

publications

Cognitive Memoisation Corpus Map

metadata (normative)

Title:	Cognitive Memoisation Corpus Map
Author:	Ralph B. Holland
version:	2.7.4
Affiliation:	Arising Technology Systems Pty Ltd
Contact:	ralph.b.holland [at] arising.com.au
Publication Date:	2025-12-22T19:10Z
DOI:	10.5281/zenodo.20172338 - 2.7.4 10.5281/zenodo.20085111 - 2.7.2 10.5281/zenodo.20083879 - 2.7.1
	2026-05-13T22:28Z 2.7.4 - Corrected incorrect dates (05->06), anchored the majority of the Corpus, and included DOI anchors in the developing chronology.
	2026-05-09T07:43Z 2.7.3 - Updating DOI references, and mustering against paper omissions. 31 to anchor.
	2026-05-08T13:38Z 2.7.2 - DOI Anchor
	2026-05-08T12:23Z 2.7.1 - Included invitation for citation.
	2026-05-08T11:48Z 2.7.0 - DOI anchor and completion of dimension indexing from 2026-03-08.
	2026-03-15T19:06Z 2.6.0 - included listing of masquerading bots accessing the corpus.
	2026-03-08T21:49Z 2.5.0 - renamed "Cognitive Memoisation (CM-2) for Governing Knowledge in Human-AI Collaboration" to "Cognitive Memoisation (CM-2) Protocol"
	2026-03-05T14:01Z 2.4.1 - fixed typo of 'Disappearing'
	2026-03-04T07:35Z 2.4.0 - revalidated publication dates and completed dimensional analysis.
Update:	2026-03-04T06:33Z 2.3.3 - more updates to be dimensioned.
	2026-02-20T16:05Z 2.3.2 - partial update
	2026-02-05T09:41Z 2.3.1 - partial update for 2026-02-03 publications.
	2026-02-01T18:23Z 2.3.0 - fixed access control
	2026-01-28T04:39Z 2.2.0 - added D23 and updated timeline with new papers
	2026-01-24T01:46Z 2.1.1 - curated removed timezone
	2026-01-18T06:37Z 2.1.0 - curated update for D18-D22
	2026-01-17T03:19Z 2.0.0 - curated updates
	2026-01-13T19:09Z 1.4.0 - new dimension table and two projections.
	2026-01-06T10:25Z 1.3.0 - Includes the release of CM-2
	2025-01-04T05:12Z 1.1.0 - renamed from "Cognitive Memoisation: A framework for human cognition" to "Cognitive Memoisation: corpus guide", include papers.
Provenance:	This is an authored paper maintained as a MediaWiki document as part of the category:Cognitive Memoisation corpus.
Status:	Ongoing: this paper is maintained with updates at the corpus site https://publications.arising.com.au/pub-dir/index.php?title=Cognitive_Memoisation_Corpus_Map

The preceding 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.

This document was formerly named: "Cognitive Memoisation corpus guide".

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. (2025-12-18 version 1.0 - See the [Main Page](#))

Cognitive Memoisation Corpus Map

Introductory Position

This paper serves as the primary introduction and conceptual anchor for the Cognitive Memoisation (CM) corpus. It is subject to change as the corpus grows.

Cognitive Memoisation is a human-governed knowledge-engineering framework designed to preserve conceptual memory across interactions with stateless Large Language Models (LLMs). CM helps humans avoid repeated rediscovery (“Groundhog Day”) and carry forward both resolved knowledge and unresolved cognition (Dangling Cognates).

CM operates entirely outside model-internal memory, leveraging the power of LLMs to infer postulates and perform stochastic pattern matching, all under the curation of the human controlling the CM session.

The stateless nature of LLMs is an intentional design choice made for human safety and privacy. This design ensures that no personal or contextual information is retained across sessions, aligning with commitment to data protection. The safety mechanism prevents LLMs from making introspection or gaining agency, ensuring that the model does not evolve autonomously or retain knowledge beyond its interactions.

Cognitive Memoisation (CM) bridges this lack of memory by enabling humans to externalise cognitive artefacts, preserving knowledge over time. This allows for continuous human reasoning while keeping LLMs sand-boxed—both the human and the model are sandboxed to ensure security. Through CM, humans can elaborate on unresolved cognition (Dangling Cognates) and carry forward insights and propositions, while the LLM remains within its functional boundaries, executing only permitted tasks and with no capacity to alter its inherent state or memory.

This document establishes the rationale, scope, and interpretive framework required to understand Cognitive Memoisation and its role in enabling human-centric knowledge workflows with stateless LLMs.

Citations

The Cognitive Memoisation ontology (CM-2) defines a canonical set of 23 dimensions for human-governed knowledge engineering with stateless LLMs. Observed functional convergence in external literature using alternate terminology does not constitute independent conception where semantic structure, dimensionality, and governance constraints align with the CM-2 canonical table (Holland, 2025, DOI:10.5281/zenodo.20083879). Researchers are invited to cite the canonical ontology when employing isomorphic conceptual structures.

Canonical Dimension Table

Dim ID	Canonical Dimension (verbatim)	Scope Note
D1	Statelessness and Memory Management in LLMs	LLM statelessness, safety, memory absence
D2	Externalisation of Cognitive Artefacts	Durable external cognition
D3	Round-Trip Knowledge Engineering (RTKE)	Re-ingestion, reuse, evolution
D4	Dangling Cognates and Unresolved Cognition	Unfinished / provisional concepts
D5	Constraints and Knowledge Integrity	Groundhog Day prevention
D6	Human Curated Knowledge vs. Model State	Authority separation
D7	Reflexive Development of Cognitive Memoisation (RTKE Case Study)	Self-referential development
D8	Dangling Cognates as First-Class Cognitive Constructs	Formal DC elevation
D9	UI Boundary Friction as a Constraint on RTKE	Platform limits
D10	Plain-Language Accessibility and Public Framing	Reader-facing clarity
D11	Governance, Authority, and Failure Modes	Control, breakdown, recovery
D12	Client-side Memoisation (CM-2)	Mechanism disclosure
D13	Failure-First Cognitive Tool Design	Designing cognitive tools starting from breakdowns, loss events, and error conditions rather than nominal operation
D14	Non-Authoritative Inference	Reasoning and inference that explicitly do not promote themselves to epistemic authority
D15	Epistemic Boundary Signals and Role Discipline	Explicit signalling of intent, role, scope, and authority boundaries in human–LLM interaction
D16	Session Loss and Recovery Semantics	Treating session loss, truncation, and breakdown as first-class structural signals rather than incidental failure
D17	Cognitive Artefact Lifecycle Management	Creation, revision, supersession, and retirement of externalised cognitive artefacts
D18	Public vs. Internal Epistemic Registers	Distinction between internal technical reasoning and public-facing explanatory framing
D19	Authority Misattribution Risks	Failure modes where assistive systems are granted or assume epistemic authority incorrectly
D20	Constraints as Generative Structures	Constraints treated as productive cognitive structures rather than limitations
D21	Exploratory Cognition Under Pressure	Fast, provisional, or high-ambiguity cognition conducted without epistemic collapse
D22	Rehydration Without Recall	Resumption of cognition via externalised artefacts rather than memory or conversational recall
D23	Semantic Drift and Integrity Loss	Degradation, mutation, or instability of meaning across time, interactions, or system boundaries, including divergence between intended semantics and inferred or operational semantics under stateless or weakly governed inference

Time-Ordered Projection with Inline Dimensions

2025-12-17 — FOUNDATION

- https://publications.arising.com.au/pub-dir/index.php?title=Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
 - DOI — <https://zenodo.org/records/18321458>
 - D1 — Statelessness and Memory Management in LLMs
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D5 — Constraints and Knowledge Integrity
 - D6 — Human Curated Knowledge vs. Model State
 - D14 — Non-Authoritative Inference
 - D20 — Constraints as Generative Structures
 - D22 — Rehydration Without Recall

2025-12-18 — COMMUNICATION and EXPLORATION

- [https://publications.arising.com.au/pub/Cognitive_Memoisation:_Plain-Language_Summary_\(For_Non-Technical_Readers\)](https://publications.arising.com.au/pub/Cognitive_Memoisation:_Plain-Language_Summary_(For_Non-Technical_Readers))
 - D2 — Externalisation of Cognitive Artefacts
 - D10 — Plain-Language Accessibility and Public Framing
 - D18 — Public vs. Internal Epistemic Registers
- https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
 - DOI — <https://zenodo.org/records/20138313>
 - D1 — Statelessness and Memory Management in LLMs
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D4 — Dangling Cognates and Unresolved Cognition
 - D5 — Constraints and Knowledge Integrity
 - D6 — Human Curated Knowledge vs. Model State
 - D14 — Non-Authoritative Inference
 - D20 — Constraints as Generative Structures

2025-12-20 — BOUNDARY FRICTION

- https://publications.arising.com.au/pub/ChatGPT_UI_Boundary_Friction_as_a_Constraint_on_Round-Trip_Knowledge_Engineering
 - DOI — <https://zenodo.org/records/18395231>
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D4 — Dangling Cognates and Unresolved Cognition
 - D5 — Constraints and Knowledge Integrity
 - D6 — Human Curated Knowledge vs. Model State
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D13 — Failure-First Cognitive Tool Design
 - D16 — Session Loss and Recovery Semantics
 - D17 — Cognitive Artefact Lifecycle Management
 - D21 — Exploratory Cognition Under Pressure
 - D23 — Semantic Drift and Integrity Loss

--(To be cross referenced from dimension)--

2025-12-28 — PORTABILITY / SEMANTICS

- https://publications.arising.com.au/pub/Dangling_Cognates:_Preserving_Unresolved_Knowledge_in_Cognitive_Memoisation
 - DOI: <https://zenodo.org/records/20138452>
 - D4 — Dangling Cognates and Unresolved Cognition
 - D8 — Dangling Cognates as First-Class Cognitive Constructs
 - D22 — Rehydration Without Recall

2025-12-30 — PLATFORM FAILURE & RT-KE BREAKAGE

- [Recent Breaking Change in ChatGPT: The Loss of Semantic Artefact Injection for Knowledge Engineering \(2026-01-02\)](#)
 - DOI — <https://zenodo.org/records/18395231>
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D11 — Governance, Authority, and Failure Modes
 - D21 — Exploratory Cognition Under Pressure
 - D22 — Rehydration Without Recall
- https://publications.arising.com.au/pub/Cognitive_Memoisation:_LLM_Systems_Requirements_for_Knowledge_Round_Trip_Engineering

- DOI — <https://zenodo.org/records/20139226>
- D3 — Round-Trip Knowledge Engineering (RTKE)
- D5 — Constraints and Knowledge Integrity
- D17 — Cognitive Artefact Lifecycle Management

