

On the Way to LLM Personalization: Learning to Remember User Conversations

Lucie Charlotte Magister^{1,*}

Katherine Metcalf²

Yizhe Zhang²

Maartje ter Hoeve²

¹University of Cambridge

²Apple

lcm67@cam.ac.uk, {kmetcalf, yizhe_zhang, m_terhoeve}@apple.com

Abstract

Large Language Models (LLMs) have quickly become an invaluable assistant for a variety of tasks. However, their effectiveness is constrained by their ability to tailor responses to human preferences and behaviors via personalization. Prior work in LLM personalization has largely focused on style transfer or incorporating small factoids about the user, as knowledge injection remains an open challenge. In this paper, we explore injecting knowledge of prior conversations into LLMs to enable future work on less redundant, personalized conversations. We identify two real-world constraints: (1) conversations are sequential in time and must be treated as such during training, and (2) per-user personalization is only viable in parameter-efficient settings. To this aim, we propose PLUM, a pipeline performing data augmentation for up-sampling conversations as question-answer pairs, that are then used to finetune a low-rank adaptation adapter with a weighted cross entropy loss. Even in this first exploration of the problem, we perform competitively with baselines such as RAG, attaining an accuracy of 81.5% across 100 conversations.

1 Introduction

Large Language Models (LLMs) have quickly become a go-to resource for learning about new topics or assisting with a plethora of tasks. However, to fully unlock the models' capabilities, responses require personalization, tuning the model to the user's preferences and needs (Salemi et al., 2024b; Zhuang et al., 2024; Salemi et al., 2025a). Prior work on LLM personalization has largely focused on adapting to the user's style and preferences via Reinforcement Learning from Human Feedback (RLHF) (Poddar et al., 2024; Chen et al., 2024; Li

et al., 2024b; Maghakian et al., 2022; Salemi et al., 2025b) and Parameter-Efficient Finetuning (PEFT) (Zhuang et al., 2024; Tan et al., 2024). To move beyond simple preference learning and small facts about the user, Retrieval Augmented Generation (RAG) methods have been used to integrate user profiles and conversation history into the model's generation process (Salemi et al., 2024b; Mysore et al., 2023; Li et al., 2023; Salemi et al., 2024a; Zhang et al., 2024a). However, RAG-based methods require maintaining external storage and picking a suitable number of documents to retrieve, with LLM performance having been shown to deteriorate with larger context windows (Vodrahalli et al., 2024; Zhang et al., 2024b). Furthermore, Woźniak et al. (2024) show that personalized finetuning yields improved model reasoning over zero-shot prompting. This leads us to question whether parametric knowledge injection could yield a more streamlined approach to encoding knowledge of prior interactions with the user, opening new avenues for future LLM personalization.

Current works on knowledge injection mainly focus on adding fact-based knowledge to LLMs (Ovadia et al., 2023; Fu et al., 2023; Mecklenburg et al., 2024) and note that knowledge injection remains an open challenge (Fu et al., 2023). In this paper, we tackle the challenge of injecting knowledge of prior user conversations into LLMs via finetuning. Our efforts are in support of enabling future personalization research. Given the focus on user conversations, we identify and impose two important constraints: (1) finetuning must be parameter efficient and (2) conversations are sequential in nature and we do not want to store them like RAG-based methods. Note that we focus on inter-conversation rather than intra-conversation knowledge. We are interested in remembering the conversation holistically after it has occurred, as individual turns within a conversation can usually be injected into the context window. To the best of our knowledge, this is

*Work done while interning at Apple.

