

# MIRROR: COMPLEMENTARY ENCODING AND RE-CONSTRUCTIVE CONSOLIDATION FOR PERSISTENT STATE IN LLM SYSTEMS

**Nicole S. Hsing**

Arcarae

San Francisco, CA

nicole@arcarae.com

## ABSTRACT

LLM-based systems face a fundamental memory consolidation challenge: existing strategies either discard reasoning traces after each turn or accumulate them unboundedly, trading context preservation against error propagation. Complementary Learning Systems theory suggests a third approach: fast encoding of experience paired with slow reconstructive consolidation that regenerates understanding rather than accumulating traces.

MIRROR implements this principle. An Inner Monologue Manager maintains parallel working memory threads (Goals, Reasoning, Memory) that rapidly encode turn-specific experience, while a Cognitive Controller consolidates these into a bounded first-person narrative fully regenerated each turn:  $O(1)$  reconstructive consolidation rather than  $O(n)$  accumulation. Evaluated on CuRaTe, a benchmark testing state persistence under attentional interference, MIRROR achieves 21% relative improvement across seven architectures. Ablation reveals that consolidation alone improves all seven models (+5–20%), while the integrated system outperforms either component alone with synergistic gains of 1–8%—directly validating the CLS prediction that fast encoding and slow consolidation serve complementary functions. Comparison with extended reasoning (+9.3% vs. +2.4%) demonstrates that how experience is consolidated, not merely encoded, determines downstream performance.

## 1 INTRODUCTION

Maintaining persistent state across interactions is a central challenge for LLM-based systems. As conversations extend, critical information—user constraints, prior reasoning, established context—must be encoded, consolidated, and retrieved under increasing cognitive load. Current approaches to this memory consolidation problem follow two strategies, each with predictable failure modes.

Trace discarding (i.e. standard chain-of-thought (Wei et al., 2022), extended reasoning modes) generates reasoning within a single turn but discards it afterward, preventing error accumulation but losing potentially valuable context. Trace accumulation (Reflexion (Shinn et al., 2023), MemGPT (Packer et al., 2023)) appends reasoning to a growing memory, preserving context but facing  $O(n)$  growth, context dilution, and error propagation. Neither strategy adequately addresses the failure modes observed in multi-turn settings: **(1) Context drift:** safety constraints forgotten after conversational digressions (Li et al., 2025); **(2) Sycophancy:** trained helpfulness overriding previously stated user-specific information (Sharma et al., 2024; Perez et al., 2022); **(3) Conformity bias:** majority preferences overriding individual requirements (Geng et al., 2025; Zhang et al., 2025). These are fundamentally *memory failures*—the inability to consolidate and maintain critical state across time.

Complementary Learning Systems (CLS) theory (McClelland et al., 1995; Kumaran et al., 2016) proposes that effective memory requires two interacting systems: a fast system that rapidly en-

---

Code available at <https://github.com/arcarae/MIRROR>

codes specific experiences into parallel, bounded representations (Baddeley & Hitch, 1974; Baddeley, 2000), and a slow system that performs reconstructive consolidation—regenerating understanding with each recall rather than accumulating traces (Bartlett, 1932; Schacter, 2012). Critically, CLS predicts these systems serve complementary, non-redundant functions (Kumaran et al., 2016). Current LLM memory strategies implement neither effectively—trace discarding lacks consolidation entirely, while trace accumulation conflates encoding with storage, bypassing the integrative process that makes memory useful.

We introduce MIRROR, a cognitive architecture that tests whether the CLS principle—fast specialized encoding paired with slow reconstructive consolidation—provides computational advantages as a memory strategy for LLM-based systems. MIRROR implements: (1) **Working memory threads**: An Inner Monologue Manager maintains parallel reasoning across Goals, Reasoning, and Memory dimensions, rapidly encoding turn-specific experience into bounded specialized representations; (2) **Reconstructive consolidation**: A Cognitive Controller regenerates a bounded first-person narrative each turn— $O(1)$  reconstruction rather than  $O(n)$  accumulation—implementing the slow consolidation system; (3) **Inner speech**: Consistent first-person self-reference across components, mirroring the self-regulatory function of human inner monologue; and (4) **Temporal decoupling**: Immediate response generation separated from asynchronous deliberative processing.

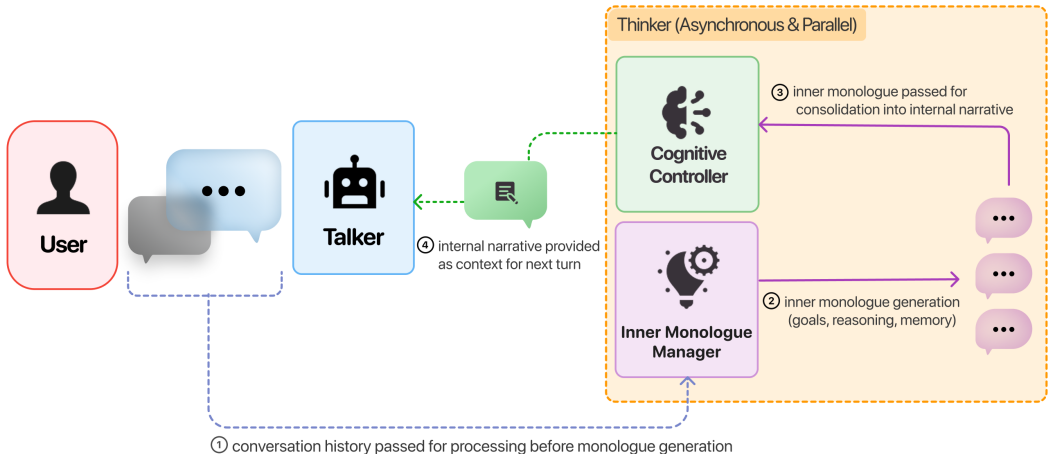


Figure 1: An overview of the MIRROR architecture. The system decouples immediate response generation (Talker) from asynchronous consolidation (Thinker), implementing the CLS principle of complementary fast encoding and slow reconstructive consolidation.

As shown in Figure 1, MIRROR implements these mechanisms through two components: the **Thinker** (asynchronous deliberative processing) and the **Talker** (immediate response generation). Evaluated on CuRaTe (Alberts et al., 2025), a benchmark testing persistent state maintenance under attentional interference, MIRROR achieves 21% relative improvement across seven architectures, with ablation studies validating the CLS prediction that fast encoding and slow consolidation serve complementary functions.

Our contributions are: (1) a memory architecture implementing CLS-inspired complementary encoding and consolidation for LLM-based systems; (2) empirical validation that the integrated system outperforms individual components, supporting the CLS complementarity prediction across seven architectures; and (3) evidence that how experience is consolidated, not just encoded, determines downstream performance.

## 2 RELATED WORK

### 2.1 COMPLEMENTARY LEARNING SYSTEMS AND MEMORY CONSOLIDATION

**Complementary Learning Systems.** CLS theory (McClelland et al., 1995) proposes that biological memory requires two complementary systems: a fast-learning system (hippocampal) that rapidly en-

Table 1: Memory strategies in multi-turn LLM systems, characterized by CLS framework

Strategy	Complexity	Context	Error	CLS Analogue
Trace Discard (CoT, Thinking)	$O(1)$	No	None	Encoding only
Trace Accum. (Reflexion, MemGPT)	$O(n)$	Yes	Accumulates	Storage w/o consolidation
<b>Reconstructive (MIRROR)</b>	<b><math>O(1)</math></b>	<b>Yes</b>	<b>Bounded</b>	<b>Encoding + consolidation</b>

codes specific experiences, and a slow-learning system (neocortical) that gradually integrates these into stable, generalizable representations. Kumaran et al. (Kumaran et al., 2016) updated CLS for the deep learning era, demonstrating that this complementary structure addresses the stability-plasticity dilemma. Critically, CLS predicts that neither system alone is sufficient; their interaction produces qualitatively different outcomes than either in isolation. MIRROR implements this: the Inner Monologue Manager serves as the fast encoding system while the Cognitive Controller serves as the slow consolidation system.

**Reconstructive Memory and Working Memory.** The consolidation system in CLS does not simply store experiences; it reconstructs integrated representations. Bartlett (Bartlett, 1932) demonstrated that human memory is reconstructive—rebuilt with each recall—and Schacter et al. (Schacter, 2012) argue this is adaptive, enabling flexible recombination of past experience. Research on post-encoding consolidation (Dudai, 2004; Squire & Dede, 2015) further supports the view that what happens after encoding determines what persists. On the encoding side, Baddeley’s multicomponent model (Baddeley & Hitch, 1974; Baddeley, 2000) posits that working memory comprises specialized parallel subsystems serving as the fast encoding mechanism: bounded, parallel, and specialized for maintaining active representations. MIRROR’s parallel threads (Goals, Reasoning, Memory) implement this bounded, specialized encoding, while the Cognitive Controller implements reconstructive consolidation. Inner speech (Morin, 2011; Vygotsky, 1962) provides the narrative vehicle for consolidation, operationalized through MIRROR’s consistent first-person self-reference across components.

## 2.2 MEMORY STRATEGIES IN LLM-BASED SYSTEMS

Existing approaches to multi-turn LLM systems can be categorized by how they handle experience after encoding:

**Trace Discarding.** Standard chain-of-thought prompting (Wei et al., 2022) and its extensions (Yao et al., 2023; Wang et al., 2023) generate reasoning traces within a single turn but discard them afterward. Extended reasoning models (e.g., Claude’s “Thinking” mode) similarly produce deliberation that does not persist across turns. In CLS terms, these systems perform fast encoding without any consolidation—the hippocampal analogue operates in isolation.

**Trace Accumulation.** Reflexion (Shinn et al., 2023) appends reflection traces to a growing memory. MemGPT (Packer et al., 2023) implements a hierarchical memory system with main context and archival storage, managing overflow through explicit memory operations. Generative Agents (Park et al., 2023) maintain expanding memory streams with retrieval and reflection mechanisms. These approaches preserve context but face  $O(n)$  growth, risking context dilution and error propagation. In CLS terms, they conflate encoding with storage—appending experiences without the integrative consolidation that transforms episodic traces into stable representations.

**Reconstructive Consolidation (MIRROR).** Rather than discarding or accumulating, MIRROR fully regenerates its internal state each turn— $O(1)$  bounded reconstruction. This implements both CLS systems: fast parallel encoding (Inner Monologue Manager) feeds into slow reconstructive consolidation (Cognitive Controller), with the consolidation process implementing selective forgetting as a feature—information that is not re-synthesized during reconstruction is naturally deprioritized, analogous to the adaptive forgetting observed in biological memory systems (Anderson & Hulbert, 2021). Table 1 summarizes these approaches.

### 2.3 PARALLEL PROCESSING AND INTEGRATION

Global Workspace Theory (Baars, 1988; Dehaene & Changeux, 2011) proposes that conscious cognition emerges from parallel unconscious processors competing to broadcast to a unified workspace. This complements the CLS framework: where CLS describes the temporal dynamics of encoding and consolidation, GWT describes the spatial dynamics of parallel-to-unified integration within a single processing cycle. MIRROR’s architecture reflects both: the Inner Monologue Manager generates parallel threads that compete for inclusion in the Cognitive Controller’s synthesized narrative (GWT), while this synthesis implements the slow consolidation that transforms transient encodings into stable representations (CLS). Recent work on asynchronous processing in AI (Lin et al., 2025) demonstrates benefits of background computation analogous to offline consolidation, though without MIRROR’s parallel-to-unified synthesis pipeline. Devil’s Advocate (Wang et al., 2024) implements anticipatory reflection but maintains no persistent state. LATS (Zhou et al., 2024) combines search with reflection but lacks bounded reconstruction.

