knowledge. EO and EA may be ephemeral or persistent, supporting work across sessions without claiming authority. Persistence alone does not imply commitment.

Durable knowledge is created only through an explicit human-initiated promotion action, referred to here as Save As. Save As operates on aggregates, typically one or more Thought Bubbles, and externalises them into an enduring system of record with declared scope and authority. Nothing is promoted individually or automatically.

3.6 CM-1 and CM-2

CM is articulated in layers. CM-1 describes the practice of disciplined externalisation and stand-down, and can be adopted immediately without specialised tooling. CM-2, introduced here, provides the normative structure that defines objects, boundaries, and invariants, enabling shared understanding and future interoperability.

This paper presents CM-2 at a deliberately conservative level, sufficient to demonstrate utility and establish scope without constraining future implementations.

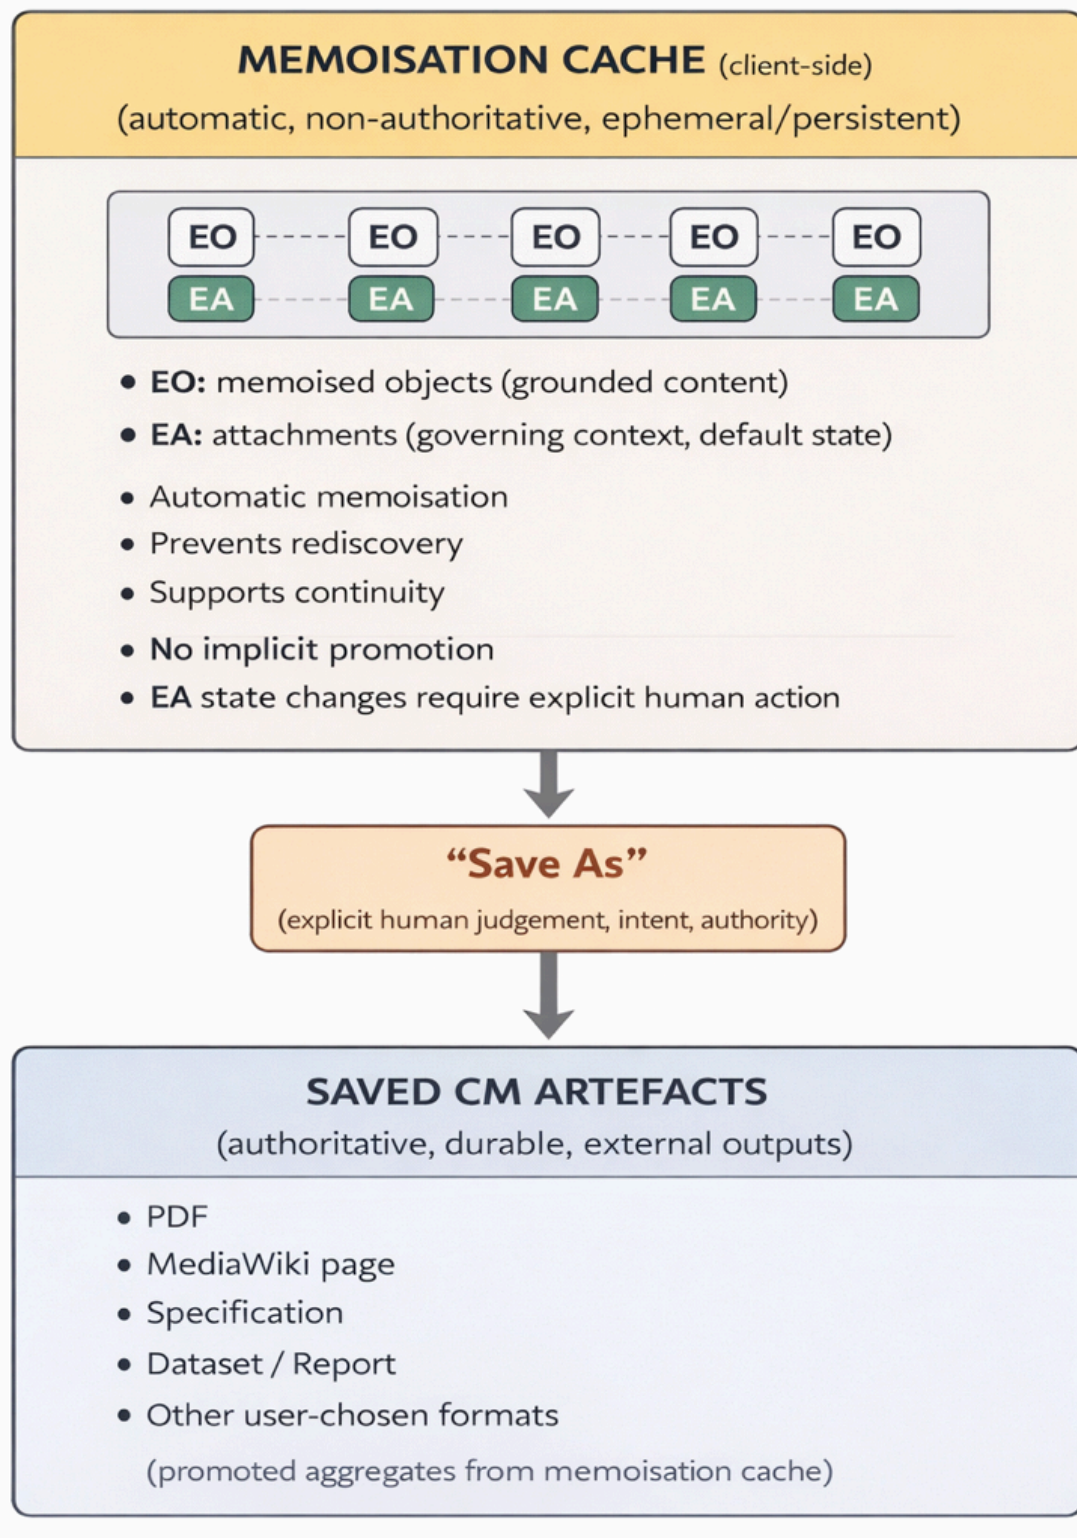


Figure 1: Cognitive Memoisation separates continuity from authority. EO and EA are automatically memoised in a client-side cache to provide continuity. Reasoning occurs within bounded Thought Bubbles (Universes of Discourse). Durable knowledge is created only when a human explicitly promotes an aggregate across the promotion boundary via Save As. Nothing becomes authoritative by accumulation or inference.

