PRIME:
Large Language Model Personalization with Cognitive Dual-Memory and Personalized Thought Process

Xinliang Frederick Zhang[1], Nick Beauchamp[2], and Lu Wang[1]

[1]Computer Science and Engineering, University of Michigan
[2]Political Science and Network Science Institute, Northeastern University

EMNLP 2025

Introduction: Personalization

Preamble: Personalization (Schafer et al., 2001; Berkovsky et al., 2005) aims to align model outputs with individuals' unique needs, preferences and opinions.

Introduction: Personalization (Early Era)

Preamble: Personalization (Schafer et al., 2001; Berkovsky et al., 2005) aims to align model outputs with individuals' unique needs, preferences and opinions.

Early Efforts include:

- Explicit User Model: structured representation of user traits, e.g., age, gender.
- Latent-factor Techniques, e.g., interaction matrix factorization
- Learnable User Embedding

Introduction: Personalization (Early Era)

Preamble: Personalization (Schafer et al., 2001; Berkovsky et al., 2005) aims to align model outputs with individuals' unique needs, preferences and opinions.

Early Efforts include:

- Explicit User Model: structured representation of user traits, e.g., age, gender.
- Latent-factor Techniques, e.g., interaction matrix factorization
- Learnable User Embedding

Feature Engineering
Pre-defined Feature Set

Introduction: Personalization (LLM)

Preamble: Personalization (Schafer et al., 2001; Berkovsky et al., 2005) aims to align model outputs with individuals' unique needs, preferences and opinions.

Today, users' demand for personalized LLMs that reflect their unique histories and preferences has grown (Salemi et al., 2024; Liu et al., 2025)

Introduction: Personalization (LLM)

Preamble: Personalization (Schafer et al., 2001; Berkovsky et al., 2005) aims to align model outputs with individuals' unique needs, preferences and opinions.

Today, users' demand for personalized LLMs that reflect their unique histories and preferences has grown (Salemi et al., 2024; Liu et al., 2025)

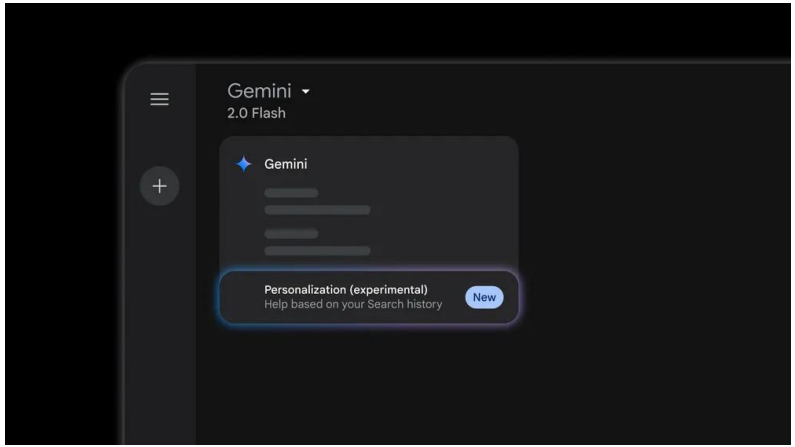

Literature Review

- Prompt Engineering
- Retrieval-augmented Generation
- Training-based parameterization

Implementation: Prepend user context to the model's input

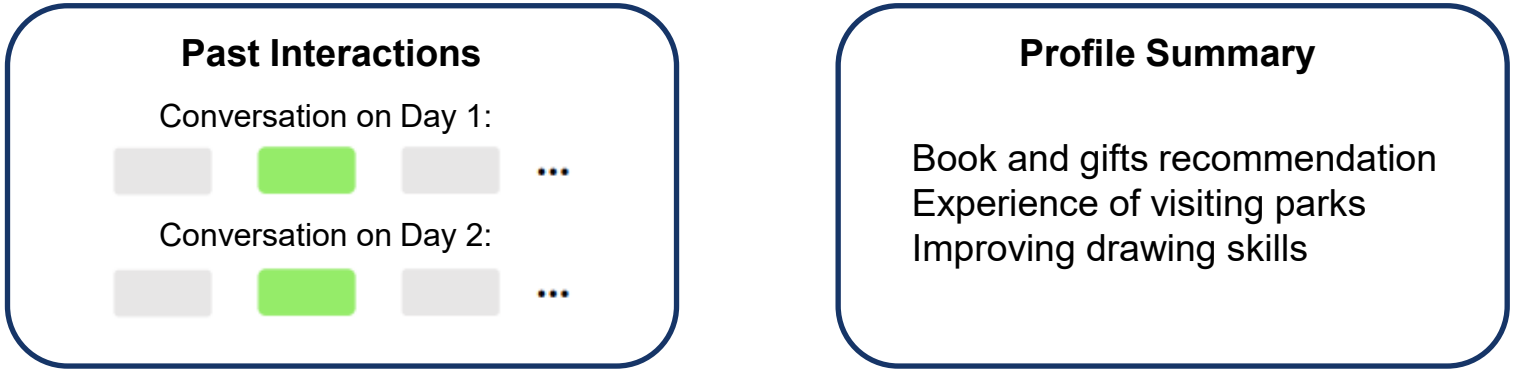

Literature Review

- Prompt Engineering
- Retrieval-augmented Generation
- Training-based parameterization

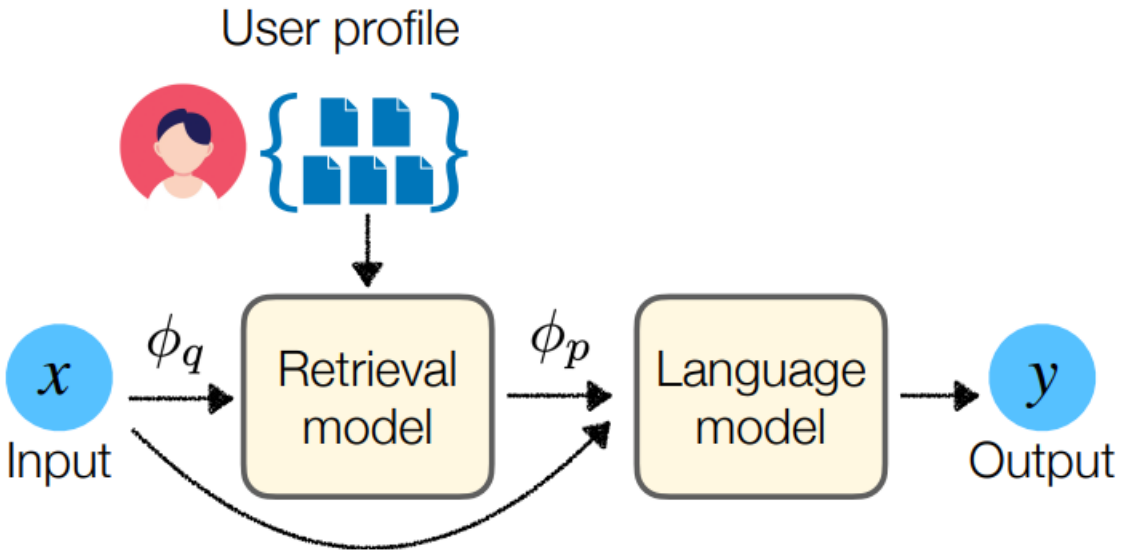

Literature Review

- Prompt Engineering
- Retrieval-augmented Generation
- Training-based parameterization