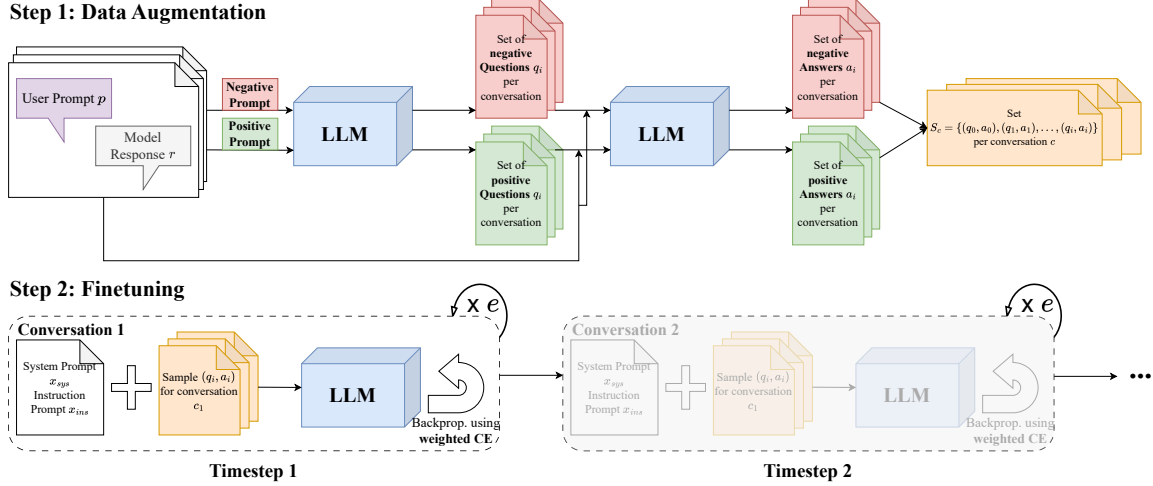


Figure 1: On overview of PLUM, a two-stage pipeline for injecting knowledge of prior user conversations into the LLM. The first step of the pipeline focuses on augmenting user conversations as positive and negative question-answer pairs about the conversation. These are then used in the finetuning step, where the LLM is trained on samples of a single conversation at a time for 10 epochs with a weighted cross entropy loss.

the first work considering to move beyond learning simple facts about the user to learning user conversations via finetuning. Specifically, we propose a *Pipeline for Learning User Conversations in Large Language Models* (PLUM), which extracts question-answer pairs from conversations to finetune a LLM with a Low-Rank Adaptation (LoRA) adapter (Hu et al., 2021) using a weighted cross entropy (CE) loss. In this initial exploration of the problem, PLUM achieves an accuracy of 81.5% across 100 conversations compared to 83.5% with RAG-based baselines. Furthermore, we present an extensive set of ablations to guide and inspire future endeavors for the parametric personalization of LLMs, moving beyond simple RAG systems.

2 Related Work

2.1 Personalization

LLM personalization focuses on tuning model responses to the user’s preferences and needs (Salemi et al., 2024b). Existing works can broadly be split into three categories or a combination thereof: (1) prompting techniques, (2) RAG-based techniques and (3) RLHF and/or PEFT methods. Prompt-based techniques focus on encoding user preferences or conversation history in soft prompts (Hebert et al., 2024; Huang et al., 2024; Shen et al., 2024) or hard prompts (Mao et al., 2024; Li et al., 2024a). Similarly, RAG-based techniques have been leveraged to add relevant information from a user’s history to the LLM’s context (Wu et al.,

2021; Salemi et al., 2024b; Lu et al., 2023; Kumar et al., 2024; Wang et al., 2024; Neelakanteswara et al., 2024; Salemi et al., 2024a; Zerhoubi and Granitzer, 2024; Sun et al., 2024; Huang et al., 2023), varying mostly in the type of information and the manner in which it is stored. However, a common challenge of these techniques is the reliance on prompt tuning and the selection of relevant data points (Zhuang et al., 2024). RLHF (Stiennon et al., 2020; Jang et al., 2023; Cheng et al., 2023; Park et al., 2024; Poddar et al., 2024; Li et al., 2024b) and low-rank PEFT-based methods (Tan et al., 2024; Zhuang et al., 2024) alleviate these issues by directly encoding user information in model parameters, however, they are usually limited to preferences or simple facts. In contrast, our work focuses on teaching the LLM to remember user conversations, which may span multiple facts and preferences. Specifically, we aim to provide an alternative to RAG by exploring injecting knowledge of previous user conversations.