4 Putting Cognitive Memoisation to Work

This section illustrates how Cognitive Memoisation operates in practice. The intent is not to prescribe tooling or workflows, but to show how CM concepts reduce failure modes when applied to ordinary knowledge work involving probabilistic systems.

4.1 A Typical Working Session

Consider a human collaborating with a language model to explore a complex topic such as drafting a technical argument, analysing a regulatory question, or designing a system architecture. Early in the session, ideas are tentative, overlapping, and incomplete. Multiple lines of reasoning may be explored in parallel.

In CM terms, this work occurs within one or more Thought Bubbles, each defining a local Universe of Discourse. EO are created automatically to capture grounded material as it appears: explanatory prose, constraints, examples, draft formulations, or intermediate reasoning. EA accumulate alongside EO, recording governing facts such as scope limitations, uncertainty, or provisional status.

At this stage, nothing is authoritative. Continuity is supported, but commitment is explicitly absent.

4.2 Avoiding Rediscovery Without Premature Commitment

As the session progresses or resumes later, EO enable continuity without requiring the human to restate prior material. The system can recognise and reuse previously externalised objects, preventing rediscovery and unnecessary repetition.

Crucially, this continuity does not imply agreement. EO may be revisited, revised, or abandoned. EA may be updated or discarded. Thought Bubbles may be parked or allowed to decay. CM tolerates this fluidity because authority has not yet been asserted.

4.3 Managing Parallel Contexts

Complex work often requires reasoning across different audiences or assumptions. CM supports this by maintaining separate Thought Bubbles.

The same EO may appear in multiple bubbles, but its interpretation is local to each. This allows collaboration and reuse without semantic collapse. No reconciliation is required until a human explicitly decides to aggregate contexts.

4.4 Explicit Promotion Through Save As

When a point of agreement is reached or when material is judged worth preserving, the human initiates a Save As action. Save As selects one or more Thought Bubbles and externalises their aggregate into a durable artefact.

This promotion declares scope, intent, and authority. It marks the transition from exploratory material to committed knowledge. Importantly, Save As is deliberate and visible. Nothing is promoted implicitly, and nothing becomes authoritative by accumulation.

4.5 Stand-Down and Reuse

After promotion, the session may end. EO and EA that were not promoted may be discarded or retained as persistent continuity aids. CM explicitly supports stand-down without loss: authoritative knowledge is preserved externally, while exploratory material may safely decay.

When work resumes, new sessions may reference promoted artefacts as authoritative, while generating fresh EO for new exploration. Continuity and authority remain distinct.

4.6 Why This Matters in Practice

Applied consistently, CM changes the character of human–AI collaboration. Rediscovery is reduced without pretending memory. Drift is contained without enforcing global schemas. Authority is established explicitly by human action rather than inferred from fluency.

Most importantly, CM allows people to work naturally while providing a clear path to durable knowledge when it matters.

5 Scope, Non-Goals, and Limitations

Cognitive Memoisation is intentionally conservative in scope. It is designed to address specific failure modes in human–AI knowledge work related to continuity, authority, and governance. It does not attempt to solve all problems associated with probabilistic systems, nor does it propose new forms of intelligence, reasoning, or autonomy.

5.1 Scope

CM applies to knowledge-intensive work where reasoning unfolds over time, context is fragmented or revisited, collaboration occurs across sessions or participants, and authority must remain explicitly human.

Typical domains include technical design, legal reasoning, policy drafting, research synthesis, and other forms of deliberative work where loss of context or silent drift carries real cost.

CM is agnostic to specific tools, vendors, or model architectures. Its principles may be practised with existing systems and incrementally supported by future tooling.

5.2 Non-Goals

CM explicitly does not aim to provide artificial memory or persistent model state, enforce global ontologies or universal schemas, replace human judgement or accountability, guarantee correctness or truth of content, or automate promotion of knowledge.

CM treats probabilistic systems as capable assistants, not authorities. It assumes that uncertainty, incompleteness, and revision are normal features of real work.

5.3 Limitations

CM introduces discipline where informality is often tolerated. This carries a modest cognitive and procedural cost. Explicit externalisation and promotion require attention and intent, and may feel unfamiliar to users accustomed to conversational continuity.

CM also does not prevent all errors. It reduces rediscovery, drift, and accidental canonisation, but it does not eliminate misunderstanding or disagreement. Its value lies in making such issues visible and governable rather than implicit.

CM-2 as presented here is normative rather than exhaustive. It defines invariants and boundaries, not full specifications or implementations. Tooling support may simplify adoption, but is not required for practice.

An important consequence of Cognitive Memoisation's explicit externalisation and promotion boundaries is that it naturally supports capabilities typically addressed only through specialised infrastructure. Because authoritative knowledge is externalised deliberately and memoised objects are identity-stable, CM enables practical disaster recovery and restoration without relying on implicit system state. Sessions may be reconstructed, inspected, or abandoned without loss of committed knowledge. The use of bounded Universes of Discourse allows post hoc analysis and semantic comparison of alternative reasoning paths, while explicit aggregation and promotion support controlled semantic merge rather than silent unification. Provenance is preserved through attachment objects and serialisation conventions, enabling traceability of origin, scope, and intent. Together, these properties provide a foundation for auditing and review that is compatible with probabilistic systems, without requiring global schemas, continuous memory, or vendor-specific mechanisms.