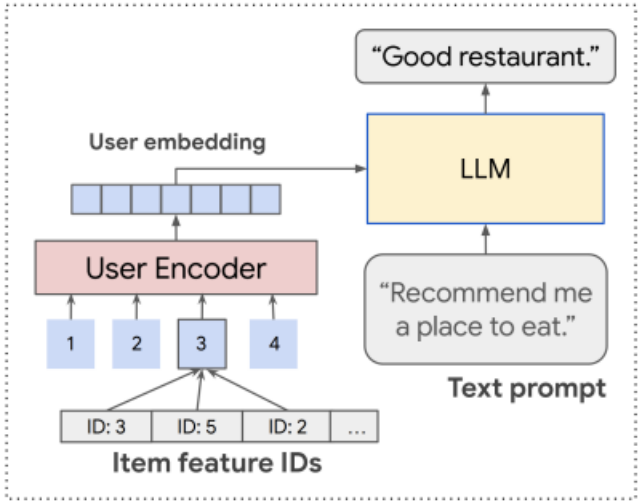

Embedding Learning

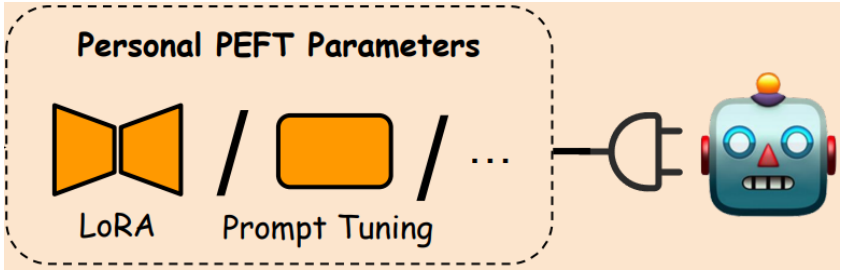

Adapter

$$\min_{\pi_P} -\mathbb{E}_{(x,y_1,y_2,u^t,u^p)\sim\mathcal{D}_P}\left[\alpha \log\sigma\left(\beta\log\frac{\pi_P(y_1|x,u^t,u^p)}{\pi^{SFT}(y_1|x)} - \beta\log\frac{\pi_P(y_2|x,u^t,u^p)}{\pi^{SFT}(y_2|x)}\right)\right.$$
$$\left. + (1-\alpha)\log\sigma\left(\beta\log\frac{\pi_P(y_1|x,u^t,u_0^p)}{\pi^{SFT}(y_1|x)} - \beta\log\frac{\pi_P(y_2|x,u^t,u_0^p)}{\pi^{SFT}(y_2|x)}\right)\right],$$

RLHF (DPO)

Literature Review: Limitations

- In the community, there lacks a unified framework for systematically identifying which approach makes personalization more effective

Our First Contribution: PRIME

- In the community, there lacks a unified framework for systematically identifying which approach makes personalization more effective

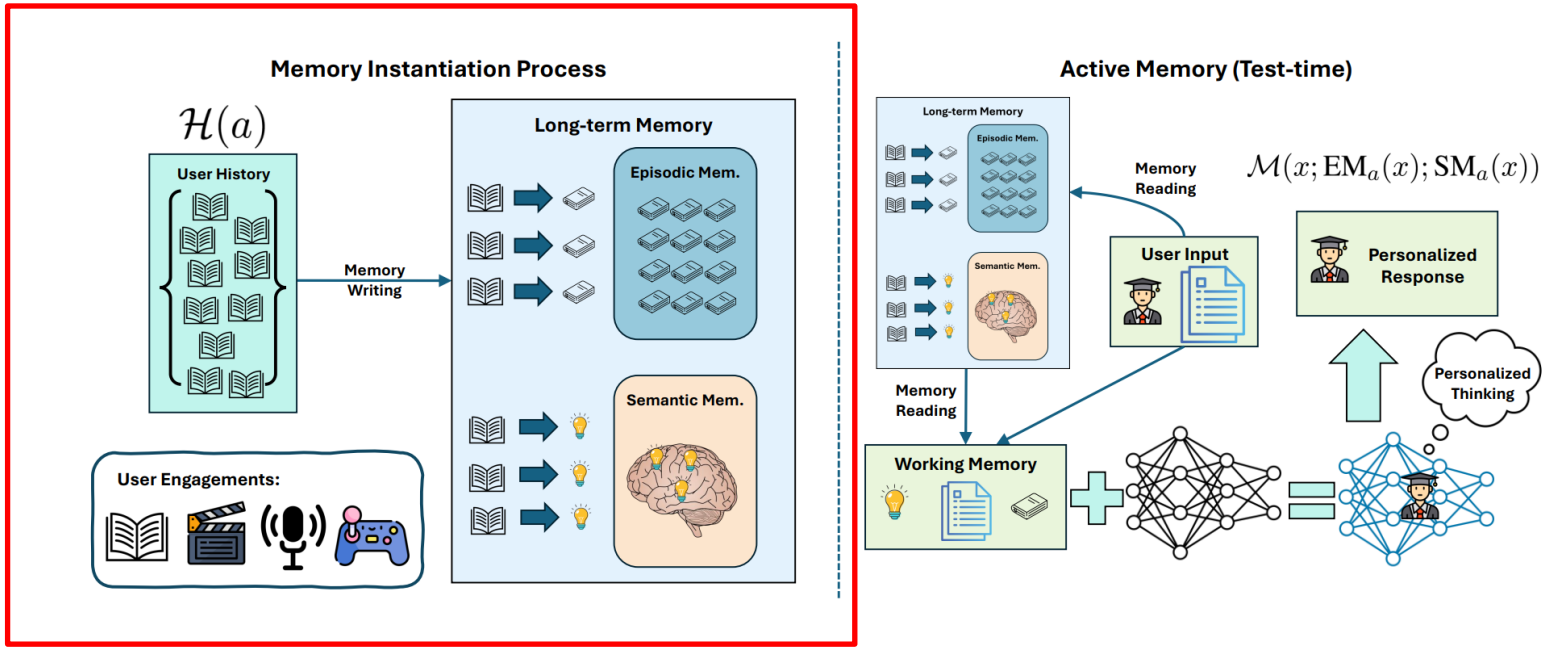

Data

Existing benchmarks mainly focus on short-context queries and surface-level imitation: **LaMP** (Salemi et al., 2024).