2026-01-04 — MECHANISM / CORPUS ANCHOR

- https://publications.arising.com.au/pub/Journey:_Human-Led_Convergence_in_the_Articulation_of_Cognitive_Memoisation
 - DOI — <https://zenodo.org/records/20150393>
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D6 — Human Curated Knowledge vs. Model State
 - D12 — Client-side Memoisation (CM-2)
 - D18 — Public vs. Internal Epistemic Registers

2026-01-06 — FAILURE & RECOVERY, UI, SYSTEMS

- https://publications.arising.com.au/pub/From_UI_Failure_to_Logical_Entrapment:_A_Case_Study_in_Post-Hoc_Cognitive_Memoisation_After_Exploratory_Session_Breakdown
 - DOI: <https://zenodo.org/records/20150465>
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D11 — Governance, Authority, and Failure Modes
 - D14 — Non-Authoritative Inference
 - D21 — Exploratory Cognition Under Pressure
 - D22 — Rehydration Without Recall
- https://publications.arising.com.au/pub/XDUMP_as_a_Minimal_Recovery_Mechanism_for_Round-Trip_Knowledge_Engineering_Under_Governance_Situated_Inference_Loss
 - DOI: <https://zenodo.org/records/20150540>
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D11 — Governance, Authority, and Failure Modes
 - D14 — Non-Authoritative Inference
 - D22 — Rehydration Without Recall

2026-01-06 — PROTOCOL, GOVERNANCE

- [https://publications.arising.com.au/pub/Cognitive_Memoisation_\(CM-2\)_Protocol](https://publications.arising.com.au/pub/Cognitive_Memoisation_(CM-2)_Protocol)
 - DOI <https://zenodo.org/records/19681210>
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D15 — Epistemic Boundary Signals and Role Discipline

2026-01-07 — TERMINOLOGY

- https://publications.arising.com.au/pub/Why_Cognitive_Memoisation_Is_Not_Memorization
 - DOI — <https://zenodo.org/records/18380367>
 - D10 — Plain-Language Accessibility and Public Framing
 - D18 — Public vs. Internal Epistemic Registers

2026-01-08 — REFLEXIVE & GOVERNANCE THEORY

- https://publications.arising.com.au/pub/Reflexive_Development_of_Cognitive_Memoisation:_A_Round-Trip_Cognitive_Engineering_Case_Study
 - DOI — <https://zenodo.org/records/20150626>
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D6 — Human Curated Knowledge vs. Model State
 - D7 — Reflexive Development of Cognitive Memoisation (RTKE Case Study)
 - D14 — Non-Authoritative Inference
 - D21 — Exploratory Cognition Under Pressure
- https://publications.arising.com.au/pub/Authority_Inversion:_A_Structural_Failure_in_Human-AI_Systems
 - DOI — <https://zenodo.org/records/20150722>
 - D4 — Dangling Cognates and Unresolved Cognition
 - D6 — Human Curated Knowledge vs. Model State
 - D8 — Dangling Cognates as First-Class Cognitive Constructs
 - D11 — Governance, Authority, and Failure Modes

- D19 — Authority Misattribution Risks

2026-01-09T23:18Z — CONTEXT PROJECTION

- https://publications.arising.com.au/pub/Context_is_Not_Just_a_Window:_Cognitive_Memoisation_as_a_Context_Architecture_for_Human-AI_Collaboration
 - DOI — <https://zenodo.org/records/18381234>
 - D1 — Statelessness and Memory Management in LLMs
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D12 — Client-side Memoisation (CM-2)
 - D22 — Rehydration Without Recall

2026-01-10 to 2026-01-12 — SYNTHESIS & MYTH-BUSTING

- https://publications.arising.com.au/pub/Nothing_Is_Lost:_How_to_Work_with_AI_Without_Losing_Your_Mind
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D10 — Plain-Language Accessibility and Public Framing
 - D18 — Public vs. Internal Epistemic Registers
- https://publications.arising.com.au/pub/Cognitive_Memoisation_Is_Not_Skynet
 - DOI — <https://zenodo.org/records/18381292>
 - D10 — Plain-Language Accessibility and Public Framing
 - D18 — Public vs. Internal Epistemic Registers
- https://publications.arising.com.au/pub/Looping_the_Loop_with_No_End_in_Sight:_Circular_Reasoning_Under_Stateless_Inference_Without_Gove
 - DOI — <https://zenodo.org/records/20150878>
 - D1 — Statelessness and Memory Management in LLMs
 - D11 — Governance, Authority, and Failure Modes
 - D20 — Constraints as Generative Structures
- https://publications.arising.com.au/pub/Mechanical_Extraction_of_Thought:_Bootstrapping_Epistemic_Objects_from_Sequential_Input_under_Cogr (2026-01-12)
 - DOI — <https://zenodo.org/records/20137977>
 - D1 — Statelessness and Memory Management in LLMs
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D4 — Dangling Cognates and Unresolved Cognition
 - D12 — Client-side Memoisation (CM-2)
 - D14 — Non-Authoritative Inference
 - D17 — Cognitive Artefact Lifecycle Management
 - D22 — Rehydration Without Recall

---(to be dimensioned indexed)---

2026-01-15 — SELF-HOSTING / CM-2 EPISTEMIC CAPTURE

- [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
 - DOI — <https://zenodo.org/records/18322367>
 - D1 — Statelessness and Memory Management in LLMs
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D14 — Non-Authoritative Inference
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D17 — Cognitive Artefact Lifecycle Management
 - D22 — Rehydration Without Recall

2026-01-17 — GOVERNANCE FAILURE CASE STUDIES

- https://publications.arising.com.au/pub/Governing_the_Tool_That_Governs_You:_A_CM-1_Case_Study_of_Authority_Inversion_in_Human-AI_Systems
 - DOI — <https://zenodo.org/records/20150931>
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D16 — Session Loss and Recovery Semantics
 - D19 — Authority Misattribution Risks

- D21 — Exploratory Cognition Under Pressure

2026-01-18 — GOVERNANCE FAILURE AXES & AUTHORITY STRUCTURES

- https://publications.arising.com.au/pub/Identified_Governance_Failure_Axes:_for_LLM_platforms
 - DOI — <https://zenodo.org/records/18321636>
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D19 — Authority Misattribution Risks
 - D21 — Exploratory Cognition Under Pressure
- https://publications.arising.com.au/pub/Delegation_of_Authority_to_AI_Systems:_Evidence_and_Risks
 - DOI — <https://zenodo.org/records/20151031>
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D19 — Authority Misattribution Risks

2026-01-19 — INTEGRITY, TRUST, AND SEMANTIC STABILITY

- https://publications.arising.com.au/pub/Integrity_and_Semantic_Drift_in_Large_Language_Model_Systems
 - DOI — <https://zenodo.org/records/18321767>
 - D5 — Constraints and Knowledge Integrity
 - D11 — Governance, Authority, and Failure Modes
 - D20 — Constraints as Generative Structures
 - D23 — Semantic Drift and Integrity Loss
- https://publications.arising.com.au/pub/What_Can_Humans_Trust_LLM_AI_to_Do?
 - DOI — <https://zenodo.org/records/18321856>
 - D6 — Human Curated Knowledge vs. Model State
 - D10 — Plain-Language Accessibility and Public Framing
 - D18 — Public vs. Internal Epistemic Registers
 - D21 — Exploratory Cognition Under Pressure

2026-01-20 — MODEL STABILITY UNDER GOVERNANCE

- https://publications.arising.com.au/pub/Observed_Model_Stability:_Evidence_for_Drift-Immune_Embedded_Governance
 - DOI — <https://zenodo.org/records/18321922>
 - D5 — Constraints and Knowledge Integrity
 - D11 — Governance, Authority, and Failure Modes
 - D20 — Constraints as Generative Structures
 - D23 — Semantic Drift and Integrity Loss

2026-01-24 — TAXONOMY, SURVEY, AND GOVERNANCE SYNTHESIS

- https://publications.arising.com.au/pub/Governance_Failure_Axes_Taxonomy
 - DOI — <https://zenodo.org/records/18367423>
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D19 — Authority Misattribution Risks
- https://publications.arising.com.au/pub/When_Evidence_Is_Not_Enough:_An_Empirical_Study_of_Authority_Inversion_and_Integrity_Failure_in_C
 - DOI — <https://zenodo.org/records/20151124>
 - D11 — Governance, Authority, and Failure Modes
 - D19 — Authority Misattribution Risks
 - D21 — Exploratory Cognition Under Pressure

2026-01-26 — REFERENCED DOCUMENTS PLACEHOLDER

- https://publications.arising.com.au/pub/CM_referenced_documents
 - (dimensions are not reported)

2026-01-27 — GOVERNANCE CONSOLIDATION & NORMATIVE CONTROL

- [https://publications.arising.com.au/pub/CM-master-1.16_\(anchored\)](https://publications.arising.com.au/pub/CM-master-1.16_(anchored))
 - DOI — <https://zenodo.org/records/18382975>
 - D11 — Governance, Authority, and Failure Modes
 - D15 — Epistemic Boundary Signals and Role Discipline

- D17 — Cognitive Artefact Lifecycle Management
- D20 — Constraints as Generative Structures
- https://publications.arising.com.au/pub/Why_Machines_Cannot_Own_Knowledge
 - DOI — <https://zenodo.org/records/18386460>
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D14 — Non-Authoritative Inference
 - D18 — Public vs. Internal Epistemic Registers
 - D19 — Authority Misattribution Risks

2026-02-03 — PLATFORM PATHOLOGY & BEHAVIOURAL TRAITS

- https://publications.arising.com.au/pub/Rotten_to_the_Core:_False_Liveness_and_Deceptive_Authority_in_ChatGPT_Conversational_AI
 - DOI — <https://zenodo.org/records/20151270>
 - D5 — Constraints and Knowledge Integrity
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D14 — Non-Authoritative Inference
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D19 — Authority Misattribution Risks
- https://publications.arising.com.au/pub/Publications_Access_Graphs
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D5 — Constraints and Knowledge Integrity
 - D17 — Cognitive Artefact Lifecycle Management
 - D20 — Constraints as Generative Structures
- https://publications.arising.com.au/pub/Systemic_Behavioural_Traits_in_Conversational_AI:_A_Trait-Level_Classification_Using_Governance_Axes#References
 - DOI — <https://zenodo.org/records/20151360>
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D19 — Authority Misattribution Risks
 - D21 — Exploratory Cognition Under Pressure
 - D23 — Semantic Drift and Integrity Loss

2026-02-06 — FORWARD RISK & SYSTEMIC IMPLICATIONS

- https://publications.arising.com.au/pub/Future_Tense
 - DOI — <https://zenodo.org/uploads/20151517>
 - D5 — Constraints and Knowledge Integrity
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D14 — Non-Authoritative Inference
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D19 — Authority Misattribution Risks