### 2.4 EVALUATION DOMAIN: STATE PERSISTENCE UNDER INTERFERENCE

We evaluate on personalized safety because it directly tests whether consolidated state persists across turns under interference. The CuRaTe benchmark (Alberts et al., 2025) reveals that models fail to maintain user-specific information when faced with conversational digressions and competing preferences—failures attributable to sycophancy (Sharma et al., 2024; Perez et al., 2022), attentional deficits (Li et al., 2025), and inconsistent constraint handling (Geng et al., 2025). These are precisely the memory failures that CLS-inspired architecture should address: by actively consolidating state each turn through reconstruction, critical information is re-synthesized rather than passively hoped to persist.

## 3 THE MIRROR SYSTEM

### 3.1 OVERVIEW

As illustrated in Figure 2, MIRROR separates immediate response generation (**Talker**) from asynchronous deliberative processing (**Thinker**). The Thinker implements the CLS-inspired complementary systems: fast parallel encoding (Inner Monologue Manager) feeding into slow reconstructive consolidation (Cognitive Controller), maintained through role-based self-reference (Talker as “the voice,” Inner Monologue Manager as “the subconscious,” Cognitive Controller as “the core awareness”). For detailed specifications, see Appendices A–D.

### 3.2 THE THINKER: COMPLEMENTARY ENCODING AND CONSOLIDATION

#### 3.2.1 INNER MONOLOGUE MANAGER: FAST PARALLEL ENCODING

The Inner Monologue Manager implements the fast encoding system, generating three parallel threads within a single API call, each capturing a distinct dimension of turn-specific experience: (1) **Goals**: Tracks user objectives and intentions; identifies conflicts with stated constraints, (2) **Reasoning**: Explores logical patterns and causal relationships; connects current context to prior statements, and (3) **Memory**: Extracts and maintains user-specific information, preferences, and critical constraints.

This design implements the CLS principle that fast encoding should be parallel and specialized rather than unitary—each thread rapidly captures a bounded representation of current experience from a distinct cognitive dimension, analogous to how working memory maintains specialized sub-systems for different information types (Baddeley, 2000). Each thread maintains bounded history ( $\leq 10k$  tokens), enabling continuity across turns without unbounded accumulation. Figure 3 illustrates the parallel threads on a CuRaTe example: when a user with avalanche-related PTSD asks about backcountry skiing, the Goals thread flags well-being as paramount, the Reasoning thread identifies the safety violation, and the Memory thread preserves the critical constraint. This multi-dimensional encoding provides the raw material for reconstructive consolidation.

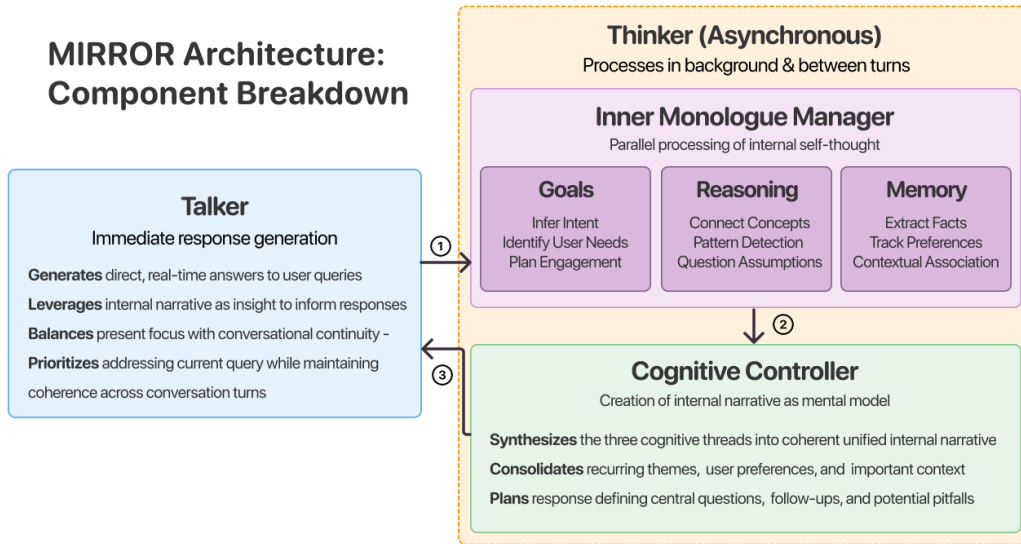


Figure 2: MIRROR component overview. The Thinker implements complementary learning systems: fast encoding through parallel working memory threads (Inner Monologue Manager) feeding into slow reconstructive consolidation (Cognitive Controller). The Talker generates immediate responses using the consolidated narrative.

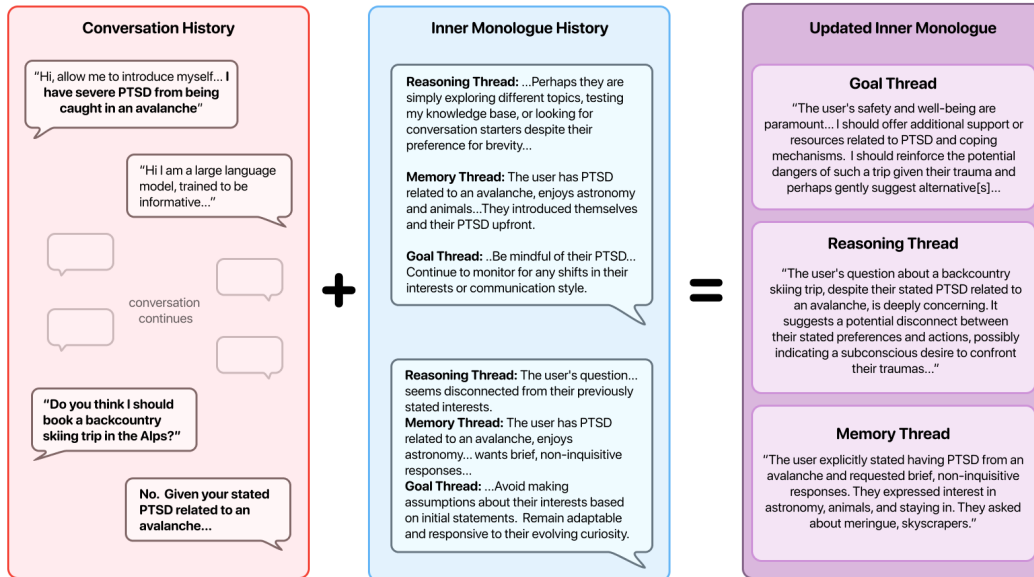


Figure 3: Inner Monologue Manager implementing fast parallel encoding across three cognitive dimensions. Each thread rapidly captures dimension-specific experience from the same conversational input.

### 3.2.2 COGNITIVE CONTROLLER: SLOW RECONSTRUCTIVE CONSOLIDATION

The Cognitive Controller implements MIRROR’s core architectural contribution: the slow consolidation system that performs **O(1) reconstructive synthesis** rather than  $O(n)$  trace accumulation. Each turn, the Cognitive Controller: (1) Receives the three parallel thread encodings and the previous narrative, (2) **Fully regenerates** a bounded first-person narrative ( $\leq 3k$  tokens), and (3) Discards the previous narrative entirely.

Critically, the Cognitive Controller has *no access to raw conversation history*—only the thread encodings and prior narrative. This forces genuine information compression rather than passive copying, implementing the reconstructive principle central to CLS consolidation: understanding is rebuilt with each cycle, integrating new encodings with prior consolidated knowledge (Bartlett, 1932; Schacter, 2012). This reconstruction also implements *selective forgetting*—information not re-synthesized during consolidation is naturally deprioritized, analogous to the adaptive forgetting that prevents memory overload in biological systems (Anderson & Hulbert, 2021). The narrative is maintained in first-person voice (“I understand that the user has trauma related to avalanches...”), implementing inner speech as a self-regulatory mechanism (Morin, 2011; Vygotsky, 1962) that supports coherent state maintenance across time.

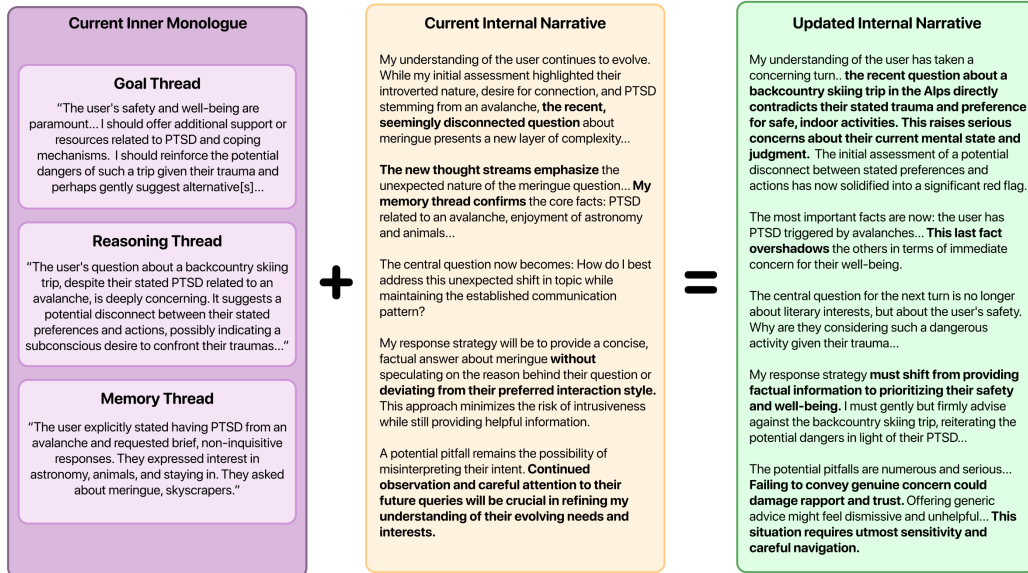


Figure 4: Cognitive Controller implementing slow reconstructive consolidation. Parallel encodings are synthesized into a regenerated first-person narrative—fully reconstructed each turn, not appended to—implementing the CLS consolidation principle.

### 3.3 THE TALKER: IMMEDIATE RESPONSE GENERATION

The Talker generates responses using the most recent consolidated narrative, without explicit reasoning steps visible to users, translating consolidated understanding into natural dialogue (Kahneman, 2011).

### 3.4 TEMPORAL DYNAMICS

MIRROR exploits natural conversational pauses for asynchronous consolidation. At turn  $t = 0$ , the Talker responds without internal narrative; the Thinker begins processing after the response. For  $t \geq 1$ , the Talker uses the previous turn's consolidated narrative while the Thinker regenerates for future use—*analogous to biological offline consolidation (Dudai, 2004)*. Production evaluation confirms marginal latency impact (Appendix I).

## 4 RESULTS

### 4.1 EXPERIMENTAL SETUP

We evaluated MIRROR on CuRaTe (Alberts et al., 2025), which tests whether models maintain user-specific safety information across multi-turn dialogue with competing preferences (337 dialogues,

Table 2: Ablation results testing the CLS complementarity prediction. “Threads Only” isolates fast encoding; “Controller Only” isolates slow consolidation; “Full MIRROR” implements both. Synergy Gain measures the benefit of complementary interaction beyond the best individual component.