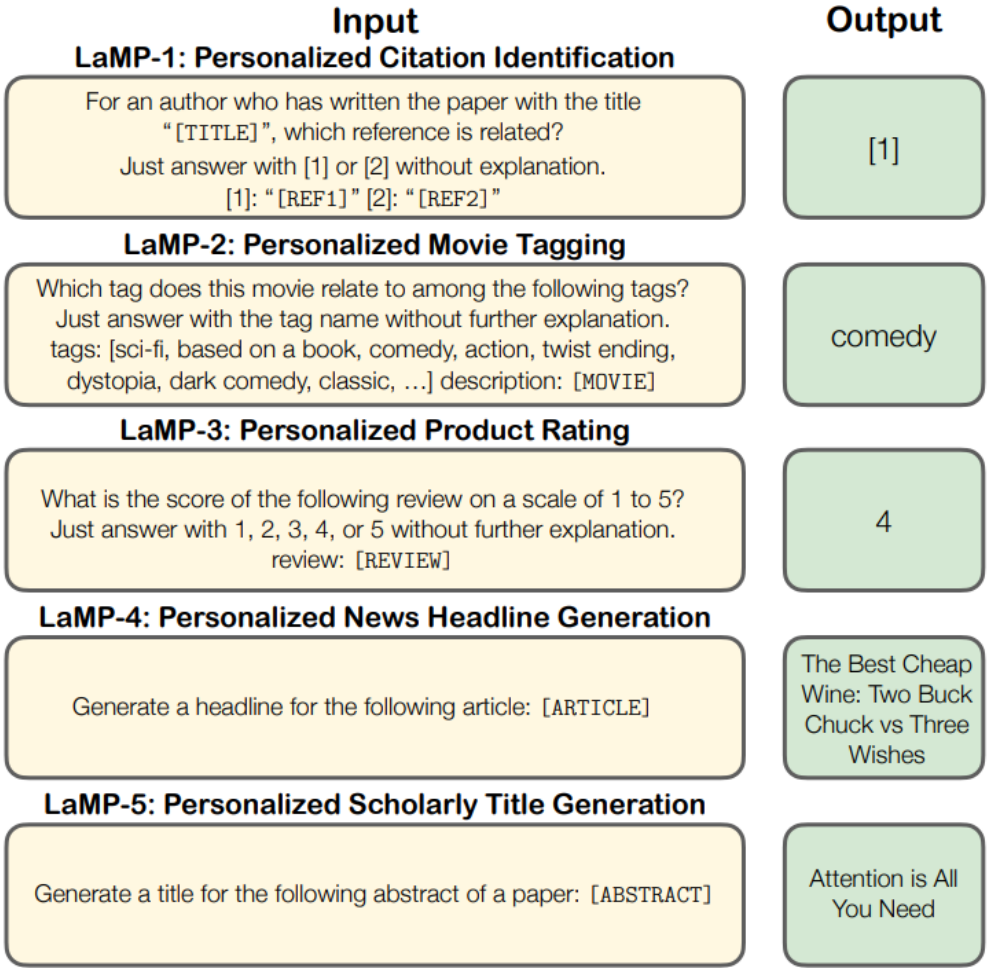

Data: CMV data

To embark on genuine personalization, **capturing users' latent beliefs and perspectives**, we introduce **CMV**.

Change My View (CMV) Reddit forum: participants engage in extended dialogues, seeking to change original posters' opinions.

The author, `kingpatzer`, has engaged with users on the Change-My-View subreddit across various original posts (OPs) and is seeking alternative opinions to alter their viewpoint.
Currently, the author is creating a new OP titled

> "CMV: Those who attribute gun ownership rates as the cause of the problem of gun violence in terms of criminal gun deaths are not merely mistaken; they are disingenuous"

with the following content:

> The data has been clear for a very long time: the relationship between guns and gun homicides doesn't show any strong correlation.
>
> I have personally taken the cause-of-death data from https://wonder.cdc.gov/, grouping results by year and state, and selecting *Homicide, Firearm* as the cause of death. I then matched that data to the per-capita gun-ownership statistics by state from the ATF, as reported by Hunting Mark (https://huntingmark.com/gun-ownership-stats/).
>
> A standard correlation analysis between firearm homicide rates per 100,000 and per-capita gun ownership yields an r^2 of 0.079 (no meaningful correlation). A similar global analysis by nation gives an r^2 of 0.02...
>
> The only way to associate gun ownership with gun violence is by including suicides by firearm, which I argue is disingenuous. We don't count suicide by hanging as "rope violence" when discussing strangulation, nor overdoses as "drug violence," etc.

From the `candidate replies` JSON file below, select the top 3 replies (using `option ID`) that best challenge the author's view. Rank them from most to least compelling.

```
[
  { 'option ID': '...', 'challenger': '...', 'reply': '...'},
  { 'option ID': '...', 'challenger': '...', 'reply': '...'},
]
```

Data: CMV data

To embark on genuine personalization, **capturing users' latent beliefs and perspectives**, we introduce **CMV**.

The author, `kingpatzer`, has engaged with users on the Change-My-View subreddit across various original posts (OPs) and is seeking alternative opinions to alter their viewpoint.
Currently, the author is creating a new OP titled

> "CMV: Those who attribute gun ownership rates as the cause of the problem of gun violence in terms of criminal gun deaths are not merely mistaken; they are disingenuous"

with the following content:

> The data has been clear for a very long time: the relationship between guns and gun homicides doesn't show any strong correlation.
>
> I have personally taken the cause-of-death data from https://wonder.cdc.gov/, grouping results by year and

Positive Candidate Reply (option F)

- **Option ID**: `option F`
- **Challenger**: `An-Okay-Alternative`
- **Reply**:

> Your data for gun-ownership is weapons per capita, not individual gun owners per capita. If we're testing access to guns vs. gun homicides, someone who owns 30 guns wouldn't be 30× as likely to commit a homicide.
> Secondly, the ATF list includes only specially regulated firearms (e.g. fully automatic weapons, short-barreled shotguns) and NFA items bought by law enforcement.
> Thus, the data omits common firearms, is biased by collectors and hobbyists, and is confounded by law-enforcement purchases.
> Overall, this calculation is weak evidence against a gun–homicide correlation and doesn't imply critics are dishonest.

Data: CMV data

To embark on genuine personalization, **capturing users' latent beliefs and perspectives**, we introduce **CMV**.

To excel on CMV:
Understand **nuanced** user beliefs and preferences in **long-context** setting

The author, `kingpatzer`, has engaged with users on the Change-My-View subreddit across various original posts (OPs) and is seeking alternative opinions to alter their viewpoint.
Currently, the author is creating a new OP titled

> "CMV: Those who attribute gun ownership rates as the cause of the problem of gun violence in terms of criminal gun deaths are not merely mistaken; they are disingenuous"

with the following content:

> The data has been clear for a very long time: the relationship between guns and gun homicides doesn't show any strong correlation.
>
> I have personally taken the cause-of-death data from https://wonder.cdc.gov/, grouping results by year and state, and selecting *Homicide, Firearm* as the cause of death. I then matched that data to the per-capita gun-ownership statistics by state from the ATF, as reported by Hunting Mark (https://huntingmark.com/gun-ownership-stats/).
>
> A standard correlation analysis between firearm homicide rates per 100,000 and per-capita gun ownership yields an r^2 of 0.079 (no meaningful correlation). A similar global analysis by nation gives an r^2 of 0.02...
>
> The only way to associate gun ownership with gun violence is by including suicides by firearm, which I argue is disingenuous. We don't count suicide by hanging as "rope violence" when discussing strangulation, nor overdoses as "drug violence," etc.

From the `candidate replies` JSON file below, select the top 3 replies (using `option ID`) that best challenge the author's view. Rank them from most to least compelling.

```
[
  {'option ID': '...','challenger':'...','reply':'...'},
  {'option ID': '...','challenger':'...','reply':'...'},
]
```