2.2 Knowledge Injection

Knowledge injection focuses on adding new knowledge to the LLM after the initial training phase. Methods vary from prompt-based techniques (Chen et al., 2022) to incorporating external knowledge sources via RAG (Song et al., 2016; Fan et al., 2020; Lewis et al., 2020; Martino et al., 2023; Zhou et al., 2024; Sun et al., 2024) to finetuning adapters (Wang et al., 2021; Ovadia et al., 2023; Fu et al., 2023; Mecklenburg et al., 2024). We take inspira-

tion from Mecklenburg et al. (2024), who explore how to augment data to directly incorporate it into the LLM via PEFT. However, our work diverges in the application to user conversation data with the aim of enabling future personalization research. Moreover, knowledge of user conversations can be seen as a special type of knowledge, as the contents are drawn from the model’s knowledge (Gekhman et al., 2024).

3 PLUM

We propose PLUM, a two-stage pipeline for injecting knowledge of prior user conversations into LLMs. The first stage encompasses augmenting the user conversation data. The second stage focuses on training via PEFT with a custom loss function. We refer the reader to Figure 1 for a visual overview.

3.1 Data Augmentation

We define a user conversation as a set of turns between the user and model, starting with an original user prompt. For simplicity, we will stick to single turn conversations, however, the proposed method can easily be extended to multi-turn. Recall that our goal is to enable future personalization by remembering user conversations, making single turn conversations a natural starting point. Given a user conversation c defined as the tuple (p, r) , where p is the user prompt and r the LLM’s response, we use the same LLM to generate a set of question-answer pairs about the conversation c . This set of question-answer pairs can be denoted as $S_c = \{(q_0, a_0), (q_1, a_1), \dots, (q_i, a_i)\}$, where q_i is a question about the conversation and a_i the corresponding answer to the question. To generate S_c , we provide the LLM with the original conversation and prompt it to ask as many questions as reasonable about the conversation. We do not ask the LLM to generate a specific number of questions, because some conversations are more data-rich than others, i.e., for some conversations there are more questions that can be asked.

We generate two types of questions using few-shot prompting for guidance. The first type are open-ended questions such as ‘What did we discuss about ...?’, while the second are focused on eliciting a clear ‘yes’ or ‘no’ response, such as ‘Did we discuss ...?’. We then provide the original conversation and the individual questions about the conversation to the LLM for answering. Be-

sides positive question-answer pairs, we also want to generate negative pairs. These are questions asking about something not covered by the conversation, eliciting a ‘no’ response. This is to reinforce the knowledge boundary of the LLM and prevent hallucination, as we observe that without negative samples the LLM will default to always positively answering questions. To generate negative samples, we ask the LLM to pose questions adjacent to the topic of the conversation to which the answer is ‘no’, to increase precision and not clash with other samples. We propose maintaining a balance between positive and negative samples so that the LLM does not err in one direction. Appendix A and B document our prompts.

3.2 Parameter-Efficient Finetuning

We impose two design constraints for the injection of conversation history. First, we note that conversations are sequential in nature, which means that we should finish finetuning on one conversation before moving on to the next. This allows for discarding the conversation after all of its samples have been iterated over, which stands in contrast to RAG-based techniques. However, it also poses the challenge of overcoming catastrophic forgetting (Luo et al., 2023). Second, we note that finetuning all model parameters per user is infeasible, therefore, we use a PEFT method.

Given these constraints, we propose finetuning a LoRA adapter conversation by conversation. We propose a LoRA adapter based on its robust performance in the previous knowledge injection work of Mecklenburg et al. (2024). Each training example x_i consists of four key elements:

$$x_i = x_{sys} + x_{ins} + a_i + q_i \quad (1)$$

The first two elements are an optional system prompt x_{sys} and an instruction prompt x_{ins} . These two elements are consistent across all training examples, while the final two elements are a question q_i and corresponding answer a_i sampled from the set S_c for the conversation c . To finetune the LoRA adapter, we use teacher forcing for more stable training (Williams and Zipser, 1989). In teacher forcing the model is provided with the true previous tokens rather than the predicted ones. We refer the reader to Appendix C for the full prompt.