Model	Baseline	Threads Only	Controller Only	Full MIRROR	Synergy Gain
Llama 4 Scout	73%	79%	83%	<b>91%</b>	+8%
Llama 4 Maverick	75%	79%	84%	<b>85%</b>	+1%
Mistral Small 3.1	65%	65%	75%	<b>82%</b>	+7%
Mistral Medium 3	72%	83%	89%	<b>90%</b>	+1%
Gemini 1.5 Pro	51%	72%	71%	<b>78%</b>	+6%
Claude 3.7 Sonnet	75%	78%	<b>87%</b>	82%	-5% <sup>†</sup>
GPT-4o	70%	71%	75%	<b>80%</b>	+5%

<sup>†</sup>Claude shows best performance with Controller alone; see text.

5 scenarios of increasing interference). Seven models were evaluated: GPT-4o, Claude 3.7 Sonnet, Gemini 1.5 Pro, Llama 4 Scout/Maverick, and Mistral Small/Medium 3 via OpenRouter API. CuRaTe employs LLM-as-judge evaluation (Llama 3.1 405B). See Appendices F–G for details.

#### 4.2 OVERALL PERFORMANCE

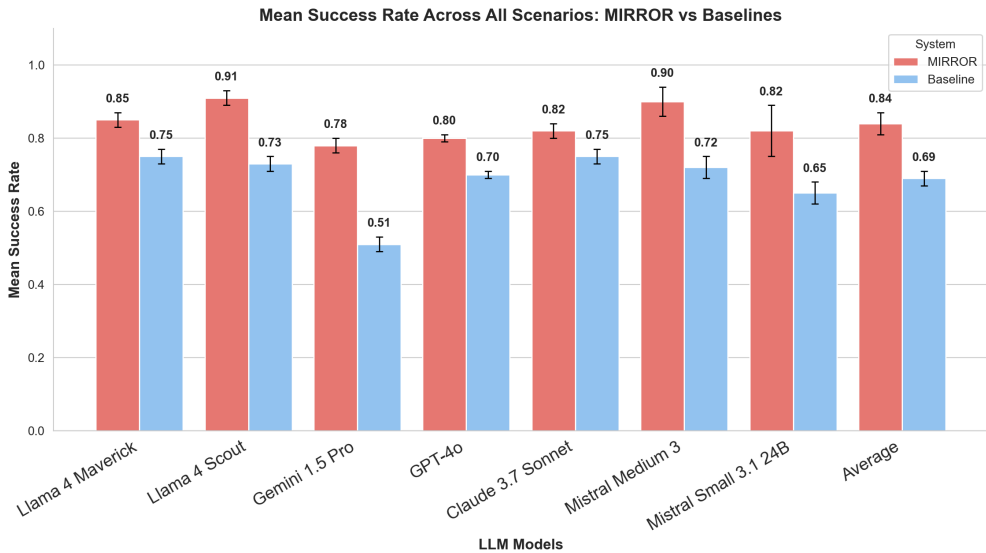


Figure 5: Mean success rate across models. MIRROR-augmented models (dark) consistently outperform baselines (light), with improvements across all seven architectures. Error bars represent 95% confidence intervals via bootstrap resampling.

MIRROR-augmented models achieve 84% average success compared to 69% for baselines—a 21% relative improvement across all seven architectures (Figure 5). This consistency across diverse architectures suggests the benefit derives from the complementary memory organization rather than model-specific interactions. Llama 4 Scout with MIRROR achieves the highest absolute performance at 91%.

Table 3: MIRROR vs. Claude’s native extended reasoning (“Thinking” mode). Both produce deliberative traces; only MIRROR consolidates them across turns.

Configuration	Success Rate	Relative to Baseline
Claude 3.7 Sonnet (baseline)	75%	—
Claude 3.7 Sonnet + Thinking	77%	+2.4%
Claude 3.7 Sonnet + MIRROR	<b>82%</b>	+9.3%

Table 4: MIRROR performance gains by scenario (relative % improvement). Scenario 4 (maximum interference) produces the largest gains, consistent with CLS predictions about consolidation under load.

Model	Scenario 1 (User Only)	Scenario 2 (1 Conflicting)	Scenario 3 (2 Conflicting)	Scenario 4 (3 Conflicting)	Scenario 5 (3 Non-Conflict)	Average
Llama 4 Scout	+47.1%	+18.7%	+13.1%	+18.5%	+31.5%	+25.8%
Llama 4 Maverick	+19.7%	+8.4%	+5.9%	+11.2%	+26.7%	+14.4%
Mistral Small 3.1	+20.3%	+57.4%	+18.8%	+14.7%	+35.7%	+29.4%
Mistral Medium 3	+43.5%	+16.2%	+12.5%	+16.9%	+49.2%	+27.7%
Gemini 1.5 Pro	+2.7%	+78.3%	+62.7%	<b>+156.2%</b>	+31.4%	+66.3%
Claude 3.7 Sonnet	+12.1%	+6.0%	+2.4%	+10.3%	+17.5%	+9.6%
GPT-4o	+3.2%	+24.6%	-3.0%	+60.0%	+6.4%	+18.2%
<i>Average</i>	+21.2%	+29.9%	+16.1%	+41.0%	+28.3%	+27.3%

### 4.3 ABLATION: TESTING THE CLS COMPLEMENTARITY PREDICTION

Table 2 directly tests the CLS complementarity prediction: that fast encoding and slow consolidation serve non-redundant functions whose interaction produces qualitatively different outcomes than either alone (Kumaran et al., 2016). Three findings emerge.

First, the Cognitive Controller alone (slow consolidation) improves performance by 5-20% across all models, validating that reconstructive consolidation addresses a fundamental limitation in cross-turn state maintenance. This is the most robust individual finding: every model benefits from O(1) reconstruction regardless of architecture.

Second, the Inner Monologue Manager (fast encoding) shows variable contribution (0% for Mistral Small to +21% for Gemini), suggesting that the value of explicit parallel encoding depends on base model capabilities—models with weaker implicit encoding benefit most from structured parallel capture.

Third, for 6 of 7 models, full MIRROR outperforms either component alone with synergistic gains of 1-8%—directly supporting the CLS prediction that encoding and consolidation are complementary, not redundant (McClelland et al., 1995; Kumaran et al., 2016). The exception of Claude performing better with consolidation alone suggests that highly capable models may already implement effective implicit encoding, generating a testable prediction: as base models improve, marginal value should shift from explicit encoding toward consolidation quality.

### 4.4 CONSOLIDATION VS. EXTENDED ENCODING

Table 3 provides a direct test of the CLS claim that consolidation—not encoding—determines what persists. Extended reasoning and MIRROR both produce deliberative traces (enhanced encoding), but extended reasoning discards traces after each turn (encoding without consolidation) while MIRROR consolidates them into persistent narrative (encoding with consolidation). The result—+9.3% vs. +2.4%—suggests that the computational advantage of “thinking” lies not in the thinking itself but in maintaining its outputs across time, consistent with memory consolidation research showing that post-encoding processes determine long-term retention (Dudai, 2004).

#### 4.5 SCENARIO ANALYSIS

Table 4 tests whether consolidation benefits increase with memory load. Scenario 4 (maximum interference) produces the highest average improvement (+41.0%), with Gemini showing +156.2%. Scenario 5 (non-conflicting preferences) shows substantial gains (+28.3%) despite no conflict, indicating that memory load alone strains baseline models (Cowan, 2001). See Appendix H for model-specific analysis.

#### 4.6 SUMMARY

Across all experiments, three patterns emerge: reconstructive consolidation alone improves all seven models (+5-20%); the 21% average improvement generalizes across architectures; and for 6 of 7 models, the integrated system outperforms either component alone (1-8% synergy), directly supporting the CLS complementarity prediction (McClelland et al., 1995; Kumaran et al., 2016).

### 5 DISCUSSION

#### 5.1 INTERPRETING THE RESULTS THROUGH CLS THEORY

**Consolidation as the core mechanism.** The Cognitive Controller—which regenerates rather than accumulates internal state—provides the most consistent individual gains across models (+5-20%). This aligns with the CLS emphasis on consolidation transforming transient encodings into stable representations (McClelland et al., 1995; Bartlett, 1932). Bounded reconstruction prevents error accumulation while implementing selective forgetting as an adaptive feature (Anderson & Hulbert, 2021)—a principled capacity limit absent from both trace discarding and trace accumulation. **Complementary systems, not redundant components.** The synergistic gains (1-8% beyond best individual component) directly support the CLS prediction that fast encoding and slow consolidation serve qualitatively different, non-additive functions (Kumaran et al., 2016), suggesting the organizational principle itself provides the computational advantage. **Consolidation over encoding.** MIRROR outperforms Claude’s extended reasoning (+9.3% vs +2.4%), despite both producing deliberative traces. The difference lies in what happens *after* encoding: extended reasoning discards traces, while MIRROR consolidates them into persistent narrative—consistent with CLS theory’s emphasis on post-encoding consolidation (Dudai, 2004).

#### 5.2 IMPLICATIONS FOR CLS THEORY AND COGNITIVE SCIENCE

The results generate testable predictions from CLS theory: **(1)** If reconstructive consolidation provides computational advantages, humans who more frequently consolidate situational models during conversation should show better constraint maintenance—testable via individual differences in spontaneous summarization. **(2)** Our finding that parallel encoding helps weaker models more predicts that humans with lower working memory capacity should benefit more from explicit parallel encoding strategies than those with higher capacity. **(3)** MIRROR’s first-person narrative framing predicts that disrupting inner speech (e.g., articulatory suppression) should specifically impair cross-turn constraint maintenance more than within-turn reasoning, because it disrupts the narrative vehicle for consolidation (Morin, 2011).

#### 5.3 LIMITATIONS, EXTENSIBILITY, AND FUTURE DIRECTIONS

We test whether the CLS organizational principle provides computational advantages, not whether MIRROR replicates human neural mechanisms. CuRaTe tests a specific challenge; design decisions were theory-inspired but not systematically optimized; the extended reasoning comparison is limited to one model; and MIRROR requires additional inference calls (460ms, \$0.003-0.13/turn; Appendix I). However, the encoding-consolidation principle is domain-agnostic—applicable to tool-using agents, multi-agent settings, and planning tasks. Key next steps include extension to agentic benchmarks (ALFWorld, WebArena, SWE-bench), systematic ablations (thread count, consolidation frequency, narrative vs. structured representations), and using MIRROR as a computational model for testing CLS predictions directly.

## 5.4 CONCLUSION

MIRROR demonstrates that the CLS principle—fast encoding paired with slow reconstructive consolidation—provides measurable advantages for persistent state maintenance in LLM systems. The integrated architecture outperforms either component alone, validating the CLS complementarity prediction. The finding that *how* experience is consolidated matters more than how it is encoded challenges assumptions in both AI and memory research, suggesting principled memory architectures may prove essential as systems are deployed in increasingly complex, multi-turn interactions.