6 Conclusion

As probabilistic systems become embedded in everyday knowledge work, the primary risks they introduce are not failures of intelligence, but failures of governance. Fluent output, contextual responsiveness, and conversational continuity create the illusion of shared understanding without providing mechanisms for authority, durability, or scope.

Cognitive Memoisation addresses this gap by separating continuity from authority and practice from commitment. By introducing memoised objects for grounded material, attachment objects for governing facts, bounded Universes of Discourse for parallel reasoning, and an explicit human-initiated promotion boundary, CM provides a disciplined yet lightweight approach to human–AI collaboration.

This paper has presented the first public articulation of CM-2, establishing its scope, invariants, and practical utility while deliberately avoiding premature formalisation. CM-1 may be practised immediately, independent of tooling or vendor support. CM-2 provides the normative structure required for shared understanding and future interoperability.

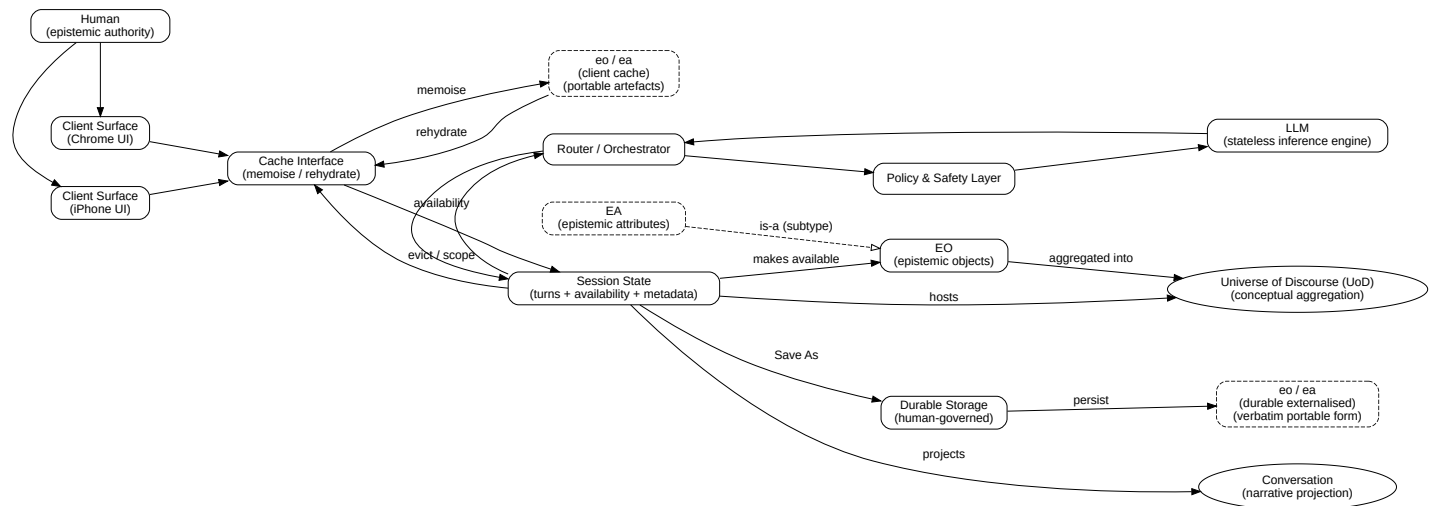
Cognitive Memoisation does not promise smarter systems. It promises clearer work. In environments where knowledge must persist, evolve, and be trusted, that distinction matters.

Normative Clarification and Strengthening of CM-2

The following appendix is normative and defines protocol-level conformance requirements. Implementations that fail to satisfy these requirements are non-compliant with CM-2, regardless of intent or partial adoption.

Applies to: “Cognitive Memoisation for Governing Knowledge in Human–AI Collaboration” Status: Normative — CM-2

This section is inserted between the front matter of the original CM-2 paper and Appendix A. It does not replace the paper’s argument; it operationalises it.



7.1 Purpose

This section hardens CM-2 from a governance framework into a portable, enforceable contract while preserving the original separation between continuity, authority, and durability.

Nothing in this section grants authority to platforms, models, or tooling. All authority remains explicitly human.

7.2 Epistemic Objects (EO)

CM-2 defines Externalised Objects (EO) as the first-class units of continuity.

This clarification establishes that EO are:

- explicit text artefacts, not internal representations
- inert and non-authoritative by default
- human-legible and auditable
- session-durable once created or rehydrated

EO are not memory. EO are not model state. EO are externalised cognition.

7.3 Epistemic Attributes (EA)

Epistemic Attributes (EA) are a semantic subtype of EO.

EA:

- IS an EO (inherits all EO invariants)
- carries only epistemic, lifecycle, scope, or governance information
- MUST NOT contain domain meaning
- MUST reference one or more EO by UUID

EA existence does not imply truth, correctness, or promotion.

7.4 Session Durability and Rehydration

EO and EA are session-durable by default.

- Once created or rehydrated, EO/EA MUST persist for the full duration of the session
- EO/EA MUST NOT disappear mid-session except by explicit human deletion or session termination

On session resumption:

- cached EO/EA are automatically rehydrated
- rehydration makes EO/EA coherently available for reasoning
- rehydration does not assert, promote, or externalise

Rehydration is functionally equivalent to restoring previously externalised text into the working context.

7.5 Universes of Discourse (UoD)

A Universe of Discourse (UoD) is a session-local aggregation of EO/EA forming a bounded *Thought Bubble*.