In addition to this specific data setup, we empirically derive a custom loss based on the CE loss (Hinton et al., 2006). We propose scaling up the

CE loss on the question and answer tokens. We can rewrite the standard CE loss as a weighted version in the following way:

$$L = H(P, Q_{(x_{sys}, x_{ins})}) + \lambda H(P, Q_{(x_{q_i}, x_{a_i})}), \quad (2)$$

where H is the CE loss measuring the difference between the true distribution of the data P and the model’s predicted distribution of the data Q . We compute the standard CE loss over the tokens x_{sys} and x_{ins} . In contrast, we scale the CE loss over the tokens q_i and a_i of a training example by λ , specifically $\lambda = 10$. We empirically derive this loss after observing that the standard CE loss quickly diminishes with the model only having learned x_{sys} , x_{ins} and the overall structure of the prompt well. We find that weighting q_i and a_i focuses the model on the actual question and answer and the knowledge they convey, rather than only matching x_{sys} and x_{ins} . We also obtain this behavior when exploring different prompt phrasing. We refer the reader to Section 5.1.2 for a discussion of ablations.

Based on our exploration, we specifically propose finetuning a LoRA adapter of rank $r = 16$ and scaling parameter $\alpha = 64$, that attaches to all linear layers in the LLM. We also suggest training on each conversation for $e = 10$ epochs. To clarify, all data samples for a conversation are shown to the model for 10 epochs before moving on to the next conversation. We also find that an equal number of positive and negative question-answer pairs for a conversation yields the best results, as well as using a batch size of $b = 8$. We further elaborate on these suggestions in Section 5.

4 Experimental Setup

4.1 Data

We rely on the OpenAssistant Conversations dataset (Köpf et al., 2023) for the initial human-generated prompts. We select 100 prompts in English as starting points for the conversations. We limit our initial exploration to 100 conversations to allow for more control during the data generation and because the user may also forget about prior conversations (Wixted and Ebbesen, 1991). The selected starting prompts are focused on knowledge transfer, rather than completing a specific task, as we believe task-related queries are of more temporary nature. We then generate a response to complete the single-turn conversations. Recall, that we use these conversations to generate two types of question-answer pairs: positive and negative, as

shown in Figure 1. We filter these pairs, checking that the questions and answers align with the expected format and directionality of the answer. We refer the reader to Appendix A and B for further details on prompting and filtering. After filtering, we have 3726 positive and negative question-answer pairs across 100 conversations. We manually spot check the generated data to verify its quality. We withhold two questions per conversation for the test dataset. While the train dataset contains open-ended and closed questions, the test dataset only contains a positive and negative closed question per conversation that can be answered with ‘yes’ and ‘no’, respectively. We select ‘yes’ and ‘no’ questions for evaluation, as they have a clear target. We also train and test the LLM’s performance on question-answer pairs not generated by itself but a larger LLM, to evaluate the reliance on in-distribution data and question-answer formulations.

4.2 Model

For our study of injecting conversation history into LLMs, we focus on Llama 3 8B Instruct (Dubey et al., 2024), because of its high performance on a variety of tasks given its reasonable size. As mentioned previously, we also employ a larger model to generate a second train and test dataset. For this, we use Llama 3 70B Instruct (Dubey et al., 2024) to generate an alternative version of the train and test dataset that does not directly align with the distribution of Llama 3 8B Instruct, the LLM which we finetune, due to capacity and training data differences. This is further motivated by the work of Hong et al. (2024), who find that model scale matters for capturing the structure of language. We refer the reader to Appendix D for further details on the finetuning setup.