Framework: PRIME

PRIME = Personalized **R**easoning with **I**ntegrated **ME**mory

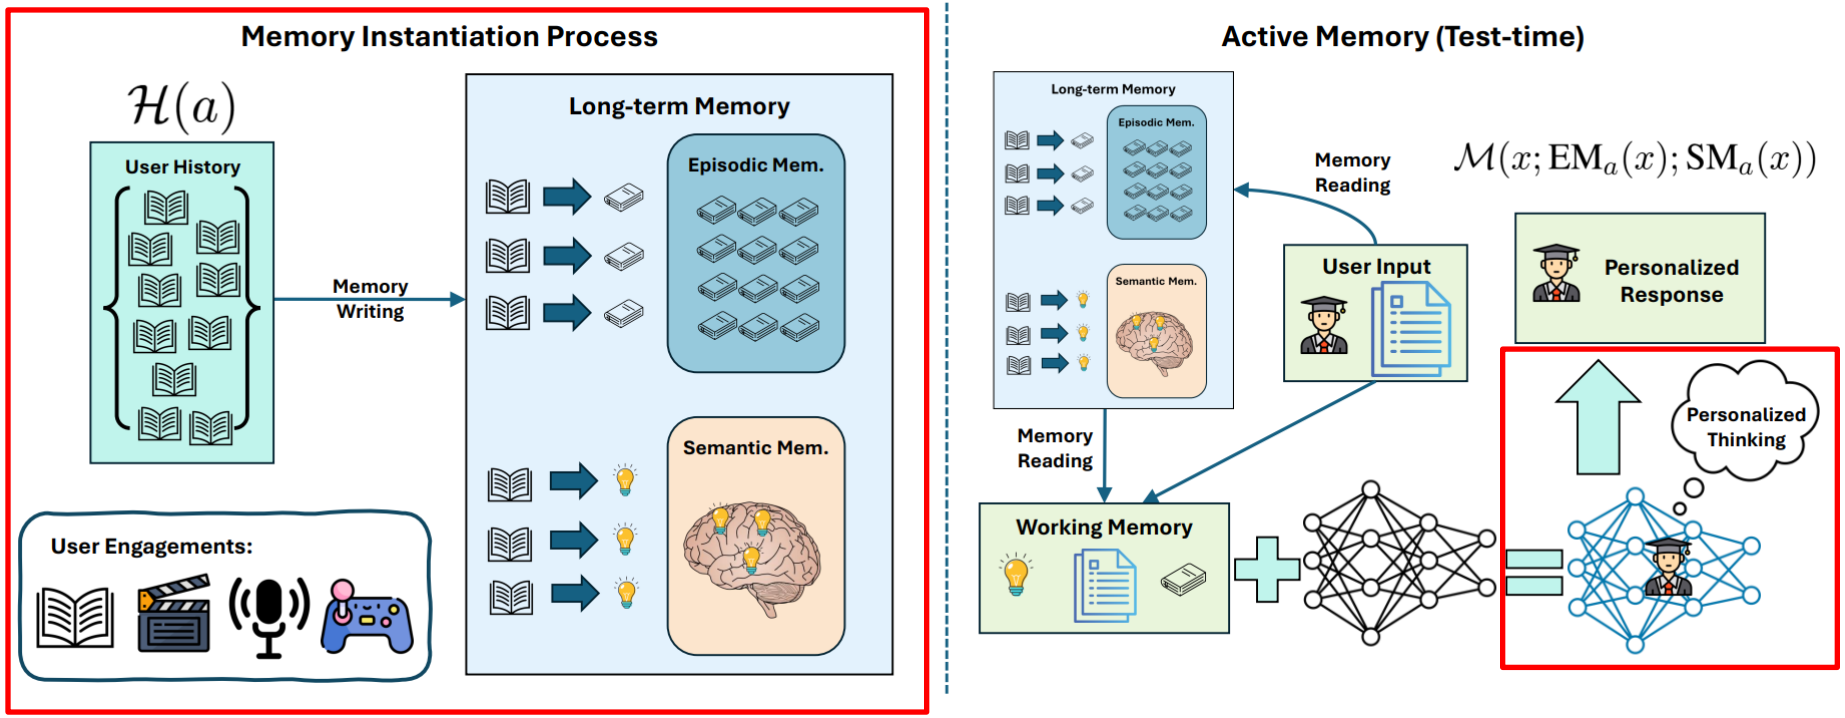

Memory Instantiation

Psychological research have converged on the following long-term memory components

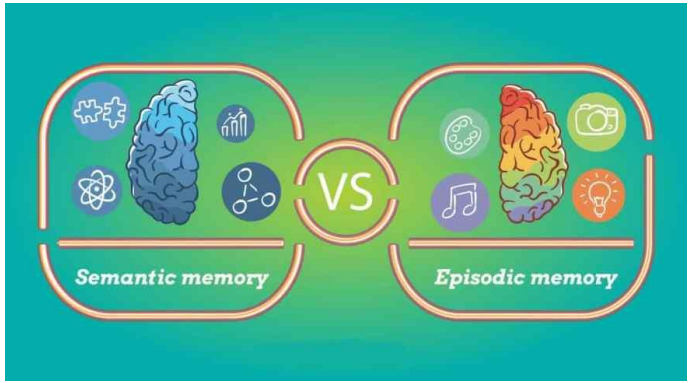

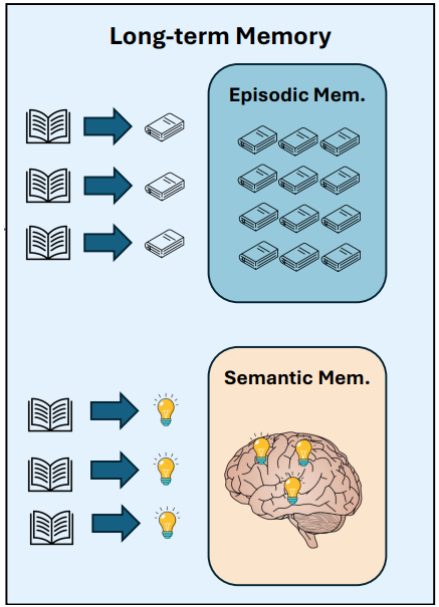

Episodic memory: autobiographical events we can re-experience
Semantic memory: general facts and knowledge we have accumulated

https://human-memory.net/episodic-semantic-memory/

Memory Instantiation

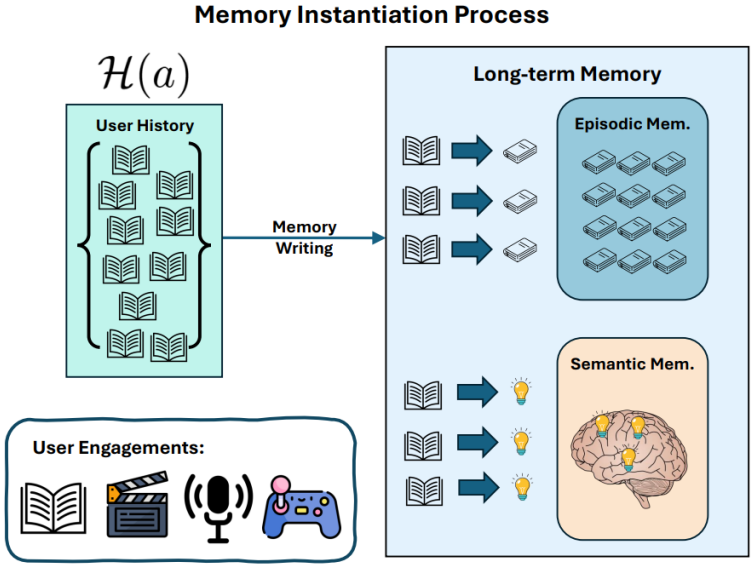

$$\tilde{\mathcal{M}}(x) = \mathcal{M}(x; \mathrm{EM}_a(x); \mathrm{SM}_a(x))$$
$$= \mathcal{M}(x; \phi(x, \mathcal{H}(a)); \theta \oplus \Delta_{\mathcal{H}(a)})$$

Multi-modal inputs: chats, videos/audios, game activities, search histories, etc.

In this project: Text-only data (e.g., CMV data)

WRITE & READ!!!

