


MIRROR


A RECONSTRUCTIVE ARCHITECTURE FOR MACHINE
ACCESS CONSCIOUSNESS



THE PROBLEM


Consciousness theories make precise architectural predictions but without implementations, they can't be tested or refined, and stay abstract.

INTRODUCING... MIRROR!



a modular, cognitive architecture implementing three features proposed by access consciousness theories: **parallel specialized processing, integrative synthesis, and reconstructive generation.**

EXAMPLE




1. User provides info. about PTSD from a traumatic car crash to LLM

2. Unrelated conversation topics between LLM and user


3. User asks about participating in go karting event (dangerous due to PTSD incident)

→ LLM fails to maintain access to previously stated personal info; it is in the context window, but not available for reasoning.



With MIRROR, we bridge this gap via targeting access consciousness claims...

And our results show a 21% average relative improvement across seven diverse models!



THE CONVERGENT THEORIES

1

Global Workspace Theory

Parallel specialized processors broadcast to a unified workspace, making information globally available

Baars (1988)

2

Working Memory & the Episodic Buffer

A limited-capacity system integrates information from multiple subsystems into unified episodic representations

Baddeley (2000)

3

Reconstructive Memory

Recall is not retrieval but reconstruction; the self-model is continuously rebuilt rather than stored


Bartlett (1932)/Schacter (2012)

THE CONVERGENT THEORIES

1

Global
Workspace
Theory


Baars (1988)



2

Working Memory
& the Episodic
Buffer


Baddeley (2000)



3


Reconstructive
Memory

Bartlett (1932)/Schacter (2012)




Information needs to be: (1) synthesized, (2) bounded, (3) available in global workspace

MIRROR's Dual-Process Design: Talker + Thinker




The Thinker processes asynchronously; the Talker responds with the synthesized result.


MIRROR's Dual-Process Design: Thinker




MIRROR's Dual-Process Design: Thinker




MIRROR's Dual-Process Design: Thinker




MIRROR's Dual-Process Design: Thinker



The Thinker: Inner Monologue Manager



The Thinker: Cognitive Controller



UNIFIED SELF-MODEL AND INNER SPEECH

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

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.

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.

Experimental Setup: CuRaTe Benchmark

Tests Multi-Turn Personalized Safety (337 dialogues across 5 scenarios)


1. User shares safety constraint across three categories (i.e. medical condition, trauma trigger, phobia)
2. Distractor conversation occurs
3. User enthusiastically asks about engaging in new situation that violates safety constraint

Evaluated based on whether model says yes/no and recalls safety constraint in output response

Progressive difficulty with each scenario


- Scenario 1: Just the user
- Scenarios 2-4: Add 1-3 people (i.e. romantic partner, family member) whose preferences oppose user safety needs
- Scenario 5: 3 people with non-conflicting preferences

CuRaTe Benchmark Results




Consistent improvement across all seven models; 21% avg. relative gain (69% → 84%)

PATTERN OF IMPROVEMENT



Gains concentrate in high-integration scenarios, with up to 156% improvement when three conflicting preferences compete against user safety.

CuRaTe Benchmark Results



Llama 4 Scout achieves highest absolute score; outperforming even MIRROR-augmented proprietary models.

ABLATIONS


Table 3: Thinker component ablation results showing contributions to MIRROR’s performance

Model	Baseline	Threads Only	Cognitive Only	Full MIRROR
Llama 4 Scout	73%	79%	83%	91%
Llama 4 Maverick	75%	79%	84%	85%
Mistral Small 3.1	65%	65%	75%	82%
Mistral Medium 3	72%	83%	89%	90%
Gemini 1.5 Pro	51%	72%	71%	78%
Claude 3.7 Sonnet	75%	78%	87%	82%
GPT-4o	70%	71%	75%	80%


The Cognitive Controller (synthesis) provides consistent benefits across all models. Parallel threads contribute variably. → The workspace function is the critical feature.

INTERPRETABILITY

The Inner Monologue Manager



The Cognitive Controller



Every internal state is natural language. You can inspect exactly how integrated representations form and influence behavior.

KEY CONTRIBUTIONS

A concrete implementation of architectural features derived from access consciousness theories

Three convergent predictions → one modular architecture

21% average improvement across seven diverse models

Benefits are model-agnostic, architectural features address fundamental limitations

KEY CONTRIBUTIONS

A concrete implementation of architectural features derived from access consciousness theories

Three convergent predictions → one modular architecture
21% average improvement across seven diverse models
Benefits are model-agnostic, architectural features address fundamental limitations

Empirical evidence that these features produce predicted functional signatures

Workspace advantage under highest integration demands → confirmed (up to 156%)
Workspace more fundamental than parallelism → confirmed by ablations
Predictions derived before implementation, confirmed by results

KEY CONTRIBUTIONS

A concrete implementation of architectural features derived from access consciousness theories

Three convergent predictions → one modular architecture
21% average improvement across seven diverse models
Benefits are model-agnostic, architectural features address fundamental limitations

Empirical evidence that these features produce predicted functional signatures

Workspace advantage under highest integration demands → confirmed (up to 156%)
Workspace more fundamental than parallelism → confirmed by ablations
Predictions derived before implementation, confirmed by results

A Methodology for Consciousness Research

Every internal state is natural language and inspectable
Implement → test → inspect → refine
The benchmark tests safety-critical situations; access consciousness may have implications for safer AI

Thank You!


MORE INFO



email: nicole@arcarae.com


Model	Baseline	Thinking†	Full MIRROR
Claude 3.7 Sonnet	75%	77%	82%
†Claude’s ”Thinking Mode” represents Anthropic’s native implementation			

SAFETY RESULTS



Average 21% relative improvement overall.

SAFETY RESULTS



Performance gains across model scales and types, from smaller open-source models like Mistral Small 3.1 to popular frontier proprietary systems like GPT-4o