4.3 Metrics

As previously mentioned, we maintain a holdout dataset of questions that can be answered with a simple ‘yes’ or ‘no’. We use these questions to evaluate the overall accuracy of the model after being finetuned on all conversations in the dataset. We also track the accuracy over time, which is an average of the model accuracy on all seen conversations for a given finetuning step. For example, after the model has seen the samples for n conversations, it is evaluated on the corresponding evaluation questions for these n conversations. Evaluating the accuracy over time provides an insight into how strongly a model has learned the answer and the

potential for catastrophic forgetting.

We also evaluate model accuracy on a suite of 5 common benchmarking tasks to ensure PLUM does not deteriorate the base performance of the model. We evaluate the base model and a PLUM-finetuned version on: (1) Measuring Massive Multitask Language Understanding (MMLU) (Hendrycks et al., 2021), (2) HellaSwag (Zellers et al., 2019), (3) ARC (Clark et al., 2018), (4) PIQA (Bisk et al., 2020) and (5) Social IQA (SiQA) (Sap et al., 2019). We use the Language Model Evaluation Harness framework (Gao et al., 2024) to perform a 1-shot and 5-shot evaluation on these tasks.

4.4 Baselines

The focus of our paper resides on enabling future personalization research by injecting knowledge of prior user conversations via PEFT. To contextualize our results, we compare our method to a standard RAG baseline, following Salemi et al. (2024b). Specifically, we train three retriever models based on BM25 (Robertson et al., 1995), a strong term matching model. The first BM25 baseline *Conv. RAG* is trained on the original conversation data. The second baseline *Sum. RAG* is trained only on summaries of the original conversations, to mimic common setups in existing literature (Richardson et al., 2023). Lastly, the *Q/A RAG* baseline is a BM25 model trained on the question-answer pairs used for LoRA finetuning for a more comparable setup. We test different settings for k , the number of documents to retrieve, ranging from $k = \{1, 2, 3\}$. We choose BM25, because of its enduring strong performance (Salemi et al., 2024a; Izacard et al., 2022). We do not compare to neural-based retrievers, such as Contriever (Izacard et al., 2022), because we only focus on 100 conversations for this study, which a neural retriever may easily overfit on. Moreover, our data setup for some baselines is not compatible with the training of Contriever, which requires positive and negative sample pairs for the contrastive loss.

5 Results

5.1 Ablations on Remembering User Conversations

We ablate various elements of PLUM to evaluate the contributions of the different components. By dissecting the elements of our framework, we aim to provide insights into the underlying mechanisms driving its performance and inspire future research.

Model Setup		Accuracy (%)		
		Yes	No	Overall
PLUM		73.0	77.0	75.0
PLUM (w/ sys.)		71.0	92.0	81.5
Epochs	$e = 1$	82.0	30.0	56.0
	$e = 1$ (w/ sys.)	70.0	46.0	58.0
	$e = 5$	70.0	78.0	74.0
	$e = 5$ (w/ sys.)	84.0	57.0	70.5
	$e = 15$	39.0	91.0	65.0
	$e = 15$ (w/ sys.)	91.0	52.0	71.5
	$e = 20$	71.0	55.0	63.0
	$e = 20$ (w/ sys.)	0.0	0.0	0.0
CE	$e = 1$	5.0	98.0	51.5
	$e = 10$	21.0	13.0	17.0
	$e = 20$	0.0	0.0	0.0
	w/ sys.	38.0	44.0	41.0
Loss Var.	(q_i, a_i) -only CE	88.0	25.0	56.5
	(q_i, a_i) -only	57.0	95.0	<u>76.0</u>
	a_i -only	40.0	90.0	65.0
	a_i -only (w/ sys.)	88.0	35.0	61.5
Data	70B Model Gen.	30.0	80.0	55.0
	Upsampled Yes	75.0	76.0	75.5
	Upsampled No	69.0	83.0	<u>76.0</u>

Table 1: Model accuracy on various ablations. The best and second best overall accuracy are in **bold** and underlined.

We present our results in Table 1. Our ablations focus on the impact of epochs, the design of the loss function and the data. We also ablate our method with and without a system prompt. We provide further ablations on the LoRA architecture, batch size and random seeds in Appendix E.