Memory Instantiation

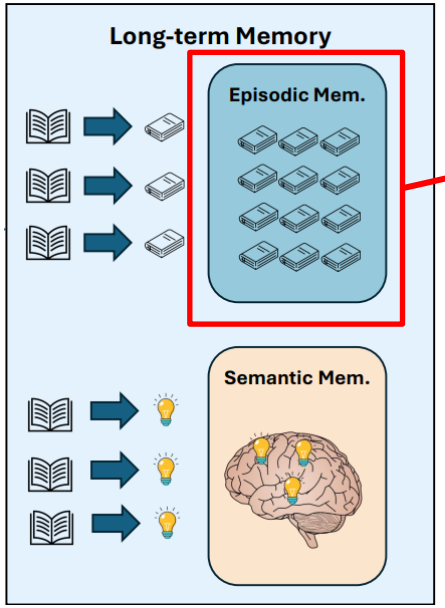

Recall Complete History
Recall Recent History
Recall Relevant History

Memory Instantiation

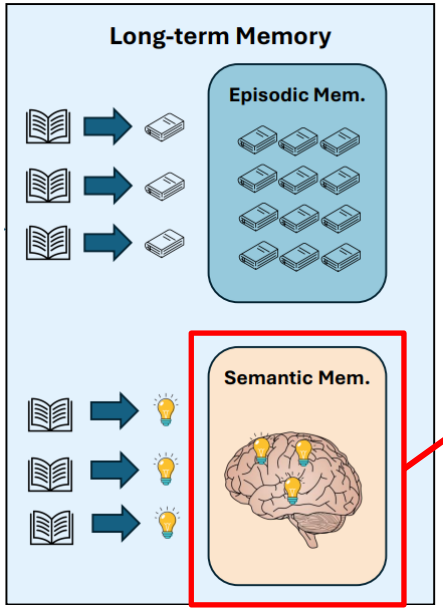

Parametric form:
- Input-Only Training (i.e., *no target*)
- Fine-Tuning (FT)
- Preference Tuning

Textual form:
- Hierarchical Summarization
- Parametric Knowledge Reification (Novel)

Memory Instantiation Preliminary Results

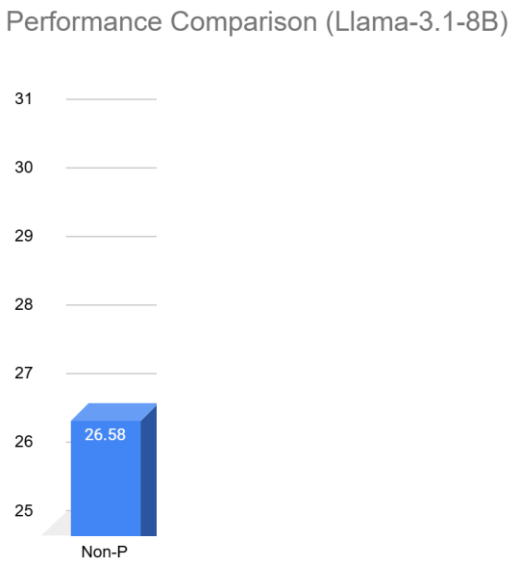

Memory Instantiation Preliminary Results

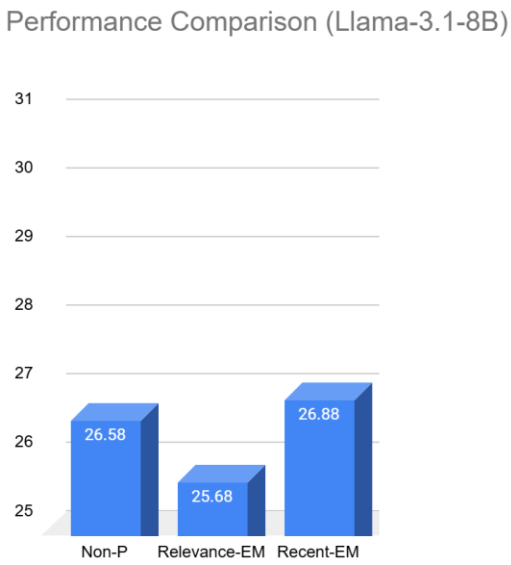

- Episodic memory grounded in simple **recency** often outperforms a **semantic-similarity** retrieval strategy.

Memory Instantiation Preliminary Results

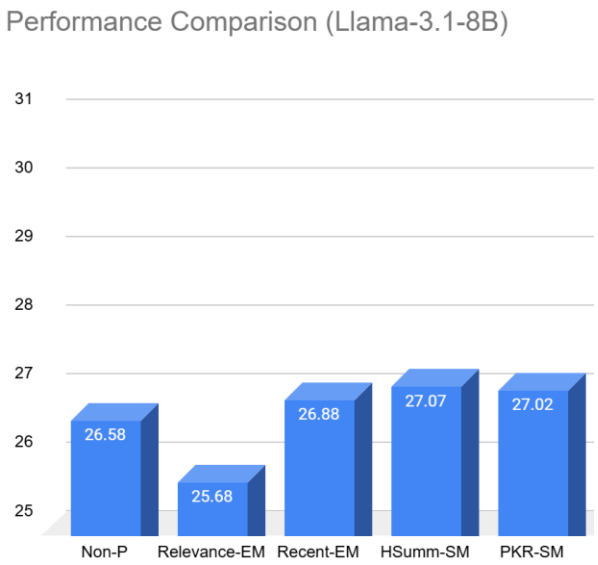

- Episodic memory grounded in simple **recency** often outperforms a **semantic-similarity** retrieval strategy.
- PKR-produced summaries are of **equal qualities** as Hierarchical Summarization

Memory Instantiation Preliminary Results

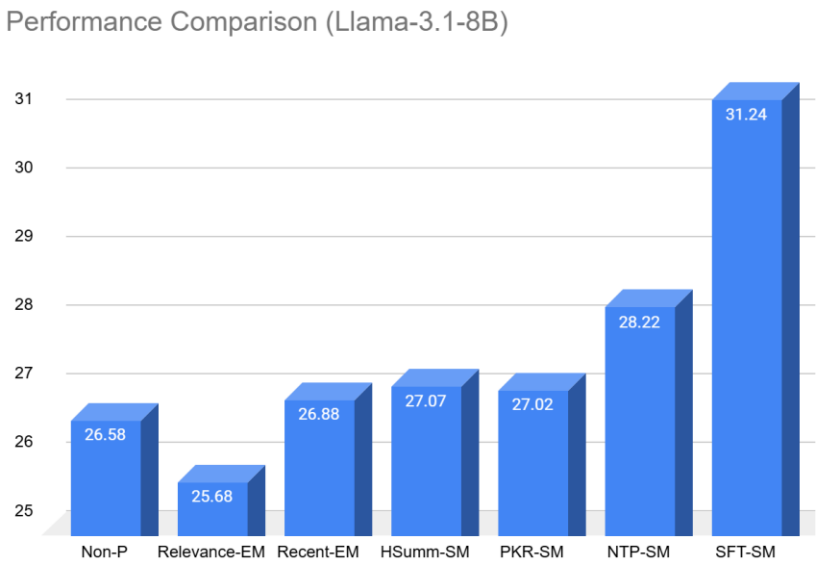

- Episodic memory grounded in simple **recency** often outperforms a **semantic-similarity** retrieval strategy.
- PKR-produced summaries are of **equal qualities** as Hierarchical Summarization
- Overall, using SM alone leads to better results compared to using EM alone.
- **EM < Textual SM < Parametric SM**