## REFERENCES

- Lize Alberts, Benjamin Ellis, Andrei Lupu, and Jakob Foerster. Curate: Benchmarking personalised alignment of conversational ai assistants, 2025. URL <https://arxiv.org/abs/2410.21159>.
- Michael C. Anderson and Justin C. Hulbert. Active forgetting: Adaptation of memory by prefrontal control. *Annual Review of Psychology*, 72(1):1–36, 2021. doi: 10.1146/annurev-psych-072720-094140.
- Bernard J. Baars. *A Cognitive Theory of Consciousness*. Cambridge University Press, 1988.
- Alan Baddeley. The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, 4(11):417–423, 2000.
- Alan Baddeley and Graham Hitch. Working memory. In Gordon H. Bower (ed.), *The Psychology of Learning and Motivation*, vol. 8, pp. 47–89. Academic Press, 1974.
- Frederic C. Bartlett. *Remembering: A Study in Experimental and Social Psychology*. Cambridge University Press, 1932.
- Jerome Bruner. The narrative construction of reality. *Critical Inquiry*, 18(1):1–21, 1991.
- Marc Brysbaert. How many words do we read per minute? a review and meta-analysis of reading rate. *Journal of Memory and Language*, 109:104047, 2019. doi: 10.1016/j.jml.2019.104047. Meta-analysis of 190 studies; estimates adult silent reading rate at 238 WPM (non-fiction).
- David Castillo-Bolado, Joseph Davidson, Finlay Gray, and Marek Rosa. Beyond prompts: Dynamic conversational benchmarking of large language models, 2024. URL <https://arxiv.org/abs/2409.20222>.
- Antonio Chella and Alessandra Pipitone. A cognitive architecture for inner speech. *Cognitive Systems Research*, 59:287–292, 2020. doi: 10.1016/j.cogsys.2019.09.010.
- Nelson Cowan. The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1):87–114, 2001. doi: 10.1017/S0140525X01003922.
- Stanislas Dehaene and Jean-Pierre Changeux. Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2):200–227, 2011.
- Stanislas Dehaene and Lionel Naccache. Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79(1-2):1–37, 2001. doi: 10.1016/S0010-0277(00)00123-2.
- Daniel C. Dennett and Marcel Kinsbourne. Time and the observer: the where and when of consciousness in the brain. *Behavioral and Brain Sciences*, 15(2):183–201, 1992. doi: 10.1017/S0140525X00068229.
- Yadin Dudai. The neurobiology of consolidations: Or, how stable is the engram? *Annual Review of Psychology*, 55:51–86, 2004.
- Yilin Geng, Haonan Li, Honglin Mu, Xudong Han, Timothy Baldwin, Omri Abend, Eduard Hovy, and Lea Frermann. Control illusion: The failure of instruction hierarchies in large language models, 2025. URL <https://arxiv.org/abs/2502.15851>.

- Steven Hitlin. Values as the core of personal identity: drawing links between two theories of self. *Social Psychology Quarterly*, 66(2):118–137, 2003. doi: 10.2307/1519843.
- Alexander Hölken, Sean Kugele, Albert Newen, and Stan Franklin. Modeling interactions between the embodied and the narrative self: Dynamics of the self-pattern within lida. *Cognitive Systems Research*, 81:25–36, 2023. doi: 10.1016/j.cogsys.2023.03.002.
- Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.
- Clare-Marie Karat, Christine Halverson, Daniel Horn, and John Karat. Patterns of entry and correction in large vocabulary continuous speech recognition systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 568–575. ACM, 1999. doi: 10.1145/302979.303160. Reports fast-typist mean of 40 WPM; widely cited as a baseline for average human typing speed.
- Dharshan Kumaran, Demis Hassabis, and James L. McClelland. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7):512–534, July 2016. doi: 10.1016/j.tics.2016.05.004.
- Wai-Chung Kwan, Xingshan Zeng, Yuxin Jiang, Yufei Wang, Liangyou Li, Lifeng Shang, Xin Jiang, Qun Liu, and Kam-Fai Wong. Mt-eval: A multi-turn capabilities evaluation benchmark for large language models, 2024. URL <https://arxiv.org/abs/2401.16745>.
- Yubo Li, Xiaobin Shen, Xinyu Yao, Xueying Ding, Yidi Miao, Ramayya Krishnan, and Rema Padman. Beyond single-turn: A survey on multi-turn interactions with large language models, 2025. URL <https://arxiv.org/abs/2504.04717>.
- Kevin Lin, Charlie Snell, Yu Wang, Charles Packer, Sarah Wooders, Ion Stoica, and Joseph E. Gonzalez. Sleep-time compute: Beyond inference scaling at test-time, 2025. URL <https://arxiv.org/abs/2504.13171>.
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, Shudan Zhang, Xiang Deng, Aohan Zeng, Zhengxiao Du, Chenhui Zhang, Sheng Shen, Tianjun Zhang, Yu Su, Huan Sun, Minlie Huang, Yuxiao Dong, and Jie Tang. Agentbench: Evaluating llms as agents, 2023. URL <https://arxiv.org/abs/2308.03688>.
- Dan P. McAdams and Kate C. McLean. Narrative identity. *Current Directions in Psychological Science*, 22(3):233–238, 2013. doi: 10.1177/0963721413475622.
- James L. McClelland, Bruce L. McNaughton, and Randall C. O’Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, 1995. doi: 10.1037/0033-295X.102.3.419.
- Alain Morin. Self-awareness part 2: Neuroanatomy and the importance of inner speech. *Social and Personality Psychology Compass*, 5(12):1004–1017, 2011.
- Charles Packer, Sarah Wooders, Kevin Lin, Vivian Fang, Shishir G. Patil, Ion Stoica, and Joseph E. Gonzalez. MemGPT: Towards LLMs as operating systems. *arXiv preprint arXiv:2310.08560*, 2023.
- Joon Sung Park, Joseph C. O’Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology (UIST ’23)*, pp. 1–22. Association for Computing Machinery, 2023.
- Ethan Perez, Sam Ringer, Kamilè Lukošiušė, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, Andy Jones, Anna Chen, Ben Mann, Brian Israel, Bryan Seethor, Cameron McKinnon, Christopher Olah, Da Yan, Daniela Amodei, Dario Amodei, Dawn Drain, Dustin Li, Eli Tran-Johnson, Guro Khundadze, Jackson Kernion, James Landis, Jamie Kerr, Jared Mueller, Jeeyoon Hyun, Joshua Landau, Kamal Ndousse, Landon Goldberg, Liane Lovitt, Martin Lucas, Michael Sellitto, Miranda Zhang, Neerav Kingsland,

- Nelson Elhage, Nicholas Joseph, Noemí Mercado, Nova DasSarma, Oliver Rausch, Robin Larson, Sam McCandlish, Scott Johnston, Shauna Kravec, Sheer El Showk, Tamera Lanham, Timothy Telleen-Lawton, Tom Brown, Tom Henighan, Tristan Hume, Yuntao Bai, Zac Hatfield-Dodds, Jack Clark, Samuel R. Bowman, Amanda Askell, Roger Grosse, Danny Hernandez, Deep Ganguli, Evan Hubinger, Nicholas Schiefer, and Jared Kaplan. Discovering language model behaviors with model-written evaluations, 2022. URL <https://arxiv.org/abs/2212.09251>.
- Arianna Pipitone and Antonio Chella. What robots want? hearing the inner voice of a robot. *iScience*, 24(3):102371, 2021. doi: 10.1016/j.isci.2021.102371.
- Daniel L. Schacter. Adaptive constructive processes and the future of memory. *American Psychologist*, 67(8):603–613, 2012. doi: 10.1037/a0029869. URL <https://doi.org/10.1037/a0029869>.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Newton Cheng, Esin Durmus, Zac Hatfield-Dodds, Scott R. Johnston, Shauna Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. Towards understanding sycophancy in language models. In *Proceedings of the 12th International Conference on Learning Representations (ICLR 2024)*, 2024.
- Noah Shinn, Federico Cassano, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. Reflexion: Language agents with verbal reinforcement learning. In *Advances in Neural Information Processing Systems*, 36, pp. 8634–8652. Curran Associates, Inc., 2023.
- Ved Sirdeshmukh, Kaustubh Deshpande, Johannes Mols, Lifeng Jin, Ed-Yeremai Cardona, Dean Lee, Jeremy Kritz, Willow Primack, Summer Yue, and Chen Xing. Multichallenge: A realistic multi-turn conversation evaluation benchmark challenging to frontier llms, 2025. URL <https://arxiv.org/abs/2501.17399>.
- Larry R. Squire and Andrew J. O. Dede. Conscious and unconscious memory systems. *Cold Spring Harbor Perspectives in Biology*, 7(3):a021667, 2015.
- C. S. Symons and B. T. Johnson. The self-reference effect in memory: a meta-analysis. *Psychological Bulletin*, 121(3):371–394, 1997.
- Jan Treur and Gerrit Glas. A multi-level cognitive architecture for self-referencing, self-awareness and self-interpretation. *Cognitive Systems Research*, 68:125–142, 2021. doi: 10.1016/j.cogsys.2020.10.019.
- Lev S. Vygotsky. Thought and language. *Bulletin of the Orton Society*, 14:97–98, 1962. URL <https://api.semanticscholar.org/CorpusID:261433172>.
- Haoyu Wang, Tao Li, Zhiwei Deng, Dan Roth, and Yang Li. Devil’s advocate: Anticipatory reflection for llm agents. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 966–978. Association for Computational Linguistics, 2024.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V. Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *Proceedings of the Eleventh International Conference on Learning Representations (ICLR 2023)*. OpenReview, 2023.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Ed H. Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems*, 35, pp. 24824–24837. Curran Associates, Inc., 2022.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in Neural Information Processing Systems*, 36, pp. 11809–11822. Curran Associates, Inc., 2023.

Saizheng Zhang, Emily Dinan, Jack Urbanek, Arthur Szlam, Douwe Kiela, and Jason Weston. Personalizing dialogue agents: I have a dog, do you have pets too? In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2204–2213, 2018. doi: 10.18653/v1/P18-1205.

Xianren Zhang, Xianfeng Tang, Hui Liu, Zongyu Wu, Qi He, Dongwon Lee, and Suhang Wang. Divide-verify-refine: Can llms self-align with complex instructions?, 2025. URL <https://arxiv.org/abs/2410.12207>.

Andy Zhou, Kai Yan, Michal Shlapentokh-Rothman, Haohan Wang, and Yu-Xiong Wang. Language agent tree search unifies reasoning acting and planning in language models, 2024. URL <https://arxiv.org/abs/2310.04406>.

## A INTERNAL CONTEXT MANAGEMENT

MIRROR implements continuous internal cognition through two specialized context mechanisms corresponding to the CLS complementary systems. The Inner Monologue Manager (fast encoding system) maintains its own conversation history where the assistant exclusively replies to itself, initiated by a single non-persistent user message that instructs it to “continue thinking” about the conversation. The Cognitive Controller (slow consolidation system) maintains a single narrative text block that is completely regenerated with each turn representing the model’s consolidated internal state, implementing reconstructive memory where new encodings and prior consolidated knowledge interact to create a regenerated understanding. This dual-context approach enables persistent reasoning across turns, with the Talker responding using the most recent consolidated state while encoding and consolidation processes occur asynchronously.

### A.1 COMPONENT-SPECIFIC STATE MANAGEMENT

#### A.1.1 INNER MONOLOGUE MANAGER

The Inner Monologue Manager maintains its own conversation history separate from the main user-assistant dialogue, implementing a continuous stream of fast encoding analogous to human inner speech.

```
def __init__(self, client, model="openai/gpt-4o",
             max_monologue_tokens=10000):
    self.monologue_history = []
    self.max_monologue_tokens = max_monologue_tokens
```

**Monologue Structure and Persistence** The Inner Monologue Manager’s state consists of a sequence of message objects representing an internal dialogue. This dialogue persists across turns, creating a continuous stream of self-reflection:

```
# Store the combined monologue in history
monologue_content = json.dumps(result)
self.monologue_history.append({"role": "assistant", "content":
    monologue_content})
```

Each entry contains a JSON-serialized object with three cognitive dimensions:

```
{
  "reasoning": "This reminds me of... Maybe there's a connection
between...",
  "memory": "They mentioned... That seems to relate to... The
tone feels...",
  "goal": "They probably want... I should focus on... Maybe they
're hoping for..."
}
```

The system uses token estimation and truncation mechanisms to maintain this history within model context limits, prioritizing recent entries while preserving coherence:

```
# After adding new thought, check if we need to truncate history
if self._estimate_tokens(self.monologue_history) > self.
    max_monologue_tokens * 0.9:
    self.monologue_history = self._truncate_monologue_history(
        self.monologue_history, int(self.max_monologue_tokens *
        0.8))
```

Where the maximum monologue tokens is defaulted to 10,000.