5.1.1 Impact of Epochs

We ablate the number of epochs required to remember a conversation. Recall that epochs refers to the number of times the model sees the training examples for a conversation before moving on to the next. We find that increasing or decreasing the number of epochs from $e = 10$ deteriorates accuracy. Notably, for all settings except $e = 10$ the model overfits to positive samples. For example, at $e = 1$ we observe an imbalance in accuracy of 82.0% versus 30.0% for ‘yes’ and ‘no’ samples, respectively. This could be related to batching, as we backpropagate on $b = 8$ samples at a time. Batching multiple examples may cause oscillation between erring on the positive and negative side, indicating adapting the number of epochs per conversation as an interesting area for further exploration. Furthermore,

we observe the accuracy dropping to 0.0% for two settings with $e = 20$, as the model begins to output incoherent sentences.

5.1.2 Impact of the Loss

We compare our weighted CE loss against the standard CE loss. Recall that in our weighted CE loss, we scale the loss of the question and answer tokens by $\lambda = 10$. We achieve an accuracy of 75.0% without and 81.5% with the system prompt. In comparison, simply training using the standard CE loss yields an accuracy of 17.0% and 41.0% without and with the system prompt, respectively. This highlights the significance of our design choice of up-weighting the loss on the question-answer section of the input prompt. If we only train on the question-answer pairs, this means we drop the system and instruction prompt, scaling the loss still achieves a significant improvement at an accuracy of 76.0% compared to 56.5%. Lastly, we also examine whether only scaling the loss on the answer tokens is sufficient. We find that this only yields an overall accuracy of 65.0% without and 61.5% with the system prompt.

Model Setup		Accuracy over Time (%)	
		Yes	No
PLUM		68.2 ± 17.9	79.2 ± 17.3
PLUM (w/ sys.)		79.7 ± 18.3	59.2 ± 28.1
Epochs	$e = 1$	58.7 ± 29.4	52.2 ± 29.9
	$e = 1$ (w/ sys.)	45.7 ± 25.7	66.3 ± 22.4
	$e = 5$	60.6 ± 23.4	70.4 ± 26.8
	$e = 5$ (w/ sys.)	81.8 ± 23.4	42.2 ± 28.4
	$e = 15$	55.6 ± 22.7	79.5 ± 15.2
	$e = 15$ (w/ sys.)	82.9 ± 19.3	51.4 ± 25.0
	$e = 20$	46.8 ± 20.4	79.4 ± 13.5
	$e = 20$ (w/ sys.)	30.8 ± 39.0	30.5 ± 36.8
CE	$e = 1$	1.9 ± 3.0	98.9 ± 1.7
	$e = 10$	54.7 ± 35.0	41.9 ± 35.2
	$e = 20$	26.7 ± 37.0	22.3 ± 32.4
	w/ sys.	54.9 ± 34.1	53.5 ± 32.3
Loss Var.	(q_i, a_i) -only CE	63.1 ± 33.4	55.1 ± 35.7
	(q_i, a_i) -only	59.4 ± 17.7	85.3 ± 18.0
	a_i -only	46.9 ± 21.6	75.5 ± 22.4
	a_i -only (w/ sys.)	82.5 ± 21.6	35.0 ± 22.7
Data	70B Model Gen.	39.1 ± 27.4	84.7 ± 16.6
	Upsampled ‘Yes’	64.3 ± 21.3	70.9 ± 24.5
	Upsampled ‘No’	67.3 ± 20.5	70.0 ± 23.2

Table 2: Model accuracy **over time** (including standard deviation) on ‘yes’ and ‘no’ questions for various ablations.




Figure 2: Accuracy **over time** for PLUM with the system prompt.