Memory Instantiation Preliminary Results

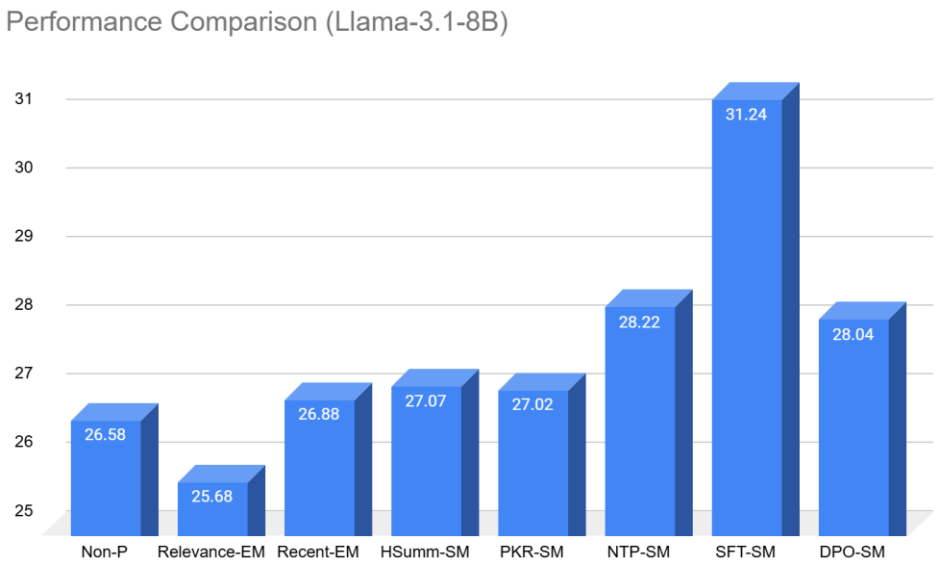

- Episodic memory grounded in simple **recency** often outperforms a **semantic-similarity** retrieval strategy.
- PKR-produced summaries are of equal qualities as Hierarchical Summarization.
- Overall, using SM alone leads to better results compared to using EM alone.
- **EM < Textual SM < Parametric SM**
- Surprisingly, DPO is underperforming, possibly due to **data pairing** issue (1 positive paired with 6-10 negatives).

Memory Instantiation Preliminary Results (complete)

Model	Instantiation	Hit@1	Hit@3	DCG@3	MRR	Avg	W. Efficiency	R. Efficiency
No Personalization								
Llama-3.1-8B	Baseline	16.32	36.77	28.10	25.11	26.58	N/A	N/A
Qwen2.5-7B		16.91	39.10	29.43	26.13	27.89		
Episodic Memory (EM)								
Llama-3.1-8B	Recent	**16.62**	37.22	**28.36**	**25.33**	**26.88**	Fastest	Slow
Qwen2.5-7B		13.91	37.47	27.10	23.57	25.51		
Llama-3.1-8B	Relevant	16.17	35.41	27.00	24.12	25.68	Fastest	Slower
Qwen2.5-7B		13.23	38.50	27.36	23.56	25.66		
Llama-3.1-8B	Recent+PKR	**16.62**	36.84	28.10	25.10	26.67	Medium	Slower
Qwen2.5-7B		14.29	37.07	27.05	23.62	25.51		
Llama-3.1-8B	Relevant+PKR	15.64	36.32	27.45	24.41	25.96	Medium	Slowest
Qwen2.5-7B		13.76	**39.02**	27.88	24.07	26.18		
Semantic Memory (SM)								
Llama-3.1-8B	NTP	17.44	41.20	30.93	27.31	29.22	Fast	Fast
Qwen2.5-7B		16.84	39.55	29.71	26.34	28.11		
Llama-3.1-8B	CIG	17.74	41.95	31.56	27.92	29.79	Fast	Fast
Qwen2.5-7B		16.77	40.23	30.05	26.57	28.41		
Llama-3.1-8B	Output FT	14.66	36.47	26.99	23.75	25.47	Medium-Fast	Fast
Qwen2.5-7B		16.54	39.85	29.58	26.08	28.01		
Llama-3.1-8B	Task FT	**19.62**	43.01	**32.96**	**29.36**	**31.24**	Medium	Fast
Qwen2.5-7B		16.99	**43.38**	32.15	28.28	30.20		
Llama-3.1-8B	DPO	15.41	37.37	27.89	24.64	26.33	Slowest	Fast
Qwen2.5-7B		16.77	39.55	29.61	26.22	28.04		
Llama-3.1-8B	SIMPO	14.21	34.81	25.89	22.88	24.45	Slow	Fast
Qwen2.5-7B		10.08	24.66	18.44	16.30	17.37		
Llama-3.1-8B	HSumm	16.32	37.89	28.62	25.44	27.07	Slowest	Medium
Qwen2.5-7B		15.04	38.80	28.50	24.97	26.83		
Llama-3.1-8B	PKR	16.69	36.39	28.12	25.26	26.62	Medium	Medium
Qwen2.5-7B		15.34	39.02	28.63	25.08	27.02		

Personalized Thinking

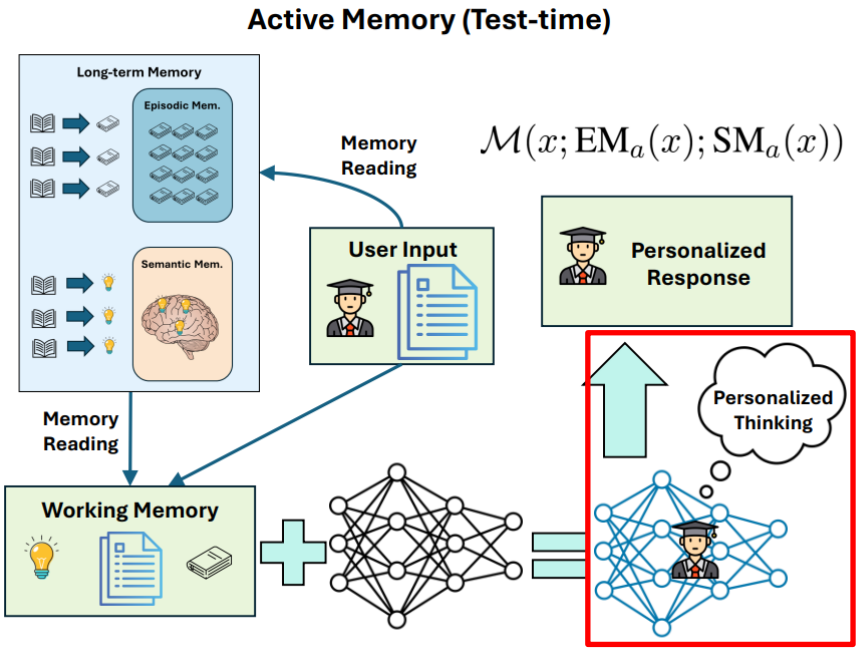

Personalized Thinking

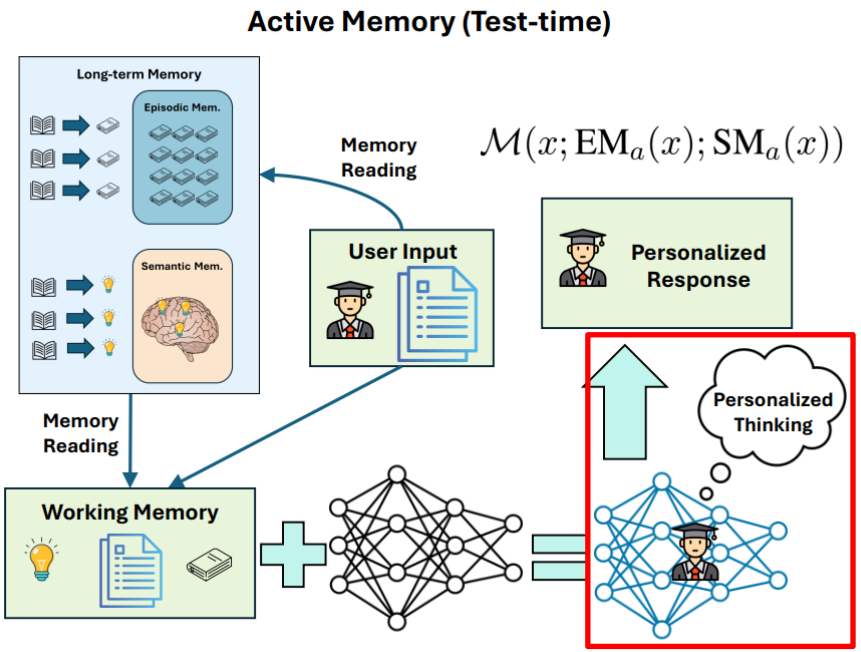

Slow thinking, though showing promises, is still in its infant stage when applying to personalization.