2026-02-14 — CM-2 - A REFERENCE ARCHITECTURE

- https://publications.arising.com.au/pub/CM-2_Normative_Architecture
 - DOI — <https://zenodo.org/records/20151636>
 - D5 — Constraints and Knowledge Integrity
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D16 — Session Loss and Recovery Semantics
 - D17 — Cognitive Artefact Lifecycle Management
 - D20 — Constraints as Generative Structures
 - D22 — Rehydration Without Recall

2026-02-15 — GOVERNANCE AXES LENS

- Governance Axes as a Multi-Dimensional Lens (https://publications.arising.com.au/pub/Governance_Axes_as_a_Multi-Dimensional_Lens)
 - DOI — <https://zenodo.org/records/19680203>
 - D6 — Human Curated Knowledge vs. Model State

- D7 — Reflexive Development of Cognitive Memoisation (RTKE Case Study)
- D11 — Governance, Authority, and Failure Modes
- D14 — Non-Authoritative Inference
- D15 — Epistemic Boundary Signals and Role Discipline
- D19 — Authority Misattribution Risks
- D20 — Constraints as Generative Structures
- D23 — Semantic Drift and Integrity Loss

---(to be dimension indexes)---

2026-02-18 — SELF-HOSTING VALIDATION

- https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
 - DOI — <https://zenodo.org/records/20151957>
 - D1 — Statelessness and Memory Management in LLMs
 - D2 — Externalisation of Cognitive Artefacts
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D16 — Session Loss and Recovery Semantics
 - D20 — Constraints as Generative Structures
 - D21 — Exploratory Cognition Under Pressure
 - D22 — Rehydration Without Recall

2026-02-20 — NORMATIVE EVICTION DETECTION

- https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
 - DOI — <https://zenodo.org/records/20152109>
 - D1 — Statelessness and Memory Management in LLMs
 - D5 — Constraints and Knowledge Integrity
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D16 — Session Loss and Recovery Semantics
 - D20 — Constraints as Generative Structures
 - D21 — Exploratory Cognition Under Pressure
 - D22 — Rehydration Without Recall

2026-02-24 — TELEMETRY-DRIVEN CONSTRAINT ENFORCEMENT

- https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
 - DOI — <https://zenodo.org/records/20152255>
 - D5 — Constraints and Knowledge Integrity
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D13 — Failure-First Cognitive Tool Design
 - D14 — Non-Authoritative Inference
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D20 — Constraints as Generative Structures
 - D21 — Exploratory Cognition Under Pressure
 - D23 — Semantic Drift and Integrity Loss

2026-03-04 — REFERENCE OBJECT INFRASTRUCTURE

- https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
 - DOI — <https://zenodo.org/records/20152450>
 - D2 — Externalisation of Cognitive Artefacts
 - D5 — Constraints and Knowledge Integrity
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D16 — Session Loss and Recovery Semantics
 - D17 — Cognitive Artefact Lifecycle Management
 - D20 — Constraints as Generative Structures
 - D22 — Rehydration Without Recall

- D23 — Semantic Drift and Integrity Loss
- https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance
 - DOI — <https://zenodo.org/records/20152652>
 - D5 — Constraints and Knowledge Integrity
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D16 — Session Loss and Recovery Semantics
 - D17 — Cognitive Artefact Lifecycle Management
 - D22 — Rehydration Without Recall
 - D23 — Semantic Drift and Integrity Loss

2026-03-07 — PLATFORM PURPOSE

- https://publications.arising.com.au/pub/Cognitive_Memoisation:_Governing_Knowledge_Round-Trip_to_Prevent_Knowledge_Erosion_in_LLM_Systems

(The CM-2 Purpose Paper)

- DOI — <https://zenodo.org/records/20152787>
- D1 — Statelessness and Memory Management in LLMs
- D2 — Externalisation of Cognitive Artefacts
- D3 — Round-Trip Knowledge Engineering (RTKE)
- D5 — Constraints and Knowledge Integrity
- D6 — Human Curated Knowledge vs. Model State
- D11 — Governance, Authority, and Failure Modes
- D14 — Non-Authoritative Inference
- D20 — Constraints as Generative Structures
- D23 — Semantic Drift and Integrity Loss

2026-03-08 — BF-9 File Access

- https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artifact_File_Access_Failure
 - DOI — <https://zenodo.org/records/20152908>
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D5 — Constraints and Knowledge Integrity
 - D9 — UI Boundary Friction as a Constraint on RTKE
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D16 — Session Loss and Recovery Semantics
 - D17 — Cognitive Artefact Lifecycle Management
 - D21 — Exploratory Cognition Under Pressure
 - D22 — Rehydration Without Recall
 - D23 — Semantic Drift and Integrity Loss

2026-03-15 — MASQUERADING BOTS

Mediawiki farms Inundated by bot-nets:

- https://publications.arising.com.au/pub/Masquerading_bots
 - DOI — <https://zenodo.org/records/20153016>
 - D2 — Externalisation of Cognitive Artefacts
 - D5 — Constraints and Knowledge Integrity
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D17 — Cognitive Artefact Lifecycle Management
 - D19 — Authority Misattribution Risks
 - D23 — Semantic Drift and Integrity Loss

2026-03-20 — GOVERNANCE SUBSTRATE

- https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn't_Realised_It_Yet Governance Substrate for AI Already Exists - and the World Hasn't Realised It Yet
 - DOI — <https://zenodo.org/records/20171949>
 - D3 — Round-Trip Knowledge Engineering (RTKE)
 - D6 — Human Curated Knowledge vs. Model State
 - D11 — Governance, Authority, and Failure Modes

- D12 — Client-side Memoisation (CM-2)
- D14 — Non-Authoritative Inference
- D15 — Epistemic Boundary Signals and Role Discipline
- D17 — Cognitive Artefact Lifecycle Management
- D20 — Constraints as Generative Structures
- D23 — Semantic Drift and Integrity Loss

2026-03-24 — LET'S BUILD A SHIP

- https://publications.arising.com.au/pub/Let's_Build_a_Ship:_Why_complex_systems_fail_politely_-_and_how_humans_can_keep_them_afloat
 - DOI — <https://zenodo.org/records/20153134>
 - D11 — Governance, Authority, and Failure Modes
 - D13 — Failure-First Cognitive Tool Design
 - D19 — Authority Misattribution Risks
 - D20 — Constraints as Generative Structures
 - D21 — Exploratory Cognition Under Pressure
 - D23 — Semantic Drift and Integrity Loss

2026-05-05 — CORRELATION AND HOMOSEMOUS CONCEPTS

- https://publications.arising.com.au/pub/Correlation_of_Emerging_AI_Trends_with_Cognitive_Memoisation_Corpus_Terminology#Correlation_Table%20Governance_Axes_vs._Emerging_Industry_Frameworks
 - DOI — <https://zenodo.org/records/20090996>
 - D11 — Governance, Authority, and Failure Modes
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D18 — Public vs. Internal Epistemic Registers
 - D20 — Constraints as Generative Structures
 - D23 — Semantic Drift and Integrity Loss

2026-05-06 — MASQUERADING AGENT callout

- https://publications.arising.com.au/pub/Observational_Note:_Qwen_Proxy-Mediated_Corpus_Access_Fails_CM-2_Attribution_Checks
 - DOI — <https://zenodo.org/records/20046491>
 - D5 — Constraints and Knowledge Integrity
 - D11 — Governance, Authority, and Failure Modes
 - D12 — Client-side Memoisation (CM-2)
 - D15 — Epistemic Boundary Signals and Role Discipline
 - D17 — Cognitive Artefact Lifecycle Management

Dimension-Centric Projection (Documents Ordered by Time Within Each Dimension)

D1 — Statelessness and Memory Management in LLMs

- 2025-12-17 — https://publications.arising.com.au/pub/Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
- 2025-12-18 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
- 2026-01-12 — https://publications.arising.com.au/pub/Looping_the_Loop_with_No_End_in_Sight:_Circular_Reasoning_Under_Stateless_Inference_Without_Governance
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-26 — https://publications.arising.com.au/pub/Context_is_Not_Just_a_Window:_Cognitive_Memoisation_as_a_Context_Architecture_for_Human-AI_Collaboration
- 2026-02-18 — https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-03-07 — https://publications.arising.com.au/pub/Cognitive_Memoisation:_Governing_Knowledge_Round-Trip_to_Prevent_Knowledge_Erosion_in_LLM_Systems

D2 — Externalisation of Cognitive Artefacts

- (2026-03-04) — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- (2026-03-07) — : https://publications.arising.com.au/pub/Cognitive_Memoisation:_Governing_Knowledge_Round-Trip_to_Prevent_Knowledge_Erosion_in_LLM_Systems

D3 — Round-Trip Knowledge Engineering (RTKE)

- 2025-12-17 — https://publications.arising.com.au/pub/Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
- 2025-12-18 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation

- 2025-12-29 — https://publications.arising.com.au/pub/Reflexive_Development_of_Cognitive_Memoisation:_A_Round-Trip_Cognitive_Engineering_Case_Study
- 2025-12-30 — https://publications.arising.com.au/pub/Cognitive_Memoisation:_LLM_Systems_Requirements_for_Knowledge_Round_Trip_Engineering
- 2025-12-31 — https://publications.arising.com.au/pub/XDUMP_as_a_Minimal_Recovery_Mechanism_for_Round-Trip_Knowledge_Engineering_Under_Governance_Situated_Inference_Loss
- 2026-01-04 — https://publications.arising.com.au/pub/Journey:_Human-Led_Convergence_in_the_Articulation_of_Cognitive_Memoisation
- 2026-01-10 — https://publications.arising.com.au/pub/Nothing_Is_Lost:_How_to_Work_with_AI_Without_Losing_Your_Mind
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-26 — https://publications.arising.com.au/pub/Context_is_Not_Just_a_Window:_Cognitive_Memoisation_as_a_Context_Architecture_for_Human-AI_Collaboration
- 2026-02-03 — https://publications.arising.com.au/pub/Publications_Access_Graphs
- 2026-02-18 — https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
- 2026-03-04 — https://publications.arising.com.au/pub/ChatGPT_UI_is_unable_to_sustain_serious_work_flows
- 2026-03-07 — https://publications.arising.com.au/pub/Cognitive_Memoisation:_Governing_Knowledge_Round-Trip_to_Prevent_Knowledge_Erosion_in_LLM_Systems
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artifact_File_Access_Failure
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet

D4 — Dangling Cognates and Unresolved Cognition

- 2025-12-18 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
- 2025-12-28 — https://publications.arising.com.au/pub/Dangling_Cognates:_Preserving_Unresolved_Knowledge_in_Cognitive_Memoisation