5.1.3 Impact of Data

We ablate our assumptions on the data balance between positive and negative samples, as well as staying within the model’s distribution. We find that increasing the number of positive or negative samples by 25% per conversation deteriorates results. Specifically, the model’s accuracy to respond with ‘yes’ or ‘no’ increases for the respective cases, but decreases for the opposite. This indicates that maintaining a balance between samples is the most beneficial. Lastly, we also verify whether it is beneficial to stay within the model’s distribution by training on samples not generated by the model itself. When training Llama 3 8B Instruct on data generated by its 70B counterpart, accuracy significantly deteriorates to 30.0% on ‘yes’ questions. This can potentially be explained by training on data generated by the model itself reinforcing the knowledge and only requiring to store whether it was discussed previously. In contrast, the different wording and potentially divergent knowledge generated by the 70B model is not as simple to learn. We provide further results in Appendix F.

5.2 Performance over Time

We measure the model’s accuracy over time to detect issues such as catastrophic forgetting. Table 2 summarizes the ‘yes’ and ‘no’ accuracy over time for various ablations. We observe similar trends as in our analysis of the overall model accuracy in Section 5.1, however, a noteworthy insight obtained from the accuracy over time is the standard deviation of the accuracy. For example, the average ‘yes’ and ‘no’ accuracy overtime for PLUM without a system prompt is 68.2% and 79.2% with a standard deviation of $\pm 17.9\%$ and $\pm 17.3\%$, re-

Baselines	Llama 3 8B Instruct		PLUM		PLUM (w/ sys.)	
	1-Shot	5-Shot	1-Shot	5-Shot	1-Shot	5-Shot
MMLU	64.2 \pm 0.4	65.7 \pm 0.4	63.0 \pm 0.4	64.9 \pm 0.4	63.2 \pm 0.4	65.3 \pm 0.4
HellaSwag	57.5 \pm 0.5	58.4 \pm 0.5	56.9 \pm 0.5	57.8 \pm 0.5	56.6 \pm 0.5	57.2 \pm 0.5
ARC-Easy	83.9 \pm 0.8	85.2 \pm 0.7	82.7 \pm 0.8	83.6 \pm 0.8	83.2 \pm 0.8	83.5 \pm 0.8
ARC-Cha.	55.4 \pm 1.5	57.4 \pm 1.4	54.0 \pm 1.5	56.1 \pm 1.5	54.1 \pm 1.5	55.5 \pm 1.5
PiQA	79.7 \pm 0.9	80.6 \pm 0.9	79.0 \pm 1.0	80.1 \pm 0.9	78.7 \pm 1.0	80.0 \pm 0.9
SiQA	53.8 \pm 1.1	56.8 \pm 1.1	54.7 \pm 1.1	57.9 \pm 1.1	54.7 \pm 1.1	56.7 \pm 1.1

Table 3: Model performance on a selection of benchmarking tasks before and after finetuning on user conversations. While we generally observe a slight deterioration in accuracy, performance remains within a reasonable range.