Personalized Thinking

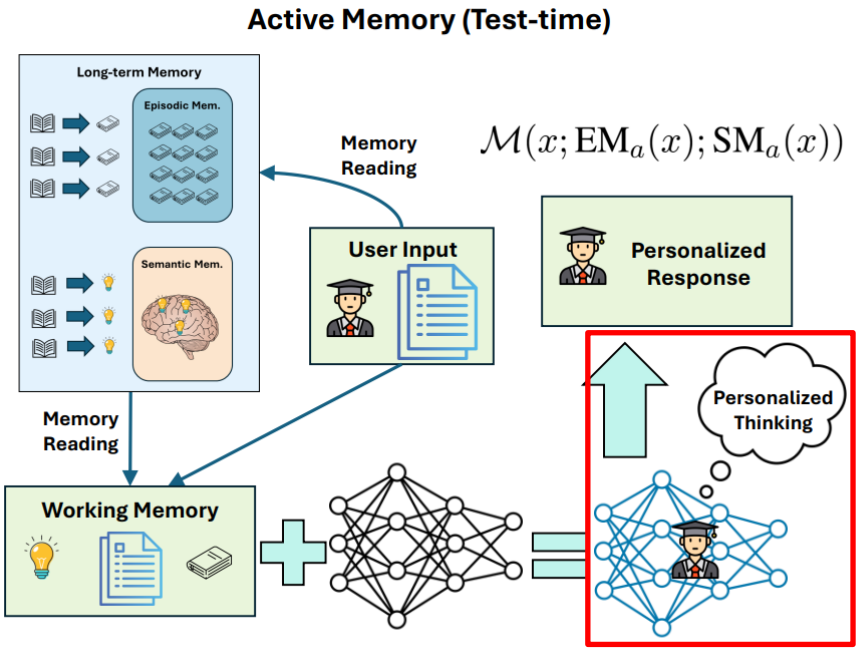

Slow thinking, though showing promises, is still in its infant stage when applying to personalization.

	Non-thinking	Thinking (R1-distill)
Llama-3.1-8B	26.58	25.18 (-5.27%)
Qwen2.5-7B	27.89	23.36 (-16.24%)
Qwen2.5-14B	30.24	31.15 (+3.01%)

Personalized Thinking

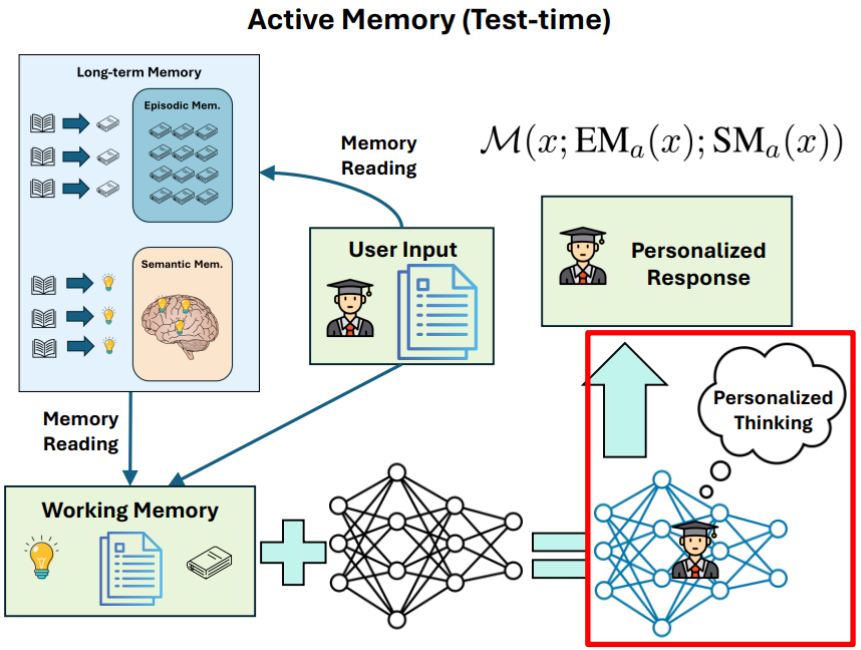

Slow thinking, though showing promises, is still in its infant stage when applying to personalization.

	Non-thinking	Thinking (R1-distill)
Llama-3.1-8B	26.58	25.18 (-5.27%)
Qwen2.5-7B	27.89	23.36 (-16.24%)
Qwen2.5-14B	30.24	31.15 (+3.01%)

Naively applying thinking models backfires!

Challenges of Realizing Personalized Thinking

- No preference data at hand for verifier training in RLHF
- No strong open-sourced thinking LLM from which we can distill personalized thinking trace
- Fast-thinking Fine-tuned LLMs cannot produce meaningful thoughts when prompted

The author, `kingpatzer`, has engaged with users on the Change-My-View subreddit across various original posts (OPs) and is seeking alternative opinions to alter their viewpoint.
Currently, the author is creating a new OP titled

> "CMV: Those who attribute gun ownership rates as the cause of the problem of gun violence in terms of criminal gun deaths are not merely mistaken; they are disingenuous"

with the following content:

> The data has been clear for a very long time: the relationship between guns and gun homicides doesn't show any strong correlation.
>
> I have personally taken the cause-of-death data from https://wonder.cdc.gov/, grouping results by year and state, and selecting *Homicide, Firearm* as the cause of death. I then matched that data to the per-capita gun-ownership statistics by state from the ATF, as reported by Hunting Mark (https://huntingmark.com/gun-ownership-stats/).
>
> A standard correlation analysis between firearm homicide rates per 100,000 and per-capita gun ownership yields an r^2 of 0.079 (no meaningful correlation). A similar global analysis by nation gives an r^2 of 0.02…
>
> The only way to associate gun ownership with gun violence is by including suicides by firearm, which I argue is disingenuous. We don't count suicide by hanging as "rope violence" when discussing strangulation, nor overdoses as "drug violence," etc.

From the `candidate replies` JSON file below, select the top 3 replies (using `option ID`) that best challenge the author's view. Rank them from most to least compelling.

```
[
  { 'option ID': '...', 'challenger': '...', 'reply': '...'},
  { 'option ID': '...', 'challenger': '...', 'reply': '...'},
]
```

["option F", "option D", "option C"]

Personalized Thinking through self-distillation

- **Step 1 (Profile Generation):** We prompt an semantic-memory-instantiated LLM, M, to generate a **profile summary**, following PKR procedure.
- **Step 2 (Review History Engagement):** We convert each historical engagement into a **query**, and prompt the same M to **answer** them.
- **Step 3 (Fast-thinking Filtering):** We apply **rejection sampling** to keep the queries the model M is able to get right.
- **Step 4 (Proxy LLM Initialization & Reasoning):** We follow **the textual semantic memory reading process** to instantiate M', where the summary is generated by M. We apply **reverse engineering** by feeding into M' the *input query and the answer, and prompt it to generate meaningful intermediate thoughts*.
- **Step 5 (Slow-thinking Filtering):** Final round of **rejection sampling** to keep reasoning traces where the final answer matches the ground truth.

Personalized Thinking (Sample Model Output)

Personalized Thinking Trace

User Experience: The author, `kingpatzer`, has engaged with users on the Change-My-View subreddit across various original posts (OPs), seeking alternative opinions to alter their viewpoint. In past discussions, they've shown openness to different perspectives and a willingness to consider data and evidence about gun ownership and gun violence. However, they can become defensive when challenged and sometimes dismiss opposing viewpoints.