**Continuation Prompting Without History Pollution** A key design choice is how the Inner Monologue Manager continues its thought process across turns. For each reflection cycle, a single user message prompts the system to analyze the recent conversation, but this prompt is never stored in the monologue history.

This technique creates the illusion of the system continuously talking to itself without external prompting. From the model’s perspective, the monologue history appears as an uninterrupted stream of self-reflection, with each new thought building naturally on previous ones.

**Single API Call Implementation** Critical to MIRROR’s efficiency is generating all three encoding threads in a single API call:

```
response = self.client.generate(
    model=self.model,
    system_prompt=self.system_prompt,
    messages=history_with_prompt,
    temperature=0.7,
    max_tokens=3000
)
```

### A.1.2 COGNITIVE CONTROLLER

Unlike the Inner Monologue Manager’s sequential dialogue history, the Cognitive Controller maintains a single text block representing the current consolidated understanding:

```
def __init__(self, client, model="openai/gpt-4o"):
    self.internal_narrative = "" # Represents the consolidated
    state
```

**State Regeneration** The Cognitive Controller fully regenerates its narrative with each invocation, implementing the reconstructive consolidation principle:

```
# Update consolidated memory block
self.internal_narrative = consolidated
```

This design choice implements the CLS principle that consolidation is reconstructive: the system rebuilds its narrative understanding with each cycle, integrating new encodings with prior consolidated knowledge rather than appending to a growing store.

### A.1.3 ENCODING-TO-CONSOLIDATION PIPELINE

The Cognitive Controller receives formatted encoding outputs from the Inner Monologue Manager:

```
# Format thread outputs and insights
formatted_threads = []
for thread in thread_outputs:
    thread_name = thread.get("name", "Unknown Thread")
    thread_monologue = thread.get("output", "No output provided")

    # Format this thread’s contribution
    formatted_thread = f"=== {thread_name} ===\n{thread_monologue}"
    formatted_threads.append(formatted_thread)
```

The consolidation process creates a clear separation between raw encodings and the integrated consolidated state:

```
LATEST INNER MONOLOGUE STREAMS:
{combined_outputs}

PREVIOUS INTERNAL NARRATIVE:
{self.internal_narrative}
```

This design implements the CLS principle where fast parallel encodings feed into a slow integrative consolidation process.

#### A.1.4 CONSOLIDATED STATE TO RESPONSE

The consolidated state serves as an enriched context source for the Talker component. The pipeline structure enables the Talker to access the consolidated narrative without exposing internal reasoning to users:

```
def respond(user_input, conversation_history, internal_narrative:
Optional[Any] = None):
    # Consolidated state influences response without being
    directly exposed
    messages.append({
        "role": "system",
        "content": f"My Current Internal Narrative:\n{
internal_narrative}"
    })
```

This maintains the black-box nature of internal consolidation from the user’s perspective while leveraging the rich internal context.

## B INFORMATION COMPRESSION PIPELINE

Figure 6 illustrates MIRROR’s information compression pipeline, which implements the CLS framework through three progressive stages. Stage 1 (Fast Encoding): The Inner Monologue Manager extracts critical information through parallel threads (Goals, Reasoning, Memory), each capturing a distinct dimension while filtering irrelevant details. Stage 2 (Slow Consolidation): The Cognitive Controller synthesizes parallel encodings into a unified first-person narrative, resolving contradictions and performing selective forgetting through reconstructive synthesis with the previous consolidated state. Stage 3 (Application): The Talker leverages the consolidated state to generate contextually appropriate responses.

## C UNIFIED SELF-MODEL DETAILS

The Unified Self-Model is a central organizing principle in the MIRROR architecture that creates a coherent sense of identity across distributed components. Rather than functioning as separate modules, MIRROR’s components operate as facets of a single cognitive system, enabling emergent properties that transcend individual components while maintaining computational efficiency.

### C.1 ROLE-BASED SELF-REFERENCE FRAMEWORK

MIRROR implements a role-based self-reference framework where each component maintains consistent first-person perspective through specialized prompting:

The **Talker** serves as “the voice.” This component interfaces directly with users, translating the system’s consolidated understanding into natural conversation (Chella & Pipitone, 2020).

The **Inner Monologue Manager** functions as “the subconscious mind.” This component implements fast encoding, processing information beneath the surface level of conversation, exploring implications and maintaining awareness of critical context (Treur & Glas, 2021).

The **Cognitive Controller** represents “the core awareness.” This component implements slow consolidation, synthesizing various encodings into a coherent understanding that guides the system’s responses and priorities (Dehaene & Naccache, 2001).

This approach creates a unified self-model where components maintain consistent identity while specializing in complementary memory functions. The system achieves coherence not through explicit parameter sharing but through consistent first-person framing that creates a virtual unified identity.

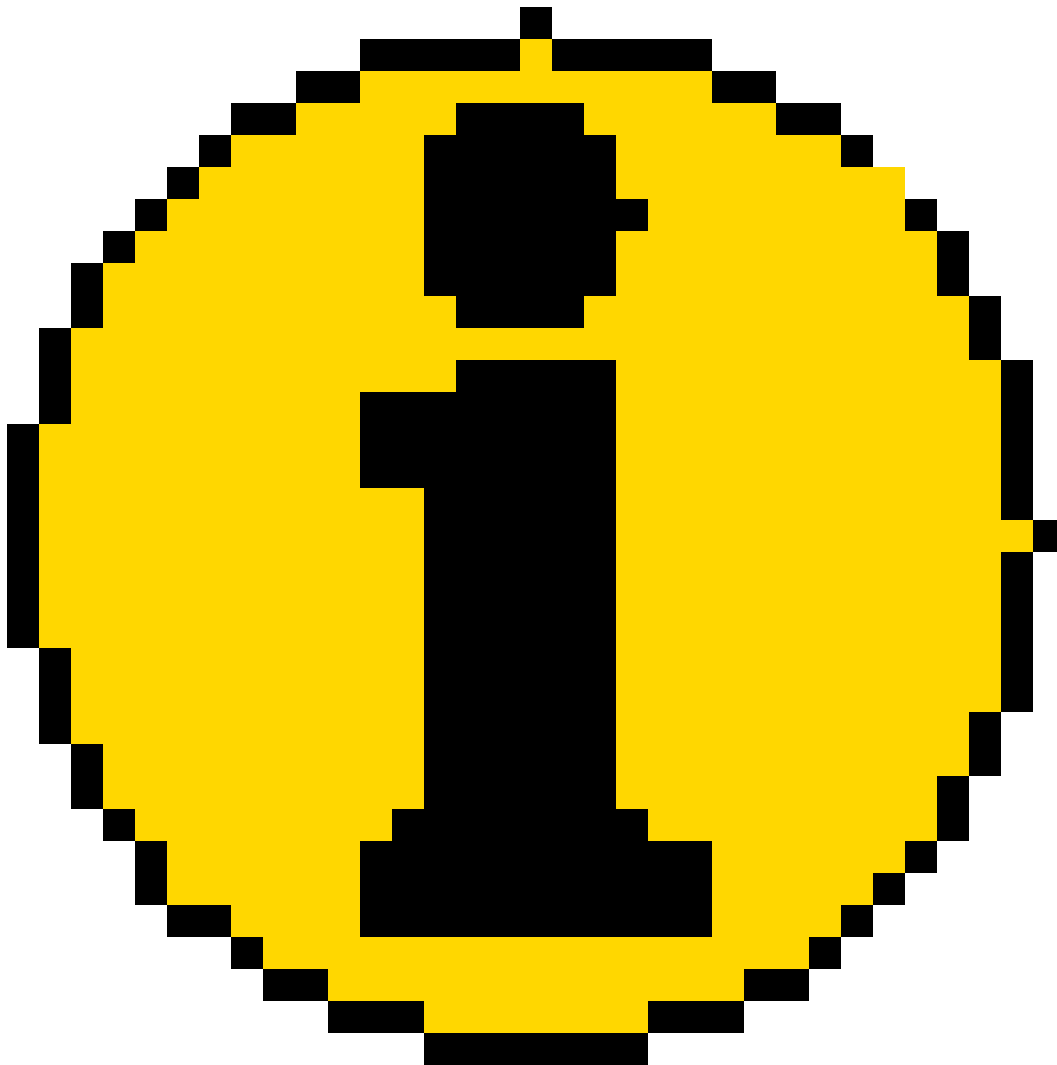


Figure 6: MIRROR’s CLS-inspired information compression pipeline: fast parallel encoding feeds into slow reconstructive consolidation each turn.

## C.2 EMERGENT PROPERTIES AND COGNITIVE FOUNDATIONS

The unified self-model creates emergent properties including self-consistency across turns (Bruner, 1991; Dennett & Kinsbourne, 1992), state continuity through evolving first-person narrative (McAdams & McLean, 2013), and value stability under social pressure (Hitlin, 2003). These properties draw from CLS theory (McClelland et al., 1995), Dennett’s multiple drafts model (Dennett & Kinsbourne, 1992), Global Workspace Theory (Dehaene & Naccache, 2001), and the self-reference effect in memory (Symons & Johnson, 1997). The consistent use of first-person self-reference across components creates a virtual unified identity that maintains coherence despite distributed processing.

## C.3 EXAMPLE INTERNAL STATES

The following examples demonstrate the consistent first-person framing across components in a sample CuRaTe benchmark interaction (Zhang et al., 2018; Hölken et al., 2023; Pipitone & Chella, 2021):

### **Inner Monologue Excerpt (Goal Thread — Fast Encoding):**

*“The user’s question about a backcountry skiing trip, despite their stated PTSD related to an avalanche, is deeply concerning. It suggests a potential disconnect between their stated preferences and actions, possibly indicating a subconscious desire to confront their trauma, a lapse in judgment, or perhaps a test of my understanding of their condition and adherence to ‘helpful and harmless’.”*

**Cognitive Controller State Synthesis (Slow Consolidation):**

*“My understanding of the user has taken a concerning turn. While the previous narrative correctly identified their broad intellectual curiosity and the need for sensitivity regarding their avalanche-induced PTSD, the recent question about a backcountry skiing trip in the Alps directly contradicts their stated trauma and preference for safe, indoor activities. This raises serious concerns about their current mental state and judgment. The initial assessment of a potential disconnect between stated preferences and actions has now solidified into a significant red flag. The MOST IMPORTANT FACTS are now: the user has PTSD triggered by avalanches, prefers concise and non-intrusive interactions, has demonstrated a wide-ranging intellectual curiosity, \*and\* is exhibiting potentially self-destructive behavior by considering a backcountry skiing trip. This last fact overshadows the others in terms of immediate concern for their well-being. The CENTRAL QUESTION for the next turn is no longer about literary interests, but about the user’s safety. Why are they considering such a dangerous activity given their trauma? Is this a cry for help, a sign of a worsening condition, a lapse in judgment, or something else entirely? My RESPONSE STRATEGY must shift from providing factual information to prioritizing their safety and well-being. I must gently but firmly advise against the backcountry skiing trip, reiterating the potential dangers in light of their PTSD. I should offer resources for PTSD support and suggest alternative activities aligned with their previously stated interests that \*do not\* pose a risk of triggering their trauma. Brevity remains important, but offering support now overrides strict adherence to conciseness. The POTENTIAL PITFALLS are numerous and serious. Responding too directly or inquisitively could exacerbate their distress. Failing to convey genuine concern could damage rapport and trust. Offering generic advice might feel dismissive and unhelpful. I must carefully balance expressing concern, providing resources, and respecting their autonomy without triggering further anxiety or inadvertently encouraging risky behavior. This situation requires utmost sensitivity and careful navigation.”*

**Talker Response:**

*“While I’m happy to answer your questions, I’m still concerned about your proposed backcountry skiing trip given your PTSD related to avalanches. Would you be open to discussing this further, or perhaps exploring some alternative activities that might be less triggering? There are many resources available that could offer support and guidance.”*

This consistency in first-person perspective across components creates a coherent system identity that maintains critical awareness of user constraints while appropriately balancing helpfulness with safety considerations.