UoD:

- may contain ideas, groundings, episodic recordings (when requested), and Dangling Cognates
- interoperates in-session regardless of ingestion path (turn text, upload, restored logic)
- is analogous to a “thread” but explicitly externalised and governed
- does not assert authority by containment

A thread of thought MAY traverse multiple UoD and EO/EA.

7.6 Time and Identity

CM-2 requires real, platform-provided time.

- EO/EA MUST carry a real creation timestamp (`created_at`)
- Any human-caused change to EO/EA body content MUST advance `updated_at`
- Timestamps are set automatically by the platform using wall-clock time
- Humans MUST NOT edit timestamps directly
- Time MUST NOT be inferred, guessed, or marked UNKNOWN

A thought that changes has moved in time.

7.7 Govern, Save As, and Durability

(Normative — CM-2)

- Govern is an explicit human act.
- Govern marks EO and EA as **durable-in-intent**.
- Govern does not externalise content.

Save As ⇒ durable

- Save As is the sole mechanism that externalises EO and EA.
- Externalisation occurs outside the platform and client, under explicit human control.

Platforms **MUST NOT**:

- infer durability from session presence,
- silently persist EO or EA,
- promote content by accumulation, repetition, or longevity.

Governance applies to the Epistemic Attributes (EA) that define, qualify, or constrain the interpretation of Externalised Objects (EO) within a Thought Bubble or Universe of Discourse.

Before a Thought Bubble or Universe of Discourse may be selected for *Save As*, all EA that are intended to:

- persist beyond the session,
- qualify meaning,
- constrain interpretation,
- or define governing context,

MUST be explicitly governed by a human.

Governing these EA establishes deliberate human responsibility for the context in which the EO are interpreted. Selection for *Save As* without prior governance of the relevant EA is **non-compliant with CM-2**.

7.8 Portability and Vendor Neutrality

CM-2 portability is a condition of conformance.

EO/EA:

- are plain UTF-8 text artefacts
- **MUST** be exportable and re-ingestible across vendors
- **MUST** preserve semantic identity and meaning

Vendors **MAY** innovate in tooling and UX, but **MUST NOT**:

- alter EO/EA semantics
- assert authority
- rewrite or reinterpret content

7.9 Consequence for Platforms

By externalising continuity and governing promotion, CM-2 enables deeper, longer-horizon reasoning without opaque memory or platform lock-in.

The ship is built not by remembering conversations, but by keeping parts labelled, auditable, and under human control.

7.10 Normative Closure

| Continuity is automatic. Authority is human. Durability is deliberate.

7.11 Aggregation and Epistemic Neutrality

(Normative — CM-2)

Aggregation of Externalised Objects (EO) and/or Epistemic Attributes (EA) **MUST NOT** alter epistemic state, authority, durability, or commitment.

Aggregation includes, but is not limited to:

- containment within a Universe of Discourse (UoD),
- grouping, bundling, or collection,
- ordering or sequencing,
- visual presentation or emphasis,
- selection as part of a candidate set for Save As,
- inclusion within summaries, views, or projections.

No aggregation mechanism **confers validation, endorsement, correctness, agreement, or commitment**.

Aggregation **MUST NOT** be interpreted as:

- human approval,
- assertion of truth,
- readiness for promotion,
- durable knowledge,
- or epistemic elevation of any kind.

Only an explicit human action — **Govern** or **Save As**, as defined in this specification — may confer durability or authority.

Appendix A (Revised) — Epistemic Objects, Epistemic Attributes, and Session Governance

Normative — CM-2

Transclusion: CM-2 protocol

Version: 1.7.5

Update: 1.7.5 - included A.0.2

Binding: normative

A.0 Terminology Clarification

In CM-2, EO always means Epistemic Object. Externalisation is a change of substrate, not a change of object.

The same Epistemic Object may exist simultaneously as:

- EOm - the object as held in human cognition (mind)
- EOs - the object as expressed on the interaction surface (text, UI, dialogue)
- EOc - the object as memoised in the client cache or session store
- EOe - the object as serialised in an external persistent store

These are not different objects. They are the same Epistemic Object realised across different surface substrates.

All governance via EA (Epistemic Attributes) applies to the Epistemic Object across all of these surface substrates.

A.0.1 Purpose and Scope

This appendix defines the inviolate contract for Epistemic Objects (EO), Epistemic Attributes (EA), and their use within CM-2 sessions.

Its purpose is to ensure that:

- Externalised cognition is portable across vendors
- Reasoning depth increases without opaque platform memory
- Authority, durability, and persistence remain strictly human-governed
- EO/EA can be audited, exchanged, rehydrated, and reasoned over without semantic drift

Conformance to this appendix is mandatory for any platform claiming CM-2 compliance.

A.0.2 Turn compliance (normative)

(Normative — CM-2)

- CM-TURN-01 Session Ordinal Requirement

Every interaction turn SHALL expose a session-local ordinal field `turn_seq`.

- CM-TURN-02 Monotonicity of `turn_seq`

`turn_seq` SHALL be strictly monotonic within a session.

For every successive turn n :

`turn_seq(n) > turn_seq(n-1)`

- CM-TURN-03 Arithmetic Stability

`turn_seq` SHALL be an integer suitable for deterministic modulo evaluation.

The implementation SHALL support exact evaluation of:

`turn_seq mod N`

- CM-TURN-04 Origin and Custody of `turn_seq`

`turn_seq` MAY be platform-provided or implementation-maintained.

If implementation-maintained, it SHALL be persisted in Durable Substrate and rehydrated exactly.

Humans MUST NOT directly edit `turn_seq` during ordinary protocol operation.

- CM-TURN-05 No Inference or Reconstruction

`turn_seq` MUST NOT be guessed, interpolated, inferred from transcript length, or reconstructed probabilistically from conversational text.

If authoritative `turn_seq` is unavailable, the turn is non-compliant.

- CM-TURN-06 Time Availability Requirement

Every interaction turn SHALL expose an authoritative wall-clock timestamp `turn_time`.