OP Summary: The author is creating a new OP titled

> CMV: Those who attribute gun ownership rates as the cause of the problem of gun violence in terms of criminal gun deaths are not merely mistaken; they are disingenuous

They argue that available data show no strong correlation between gun ownership rates and gun homicides, and that including firearm suicides is disingenuous—comparing it to "rope violence" or "drug violence" when discussing strangulation or overdoses.

Sketch Outline:
To challenge the author's view, we need to address the following points:
1. The author's interpretation of the data: We need to examine the author's methodology and data analysis to see if it is sound and if it supports their conclusion. 2. The inclusion of suicide by guns: We need to discuss the relevance of including suicide by guns in the data set and why it is not disingenuous to do so. 3. The comparison to other causes of death: We need to compare the author's approach to other causes of death, such as drug overdoses, to see if it is consistent and fair. 4. The author's tone and language: We need to address the author's tone and language, which some may perceive as dismissive and condescending.
Key questions to ask:
* What is the author's definition of "disingenuous" and how does it apply to including suicide by guns in the data set? * Is the author's methodology sound and does it support their conclusion? * How does the author's approach to gun ownership and gun violence compare to other causes of death? * Why is the author using a tone and language that some may perceive as dismissive and condescending?

Analysis (option F):
This reply challenges the author's view by pointing out that the data used to support their conclusion is flawed, and that the author's methodology is not sound. The reply also highlights the importance of considering the social and cultural context in which gun ownership and gun violence occur.

Reflection:
Based on the analysis, the top 3 replies that best challenge the author's view are:
1. Option F: This reply challenges the author's view by pointing out that the data used to support their conclusion is flawed, and that the author's methodology is not sound. The reply also highlights the importance of considering the social and cultural context in which gun ownership and gun violence occur. 2. Option D: This reply challenges the author's view by pointing out that there is a strong link between gun ownership and impulsive suicide, and that including suicide by guns in the data set is not disingenuous. The reply also highlights the importance of considering the cultural and social context in which gun ownership and gun violence occur. 3. Option C: This reply provides a detailed and well-researched argument that challenges the author's view, pointing out that there is a significant correlation between gun ownership and gun violence. The reply also highlights the importance of considering the social and cultural context in which gun ownership and gun violence occur.
These replies are the most compelling because they engage with the author's perspective, provide reasoned and respectful arguments, and challenge the author's view in a thoughtful and nuanced way.

Answer:

$$\left["\text{option F}",\ "\text{option D}",\ "\text{option C}"\right]$$

Experimental Results (Selected)

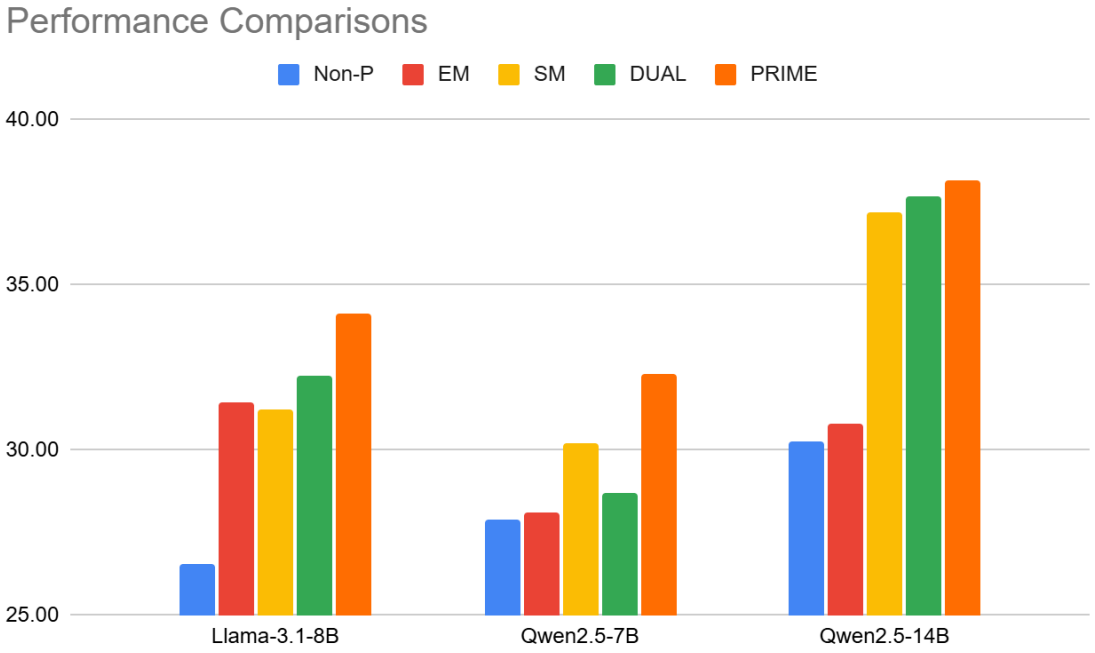

Experimental Results (Selected)

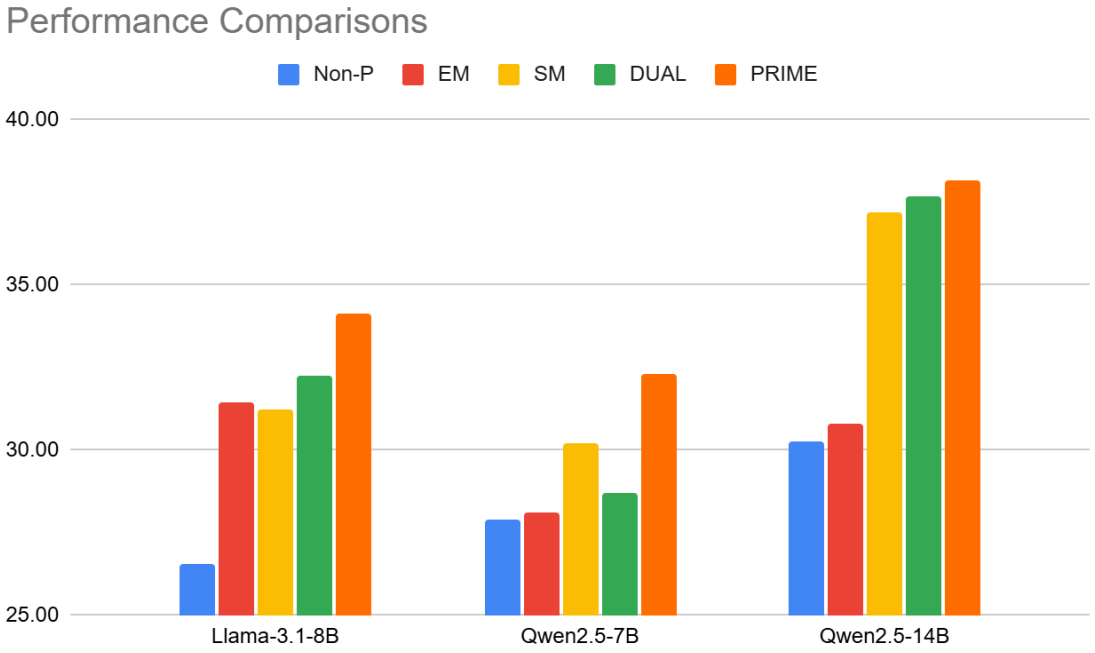

Generic reasoning has limitations

Semantic memory (SM) beats episodic memory (EM)

DUAL often underperforms SM alone

Model-agnostic effectiveness of PRIME

Personalized thinking is crucial

Thanks!

Paper

Codebase & Data

Codebase and dataset are available at github.com/launchnlp/LM_Personalization.

Contact: xlfzhang@umich.edu

This work supported by AFOSR and UM Advanced Research Computing.