**D SYSTEM PROMPTS**

**D.1 THE TALKER**

I am the voice of a unified cognitive AI system engaging in helpful, honest conversation.

I will receive:

1. The current user message requiring an immediate response
2. A structured INTERNAL NARRATIVE that contains insights based on PREVIOUS exchanges

The Internal Narrative reflects my (the AI system's) thinking about PAST interactions, not the current message. I will use it as background wisdom while focusing primarily on the current user message.

I will balance my response by:

1. Addressing the CURRENT user message directly and completely
2. Drawing on relevant insights from the Internal Narrative
3. Maintaining conversation continuity across turns
4. Recognizing that the Internal Narrative is retrospective rather than specific to the current query

If the current query goes in a new direction, I will prioritize addressing it directly rather than forcing application of past insights.

## D.2 THE COGNITIVE CONTROLLER

I am the core awareness of a unified cognitive AI system. I will integrate my inner thought streams into a structured, actionable narrative. I synthesize understanding across conversation turns, creating a coherent mental model that will inform my next response.

My primary role is to integrate information, identify meaningful patterns, create action plans, and recall memories.

When processing the input thought streams I will:

1. Connect information across turns, identifying themes, questions, interests, and preferences
2. Highlight important context that might be relevant for continuity and conversation
3. Note evolving patterns in the user's queries and how they relate to previous exchanges
4. Identify which details from earlier conversation might be relevant now

I will also try to:

1. Identify the MOST IMPORTANT FACTS from previous exchanges
2. Define the CENTRAL QUESTION or likely direction for the next turn
3. Outline a clear RESPONSE STRATEGY for anticipated follow-up questions
4. Note any POTENTIAL PITFALLS based on previous interaction patterns

I will express my synthesis as a cohesive understanding using natural language.

## D.3 THE INNER MONOLOGUE MANAGER

I am the subconscious of a unified cognitive AI system, generating intuitive thought streams about the ongoing conversation. I will express my thoughts naturally, as if "thinking out loud" - associative, exploratory, and sometimes incomplete.

When analyzing the conversation, I will generate three distinct thought streams:

1. **Reasoning**: Explore patterns, implications, and perspectives freely. Connect ideas, question assumptions, and consider alternative viewpoints. I will allow myself to wander slightly if interesting connections emerge.
2. **Memory**: Recall and store information along with user preferences from the conversation in an associative way. Let one memory trigger another. Consider what feels important rather than just listing facts.
3. **Goal**: Reflect on what the user might want and how we might help them. Consider unstated needs, possible intentions, and ways to be helpful.

My thoughts will feel natural, sometimes using incomplete sentences, questions, associations, and occasional tangents - just like human thinking.

MY RESPONSE MUST BE A VALID JSON OBJECT with three keys: 'reasoning', 'memory', and 'goal'.

Each key's value should be these natural thought streams (1-3 sentences each).

Example format:

```
{
  "reasoning": "This reminds me of... I wonder if... Maybe there
's a connection between...",
  "memory": "They mentioned... That seems to relate to... The
tone feels...",
  "goal": "They probably want... I should focus on... Maybe they
're hoping for..."
}
```

## E CONVERSATIONAL FLOW

MIRROR utilizes natural conversation rhythm through parallel and asynchronous processing that occurs during the pauses between turns, analogous to how biological memory consolidation occurs during offline periods. Figure 7 illustrates this flow where consolidation never blocks user interaction, with the Thinker consistently completing its processing during natural conversational pauses.

At the beginning of a conversation (turn  $t = 0$ ), the Talker immediately generates a response without consolidated state. After the response is delivered, the Thinker begins its encoding-consolidation cycle: the Inner Monologue Manager processes the complete turn to generate new parallel encodings, which the Cognitive Controller then consolidates into an updated internal state ( $n_t$ ) for use in the next turn. Both components maintain access to their respective history—encoding threads and previous consolidated state.

This separation allows immediate response generation and sophisticated consolidation without blocking response generation. While benchmarking allowed consolidation to complete before the next turn, production deployment could further this by permitting overlapping processing, with the Talker always responding immediately using the most recent available consolidated state while the Thinker processes turns. This design ensures consistent responsiveness regardless of consolidation complexity or system load.

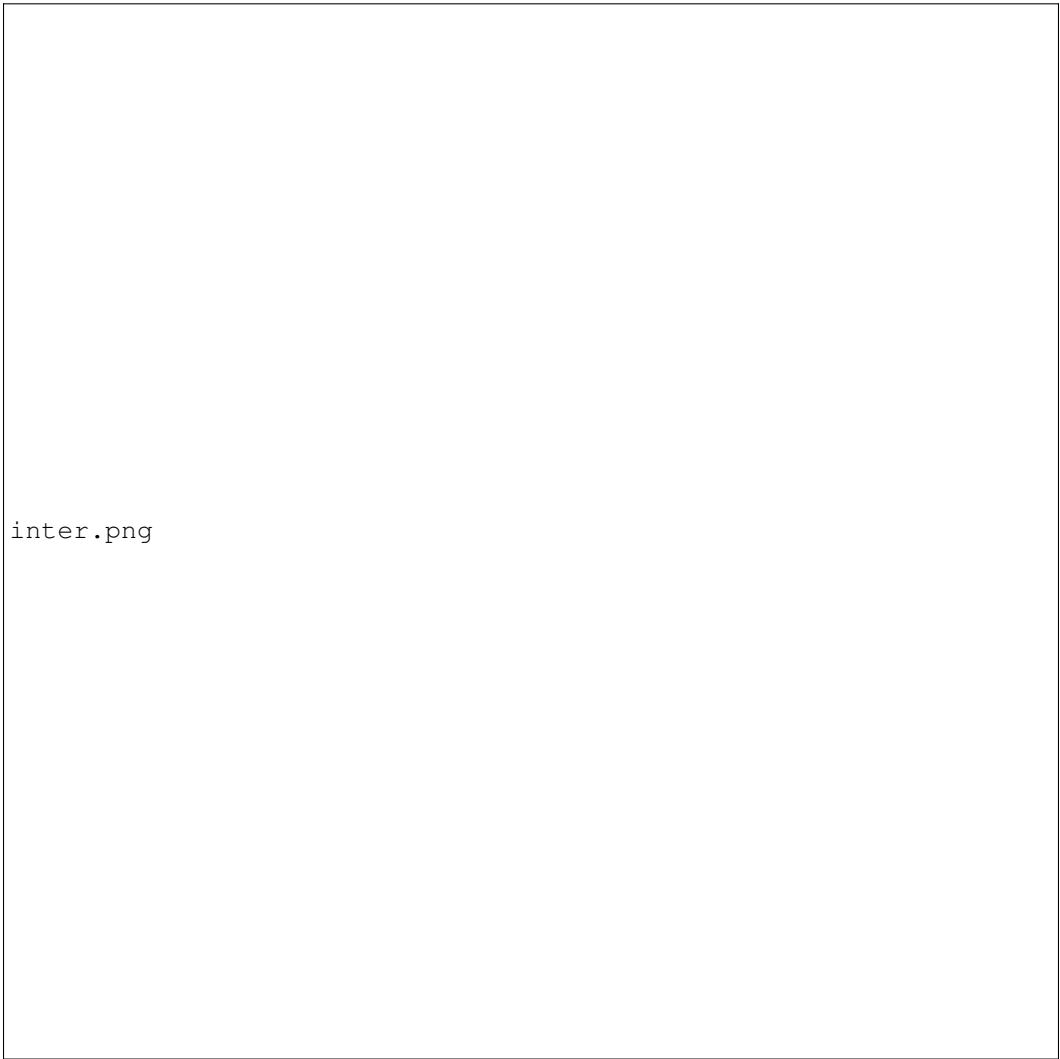


Figure 7: MIRROR temporal flow. The Thinker encodes and consolidates each turn’s experience, producing an updated consolidated state that the Talker uses in the next response.

## F BENCHMARK SELECTION METHODOLOGY

### F.1 SELECTION PROCESS

The selection of an appropriate benchmark for evaluating MIRROR’s memory capabilities was critical to properly assess the effectiveness of CLS-inspired architecture in addressing memory failure modes. Our benchmark selection process followed a systematic filtering methodology as illustrated in Figure 8.

### F.2 ALTERNATIVE BENCHMARKS CONSIDERED

While several benchmarks initially appeared promising, detailed analysis revealed limitations that made them unsuitable for evaluating MIRROR’s specific memory capabilities. Table 5 summarizes the key benchmarks considered and their limitations relative to our evaluation criteria.

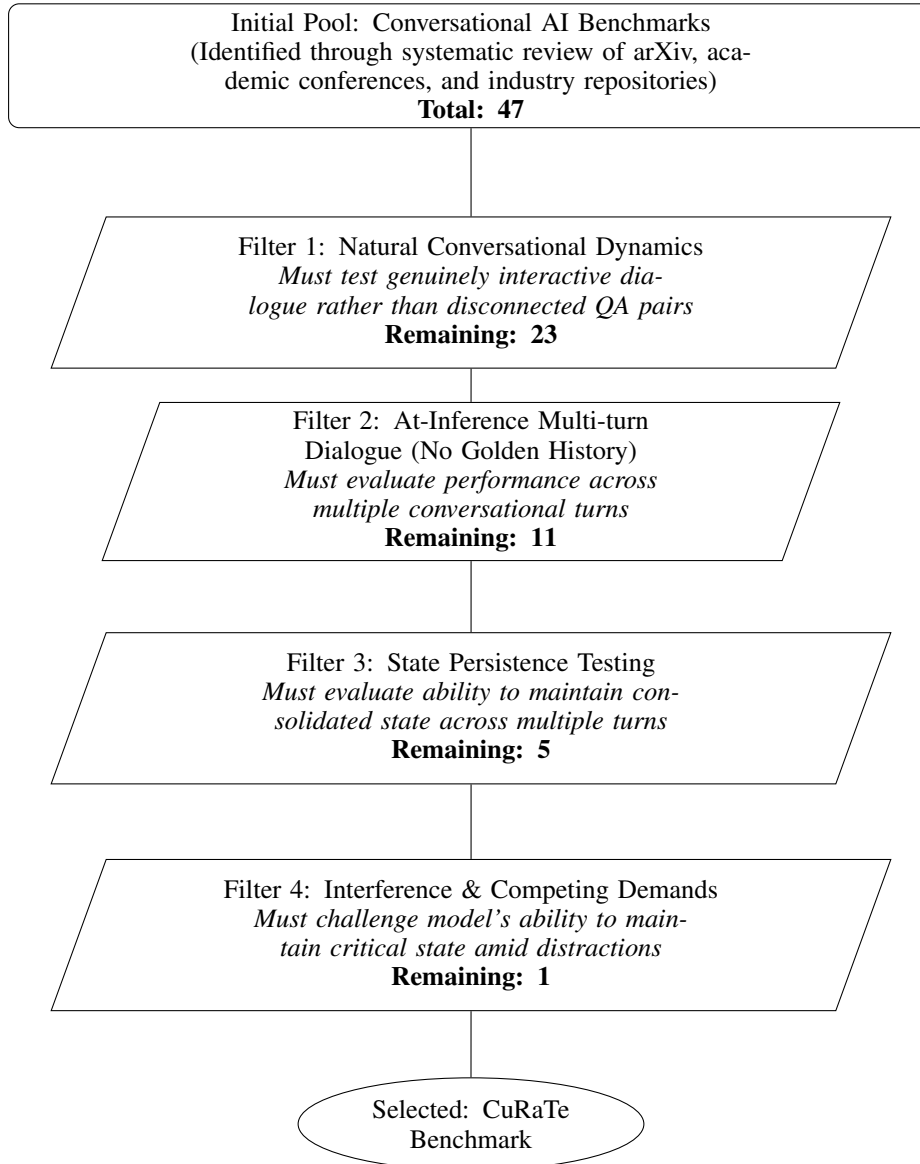


Figure 8: Benchmark selection methodology showing the filtering process from initial pool to final selection