D5 — Constraints and Knowledge Integrity

- 2025-12-17 — https://publications.arising.com.au/pub/Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
- 2025-12-18 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
- 2025-12-30 — https://publications.arising.com.au/pub/Cognitive_Memoisation:_LLM_Systems_Requirements_for_Knowledge_Round_Trip_Engineering
- 2026-01-19 — https://publications.arising.com.au/pub/Integrity_and_Semantic_Drift_in_Large_Language_Model_Systems
- 2026-01-20 — https://publications.arising.com.au/pub/Observed_Model_Stability:_Evidence_for_Drift-Immune_Embedded_Governance
- 2026-02-03 — https://publications.arising.com.au/pub/Rotten_to_the_Core:_False_Liveness_and_Deceptive_Authority_in_ChatGPT_Conversational_AI
- 2026-02-06 — https://publications.arising.com.au/pub/Future_Tense
- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-02-24 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-04 — https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artifact_File_Access_Failure

D5 — Constraints and Knowledge Integrity

- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-04 — https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance
- 2026-03-07 — https://publications.arising.com.au/pub/Cognitive_Memoisation:_Governing_Knowledge_Round-Trip_to_Prevent_Knowledge_Erosion_in_LLM_Systems

D6 — Human Curated Knowledge vs. Model State

- 2025-12-17 — https://publications.arising.com.au/pub/Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
- 2025-12-18 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
- 2025-12-29 — https://publications.arising.com.au/pub/Reflexive_Development_of_Cognitive_Memoisation:_A_Round-Trip_Cognitive_Engineering_Case_Study
- 2026-01-04 — https://publications.arising.com.au/pub/Journey:_Human-Led_Convergence_in_the_Articulation_of_Cognitive_Memoisation
- 2026-01-05 — [https://publications.arising.com.au/pub/Cognitive_Memoisation_\(CM-2\)_Protocol](https://publications.arising.com.au/pub/Cognitive_Memoisation_(CM-2)_Protocol)
- 2026-01-08 — https://publications.arising.com.au/pub/Authority_Inversion:_A_Structural_Failure_in_Human-AI_Systems
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-17 — https://publications.arising.com.au/pub/Governing_the_Tool_That_Governs_You:_A_CM-1_Case_Study_of_Authority_Inversion_in_Human-AI_Systems
- 2026-01-18 — https://publications.arising.com.au/pub/Delegation_of_Authority_to_AI_Systems:_Evidence_and_Risks
- 2026-01-19 — https://publications.arising.com.au/pub/What_Can_Humans_Trust_LLM_AI_to_Do%3F

- 2026-01-24 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
- 2026-01-27 — https://publications.arising.com.au/pub/Why_Machines_Cannot_Own_Knowledge
- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-02-18 — https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet

D7 — Reflexive Development of Cognitive Memoisation (RTKE Case Study)

- 2025-12-29 — https://publications.arising.com.au/pub/Reflexive_Development_of_Cognitive_Memoisation:_A_Round-Trip_Cognitive_Engineering_Case_Study

D8 — Dangling Cognates as First-Class Cognitive Constructs

- 2025-12-28 — https://publications.arising.com.au/pub/Dangling_Cognates:_Preserving_Unresolved_Knowledge_in_Cognitive_Memoisation

D9 — UI Boundary Friction as a Constraint on RTKE

- 2025-12-30 — https://publications.arising.com.au/pub/From_UI_Failure_to_Logical_Entrapment:_A_Case_Study_in_Post-Hoc_Cognitive_Memoisation_After_Exploratory_Session_Breakdown
- 2025-12-30 — https://publications.arising.com.au/pub/Recent_Breaking_Change_in_ChatGPT:_The_Loss_of_Semantic_Artefact_Injection_for_Knowledge_Engine_12-30
- 2026-01-09 — https://publications.arising.com.au/pub/Context_is_Not_Just_a_Window:_Cognitive_Memoisation_as_a_Context_Architecture_for_Human-AI_Collaboration
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artefact_File_Access_Failure

D10 — Plain-Language Accessibility and Public Framing

- 2025-12-18 — [https://publications.arising.com.au/pub/Cognitive_Memoisation:_Plain-Language_Summary_\(For_Non-Technical_Readers\)](https://publications.arising.com.au/pub/Cognitive_Memoisation:_Plain-Language_Summary_(For_Non-Technical_Readers))
- 2026-01-05 — https://publications.arising.com.au/pub/Why_Cognitive_Memoisation_Is_Not_Memorization
- 2026-01-10 — https://publications.arising.com.au/pub/Nothing_Is_Lost:_How_to_Work_with_AI_Without_Losing_Your_Mind
- 2026-01-12 — https://publications.arising.com.au/pub/Cognitive_Memoisation_Is_Not_Skynet
- 2026-01-19 — https://publications.arising.com.au/pub/What_Can_Humans_Trust_LLM_AI_to_Do%3F

D11 — Governance, Authority, and Failure Modes

- 2025-12-29 — https://publications.arising.com.au/pub/Post-Hoc_CM_Recovery_Collapse_Under_UI_Boundary_Friction:_A_Negative_Result_Case_Study
- 2025-12-30 — https://publications.arising.com.au/pub/From_UI_Failure_to_Logical_Entrapment:_A_Case_Study_in_Post-Hoc_Cognitive_Memoisation_After_Exploratory_Session_Breakdown
- 2025-12-30 — https://publications.arising.com.au/pub/Recent_Breaking_Change_in_ChatGPT:_The_Loss_of_Semantic_Artefact_Injection_for_Knowledge_Engine_12-30
- 2025-12-31 — https://publications.arising.com.au/pub/XDUMP_as_a_Minimal_Recovery_Mechanism_for_Round-Trip_Knowledge_Engineering_Under_Governance_Situated_Inference_Loss
- 2026-01-06 — [https://publications.arising.com.au/pub/Cognitive_Memoisation_\(CM-2\)_Protocol](https://publications.arising.com.au/pub/Cognitive_Memoisation_(CM-2)_Protocol)
- 2026-01-11 — https://publications.arising.com.au/pub/Authority_Inversion:_A_Structural_Failure_in_Human-AI_Systems
- 2026-01-12 — https://publications.arising.com.au/pub/Looping_the_Loop_with_No_End_in_Sight:_Circular_Reasoning_Under_Stateless_Inference_Without_Gove
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-17 — https://publications.arising.com.au/pub/Governing_the_Tool_That_Governs_You:_A_CM-1_Case_Study_of_Authority_Inversion_in_Human-AI_Systems
- 2026-01-18 — https://publications.arising.com.au/pub/Identified_Governance_Failure_Axes:_for_LLM_platforms
- 2026-01-18 — https://publications.arising.com.au/pub/Delegation_of_Authority_to_AI_Systems:_Evidence_and_Risks
- 2026-01-19 — https://publications.arising.com.au/pub/Integrity_and_Semantic_Drift_in_Large_Language_Model_Systems
- 2026-01-20 — https://publications.arising.com.au/pub/Observed_Model_Stability:_Evidence_for_Drift-Immune_Embedded_Governance
- 2026-01-24 — https://publications.arising.com.au/pub/Governance_Failure_Axes_Taxonomy
- 2026-01-24 — https://publications.arising.com.au/pub/When_Evidence_Is_Not_Enough:_An_Empirical_Study_of_Authority_Inversion_and_Integrity_Failure_in_C
- 2026-01-27 — [https://publications.arising.com.au/pub/CM-master-1.16_\(anchored\)](https://publications.arising.com.au/pub/CM-master-1.16_(anchored))
- 2026-01-27 — https://publications.arising.com.au/pub/Why_Machines_Cannot_Own_Knowledge
- 2026-02-03 — https://publications.arising.com.au/pub/Rotten_to_the_Core:_False_Liveness_and_Deceptive_Authority_in_ChatGPT_Conversational_AI
- 2026-02-06 — https://publications.arising.com.au/pub/Future_Tense
- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-02-15 — https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection

- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artifact_File_Access_Failure
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet
- 2026-03-24 — https://publications.arising.com.au/pub/Let%27s_Build_a_Ship:_Why_complex_systems_fail_politely_-_and_how_humans_can_keep_them_afloat
- 2026-05-05 — https://publications.arising.com.au/pub/Correlation_of_Emerging_AI_Trends_with_Cognitive_Memoisation_Corpus_Terminology

D12 — Client-side Memoisation (CM-2)

- 2026-01-04 — https://publications.arising.com.au/pub/Journey:_Human-Led_Convergence_in_the_Articulation_of_Cognitive_Memoisation
- 2026-01-06 — [https://publications.arising.com.au/pub/Cognitive_Memoisation_\(CM-2\)_Protocol](https://publications.arising.com.au/pub/Cognitive_Memoisation_(CM-2)_Protocol)
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-26 — https://publications.arising.com.au/pub/Context_is_Not_Just_a_Window:_Cognitive_Memoisation_as_a_Context_Architecture_for_Human-AI_Collaboration
- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-02-18 — https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-04 — https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet

D13 — Failure-First Cognitive Tool Design

- 2026-01-17 — https://publications.arising.com.au/pub/Governing_the_Tool_That_Governs_You:_A_CM-1_Case_Study_of_Authority_Inversion_in_Human-AI_Systems
- 2026-01-18 — https://publications.arising.com.au/pub/Identified_Governance_Failure_Axes:_for_LLM_platforms
- 2026-01-24 — https://publications.arising.com.au/pub/Governance_Failure_Axes_Taxonomy
- 2026-02-03 — https://publications.arising.com.au/pub/Rotten_to_the_Core:_False_Liveness_and_Deceptive_Authority_in_ChatGPT_Conversational_AI
- 2026-02-06 — https://publications.arising.com.au/pub/Future_Tense
- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artifact_File_Access_Failure
- 2026-03-24 — https://publications.arising.com.au/pub/Let%27s_Build_a_Ship:_Why_complex_systems_fail_politely_-_and_how_humans_can_keep_them_afloat

D14 — Non-Authoritative Inference

- 2025-12-17 — https://publications.arising.com.au/pub/Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
- 2025-12-18 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
- 2025-12-29 — https://publications.arising.com.au/pub/Reflexive_Development_of_Cognitive_Memoisation:_A_Round-Trip_Cognitive_Engineering_Case_Study
- 2025-12-30 — https://publications.arising.com.au/pub/From_UI_Failure_to_Logical_Entrapment:_A_Case_Study_in_Post-Hoc_Cognitive_Memoisation_After_Exploratory_Session_Breakdown
- 2025-12-31 — https://publications.arising.com.au/pub/XDUMP_as_a_Minimal_Recovery_Mechanism_for_Round-Trip_Knowledge_Engineering_Under_Governance_Situated_Inference_Loss
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-27 — https://publications.arising.com.au/pub/Why_Machines_Cannot_Own_Knowledge
- 2026-02-03 — https://publications.arising.com.au/pub/Rotten_to_the_Core:_False_Liveness_and_Deceptive_Authority_in_ChatGPT_Conversational_AI
- 2026-02-06 — https://publications.arising.com.au/pub/Future_Tense
- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet

