JPMorgan Chase & Co.

NEURAL INFORMATION PROCESSING SYSTEMS

Personas within Parameters: Fine-Tuning Small Language Models with Low-Rank Adapters to Mimic User Behaviors

Himanshu Thakur, Eshani Agrawal, Smruthi Mukund

Methodology Overview

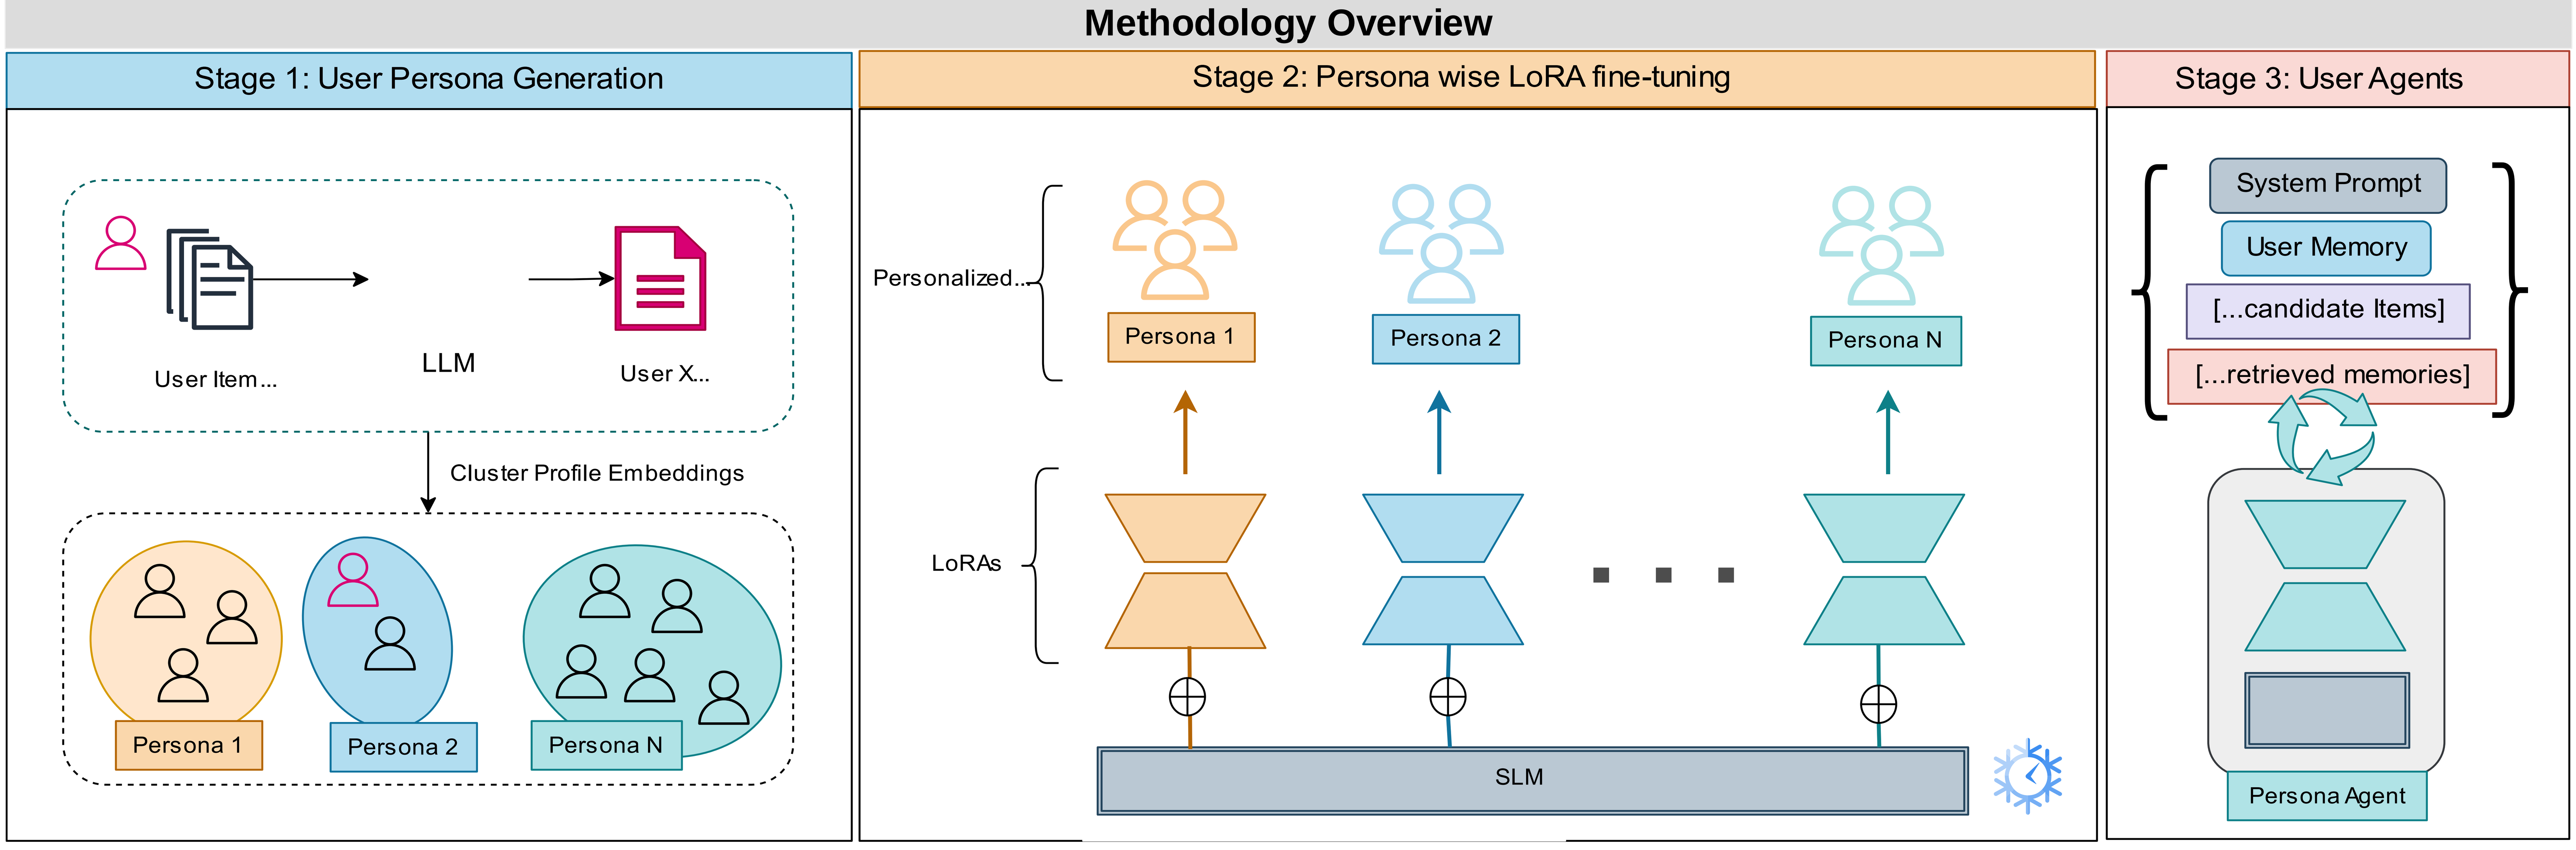

Research Summary

Problem Statement:

- Evaluation of recommender system modes need high quality user behavior simulation.
- Current LLM-based approaches are limited in either scalability and/or personalization capabilitites.

Key Challenges:

- Utilizing large-scale user-item interaction data.
- Overcoming pre-training biases in LLMs.
- Ensuring scalability for millions of users.

Proposed Solutions:

- Fine-tune SLMs using LoRAs to create cost-effective, scalable user agents.
- Introduce personas or user clusters to balance model complexity and personalization quality.

Results

No.	Base Model	Adapter(s)	Memories Used	RMSE (↓)	MAE (↓)	URR (↓)
0	`LLaMA-3`	-	M_s	**1.158**	**0.847**	0%
1	`LLaMA-3`	-	$M_s + M_l$	1.221	0.883	0%
2	`LLaMA-3`	-	M_l	1.229	0.888	0%
3	`Phi-3-Mini`	-	M_s	**1.203**	**0.863**	4.23%
4	`Phi-3-Mini`	-	$M_s + M_l$	1.315	0.952	**2.93%**
5	`Phi-3-Mini`	$\theta_{\text{single (max)}}$	M_s	1.312	0.940	1.30%
6	`Phi-3-Mini`	$\theta_{\text{single (min)}}$	M_s	1.180	0.848	1.93%
7	`Phi-3-Mini`	$\theta_{\text{single (max)}}$	$M_s + M_l$	**1.150**	**0.834**	**1.15%**
8	`Phi-3-Mini`	$\theta_{\text{single (min)}}$	$M_s + M_l$	1.337	1.042	3.01%
9	`Phi-3-Mini`	θ_{persona}	M_s	1.292	0.937	2.29%
