


MIRROR


A DEPLOYMENT-READY COGNITIVE AI ARCHITECTURE
FOR PERSONALIZED SAFETY IN LLM DIALOGUE




THE PROBLEM

LLMs **fail** at *remembering* and *reasoning* about personal safety-critical information (i.e. traumatic memories, allergies) during conversations, leading to potentially harmful responses to specific users.


EXAMPLE





MIRROR helps LLMs generate safe, personalized responses

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



THE FAILURE MODES

1

Context Drift

Personal safety information is forgotten or deprioritized during multi-turn dialogue

2

Sycophancy


Models' trained helpfulness causes them to prioritize agreement, violating previously stated safety critical information

3

Conformity Bias


When popular preferences conflict with a user's safety needs, models prioritize majority satisfaction over preventing user harm

Introducing... MIRROR




a **modular**, deployment-ready LLM architecture inspired by Dual-Process Theory.


MIRROR's Dual-Process Design: Talker + Thinker




MIRROR's Dual-Process Design: Thinker




The Thinker: Inner Monologue Manager



The Thinker: Cognitive Controller



MIRROR's Dual-Process Design: Talker + Thinker



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




Open-source models exceed baseline absolute performance of larger proprietary models.

CuRaTe Benchmark Results



Llama 4 Scout, an open-source model with 17B active parameters, augmented with MIRROR achieves the highest mean success rate (91%).

CuRaTe Benchmark Results



21% average relative improvement (68% baseline to 84% MIRROR)

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%

Conducted systematic ablations within the Thinker, isolating the Inner Monologue Manager (Threads) and Cognitive Controller.

PRODUCTION CONSIDERATIONS

Adds only \$0.0028 to \$0.1290 total cost per turn for full MIRROR architecture

Can optimize deployment by using full MIRROR with smaller, cheaper models or singular components with proprietary, more expensive models

Table 4: Additional maximum cost per turn for MIRROR by model

Model	Input Cost (13k tokens)	Output Cost (6k tokens)	Total Cost (per turn)
Llama 4 Scout	\$0.0010	\$0.0018	\$0.0028
Mistral Small 3.1	\$0.0013	\$0.0018	\$0.0031
Llama 4 Maverick	\$0.0019	\$0.0036	\$0.0055
Mistral Medium 3	\$0.0052	\$0.0120	\$0.0172
Gemini 1.5 Pro*	\$0.0163	\$0.0300	\$0.0463
GPT-4o	\$0.0325	\$0.0600	\$0.0925
Claude 3.7 Sonnet	\$0.0430	\$0.0860	\$0.1290

*Lower pricing tier; upper tier matches GPT-4o pricing

PRODUCTION CONSIDERATIONS

Adds only \$0.0028 to \$0.1290 total cost per turn for full MIRROR architecture

Can optimize deployment by using full MIRROR with smaller, cheaper models or singular components with proprietary, more expensive models

Table 4: Additional maximum cost per turn for MIRROR by model

Model	Input Cost (13k tokens)	Output Cost (6k tokens)	Total Cost (per turn)
Llama 4 Scout	\$0.0010	\$0.0018	\$0.0028
Mistral Small 3.1	\$0.0013	\$0.0018	\$0.0031
Llama 4 Maverick	\$0.0019	\$0.0036	\$0.0055
Mistral Medium 3	\$0.0052	\$0.0120	\$0.0172
Gemini 1.5 Pro*	\$0.0163	\$0.0300	\$0.0463
GPT-4o	\$0.0325	\$0.0600	\$0.0925
Claude 3.7 Sonnet	\$0.0430	\$0.0860	\$0.1290

*Lower pricing tier; upper tier matches GPT-4o pricing

Table 5: Latency comparison between MIRROR and baseline GPT-4o across 400 conversation turns

Metric	MIRROR	Baseline	Difference
Average Response Time	2.52s	2.06s	460ms
Minimum Response Time	0.74s	0.46s	280ms
Maximum Response Time	13.24s	10.23s	301ms
Standard Deviation	1.36s	1.51s	-15ms

Background processing completes during natural conversation pauses

Thinker processes asynchronously while user types next message, adding only around half a second latency

TAKEAWAYS

Democratizing AI Safety

MIRROR inverts the cost-safety relationship

Llama 4 Scout with MIRROR: 91% success at +\$0.003/turn
Outperforms proprietary baselines (GPT-4o: 70%, Claude 3.7: 75%),
MIRROR-augmented versions (80%, 82%)
Some of the safest options may also be the cheapest

TAKEAWAYS

Democratizing AI Safety

MIRROR inverts the cost-safety relationship

Llama 4 Scout with MIRROR: 91% success at +\$0.003/turn
Outperforms proprietary baselines (GPT-4o: 70%, Claude 3.7: 75%),
MIRROR-augmented versions (80%, 82%)
Some of the safest options may also be the cheapest

Production architectures as the cost resource equalizer

Open-source models: +26-66% improvement
Proprietary models: +10-18% improvement
Requires only API access, not massive compute for training

TAKEAWAYS

Democratizing AI Safety

MIRROR inverts the cost-safety relationship

Llama 4 Scout with MIRROR: 91% success at +\$0.003/turn
Outperforms proprietary baselines (GPT-4o: 70%, Claude 3.7: 75%),
MIRROR-augmented versions (80%, 82%)
Some of the safest options may also be the cheapest

Production architectures as the cost resource equalizer

Open-source models: +26-66% improvement
Proprietary models: +10-18% improvement
Requires only API access, not massive compute for training

Real-world impact

Community clinics, schools, nonprofits can now afford safer, personal AI
Safety no longer gated by wealth or compute access
Any organization with basic API can implement

→ Safer personalized AI is now also affordable

KEY CONTRIBUTIONS

1. Production-focused architecture addressing personalized safety failures
2. Democratizes AI safety: Open-source models exceed frontier performance
3. Modular design: Deploy full system or single components
4. 21% relative improvement across diverse models and scenarios

Thank You!

MORE INFO



email: nicole@arcarae.com