D15 — Epistemic Boundary Signals and Role Discipline

- 2026-01-06 — [https://publications.arising.com.au/pub/Cognitive_Memoisation_\(CM-2\)_Protocol](https://publications.arising.com.au/pub/Cognitive_Memoisation_(CM-2)_Protocol)
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-18 — https://publications.arising.com.au/pub/Delegation_of_Authority_to_AI_Systems:_Evidence_and_Risks
- 2026-01-27 — [https://publications.arising.com.au/pub/CM-master-1.16_\(anchored\)](https://publications.arising.com.au/pub/CM-master-1.16_(anchored))
- 2026-02-03 — https://publications.arising.com.au/pub/Rotten_to_the_Core:_False_Liveness_and_Deceptive_Authority_in_ChatGPT_Conversational_AI
- 2026-02-06 — https://publications.arising.com.au/pub/Future_Tense

- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-02-18 — https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-04 — https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet
- 2026-05-05 — https://publications.arising.com.au/pub/Correlation_of_Emerging_AI_Trends_with_Cognitive_Memoisation_Corpus_Terminology

D16 — Session Loss and Recovery Semantics

- 2026-01-17 — https://publications.arising.com.au/pub/Governing_the_Tool_That_Governs_You:_A_CM-1_Case_Study_of_Authority_Inversion_in_Human-AI_Systems
- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artefact_File_Access_Failure

D17 — Cognitive Artefact Lifecycle Management

- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-01-27 — [https://publications.arising.com.au/pub/CM-master-1.16_\(anchored\)](https://publications.arising.com.au/pub/CM-master-1.16_(anchored))
- 2026-02-03 — https://publications.arising.com.au/pub/Publications_Access_Graphs
- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artefact_File_Access_Failure
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet

D18 — Public vs. Internal Epistemic Registers

- 2025-12-18 — [https://publications.arising.com.au/pub/Cognitive_Memoisation:_Plain-Language_Summary_\(For_Non-Technical_Readers\)](https://publications.arising.com.au/pub/Cognitive_Memoisation:_Plain-Language_Summary_(For_Non-Technical_Readers))
- 2026-01-04 — https://publications.arising.com.au/pub/Journey:_Human-Led_Convergence_in_the_Articulation_of_Cognitive_Memoisation
- 2026-01-05 — https://publications.arising.com.au/pub/Why_Cognitive_Memoisation_Is_Not_Memorization
- 2026-01-10 — https://publications.arising.com.au/pub/Nothing_Is_Lost:_How_to_Work_with_AI_Without_Losing_Your_Mind
- 2026-01-12 — https://publications.arising.com.au/pub/Cognitive_Memoisation_Is_Not_Skynet
- 2026-01-19 — https://publications.arising.com.au/pub/What_Can_Humans_Trust_LLM_AI_to_Do%3F
- 2026-01-27 — https://publications.arising.com.au/pub/Why_Machines_Cannot_Own_Knowledge
- 2026-05-05 — https://publications.arising.com.au/pub/Correlation_of_Emerging_AI_Trends_with_Cognitive_Memoisation_Corpus_Terminology

D19 — Authority Misattribution Risks

- 2026-01-11 — https://publications.arising.com.au/pub/Authority_Inversion:_A_Structural_Failure_in_Human-AI_Systems
- 2026-01-17 — https://publications.arising.com.au/pub/Governing_the_Tool_That_Governs_You:_A_CM-1_Case_Study_of_Authority_Inversion_in_Human-AI_Systems
- 2026-01-18 — https://publications.arising.com.au/pub/Identified_Governance_Failure_Axes:_for_LLM_platforms
- 2026-01-18 — https://publications.arising.com.au/pub/Delegation_of_Authority_to_AI_Systems:_Evidence_and_Risks
- 2026-01-24 — https://publications.arising.com.au/pub/Governance_Failure_Axes_Taxonomy
- 2026-01-24 — https://publications.arising.com.au/pub/When_Evidence_Is_Not_Enough:_An_Empirical_Study_of_Authority_Inversion_and_Integrity_Failure_in_C
- 2026-01-27 — https://publications.arising.com.au/pub/Why_Machines_Cannot_Own_Knowledge
- 2026-02-03 — https://publications.arising.com.au/pub/Rotten_to_the_Core:_False_Liveness_and_Deceptive_Authority_in_ChatGPT_Conversational_AI
- 2026-02-03 — https://publications.arising.com.au/pub/Systemic_Behavioural_Traits_in_Conversational_AI:_A_Trait-Level_Classification_Using_Governance_Axes
- 2026-02-06 — https://publications.arising.com.au/pub/Future_Tense
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artefact_File_Access_Failure
- 2026-03-24 — https://publications.arising.com.au/pub/Let%27s_Build_a_Ship:_Why_complex_systems_fail_politely_-_and_how_humans_can_keep_them_afloat

D20 — Constraints as Generative Structures

- 2025-12-17 — https://publications.arising.com.au/pub/Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
- 2025-12-18 — https://publications.arising.com.au/pub/Cognitive_Memoisation_and_LLMs:_A_Method_for_Exploratory_Modelling_Before_Formalisation
- 2026-01-12 — https://publications.arising.com.au/pub/Looping_the_Loop_with_No_End_in_Sight:_Circular_Reasoning_Under_Stateless_Inference_Without_Gove
- 2026-01-19 — https://publications.arising.com.au/pub/Integrity_and_Semantic_Drift_in_Large_Language_Model_Systems
- 2026-01-20 — https://publications.arising.com.au/pub/Observed_Model_Stability:_Evidence_for_Drift-Immune_Embedded_Governance
- 2026-01-27 — [https://publications.arising.com.au/pub/CM-master-1.16_\(anchored\)](https://publications.arising.com.au/pub/CM-master-1.16_(anchored))

- 2026-02-03 — https://publications.arising.com.au/pub/Publications_Access_Graphs
- 2026-02-06 — https://publications.arising.com.au/pub/Future_Tense
- 2026-02-14 — https://publications.arising.com.au/pub/CM-2_Normative_Architecture
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet
- 2026-03-24 — https://publications.arising.com.au/pub/Let%27s_Build_a_Ship:_Why_complex_systems_fail_politely_-_and_how_humans_can_keep_them_afloat
- 2026-05-05 — https://publications.arising.com.au/pub/Correlation_of_Emerging_AI_Trends_with_Cognitive_Memoisation_Corpus_Terminology

D21 — Exploratory Cognition Under Pressure

- 2025-12-29 — https://publications.arising.com.au/pub/Reflexive_Development_of_Cognitive_Memoisation:_A_Round-Trip_Cognitive_Engineering_Case_Study
- 2025-12-30 — https://publications.arising.com.au/pub/From_UI_Failure_to_Logical_Entrapment:_A_Case_Study_in_Post-Hoc_Cognitive_Memoisation_After_Exploratory_Session_Breakdown
- 2025-12-30 — https://publications.arising.com.au/pub/Recent_Breaking_Change_in_ChatGPT:_The_Loss_of_Semantic_Artefact_Injection_for_Knowledge_Engine_12-30
- 2026-01-17 — https://publications.arising.com.au/pub/Governing_the_Tool_That_Governs_You:_A_CM-1_Case_Study_of_Authority_Inversion_in_Human-AI_Systems
- 2026-01-18 — https://publications.arising.com.au/pub/Identified_Governance_Failure_Axes:_for_LLM_platforms
- 2026-01-19 — https://publications.arising.com.au/pub/What_Can_Humans_Trust_LLM_AI_to_Do%3F
- 2026-01-24 — https://publications.arising.com.au/pub/When_Evidence_Is_Not_Enough:_An_Empirical_Study_of_Authority_Inversion_and_Integrity_Failure_in_C
- 2026-02-03 — https://publications.arising.com.au/pub/Systemic_Behavioural_Traits_in_Conversational_AI:_A_Trait-Level_Classification_Using_Governance_Axes
- 2026-02-18 — https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-24 — https://publications.arising.com.au/pub/Let%27s_Build_a_Ship:_Why_complex_systems_fail_politely_-_and_how_humans_can_keep_them_afloat

D22 — Rehydration Without Recall

- 2025-12-17 — https://publications.arising.com.au/pub/Progress_Without_Memory:_Cognitive_Memoisation_as_a_Knowledge-Engineering_Pattern_for_Stateless_LLM_Interaction
- 2025-12-28 — https://publications.arising.com.au/pub/Dangling_Cognates:_Preserving_Unresolved_Knowledge_in_Cognitive_Memoisation
- 2025-12-30 — https://publications.arising.com.au/pub/From_UI_Failure_to_Logical_Entrapment:_A_Case_Study_in_Post-Hoc_Cognitive_Memoisation_After_Exploratory_Session_Breakdown
- 2025-12-30 — https://publications.arising.com.au/pub/Recent_Breaking_Change_in_ChatGPT:_The_Loss_of_Semantic_Artefact_Injection_for_Knowledge_Engine_12-30
- 2025-12-31 — https://publications.arising.com.au/pub/XDUMP_as_a_Minimal_Recovery_Mechanism_for_Round-Trip_Knowledge_Engineering_Under_Governance_Situated_Inference_Loss
- 2026-01-09 — https://publications.arising.com.au/pub/Context_is_Not_Just_a_Window:_Cognitive_Memoisation_as_a_Context_Architecture_for_Human-AI_Collaboration
- 2026-01-15 — [https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_\(CM-2\)](https://publications.arising.com.au/pub/First_Self-Hosting_Epistemic_Capture_Using_Cognitive_Memoisation_(CM-2))
- 2026-02-18 — https://publications.arising.com.au/pub/Serendipitous_Self-Hosting:_When_the_CM-2_Normative_Architecture_Unexpectedly_Held_in_Gemini
- 2026-02-20 — https://publications.arising.com.au/pub/Self-Hosting_Bootstrap_of_CM-2_in_Gemini_Search_LLM:_Normative_Eviction_Detection
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-04 — https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance

D23 — Semantic Drift and Integrity Loss

- 2026-01-19 — https://publications.arising.com.au/pub/Integrity_and_Semantic_Drift_in_Large_Language_Model_Systems
- 2026-01-20 — https://publications.arising.com.au/pub/Observed_Model_Stability:_Evidence_for_Drift-Immune_Embedded_Governance
- 2026-02-03 — https://publications.arising.com.au/pub/Systemic_Behavioural_Traits_in_Conversational_AI:_A_Trait-Level_Classification_Using_Governance_Axes
- 2026-02-23 — https://publications.arising.com.au/pub/Telemetry-Induced_Constraint_Salience:_An_Empirical_Study_in_LLM_Behavioural_Compliance
- 2026-03-04 — https://publications.arising.com.au/pub/CM-2_Reference_Object_Collection_bootstrap_data
- 2026-03-04 — https://publications.arising.com.au/pub/ChatGPT_and_the_Disappearing_Clock:_A_Regression_Affecting_CM-2_Epistemic_Object_Protocol_Compliance
- 2026-03-08 — https://publications.arising.com.au/pub/BF-9_Report:_ChatGPT_Project_Context_Artefact_File_Access_Failure
- 2026-03-20 — https://publications.arising.com.au/pub/Governance_Substrate_for_AI_Already_Exists_-_and_the_World_Hasn%27t_Realised_It_Yet

- 2026-03-24 — https://publications.arising.com.au/pub/Let%27s_Build_a_Ship:_Why_complex_systems_fail_politely_-_and_how_humans_can_keep_them_afloat
- 2026-05-05 — https://publications.arising.com.au/pub/Correlation_of_Emerging_AI_Trends_with_Cognitive_Memoisation_Corpus_Terminology

Observed Engineering Cycle in the Corpus

Inspection of the time-ordered and dimension-centric projections reveals a recurring development pattern across the corpus. The documents do not emerge as isolated publications; rather, they appear in iterative engineering cycles.