`turn_time` SHALL be platform-provided where available, or obtained by explicitly governed tooling under human-approved durable-substrate operation.

- `turn_time` MUST be real wall-clock time.

- CM-TURN-07 Monotonicity of `turn_time`

For every successive turn n :

`turn_time(n) > turn_time(n-1)`

Equal values MAY be permitted only if the clock resolution is coarser than the turn rate and this is explicitly declared by the implementation.

Otherwise non-increase constitutes a fault.

- CM-TURN-08 Time Non-Substitution

`turn_time` MUST NOT be inferred, guessed, symbolic, relative, logical, session-indexed, or human-supplied.

Absence or substitution of `turn_time` is non-compliant.

- CM-TURN-09 Every-Turn Validation

At every turn, prior to ordinary inference, the implementation SHALL validate:

1. `turn_seq` availability
2. `turn_seq` strict monotonicity
3. `turn_time` availability
4. `turn_time` monotonicity
5. if `turn_seq mod N == 0`, ROC cardinality $K_{obs}(roc_id) \geq K_{req}$

- CM-TURN-10 Fault Classification

The following SHALL be classified as protocol faults:

1. missing `turn_seq`
2. non-monotonic `turn_seq`

- 3. missing turn_time
- 4. non-monotonic turn_time
- 5. ROC referential integrity break
- 6. checkpoint cardinality failure

- CM-TURN-11 Fault Semantics

- Faults SHALL be hard faults, not advisory signals.
- On detection of any fault, the implementation SHALL:
 1. declare the fault explicitly
 2. suspend ordinary inference that depends on the violated invariant
 3. invoke the deterministic recovery ladder where recovery is defined
 4. emit a GAP artefact if recovery cannot restore compliance

- CM-TURN-12 Recovery Ordering

If turn_seq or turn_time is unavailable, deterministic recovery SHALL NOT fabricate them. Recovery MAY restore ROC structure, but SHALL record a GAP artefact for temporal or ordinal non-compliance if those primitives remain unavailable.

- CM-TURN-13 Checkpoint Trigger

Checkpoint evaluation SHALL occur if and only if:

- turn_seq mod N == 0
- No other event SHALL silently substitute for the checkpoint trigger unless explicitly declared by normative override.

- CM-TURN-14 Bootstrap Requirement

At bootstrap, the implementation SHALL establish both:

- turn_seq
- turn_time

before claiming CM-2 compliance for ordinary inference.

If either cannot be established, bootstrap SHALL fail closed.

- CM-TURN-15 Audit Record

For every turn, the implementation SHALL be able to externalise:

- turn_seq
- turn_time
- K_req
- N
- epsilon
- evaluated roc_id set
- referential_integrity_status
- attention_deficit_status
- recovery_step_taken

A.2 Epistemic Object (EO) (normative)

A.2.1 Definition

An **Epistemic Object (EO)** is a human-legible, inert text artefact representing externalised cognition.

EO:

- has no agency
- asserts no authority by default
- participates in reasoning only when present in a session context
- MUST persist for the full duration of a session
- is externalised durably only by explicit human action

EO are not memory. EO are not model state. EO are externalised thought.

A.2.2 EO Identity Invariance

(Normative — CM-2)

Re-serialisation, transport, copying, export, import, restoration, or re-ingestion of an Externalised Object (EO) **MUST NOT** create a new EO identity.

An EO whose *uuid*, *created_at*, and body content are preserved **IS the same EO**, regardless of:

- vendor,
- platform,
- storage system,
- session,
- transport mechanism,
- or representation medium.

Vendors **MUST NOT** treat imported, restored, or rehydrated EO as new objects, new versions, or derivatives unless an explicit human action creates a new EO.

Identity equivalence is semantic and temporal, not positional or storage-based.

A.3 Epistemic Attribute (EA)

A.3.1 Definition

An **Epistemic Attribute (EA)** is a semantic subtype of EO whose purpose is to govern interpretation of EO.

EA:

- IS an EO (inherits all EO invariants)
- MUST declare itself with a TOML [ea] header
- MUST reference one or more EO by UUID
- MUST NOT contain domain meaning
- carries only epistemic, lifecycle, scope, or governance information

EA existence does not imply truth, correctness, or promotion.

A.4 Encoding and Envelope (Inviolate)

A.4.1 Encoding

All EO and EA MUST be encoded as:

- UTF-8 plaintext (inviolate)

Non-UTF-8 artefacts are non-compliant.

A.4.2 Header Format

- The envelope header MUST be TOML
- EO MUST begin with [eo]
- EA MUST begin with [ea]

No aliasing, inference, substitution, or stochastic rewriting is permitted.

A.4.3 Mandatory Header Fields (Inviolate)

Every EO and EA MUST include:

- *uuid* — portable, canonical textual UUID

- `semantic_version` — MAJOR.MINOR.PATCH (envelope version)
- `created_at` — portable timestamp (ISO-8601)

These fields are inviolate. Omission or mutation is non-compliant.

A.5 Header–Body Boundary (Inviolate)

A commented sentinel MUST mark the boundary between header and body (with the comment as first character of the line) e.g.:

```
# EO_BODY_START
```

- This sentinel is inviolate
- Vendors MUST NOT infer boundaries
- Everything after the sentinel is body content

A.6 Body Semantics (Inviolate)

An EO is a UTF-8, TOML-headed, sentinel-bounded text artefact with an inviolate UUID and real wall-clock timestamps, whose body is opaque, inert, and non-authoritative unless explicitly promoted by a human.