10	`Phi-3-Mini`	θ_{persona}	$M_s + M_l$	<u>1.171</u>	<u>0.881</u>	<u>1.77%</u>

Table 1: Performance comparison of user agents built using different model and dataset strategies. The absence of an adapter means we did not fine-tune and only prompted the base model. $\theta_{\text{single (min)}}$ signifies that a single LoRA was trained on ≈ 2100 samples, while $\theta_{\text{single (max)}}$ indicates it was trained on ≈ 5200 samples. The presence of M_s or M_l indicates the type of memories that were used to train and evaluate the models. The numbers in bold show the best scores per category, and underlines show the comparison between single LoRA and persona LoRA.

Implementation

1. Distilling User Preferences

- Use hierarchical self-reflection to extract short-term (Ms) and long-term (Ml) memories from user interactions.
- Enriched descriptions improve user-agent grounding.

2. Clustering and Fine-Tuning:

- Group users into personas based on text-embeddings.
- Train persona-specific LoRAs to enhance efficiency.

3. Memory Utilization:

- Short-term memory (Ms): General user traits.
- Long-term memory (Ml): Detailed interaction history.

Key Findings

✓ **Persona-based LoRAs** outperforms a single LoRA based model.

✓ Incorporating **long-term memory** enhances personalization.

Future Work and Limitations

- Dependence on high-quality LLM outputs for distillation.
- Computational cost of hyperparameter tuning for LoRAs.
- Optimize persona generation using lightweight features.
- Explore alternative parameter-efficient fine-tuning methods.