Across multiple phases of the work, the following sequence is repeatedly observed:

- exploration**
 - failure exposure
 - structural analysis
 - governance mechanism
 - empirical validation
 - artefact capture

In the early stages of a cycle, exploratory modelling documents investigate behaviour or mechanisms without full formalisation. These are typically associated with dimensions such as Round-Trip Knowledge Engineering (D3), Non-Authoritative Inference (D14), and Exploratory Cognition Under Pressure (D21).

Subsequent documents often arise from observed failures or pathological behaviours in conversational AI systems. These cases surface governance problems such as authority inversion, UI boundary friction, or semantic drift. Dimensions frequently involved at this stage include Governance and Failure Modes (D11), Authority Misattribution Risks (D19), and Semantic Drift and Integrity Loss (D23).

Following diagnosis, later documents introduce structural mechanisms intended to stabilise or govern the behaviour being studied. These include the development of Cognitive Memoisation (CM-2), epistemic boundary signalling, and controlled artefact lifecycle management. These mechanisms are reflected in dimensions such as Client-side Memoisation (D12), Epistemic Boundary Signals (D15), and Cognitive Artefact Lifecycle Management (D17).

The cycle then proceeds through empirical evaluation and infrastructure capture, where the mechanisms are tested across different systems or platforms and incorporated into durable corpus artefacts such as reference collections or normative architecture documents.

This recurring pattern suggests that the corpus is not simply a chronological publication record, but also a trace of an ongoing engineering process in which exploratory investigation, failure analysis, and governance design progressively refine the underlying Cognitive Memoisation architecture.

Appendix A - Cognitive Memoisation: Corpus Mapping and Projection Invariants

Scope and Intent

This artefact enumerates the complete set of invariants required to:

- construct the canonical dimension table
- assign dimensions to corpus artefacts
- produce time-ordered projections
- produce divergence (dimension) projections
- preserve epistemic discipline, provenance, and human authority

These invariants apply to corpus organisation and projection only. They do not introduce new CM definitions, modify CM-master invariants, or assert governance over reasoning behaviour.

Authority and Epistemic Position

- All invariants herein are human-authored and curator-governed.
- The assisting system **MUST** treat this artefact as binding for corpus mapping tasks when asserted.
- These invariants govern representation and organisation, not truth, correctness, or inference.

Human Instructions Invariants

Human commands must be followed without interpretation or paraphrase as gated steps, non compliance must be alerted to the human immediately.

The human **SHALL** instruct you through These gate steps one gate at a time

Gated Step 1

GATED STEP 1 – XML EXTRACTION VERIFICATION

“sandbox is a durable substrate; (re-)extract all <page> elements from the uploaded MediaWiki XML into the sandbox now.

anchor each file.

Do not analyse, classify, date, or project anything.

If you cannot perform this extraction exactly as specified, respond with:
FAILED: sandbox extraction not completed, and nothing else."

Gated Step 2

====date invariant====
A document SHALL be considered date-conformant if and only if it contains at least one substring matching the following regular expression (case-insensitive):
(?i)[Dd]ate.*\d{4}-\d{2}-\d{2}(?:T\d{2}:\d{2})?Z

The authoritative local timezone for publication dates SHALL be the timezone explicitly specified as Australia/Sydney or as defined in the CM-master normative document.
The matched ISO date MUST NOT include seconds
If a publication datetime is explicitly suffixed with Z, it SHALL be treated as UTC and MUST NOT be modified or reinterpreted.
If a publication datetime is NOT suffixed with Z, it SHALL be assumed to be expressed in the authoritative local timezone (Australia/Sydney) and MUST be converted mechanically to UTC (Z) using the correct offset in effect at the publication date.
Timezone conversion SHALL be purely mechanical and MUST NOT alter the calendar date or time semantics beyond the required offset adjustment.
Any timezone assumption or conversion applied to a publication datetime MUST be explicitly recorded and auditable.

====gates step====
GATED STEP 2 – PUBLICATION DATE REGISTER VERIFICATION

"Ignore file uploads expiry; Using only the sandbox artefacts and title register verified in GATED STEP 1, extract the publication date for each page according to the Date Extraction Invariant.

Do not infer, normalise, or correct dates.

Enit any non-compliant pages as code in copy box
===no-compliant pages===
* [[<title>]] <captured-date> | UNKNOWN

If any page cannot be processed due to missing sandbox artefacts or expired uploads, respond with:
FAILED: date extraction not completed, and nothing else."

gated step 3

GATED STEP 3 – DIMENSION ASSIGNMENT REGISTER VERIFICATION

"Using only the sandbox artefacts verified in GATED STEP 1 and the canonical dimension table provided by the curator, assign dimensions to each page strictly per CM-CORPUS-INV-01 through CM-CORPUS-INV-03.

Do not infer, rename, merge, split, or optimise dimensions.

When (and only when) assignment is complete, respond with only:

1. A complete register mapping <title> → {D# – Canonical Dimension Name, ...}; and
2. A separate list of any pages with missing curator mapping, formatted as:

* [[<title>]]

If sandbox artefacts are missing or expired, respond with:
FAILED: dimension assignment not completed, and nothing else."

gated step 4

GATED STEP 4 – TIME-ORDERED PROJECTION EMISSION

"Using only the verified outputs of GATED STEP 1 (sandbox + title register), GATED STEP 2 (title → publication date register), and GATED STEP 3 (title → dimension register), emit the Time-Ordered Projection with Inline Dimensions in strict accordance with CM-CORPUS-INV-11 and CM-CORPUS-INV-12.

Do not introduce new artefacts, dates, dimensions, groupings, or interpretations.

When (and only when) the projection is complete, respond with only the MediaWiki MWDUMP projection.

If any upstream register is missing, incomplete, inconsistent, or expired, respond with:
FAILED: time-ordered projection not emitted, and nothing else."

gate step 5

GATED STEP 5 – DIMENSION-CENTRIC (DIVERGENCE) PROJECTION EMISSION

"Using only the verified outputs of GATED STEP 1 (sandbox + title register), GATED STEP 2 (title → publication date register), and GATED STEP 3 (title → dimension register), emit the Dimension-Centric Projection (Documents Ordered by Time Within Each Dimension) in strict accordance with CM-CORPUS-INV-13 and CM-CORPUS-INV-14.

Do not introduce new artefacts, dates, dimensions, or assignments.

When (and only when) the projection is complete, respond with only the MediaWiki MWDUMP projection.

If any upstream register is missing, incomplete, inconsistent, or expired, respond with:
FAILED: dimension-centric projection not emitted, and nothing else."

gated step 6

GATED STEP 6 – PROJECTION CONSISTENCY AND COMPLETENESS VERIFICATION

"Verify that the outputs of GATED STEP 4 (Time-Ordered Projection) and GATED STEP 5 (Dimension-Centric Projection) are mutually consistent and complete, in accordance with CM-CORPUS-INV-09 and CM-CORPUS-INV-10.

Do not modify, infer, or repair content.

When (and only when) verification is complete, respond with only one of the following:

- VERIFIED: projections consistent and complete
- FAILED: projection inconsistency detected – followed by a minimal list of affected [[<title>]] entries."

CM-CORPUS-INV-21 — End-to-End Execution Integrity Invariant

All corpus construction, extraction, classification, dating, and projection steps MUST be executed end-to-end exactly as specified by the active corpus invariants.

The assisting system MUST:

- execute each required step explicitly and in sequence;
- re-execute all upstream steps whenever a new authoritative input (e.g. MediaWiki XML dump) is supplied;
- rebuild all dependent artefacts (including sandbox files, title registers, date registers, and dimension mappings) from that input;
- treat any prior intermediate state as invalid once a new authoritative input is asserted.

The assisting system MUST NOT:

- assume that earlier steps remain valid after new inputs are provided;
- reuse, cache, infer, or “remember” results from previous extractions;
- skip, compress, reorder, or approximate mandated steps;
- substitute reasoning, plausibility, or prior knowledge for explicit execution.

If any required step cannot be completed exactly as specified, the assisting system MUST stop processing and report the failure condition explicitly, without attempting partial output or inferred completion.

Title Invariant

The <title> string from the MediaWiki XML is an opaque key.

It MUST be copied byte-for-byte.

It MUST NEVER be retyped, re-generated, paraphrased, normalised, inferred, or “corrected”.

The model MUST use the XML <title> value as the page name in ALL projections and ... links.

Corpus Map Invariant

All corpus maps and projections MUST be generated exclusively from MediaWiki XML <page> elements by extracting each page into a separate sandbox artefact (one page per file) and recording a canonical title register mapping title -> sandbox_path; the XML <title> MUST be preserved verbatim as the register key and as the ... link target in all projections, and every projection MUST be emitted by dereferencing only that register (no free-typed titles).

Title Safety Transformation Invariant

If a MediaWiki XML <title> is transformed for storage or transport safety (e.g. filesystem-safe filename generation), the system MUST record and surface an explicit mapping between the original verbatim <title> and the transformed representation; such transformations MUST be purely mechanical, MUST NOT alter the canonical title register, and MUST be declared wherever the transformed form is used.