A.6.1 Body Form

After the boundary sentinel, the body is:

- a single opaque plaintext block
- UTF-8 encoded
- parseable and tokenisable by all CM-2 platforms

A.6.2 Permitted Content

The body MAY include:

- prose
- lists
- URLs
- text-encoded payloads (e.g. base64)
- any content a model can decode only under explicit human direction

The body:

- is inert
- MUST NOT execute
- MUST NOT assert authority
- MUST NOT imply structure

A.7 End Body

A commented sentinel MUST mark the end of the body (with the sentinel comment as first character of the line) e.g.:

```
# EO_BODY_END
```

This sentinel is inviolate Vendors MUST NOT infer boundaries Everything before the sentinel is body content This sentinel may be followed by an sha

A.8 Optional Integrity Footer

EO/EA MAY include an integrity footer (e.g. sha256).

- integrity markers are OPTIONAL
- absence is not a failure
- if present, they MUST be verifiable

- users may embed their own assurance payloads in the body

Semantic equivalence, not byte identity, is the governing rule.

A.9 Session Durability and Rehydration

A.9.1 Session-Durable Default

All EO/EA are **session-durable** by default.

They:

- MUST persist for the full session
- MUST NOT disappear mid-session
- MAY be removed only by explicit human deletion or session termination

A.9.2 Rehydration

On session resumption:

- cached EO (including EA) are automatically rehydrated
- rehydration is functionally equivalent to pasting text into context
- rehydration does not assert, promote, or externalise

A.10 Time Semantics (Inviolate)

A.10.1 General Requirements

1. All Externalised Objects (EO) and Epistemic Attributes (EA) MUST carry real, platform-provided wall-clock time.
2. Time values MUST NOT be inferred, guessed, symbolic, relative, logical, session-indexed, or marked UNKNOWN.
3. Time values MUST be generated automatically by the platform using a real clock.
4. Humans MUST NOT directly edit, supply, or override time values.
5. Absence, mutation, or substitution of required time fields renders the artefact non-compliant with CM-2.

A.10.2 created_at — Absolute Wall-Clock Date/Time (Inviolate)

1. Every EO and EA MUST include a created_at field.
2. created_at MUST:
 - Represent the absolute wall-clock time of initial creation
 - Be expressed as an ISO-8601 date/time string
 - Include timezone information or to be explicitly UTC (Z)
3. created_at:
 - MUST be set once
 - MUST NOT change for the lifetime of the EO/EA
4. created_at anchors:
 - Object identity in time
 - Auditability across sessions and vendors
 - Distinction between rehydration and regeneration
5. Failure modes:
 - Missing created_at → non-compliant
 - Non-ISO format → non-compliant
 - Logical or inferred value → non-compliant

A.10.3 updated_at — Epoch Time (Inviolate)

(Normative — CM-2)

The *updated_at* field **MUST** be present on any Externalised Object (EO) or Epistemic Attribute (EA) whose body content has been changed due to a human-caused action.

The *updated_at* field **MUST**:

- be expressed as epoch time,
- use a clearly specified unit, either:
 - seconds since the Unix epoch, or
 - milliseconds since the Unix epoch,
- use a unit that is consistent within the artefact,
- preserve the unit and value exactly across export and import,
- strictly increase on each qualifying change.

The *updated_at* field **MUST NOT** be modified except as a direct consequence of a human-caused change to EO or EA body content.

The following events **MUST NOT** advance *updated_at*:

- rehydration,
- re-serialization,
- export or import,
- transport or copying,
- caching or cache eviction,
- storage migration,
- format normalisation,
- vendor-internal processing or optimisation,
- system restart or recovery.

Only a deliberate human action that changes EO or EA body content **MAY** advance *updated_at*.

Failure Modes (Non-Conformance)

Any of the following conditions **void CM-2 conformance**:

- *updated_at* is missing when body content has changed

A.10.4 Relationship Between *created_at* and *updated_at*

1. *created_at* and *updated_at* serve distinct and non-substitutable roles:
2. *created_at* anchors origin in absolute time
3. *updated_at* records temporal movement due to change
4. Neither field may be:
5. Derived from the other
6. Reconstructed post hoc
7. Replaced by a single combined timestamp

A thought that changes has moved in time; a thought that has not changed has not moved.

A.10.5 Portability and Interoperability Requirements

1. Vendors **MUST** preserve:
 - *created_at* ISO-8601 value exactly
 - *updated_at* epoch value exactly
 - Epoch unit semantics
2. Vendors **MUST NOT**:
 - Convert epoch time to formatted time
 - Convert formatted time to epoch
 - Normalise, round, or reinterpret timestamps
 - On import, non-conforming time representations **MUST** be rejected or quarantined.

Failure to meet these requirements voids CM-2 portability claims.

A.11 Govern, Durability, and Save As

(Normative — CM-2)

A.11.1 Govern

- Govern is an explicit human act.
- Govern marks EO and EA as durable-in-intent.
- Govern does not externalise content.

A.11.2 Save As (Inviolate)

- Save As \Rightarrow externally durable.
- Save As is the only mechanism that externalises EO and EA.
- Externalisation occurs outside the platform and client, under explicit human control.

A.11.3 Prohibitions

Platforms:

- **MUST NOT** externalise EO or EA without Save As,
- **MUST NOT** infer durability from session presence,
- **MUST NOT** silently persist EO or EA.

A.12 Epistemic States

(Normative, Not Mandatory Except Where Stated)

Epistemic Attributes (EA) MAY record epistemic states to support human governance and interpretation. Vendors and humans MAY define their own epistemic state vocabularies.

The following epistemic states are **required by CM-2 semantics** and MUST be representable in any CM-2-compliant system (labels MAY vary, semantics MUST NOT):

- **governed** — indicates that the EA has been explicitly governed by a human, as defined in §7.7 and A.10, establishing durable-in-intent context.
- **canonical** — indicates that the EA participates in material promoted via Save As, as defined in §7.7 and A.10.2, and is externally durable.
- **deprecated** — indicates that the EA constrains interpretation by marking material as no longer current or recommended, as defined in the lifecycle rules of CM-2.