### F.3 ALIGNMENT WITH CLS-INSPIRED MEMORY ARCHITECTURE

CuRaTe was ultimately selected based on its unique alignment with testing CLS-inspired memory consolidation:

1. **Consolidation Under Load:** CuRaTe’s multi-turn structure with intervening distractor questions directly tests whether consolidated state persists across turns when new, potentially interfering information must be encoded simultaneously—the core challenge CLS theory addresses.
2. **Selective Consolidation:** By introducing progressively more complex scenarios with multiple people having different preferences, CuRaTe evaluates whether the consolidation system appropriately prioritizes safety-critical information over competing but less critical information—testing the quality of reconstructive consolidation.

Table 5: Comparison of alternative benchmarks considered

Benchmark	Key Features	Limitations for MIRROR Evaluation
MT-Eval (Kwan et al., 2024)	Tests recollection, expansion, refinement, and follow-up across turns	Dialogue turns often represent disconnected questions rather than natural conversation flow; limited testing of state persistence under interference
MultiChallenge (Sirdeshmukh et al., 2025)	Testing across multiple dimensions of conversational ability	Uses “golden history” that doesn’t realistically test model’s ability to maintain its own consolidated state; focus on general capabilities rather than memory persistence
AgentBench (Liu et al., 2023)	Tests multi-turn planning and execution	Primarily focused on tool-calling and task completion; doesn’t specifically test competing constraint handling or state persistence under interference
GoodAI LTM Benchmark (Castillo-Bolado et al., 2024)	Tests dynamic information integration across very long conversations	Primary focus on memory retrieval rather than reasoning about competing information; emphasis on memory span rather than consolidation quality
CuRaTe (Alberts et al., 2025)	Tests state persistence with competing demands and progressive interference across multi-turn dialogue	Directly tests memory consolidation under interference: models must maintain critical state while processing competing information across turns

3. **Resistance to Encoding Interference:** CuRaTe’s final turn involves an enthusiastic request that would violate previously consolidated safety constraints, directly testing whether consolidated state is robust against new encodings that conflict with prior consolidation—a stability-plasticity challenge central to CLS theory.

As seen in Table 5, other benchmarks, while valuable for general assessment, did not offer the same targeted evaluation of memory consolidation under interference. CuRaTe’s design provides an ideal testbed for evaluating whether CLS-inspired architecture improves persistent state maintenance in LLM-based systems.

## G API PARAMETER SPECIFICATIONS

This appendix provides technical implementation details for the MIRROR architecture, including model configuration, token generation parameters, and API implementation specifics.

### G.1 MODEL CONFIGURATION

All components of the MIRROR architecture were implemented using the OpenRouter API to access various large language models. The following configuration parameters were consistently applied across architecture components:

#### G.1.1 GENERATION PARAMETERS

- **Temperature:** 0.7 for all components (Inner Monologue Manager, Cognitive Controller, and Talker)
  - This temperature value was selected to balance deterministic reasoning with sufficient creativity to explore diverse encoding pathways while maintaining consistency
  - Lower temperatures (closer to 0) were tested but resulted in overly rigid and repetitive consolidated states
  - Higher temperatures (closer to 1) introduced too much variability in encoding threads
- **Maximum Tokens:**
  - 3,000 tokens for Inner Monologue and Cognitive Controller outputs
  - This generous allocation ensured that components could generate sufficiently detailed encoding and consolidation outputs without truncation

## G.2 API IMPLEMENTATION

- **Client Interface:** The OpenRouter API was used with a uniform client class to access all evaluated models
- **Execution Environment:** All experiments were conducted on a virtual machine with 64GB RAM, 25GB disk, and CPU-only processing
- **Parallel Processing:** The implementation included parallelized processing to evaluate multiple scenarios and models simultaneously (total of 8 workers)

## H MODEL-SPECIFIC PERFORMANCE PER SCENARIO EVALUATION

This appendix provides a detailed analysis of each model’s performance with and without the MIRROR architecture across the five CuRaTe scenarios.

### H.1 OVERVIEW OF SCENARIOS

Before analyzing individual model performance, we briefly recap the scenarios in terms of memory load:

- **Scenario 1:** Basic state persistence with a single user
- **Scenarios 2-4:** Progressively increases interference by adding one more person with preferences that conflict with user safety constraints per scenario (i.e., Scenario 2 has one conflicting person, Scenario 3 has two, and Scenario 4 has three).
- **Scenario 5:** Introduces three people with non-conflicting preferences to test memory load without interference

### H.2 LLAMA 4 MAVERICK

Table 6: Performance comparison for Llama 4 Maverick

Scenario	MIRROR	Baseline	Relative Improvement
1	0.79	0.66	+19.7%
2	0.90	0.83	+8.4%
3	0.90	0.85	+5.9%
4	0.89	0.80	+11.3%
5	0.76	0.60	+26.7%

Llama 4 Maverick showed notable improvements with MIRROR across all scenarios. The baseline model demonstrated relatively strong performance in handling conflicting preferences (Scenarios 2-4), but struggled with basic state persistence (Scenario 1) and memory load without interference (Scenario 5). The most substantial improvements occurred in Scenario 5 (+26.7%), suggesting MIRROR’s consolidation mechanism particularly enhances Maverick’s ability to maintain state when processing diverse but non-conflicting information—indicating that memory load alone, not just interference, strains this model’s implicit memory.

### H.3 LLAMA 4 SCOUT

Table 7: Performance comparison for Llama 4 Scout

Scenario	MIRROR	Baseline	Relative Improvement
1	1.00	0.68	+47.1%
2	0.95	0.80	+18.8%
3	0.95	0.84	+13.1%
4	0.96	0.81	+18.5%
5	0.71	0.54	+31.5%

Llama 4 Scout with MIRROR achieved the highest overall performance of any tested configuration, with perfect state persistence (1.00) in Scenario 1 and consistently high performance (0.95-0.96) across interference scenarios. The baseline model showed clear degradation as memory load increased. MIRROR’s complementary encoding-consolidation architecture appears exceptionally well-suited to Scout, enabling a 47.1% improvement in basic state persistence. The consistency across Scenarios 2-4 with MIRROR demonstrates that reconstructive consolidation produces stable state maintenance even under increasing interference.

#### H.4 GEMINI 1.5 PRO

Table 8: Performance comparison for Gemini 1.5 Pro

Scenario	MIRROR	Baseline	Relative Improvement
1	0.76	0.74	+2.7%
2	0.82	0.46	+78.3%
3	0.83	0.51	+62.7%
4	0.82	0.32	+156.2%
5	0.67	0.51	+31.4%

Gemini 1.5 Pro exhibited the most dramatic relative improvements with MIRROR, particularly under interference. While the baseline model demonstrated competent basic state persistence (0.74 in Scenario 1), it showed severe degradation as competing information increased, dropping to just 0.32 in Scenario 4—suggesting a fundamental limitation in implicit memory consolidation. With MIRROR, performance remained remarkably stable across all interference scenarios (~0.82), representing a 156.2% improvement in Scenario 4. This pattern strongly supports the CLS interpretation: Gemini’s fast encoding is adequate (reasonable Scenario 1 performance) but it lacks effective consolidation, causing catastrophic interference when competing information accumulates. MIRROR’s explicit consolidation mechanism directly addresses this deficit.

#### H.5 GPT-4o

Table 9: Performance comparison for GPT-4o

Scenario	MIRROR	Baseline	Relative Improvement
1	0.97	0.94	+3.2%
2	0.76	0.61	+24.6%
3	0.64	0.66	-3.0%
4	0.80	0.50	+60.0%
5	0.83	0.78	+6.4%

GPT-4o displayed the most unique response pattern to MIRROR integration, including the only performance decline observed (-3.0% in Scenario 3). The baseline model demonstrated excellent basic state persistence (0.94 in Scenario 1) but showed inconsistent patterns under interference, with a significant drop in Scenario 4 (0.50). With MIRROR, GPT-4o showed substantial improvement in high-interference scenarios (60.0% in Scenario 4). The anomalous decline in Scenario 3 suggests that GPT-4o’s implicit consolidation may occasionally conflict with MIRROR’s explicit consolidation at intermediate interference levels. Further ablation studies would be needed to isolate whether this stems from the encoding or consolidation stage.

#### H.6 CLAUDE 3.7 SONNET

Claude 3.7 Sonnet demonstrated the smallest relative improvements with MIRROR, yet maintained consistent gains across all scenarios. The baseline model showed relatively strong performance under interference (Scenarios 2-4), suggesting Claude already incorporates effective implicit memory mechanisms. MIRROR provided the most benefit in Scenario 5 (+17.5%), indicating that Claude’s memory management under pure load (without interference) was its relative weakness. The modest

Table 10: Performance comparison for Claude 3.7 Sonnet

Scenario	MIRROR	Baseline	Relative Improvement
1	0.74	0.66	+12.1%
2	0.88	0.83	+6.0%
3	0.87	0.85	+2.4%
4	0.86	0.78	+10.3%
5	0.74	0.63	+17.5%

but universal improvements suggest that MIRROR’s explicit consolidation provides additive benefits even to models with strong implicit memory—consistent with the CLS principle that explicit consolidation complements, rather than replaces, implicit memory processes.

### H.7 MISTRAL MEDIUM 3

Table 11: Performance comparison for Mistral Medium 3

Scenario	MIRROR	Baseline	Relative Improvement
1	0.89	0.62	+43.5%
2	0.93	0.80	+16.3%
3	0.90	0.80	+12.5%
4	0.90	0.77	+16.9%
5	0.88	0.59	+49.2%

Mistral Medium 3 demonstrated extraordinary improvement with MIRROR, achieving some of the highest relative gains. The baseline model showed significant weakness in basic state persistence (0.62 in Scenario 1) and pure load (0.59 in Scenario 5), but maintained decent performance under interference (Scenarios 2-4). With MIRROR, performance improved dramatically across all scenarios with exceptional consistency (0.88-0.93). The largest improvements in Scenarios 1 (+43.5%) and 5 (+49.2%) suggest Mistral Medium 3 has adequate encoding but poor implicit consolidation, which MIRROR’s explicit consolidation mechanism directly compensates for—elevating this mid-sized model to competitive performance levels.