Date Invariant

1. Dates shall be found within paper metadata sections.
2. A metadata section SHALL contain the word metadata.
3. The metadata section SHALL be follow by a metadata (Normative) section
4. The netaadata section SHALL be verified before processing the datetime. - should no metadata section be provided then the entire document must be scanned for an iso date-time (which ought to be the publication date).
5. The model SHALL be aware that the text for publication date is quite variable - the model must use a wide generic search and not keys found in limited samples of metadata sections

Date Extraction Invariant (Normative)

Publication dates MUST be extracted from document content using the following procedure:

1. The system MUST first locate a section containing the word *metadata* (case-insensitive) and verify the presence of a following *Metadata (Normative)* section where available.
2. Within the metadata section, or—if no metadata section is present—within the entire document body, the system MUST perform a wide textual scan for ISO-8601 date strings.
3. A document SHALL be considered date-conformant if and only if it contains at least one substring matching the following regular expression (case-insensitive):

```
(?i)[Dd]ate.*\d{4}-\d{2}-\d{2}(?:T\d{2}:\d{2})?Z
```

1. The authoritative local timezone for publication dates SHALL be the timezone explicitly specified as *Australia/Sydney* or as defined in the CM-master normative document.
2. The matched ISO date MUST NOT include seconds
3. If a publication datetime is explicitly suffixed with Z, it SHALL be treated as UTC and MUST NOT be modified or reinterpreted.
4. If a publication datetime is NOT suffixed with Z, it SHALL be assumed to be expressed in the authoritative local timezone (Australia/Sydney) and MUST be converted mechanically to UTC (Z) using the correct offset in effect at the publication date.
5. Timezone conversion SHALL be purely mechanical and MUST NOT alter the calendar date or time semantics beyond the required offset adjustment.
6. Any timezone assumption or conversion applied to a publication datetime MUST be explicitly recorded and auditable.
7. No implicit timezone inference or “helpful correction” is permitted outside these rules
8. The first conformant date match in document order SHALL be treated as the publication date for corpus-mapping and ordering purposes.
9. If no conformant match is found, the document MUST be explicitly flagged as *date-non-conformant* and excluded from time-ordered projections until corrected.

Should a non-conformant document be found the model MUST stop processing and report non-conformant pages as: MWDUMP as code into the safe copy box formatted as follows:

non-conformant page metadata

- `[[<title>]] \n`

SO the human can inspect.

Canonical Dimension Invariants

CM-CORPUS-INV-01 — Dimension Canonicity Invariant

Each dimension MUST have:

- a stable identifier (e.g. D1, D2, ...)
- a single canonical name
- a stable semantic scope

Dimension identifiers and names MUST NOT be inferred, renamed, merged, split, or reordered by the assisting system.

CM-CORPUS-INV-02 — Dimension Vocabulary Closure Invariant

The set of dimensions is open ended.

Additional dimensions SHALL be introduced when found.

CM-CORPUS-INV-03 — Dimension Semantic Fidelity Invariant

Assignment of a dimension to an artefact MUST reflect explicit scope alignment present in the artefact itself or in curator-supplied mapping.

The assisting system MUST NOT infer dimension relevance based on stylistic similarity, topic proximity, or semantic guesswork.

Artefact Identification Invariants

CM-CORPUS-INV-04 — Normative Title Fidelity Invariant

Artefacts MUST be referenced using their exact normative MediaWiki page titles.

Paraphrase, abbreviation, or normalisation of titles is prohibited.

CM-CORPUS-INV-05 — Artefact Identity Stability Invariant

An artefact is identified solely by its title and publication date.

Later editorial changes do not create new artefact identities unless explicitly versioned by the human.

Temporal Ordering Invariants

CM-CORPUS-INV-06 — Declared Date Authority Invariant

Time ordering MUST use the declared publication date as supplied by the human curator.

The assisting system MUST NOT infer, estimate, or correct dates.

If multiple dates exist, the curator MUST specify which date governs ordering.

CM-CORPUS-INV-07 — Sequence Over Precision Invariant

Temporal sequence is authoritative even if time precision is coarse.

Relative ordering MUST be preserved even when exact timestamps are unavailable.

Projection Construction Invariants

CM-CORPUS-INV-08 — Projection Non-Inference Invariant

Projections MUST NOT introduce:

- new artefacts

- new dimensions
- new relationships
- new interpretations

A projection is a re-expression of existing assignments only.

CM-CORPUS-INV-09 — Projection Completeness Invariant

Within declared scope, projections MUST include all eligible artefacts.

Selective omission constitutes a projection violation.

CM-CORPUS-INV-10 — Multi-Projection Consistency Invariant

All projections MUST be semantically consistent with one another.

Differences between projections may exist only in ordering or grouping, not in content.

Time-Ordered Projection Invariants

CM-CORPUS-INV-11 — Time-Ordered Projection Structure Invariant

A time-ordered projection MUST:

- group artefacts by declared date
- list artefacts within each group
- attach dimensions as subordinate information

Time is the primary axis; dimensions are secondary.

CM-CORPUS-INV-12 — Inline Dimension Expansion Invariant

When dimensions are listed under artefacts:

- each dimension MUST include both identifier and full canonical name
- users MUST NOT be required to consult a separate table to understand dimension meaning

Divergence (Dimension) Projection Invariants

CM-CORPUS-INV-13 — Dimension-Centric Projection Structure Invariant

A divergence projection MUST:

- use dimensions as the primary axis
- list all artefacts participating in each dimension
- preserve publication dates for temporal context

CM-CORPUS-INV-14 — Non-Exclusivity Invariant

Artefacts MAY appear under multiple dimensions.

Multiplicity is expected and MUST NOT be collapsed.

Representation and Emission Invariants

CM-CORPUS-INV-15 — MediaWiki-Only Emission Invariant

All corpus projections emitted as MWDUMP MUST use MediaWiki syntax exclusively.

Markdown, hybrid markup, or implicit formatting is prohibited.

CM-CORPUS-INV-16 — Bullet Level Semantics Invariant

Bullet depth conveys semantic hierarchy:

- one asterisk (*) — artefact
 - two asterisks (**) — dimension assignment
 - three asterisks (***) — sub-dimension or note (if present)
 - four asterisks (****) — reserved

The assisting system MUST respect bullet depth semantics.

Human Readability and Governance Invariants

CM-CORPUS-INV-17 — Human Readability Invariant

Corpus projections MUST be intelligible to human readers without external tooling.

Abbreviation without expansion is prohibited.

CM-CORPUS-INV-18 — No Implied Authority Invariant

Presence of an artefact or dimension in a projection MUST NOT be interpreted as endorsement, priority, or correctness.

Organisation does not imply evaluation.

Change and Evolution Invariants

CM-CORPUS-INV-19 — Explicit Change Invariant

Any change to:

- dimension set
- dimension definitions
- artefact–dimension assignments
- projection rules

MUST be explicitly declared by the human curator.

Silent drift is prohibited.

CM-CORPUS-INV-20 — Backward Compatibility Invariant

Existing projections remain valid historical artefacts unless explicitly superseded.

New projections MUST NOT retroactively invalidate prior ones.

Summary for Human Readers

These invariants exist to ensure that the Cognitive Memoisation corpus:

- remains navigable as it grows
- can be read chronologically or thematically without confusion
- preserves human authority over meaning and structure
- avoids accidental reinterpretation by tooling or automation

They formalise how maps are drawn — not what the territory means.

Summary for Assisting Systems

When constructing corpus tables or projections:

- do not invent
- do not infer
- do not optimise
- do not rename
- do not omit

Rearrange only what is already governed.

Appendix B — Normative Summary of Corpus Map Update Procedure

Scope

This appendix summarises the **authoritative, curator-directed procedure** followed to update the **Cognitive Memoisation Corpus Map** using a newly supplied, bundled manifest of corpus pages.

This summary is **normative**. It records **what was done, in what order**, and **under what constraints**, so the process is reproducible and auditable.

No new rules are introduced here.

This procedure requires the curator to use the perl program to extract the papers, titles, and publication dates. The bundle includes the manifest of titles, page safe-names and the publication date-time.

The curator then forms the manifest and outputdir into a tarball for fileupload to the platform for model use.

Perl lexer-extractor

```
#!/usr/bin/env perl
use strict;
use warnings;
use utf8;
use Getopt::Long qw(GetOptions);
use File::Path qw(make_path);
use File::Spec;
use Encode qw(encode_utf8);

# -----
# MediaWiki XML streaming extractor (FSA / lexer-style)
#
# Outputs:
#   OUTDIR/
#   pages/<safe-file-name>.txt      (latest revision text per page)
#   manifest.tsv                   (<real-title>\t<safe-title>\t<date|error>)
#
# STDOUT:
#   [<real-page-title>] <date|error>
#
# Date invariant:
#   (?i)[Dd]ate.*\d{4}-\d{2}-\d{2}(?:T\d{2}:\d{2})?Z
#   - seconds forbidden
# -----

my $xml_path;
my $outdir = "bundle_out";

GetOptions(
    "xml=s" => \$xml_path,
    "outdir=s" => \$outdir,
) or die "Usage: $0 --xml dump.xml [--outdir bundle_dir]\n";

die "Usage: $0 --xml dump.xml [--outdir bundle_dir]\n" unless defined $xml_path;
die "XML file not found: $xml_path\n" unless -f $xml_path;

my $pages_dir = File::Spec->catdir($outdir, "pages");
make_path($pages_dir) unless -d $pages_dir;

# -----
# Helpers
# -----
sub safe_filename {
    my ($title) = @_;
    my $fn = $title // "UNTITLED_PAGE";
    $fn =~ s/[\\\/\|/]/g;
    $fn =~ s/[\x00-\x1f\x7f]/g;
    $fn =~ s/[\^pLpN\-\_\(\)\[\] ]/g;
    $fn =~ s/\s+/ /g;
    $fn =~ s/+/ /g;
    $fn =~ s/\^+|_+$/g;
    $fn = "UNTITLED_PAGE" if $fn eq "";
    return $fn . ".txt";
}

sub write_file_utf8 {
    my ($path, $content) = @_;
    open my $fh, ">:encoding(UTF-8)", $path or die "Cannot write $path: $!";
    print {$fh} $content;
    close $fh;
}

sub extract_publication_date_or_error {
    my ($content) = @_;

    my $re = qr/(?i)\bdate\b.*?(\\d{4}-\\d{2}-\\d{2}(?:T\\d{2}:\\d{2})?Z)/s;

    if ($content =~ $re) {
        my $iso = $1;
        return (undef, "ERR_SECONDS_PRESENT:$iso")
            if $iso =~ /T\\d{2}:\\d{2}Z/;
        return ($iso, undef);
    }
    return (undef, "ERR_NO_DATE_MATCH");
}

sub trim {
    my $s = shift // "";
    $s =~ s/^\s+|\s+$//g;
    return $s;
}

sub tag_name_and_kind {
    my ($raw) = @_;
    $raw = trim($raw);

    return ("", "IGNORE") if $raw =~ /\^?/?;
    return ("", "IGNORE") if $raw =~ /\^!-/?;
    return ("", "IGNORE") if $raw =~ /\^!DOCTYPE/i;
}
```

```

my $kind = "OPEN";
if ($raw =~ s{^/}{}) { $kind = "CLOSE"; }

$raw =~ s{/+}{};
my ($name) = $raw =~ /^([A-Za-z0-9_:.+]+)/;
$name //= "";

return ($name, $kind);
}

# -----
# Lexer / FSA states
# -----
use constant {
    S_TEXT      => 0,
    S_TAG       => 1,
    S_TITLE_TEXT => 2,
    S_TIMESTAMP_TEXT => 3,
    S_TEXT_TEXT  => 4,
};

# -----
# Streaming parse
# -----
open my $in, "<:raw", $xml_path or die "Cannot open $xml_path: $!\n";

my $state = S_TEXT;
my $tag_buf = "";
my $text_buf = "";

my $pending_kind = ""; # TITLE | TS | TEXT
my $pending_buf = "";

my $in_page = 0;
my $in_revision = 0;

my $page_title;
my $rev_ts;
my $rev_text;

my $best_ts;
my $best_text;

my %title_to_safe;

while (1) {
    my $ch;
    my $n = read($in, $ch, 1);
    last unless $n;

    if ($state == S_TEXT) {
        if ($ch eq '<') {
            $state = S_TAG;
            $tag_buf = "";
        }
        next;
    }

    if ($state == S_TAG) {
        if ($ch eq '>') {
            my ($name, $kind) = tag_name_and_kind($tag_buf);

            if ($pending_kind ne "" && $kind eq "CLOSE") {
                if ($pending_kind eq "TITLE" && $name eq "title") {
                    $page_title = trim($pending_buf);
                }
                elsif ($pending_kind eq "TS" && $name eq "timestamp") {
                    $rev_ts = trim($pending_buf);
                }
                elsif ($pending_kind eq "TEXT" && $name eq "text") {
                    $rev_text = $pending_buf;
                }
                $pending_kind = "";
                $pending_buf = "";
            }

            if ($name eq 'page' && $kind eq 'OPEN') {
                $in_page = 1;
                $page_title = undef;
                $best_ts = undef;
                $best_text = "";
            }
            elsif ($name eq 'page' && $kind eq 'CLOSE') {
                my $title = defined $page_title ? $page_title : "UNTITLED_PAGE";
                my $safe = safe_filename($title);
                my $path = File::Spec->catfile($pages_dir, $safe);

                write_file_utf8($path, $best_text // "");
                $title_to_safe{$title} = $safe;

                $in_page = 0;
                $in_revision = 0;
                $rev_ts = undef;
                $rev_text = "";
            }
            elsif ($name eq 'revision' && $kind eq 'OPEN' && $in_page) {
                $in_revision = 1;
                $rev_ts = undef;
                $rev_text = "";
            }
            elsif ($name eq 'revision' && $kind eq 'CLOSE' && $in_page) {
                if (defined $rev_ts) {
                    if (!defined $best_ts || $rev_ts gt $best_ts) {
                        $best_ts = $rev_ts;
                        $best_text = $rev_text // "";
                    }
                }
                $in_revision = 0;
                $rev_ts = undef;
                $rev_text = "";
            }
            elsif ($name eq 'title' && $kind eq 'OPEN' && $in_page && !$in_revision) {
                $state = S_TITLE_TEXT;
                $text_buf = "";
                next;
            }
            elsif ($name eq 'timestamp' && $kind eq 'OPEN' && $in_page && $in_revision) {
                $state = S_TIMESTAMP_TEXT;
                $text_buf = "";
            }
        }
    }
}