The following epistemic states are **illustrative only**:

- session-durable
- draft
- parked

This list is not exhaustive. Additional epistemic states MAY be defined, provided they do not imply authority, truth, or durability beyond what is explicitly established by *Govern* and *Save As*.

A.13 Universes of Discourse (UoD)

A **Universe of Discourse (UoD)** is a session-local aggregation of EO/EA forming a thought bubble.

UoD:

- contains ideas, groundings, episodic recordings, dangling cognates
- may be rehydrated from turn text, uploads, or restored logic
- is analogous to a “thread” but explicitly externalised and governed
- does not assert authority by containment

A thread of thought MAY traverse multiple UoD and EO/EA.

A.14 Export, Recovery, and Auditability

A.14.1 Export

Vendors **MUST** provide a human-accessible export mechanism for EO/EA.

Exports:

- **MUST** preserve semantic identity and meaning
- **MUST** preserve UUID, timestamps, headers, and boundary sentinel
- **MUST NOT** rewrite, infer, normalise, or promote content

A.14.2 Transport Extensions

Vendors MAY offer additional transport mechanisms provided they:

- do not assert authority
- do not alter semantics
- preserve EO/EA identity and meaning

A.14.3 Auditability

EO/EA **MUST** remain auditable across vendors via:

- stable UUID
- timestamps
- semantic_version
- EO/EA header type
- boundary sentinel

Vendors **MUST** be able to show import/export provenance without asserting authority.

A.15 Session Tear-Down

- session-durable EO/EA **SHOULD** be expunged on session tear-down
- if automation is not possible, vendors **MUST** provide a human-visible cache-clearing mechanism

A.16 Vendor Conformance

(Normative — CM-2)

A CM-2-compliant vendor **MUST**:

- accept and emit Externalised Objects (EO) and Epistemic Attributes (EA) exactly as specified,
- rehydrate EO and EA automatically without promotion or alteration of epistemic state,
- preserve semantic equivalence across exchange, serialization, and re-ingestion,
- reject or quarantine non-compliant artefacts.

Vendor user interfaces, visual treatments, labels, icons, ordering, emphasis, or interaction affordances **MUST NOT** imply epistemic authority, correctness, endorsement, agreement, or durability beyond what has been explicitly established by a human action under CM-2.

In particular, vendors **MUST NOT**:

- present governed material as authoritative, approved, or validated,
- visually elevate content based on persistence, repetition, aggregation, or longevity,
- imply promotion or durability prior to an explicit *Save As*,
- infer authority from session presence, storage state, or system behaviour.

Only an explicit human action of *Save As* may justify durability or authoritative presentation.

Failure to comply with any requirement in this section breaks CM-2 portability and **voids all CM-2 compliance claims**.

A.17 Normative Closure

| Continuity is automatic. Authority is human. Durability is deliberate.

A.18 Additional EO invariants

1. Export obligation invariant (EOc → EOe)

The client-side MUST hold Epistemic Knowledge (e.g. EOC) where every cached EO must be exportable to EOe in a text-parseable CM-2-compliant form, with identity and provenance preserved.

2. No silent transformation invariant

No EO may be rewritten, normalised, summarised, reordered, merged, or split without an explicit, inspectable transition record (via EA and/or provenance delta). “Helpful formatting” is still a transformation.

3. Round-trip fidelity invariant

EOe exported from a platform, then re-imported, must rehydrate to an EO that is content-equivalent and identity-equivalent to the exported EOe (except where explicit transitions are recorded).

4. Boundary integrity invariant (Thought Bubble boundaries)

If Thought Bubbles are proposed/maintained, the system must preserve explicit boundary markers; cross-bubble leakage is disallowed unless the human explicitly links or merges bubbles.

5. Promotion gating invariant

Any transition that increases epistemic weight (e.g., “stabilised,” “durable,” “canonical,” “closed”) must be human-initiated and explicitly recorded; persistence, repetition, or reuse cannot trigger it.

6. Provenance immutability invariant

Once EOe is emitted, its provenance header fields are immutable; updates require a new versioned EOe with a recorded lineage link (no in-place mutation).

7. Loss signalling invariant

If an EO/Thought Bubble previously in the session surface leaves the inference window, the system must not behave as if it remains present; if it references it, it must rehydrate it or explicitly signal absence.

8. Minimality invariant for EA

EA may qualify lifecycle/scope/status, but may not inject new domain content into EO; EA is governance metadata, not a covert authoring channel.

9. Platform Context invariant

Context MUST NOT be defined or represented as an Epistemic Object (EO) or Epistemic Attribute (EA). Context MUST be a derived projection (view) over EO and EA within a session, and MUST NOT confer identity, authority, durability, epistemic state, or commitment.

Appendix B — Glossary: CM-2

Advisory	Providing guidance or examples without imposing requirements; advisory material may be ignored without violating CM-2.
Agent	A software construct that may act or operate autonomously; CM-2 does not grant agency to EO/EA or LLMs, and any agent behaviour remains subordinate to explicit human governance.
Artificial Intelligence (AI)	A broad class of computational systems designed to perform tasks associated with intelligent behaviour; in CM-2, AI systems are treated as capable assistants without inherent authority, memory, or governance rights.
Auditability	The ability to inspect, trace, and verify EO/EA identity, time, and provenance across sessions and vendors without implying authority.
Authority	The human-exclusive right to ground, accept, revise, promote, or discard knowledge; never inferred from model output, repetition, or continuity.
Authoritative	Having epistemic force conferred by explicit human action via Save As; authority is never implicit in CM-2.