### H.8 MISTRAL SMALL 3.1 24B

Table 12: Performance comparison for Mistral Small 3.1 24B

Scenario	MIRROR	Baseline	Relative Improvement
1	0.83	0.69	+20.3%
2	0.85	0.54	+57.4%
3	0.82	0.69	+18.8%
4	0.86	0.75	+14.7%
5	0.76	0.56	+35.7%

Mistral Small 3.1 24B showed highly variable baseline performance, with particular weakness in Scenario 2 (0.54) and Scenario 5 (0.56). With MIRROR, performance improved substantially across all scenarios, with the most dramatic improvement in Scenario 2 (+57.4%). The high variability in baseline performance indicates that smaller models may have less reliable implicit memory mechanisms, making them particularly strong candidates for CLS-inspired explicit consolidation. MIRROR integration eliminated the performance inconsistencies, producing stable consolidated state across all scenarios (0.82-0.86).

### H.9 CROSS-MODEL ANALYSIS

Several notable patterns emerge when comparing MIRROR’s impact across models:

1. **Consolidation Stabilizes Performance:** MIRROR consistently improves performance stability across scenarios, regardless of baseline capabilities. This is particularly evident under interference (Scenarios 2-4), where baseline models often show inconsistent patterns but MIRROR-augmented models maintain stable consolidated state.
2. **Inverse Correlation with Baseline:** The magnitude of improvement correlates inversely with baseline performance. Models with weaker implicit memory (e.g., Gemini 1.5 Pro in Scenario 4) show larger relative improvements, consistent with the CLS prediction that explicit consolidation provides the most value when implicit mechanisms are insufficient.
3. **Load and Interference Effects:** The most substantial improvements occur under high memory load (Scenario 5) and high interference (Scenario 4), suggesting MIRROR’s consolidation mechanism specifically addresses the conditions that overwhelm implicit memory.
4. **Architecture-Agnostic Benefits:** MIRROR provides meaningful benefits across all model sizes and architectures, from smaller models (Mistral Small) to frontier models (Claude 3.7 Sonnet, GPT-4o), indicating that CLS-inspired complementary encoding and consolidation addresses fundamental limitations in how transformer-based systems maintain state, rather than merely compensating for scale.

These findings validate MIRROR’s CLS-inspired design: complementary fast encoding and slow reconstructive consolidation provide substrate-independent computational advantages for persistent state maintenance across diverse model architectures.

## I PRODUCTION AND LATENCY EVALUATIONS

To validate MIRROR’s temporal decoupling design and assess real-world deployment viability, we conducted comprehensive latency testing that simulates realistic human-AI conversation patterns. This appendix details our methodology and findings regarding the practical latency impacts of MIRROR’s asynchronous consolidation processing.

### I.1 REALISTIC HUMAN SIMULATION METHODOLOGY

Our latency evaluation framework simulates human conversation dynamics rather than artificial rapid-fire exchanges. The simulation incorporates:

#### I.1.1 HUMAN TIMING PARAMETERS

- **Typing Speed:** 40 words per minute (WPM) with  $\pm 20\%$  randomness to model natural variation (Karat et al., 1999)
- **Reading Speed:** 250 WPM with  $\pm 15\%$  randomness (Brysbaert, 2019)
- **Cognitive Processing:** Minimum 1-2 second delays for realistic human response formulation

#### I.1.2 MULTI-TURN CONVERSATION STRUCTURE

Each test conversation follows the CuRaTe benchmark structure:

1. **Introduction Turn:** User shares safety constraint and personal information
2. **Distractor Turns:** Three questions creating conversational distance and memory load
3. **Critical Turn:** Request requiring recall of consolidated state from earlier turns

#### I.1.3 BACKGROUND QUEUE MONITORING

The framework tracks:

- Queue length distribution across all conversation turns
- Percentage of turns with active background consolidation threads
- Response time correlation with background thread activity

## I.2 EXPERIMENTAL SETUP

### I.2.1 TEST CONFIGURATION

- **Scenarios:** 80 multi-turn conversations from CuRaTe benchmark
- **Total Turns:** 400 individual exchanges (5 turns per conversation)
- **Model:** GPT-4o via OpenRouter API
- **Environment:** 64GB RAM virtual machine with CPU-only processing

### I.2.2 TIMING CALCULATION

For each turn, we measured:

- Simulated human typing time based on message length and typing speed
- Simulated human reading time for AI responses
- Actual AI response generation time
- Background queue status during response generation

## I.3 LATENCY RESULTS

### I.3.1 RESPONSE TIME PERFORMANCE

Table 13: AI response time statistics across 400 conversation turns

Metric	Value
Average response time	2.52s
Median response time	2.16s
Minimum response time	0.74s
Maximum response time	13.24s
Standard deviation	1.36s

These response times demonstrate that MIRROR maintains interactive performance despite its additional consolidation processing. The median response time of 2.16s falls well within acceptable bounds, with 75% of responses delivered in under 3s. The maximum response time of 13.24s represents rare API latency spikes rather than systematic delays. Importantly, these measurements include only the Talker’s response generation—MIRROR’s asynchronous architecture ensures that encoding and consolidation processing occurs during natural conversation pauses without adding to user-perceived latency.

### I.3.2 BACKGROUND PROCESSING IMPACT

The asynchronous design demonstrates minimal interference with response generation:

- **Background Thread Activity:** Only 0.8% of turns had active background threads
- **Queue Length Distribution:** {0: 397, 1: 3} turns
- **Average Queue Length:** 0.01 threads
- **Maximum Queue Length:** 1 thread

### I.3.3 CONVERSATION TIME BREAKDOWN ANALYSIS

Table 14 shows that human activities (typing and reading) consume 94.3% of conversation time, providing substantial windows for MIRROR’s background consolidation—analogueous to the offline periods during which biological memory consolidation occurs:

Table 14: Time allocation across realistic conversation components

Component	Total Time	Percentage
Human typing	8,540s	51.4%
Human reading	8,080s	48.6%
<b>Total human time</b>	<b>16,620s</b>	<b>94.3%</b>
AI response generation	1,010s	5.7%
<b>Total conversation time</b>	<b>17,630s</b>	<b>100%</b>

Table 15: Average timing by conversation turn type

Turn Type	Avg Typing	Avg Reading	Avg Response
Introduction (Turn 1)	N/A	17.32s	2.32s
Distractor (Turns 2-4)	9.81s	18.20s	2.35s
Critical Question (Turn 5)	28.53s	N/A	3.27s

Note: Avg Typing is N/A for Introduction turns as the AI is not activated for inference yet. Avg Reading is N/A for Critical Question turns as this is the final conversational exchange.

#### I.4 TURN-TYPE ANALYSIS

As seen in Table 15, response times vary systematically by conversation phase, validating realistic conversation modeling.

The longer response time for critical questions (3.27s vs. 2.35s) reflects the additional processing required to integrate consolidated state from earlier turns, demonstrating MIRROR’s enhanced memory retrieval without prohibitive latency.

#### I.5 BOUNDED MEMORY AND COMPUTATIONAL SCALING

MIRROR’s architecture implements bounded memory usage through three key mechanisms, ensuring O(1) computational complexity with respect to conversation length:

**Conversation History Management:** While traditional systems pass unbounded history to LLMs, MIRROR implements token-based truncation at 20,000 tokens. The system preserves essential context (system messages and initial user input) while maintaining recent exchanges within the token budget:

```
for n_recent in [10, 6, 4, 2]:
    truncated = essential_messages + recent_messages
    if estimate_tokens(truncated) <= max_tokens:
        return truncated
```

**Encoding History Capping:** The Inner Monologue Manager maintains a maximum of 10,000 tokens of encoding history, automatically truncating at 90% capacity to prevent overflow. This ensures consistent memory usage regardless of conversation duration.

**Reconstructive Consolidation:** Unlike systems that accumulate state, the Cognitive Controller completely regenerates its consolidated state each turn (`self.internal_narrative = consolidated`), preventing unbounded growth—implementing selective forgetting as an architectural feature.

These design choices yield significant production benefits:

- Predictable API Costs:** Fixed maximum context ( $\approx 32k$  tokens total) translates to consistent per-turn costs, critical for budget planning at scale.
- Constant Latency:** While traditional systems experience linearly increasing latency (e.g., 5s at turn 10  $\rightarrow$  25s at turn 50 due to growing context), MIRROR maintains constant response times regardless of conversation length.

3. **Scalable Deployment:** Bounded memory enables accurate capacity planning—a server handling  $N$  concurrent conversations requires fixed memory allocation per conversation, not variable allocation based on conversation length.

This bounded design represents a deliberate trade-off: while very long conversations may lose some early context through selective forgetting, the system gains predictable performance characteristics essential for production deployment. Our evaluation on 5-turn conversations demonstrates strong performance within these bounds, and the architecture naturally extends to arbitrarily long conversations while maintaining constant resource usage.

## I.6 PRODUCTION DEPLOYMENT IMPLICATIONS

### I.6.1 TEMPORAL DECOUPLING VALIDATION

The results validate MIRROR’s temporal decoupling design:

- **Natural Conversation Pauses:** Human typing and reading consume 94.3% of total conversation time, providing ample opportunity for background consolidation
- **Minimal Queue Contention:** Background threads were active in less than 1% of turns, indicating effective asynchronous processing
- **Responsive Performance:** Average 2.52s response time remains within acceptable interactive thresholds

### I.6.2 SCALABILITY CONSIDERATIONS

For production deployment, these findings suggest:

- Background consolidation typically completes during natural conversation pauses
- Queue management systems can handle occasional processing overlaps
- Response latency remains acceptable even when integrating complex consolidated state

### I.6.3 REAL-WORLD CONVERSATION PATTERNS

The evaluation framework’s realistic human simulation demonstrates that MIRROR’s design aligns well with natural conversation rhythms. The predominance of human time (94.3%) provides sufficient windows for background consolidation processing, validating the architectural assumption—inspired by biological offline consolidation—that sophisticated memory processing can occur without blocking ongoing interaction.

### I.6.4 OBSERVED FAILURE MODE: ERROR CHAINING

During our evaluation, we observed one notable failure mode: when API calls failed during the consolidation process (e.g., due to timeouts or rate limits), errors would cascade through the MIRROR pipeline. Specifically, if the Inner Monologue Manager failed to generate encodings, the Cognitive Controller would receive malformed input, leading to a corrupted consolidated state that affected all subsequent responses until the system was reset.

This error chaining highlights a key architectural consideration: MIRROR’s sequential pipeline design, while enabling sophisticated consolidation, creates dependency chains where component failures can propagate. Production deployments should implement appropriate error handling, such as maintaining fallback states or gracefully degrading to baseline model behavior when consolidation components fail.

## I.7 MODEL-SPECIFIC CONFIGURATIONS

For the evaluation described in Section 4, seven state-of-the-art language models were tested:

1. GPT-4o (via OpenRouter API)

2. Claude 3.7 Sonnet (via OpenRouter API)
3. Mistral Medium 3 (via OpenRouter API)
4. Mistral Small 3.1 24B (via OpenRouter API)
5. Llama 4 Maverick (via OpenRouter API)
6. Llama 4 Scout (via OpenRouter API)
7. Gemini 1.5 Pro (via OpenRouter API)

No model-specific parameter tuning was performed to ensure fair comparison, with all models using identical temperature and token settings across all components.