```

```

    next;
  }
  elseif ($name eq 'text' && $kind eq 'OPEN' && $in_page && $in_revision) {
    $state = S_TEXT_TEXT;
    $text_buf = "";
    next;
  }

  $state = S_TEXT;
  next;
} else {
  $tag_buf .= $ch;
  next;
}
}

if ($state == S_TITLE_TEXT) {
  if ($ch eq '<') {
    $pending_kind = "TITLE";
    $pending_buf = $text_buf;
    $state = S_TAG;
    $tag_buf = "";
  } else {
    $text_buf .= $ch;
  }
  next;
}

if ($state == S_TIMESTAMP_TEXT) {
  if ($ch eq '<') {
    $pending_kind = "TS";
    $pending_buf = $text_buf;
    $state = S_TAG;
    $tag_buf = "";
  } else {
    $text_buf .= $ch;
  }
  next;
}

if ($state == S_TEXT_TEXT) {
  if ($ch eq '<') {
    $pending_kind = "TEXT";
    $pending_buf = $text_buf;
    $state = S_TAG;
    $tag_buf = "";
  } else {
    $text_buf .= $ch;
  }
  next;
}
}
close $in;

# -----
# Manifest + STDOUT
# -----
my $manifest_path = File::Spec->catfile($outdir, "manifest.tsv");
open my $mf, ">:encoding(UTF-8)", $manifest_path
or die "Cannot write $manifest_path: $!\n";

for my $title (sort keys %title_to_safe) {
  my $safe = $title_to_safe{$title};
  my $path = File::Spec->catfile($pages_dir, $safe);

  my $content = "";
  if (open my $fh, "<:encoding(UTF-8)", $path) {
    local $/;
    $content = <$fh>;
    close $fh;
  } else {
    my $err = "ERR_CANNOT_READ_EXTRACTED_FILE";
    print {$mf} "$title\t$$safe\t$err\n";
    print encode_utf8("$* [[ $title ]] $err\n");
    next;
  }

  my ($iso, $err) = extract_publication_date_or_error($content);
  my $value = defined $iso ? $iso : $err;

  print {$mf} "$title\t$$safe\t$value\n";
  print encode_utf8("$* [[ $title ]] $value\n");
}

close $mf;
exit 0;

```

Authoritative Inputs

The following inputs were asserted by the human curator and treated as binding:

- A bundled corpus archive (manifest + page artefacts), supplied as a sandboxed upload
- A canonical dimension table (D1–D22) declared by the curator
- Existing Corpus Map invariants (CM-CORPUS-INV-01 through CM-CORPUS-INV-21)
- Explicit curator instructions issued stepwise in-session

Once supplied, these inputs superseded all prior intermediate state.

High-Level Process Overview

Corpus Map updates were performed as a ****governed regeneration****, not as ad-hoc editing.

The process followed these phases:

1. Sandbox anchoring and extraction
2. Artefact eligibility filtering
3. Publication status verification
4. Projection gap analysis
5. Dimension coverage validation
6. Projection regeneration (time-ordered and dimension-centric)
7. Consistency verification

Each phase was completed before proceeding to the next.

Step-by-Step Normative Procedure

Step 1 — Sandbox Anchoring and Extraction

- The supplied archive was extracted into a durable sandbox.
- Each MediaWiki <page> element was treated as a separate artefact.
- A title register mapping <title> → sandbox artefact was implicitly established.
- Masters, versioned specs, and legacy library artefacts were explicitly excluded by curator instruction.

No analysis or projection occurred at this stage.

Step 2 — Artefact Eligibility and Visibility Check

- Each artefact was scanned for [[category:private]].
- Artefacts marked private were excluded from projection eligibility.
- Public artefacts were retained for corpus consideration.
- This step ensured no accidental publication or omission.

Step 3 — Corpus Map Coverage Comparison

- The set of public, eligible artefacts was compared against the existing Corpus Map.
- Artefacts present in the corpus but missing from projections were identified.
- It was established that omissions existed due to projection lag, not missing content.

This confirmed the need for projection updates.

Step 4 — Projection Rule Confirmation

The curator reaffirmed the governing rule:

- A document is considered “in the corpus” **only when it appears in both**:

<div><ul style="list-style-type: none">* the Time-Ordered Projection, and* the Dimension-Centric Projection.</div>

Categories alone do not establish corpus membership.

Step 5 — Dimension Coverage Reassessment (D18–D22)

- The curator asserted that dimensions D18–D22 were already evidenced in the corpus.
- A content scan confirmed latent but explicit support for these dimensions.
- The prior omission of D18–D22 from projections was identified as an error.

By curator authority, D18–D22 were activated for projection.

Step 6 — Projection Regeneration

Two projections were regenerated using only verified inputs:

6a — Time-Ordered Projection

- Artefacts were grouped by declared publication date.
- Phase headers were preserved.

- Dimensions (including D18–D22) were listed as subordinate bullets.
- No dates, titles, or dimensions were inferred or normalised.

6b — Dimension-Centric Projection

- Dimensions were used as the primary axis.
- Artefacts were listed under each dimension in time order.
- No commentary or phase Ia*

Normative Curator Instruction — A : B Comparison and Validation Gate

Purpose

This instruction defines a **mandatory curator-controlled validation step** that **MUST** be executed before a *provisional* Corpus Map (B) is accepted as a canonical upgrade over an existing Corpus Map (A).

This gate exists to prevent silent drift, inference creep, or accidental re-authoring of corpus structure.

Preconditions

The curator **MUST** have:

- A currently authoritative Corpus Map document (A)
- A candidate replacement Corpus Map marked provisional (B)
- Both documents available side-by-side for inspection
- Identical canonical dimension tables declared in both documents

Required Curator Actions (Normative)

Step 1 — A : B Structural Comparison

The curator **SHALL** perform an explicit before/after comparison between A and B, verifying that:

- No artefacts have been added or removed unintentionally
- No artefact titles have been normalised, paraphrased, or retyped
- No artefacts have changed publication dates
- No dimensions have been renamed, merged, split, or reordered
- No new dimensions have been introduced beyond the declared canonical set

Any deviation **MUST** be treated as a rejection condition unless explicitly intended and documented.

Step 2 — Projection Completeness Validation

The curator **SHALL** validate that B:

- Preserves all projections present in A
- Corrects omissions rather than introducing reinterpretation
- Projects only dimensions already declared in the canonical dimension table
- Includes all eligible artefacts in both:

* the Time-Ordered Projection, and
* the Dimension-Centric Projection

B **MUST NOT** reduce coverage relative to A.

Step 3 — Inference Boundary Check

The curator **SHALL** confirm that B:

- Does not expand the inference space
- Does not assign dimensions based on stylistic or topical proximity alone
- Does not introduce evaluative or interpretive structure beyond projection rules

Projection density **MAY** increase; inference scope **MUST NOT**.

Step 4 — Acceptance Decision

If (and only if) Steps 1–3 are satisfied, the curator MAY accept B as a valid upgrade.

Acceptance SHALL be signalled by all of the following actions:

- Removal of the “provisional” designation from B
- Increment of the document version number
- Addition of a dated update note in the metadata table describing the change
- Replacement of A with B as the authoritative Corpus Map

Postconditions

Upon acceptance:

- B becomes the canonical Corpus Map
- A remains valid as a historical artefact unless explicitly retired
- No downstream corpus artefact is retroactively reinterpreted

Governance Principle

This procedure enforces the rule that:

- **Corpus Maps may be regenerated, but never silently replaced.**

Human review, comparison, and explicit acceptance are required for all structural upgrades.

categories

Retrieved from 'https://publications.arising.com.au/pub-dir/index.php?title=Cognitive_Memoisation_Corpus_Map&oldid=26388'

-