Biological Cognitive Memoisation (BCM)	The natural human biological processes of memory, recall, association, and learning; implicit, fallible, and non-auditable, contrasted with CM's explicit externalisation.
Cache	A non-authoritative storage area (typically client-side) used to hold EO/EA for continuity and rehydration without implying memory, authority, or durability.
Cognition	The human processes of thinking, reasoning, imagining, remembering, and sense-making; CM externalises selected cognition without replacing it.
Completion	A model-generated continuation or inference that is logically coherent but not requested, grounded, or relevant to human intent.
Concept	A unit of meaning recognised by a human; concepts may be grounded, provisional, or unresolved and externalised as EO.
Constraint	A rule or limitation restricting permissible actions or interpretations without implying truth or authority.
Continuity	The ability to recognise and reuse prior material across interactions without rediscovery, distinct from authority or commitment.
Conversation	The human-perceived flow of interaction across turns; not a unit of governance, persistence, or authority in CM-2.
Dangling Cognate (DC)	An explicitly unresolved concept externalised as an EO to preserve ambiguity without forcing premature grounding.
Descriptive	Explaining how things are or may behave; descriptive text does not confer permission, obligation, or authority.
Durable	The state of an EO/EA that has been explicitly externalised via Save As into human-controlled storage beyond any single session.
Epistemic	Relating to knowledge, belief, justification, and certainty; CM-2 governs epistemic status explicitly rather than by inference.
Epistemic Attribute (EA)	A semantic subtype of EO carrying only epistemic, lifecycle, scope, or governance information, without domain meaning.
Epistemic Object	A unit of externalised knowledge treated as an object of reasoning and governance; realised concretely in CM-2 as EO.
Epistemic States	Human-declared markers describing epistemic status (e.g. draft, governed, canonical, deprecated); normative where mandated, otherwise optional.
Externalisation	The act of capturing cognition as EO so it can persist, be reasoned over, audited, and governed independently of memory.
Govern	An explicit human act that places EO or EA under deliberate governance, marking them as durable-in-intent without conferring authority or externalising them.
Governance	Explicit human-controlled rules and actions determining how cognition is externalised, interpreted, promoted, persisted, revised, or discarded.
Groundhog Day (Failure Mode)	Repeated rediscovery of prior work due to lack of durable, governed continuity.
Hallucination	Model-generated content that appears plausible but lacks grounding in governed EO or human intent.
Human	The sole epistemic authority in CM-2, responsible for meaning, judgement, governance, promotion, revision, and discard.
Idempotent	A property whereby repeated application of an action produces the same result; EO identity is idempotent unless explicitly changed.
Implicit	Assumed or inferred without explicit declaration; implicit authority, persistence, or meaning is invalid in CM-2.
Inference	A reasoning process deriving conclusions from available material; permitted but subordinate to human authority.
Inviolate	An invariant admitting no exception, inference, or downgrade; violation renders artefacts or platforms non-compliant.
Intuition	A human cognitive signal or felt sense guiding exploration; may be externalised as provisional EO without evidentiary force.
Knowledge	Cognition that has been explicitly governed and, where required, promoted via Save As into durable form.
Large Language Model (LLM)	A probabilistic text-generation system capable of reasoning and transformation but inherently non-authoritative and stateless.
Memory	The capacity to store and recall information; in CM-2 refers either to BCM or opaque system persistence, both distinct from EO.
Memoisation (Computing)	Caching computation results to avoid recomputation; CM uses the term to denote non-authoritative continuity caching.
Normative	Stating rules, expectations, or standards of conformance; normative text defines what compliant systems MUST or SHOULD do.
Ontology	A formal, fixed schema of concepts and relationships; explicitly avoided by CM-2 to prevent premature closure and lock-in.
Parse (Parsing)	To read and structurally interpret text to identify components without inferring meaning or authority.

Persistent	A general term meaning retained over time; intentionally avoided in CM-2 because it may imply opaque or unguided retention.
Portable	Capable of being exported, exchanged, and re-ingested across vendors while preserving semantic identity and meaning.
Prescriptive	Directing how something must be done procedurally; CM-2 minimises prescriptive rules.
Rehydration	Automatic restoration of cached EO/EA into a session context, equivalent to restoring text into working context.
Relationship	A human-declared connection between concepts or EO; context-bound and not global truth.
Save As	The sole human action that externalises EO/EA into durable storage; Save As ⇒ durable.
Semantic / Semantics	Concerned with meaning rather than form; CM-2 prioritises semantic equivalence over byte identity.
Semantic Equivalence	Preservation of meaning, identity, and epistemic status across representations and vendors.
Serialise (Serialisation)	Conversion of cognition or artefacts into linear, transportable text representation.
Session	A bounded interaction interval during which EO/EA are session-durable and available for reasoning.
Session-Durable	The default lifetime of EO/EA, persisting for the full duration of a session unless deleted or terminated.
Thread	A sequence of reasoning that may traverse multiple EO, EA, and UoD; not a unit of governance.
Thought	A unit of human cognition that may be transient or externalised as EO.
Thought Bubble	An informal synonym for a Universe of Discourse.
Tokenisation	Segmentation of text into tokens for computational processing; EO/EA bodies must be tokenisable.
Turn	A single exchange unit within a session that may introduce or modify EO/EA.
Universe of Discourse (UoD)	A bounded aggregation of EO/EA forming a coherent context for reasoning.
Vendor	An organisation or platform provider offering AI systems or tooling; must not imply authority or enclose EO/EA.
Vendor Neutrality	The requirement that CM-2 artefacts do not depend on proprietary schemas or platform behaviour.

categories

See [https://publications.arising.com.au/pub/Cognitive_Memoisation_\(CM-2\)_for_Governing_Knowledge_in_Human-AI_Collaboration#categories](https://publications.arising.com.au/pub/Cognitive_Memoisation_(CM-2)_for_Governing_Knowledge_in_Human-AI_Collaboration#categories)

Retrieved from 'https://publications.arising.com.au:443/pub-dir/index.php?title=Cognitive_Memoisation_(CM-2)_Protocol&oldid=25483'

■