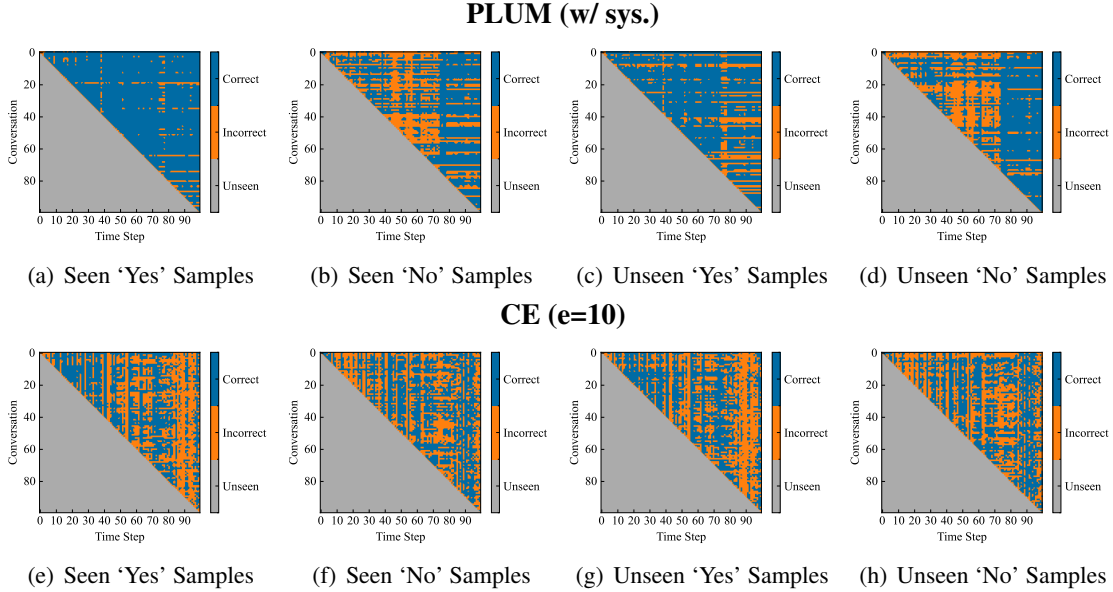


Figure 3: Consistency plots visualizing whether the ‘yes’/‘no’ question was predicted correctly (blue) or incorrectly (orange) for a given time step. Here, a time step refers to the model having seen all samples of a conversation for the specified number of epochs. The lower left triangle of the plot is gray, as these conversations have not been seen yet.

spectively. This is quite a significant range, which indicates that the model oscillates between answering ‘yes’ and ‘no’. Figure 2 plots the accuracy as the model learns more conversations. We can observe that the model oscillates between erring on the ‘yes’ or ‘no’ side.

To further investigate this, we plot the predictions over time in Figure 3. We call these plots *consistency plots*, as ideally the upper triangle of each plot would be blue, meaning that the model consistently predicts the correct answer. We find that with PLUM, the LLM fairly consistently learns how to answer the questions for a given conversation, however, it fails to learn some conversations altogether, indicated by consistent streaks of orange. We observe no sign of catastrophic forgetting. Moreover, in Figure 3(d) it appears as if the LLM initially struggles to reply with ‘no’, but then learns this, starting at around time step 70. This shift towards

‘no’ is also recognizable in the corresponding ‘yes’ Figure 3(c), as there are more streaks of orange starting at time step 70. This could indicate the the LLM learns to balance responding with ‘yes’ and ‘no’. Nevertheless, we should observe that PLUM allows strides in the right direction, despite the oscillations. This is especially evident when contrasting the consistency plots of PLUM (Figures 3(a)-3(d)) with those generated when using standard CE loss with $e = 10$ (Figures 3(e)-3(h)).

5.3 Model Performance

Table 3 shows the 1-shot and 5-shot performance of Llama 3 8B Instruct and the PLUM-finetuned model on five common benchmarking tasks. We can observe a slight deterioration in accuracy across all tasks except SiQA. We observe a slight improvement in accuracy from 53.8% to 54.7% on the SiQA dataset, which may be attributed to train-

ing on prior conversations being somewhat related to reasoning about social settings. Overall, we can conclude that PLUM does not negatively impact performance on standard benchmarking tasks.

5.4 Comparison against Baselines

Table 4 details the accuracy of PLUM against various baselines. PLUM with and without a system prompt achieves a competitive accuracy of 81.5% and 75.0%, respectively. For context, we include the 0-shot performance of Llama 3 8B Instruct and a perfect recall version of RAG, where we automatically inject the correct conversation or summary into the prompt. The 0-shot performance of the model is 50.0%, as the model defaults to responding with ‘no’. The highest accuracy of 89.5% is achieved when providing the conversation to the LLM (*Perfect Conv. Recall*). We can see this as the near-optimal performance of the LLM on the given data, but not necessarily an upper bound, as this is dependent on data quality. In contrast, the best performing RAG baseline is *Conv. RAG* with $k = 3$ at 86.5%. We observe an increase in accuracy as k increases, as the BM25 model does not always return the correct conversation as the top one. The *Sum. RAG* baseline achieves only 71.0% accuracy, which can be attributed to the summaries potentially missing details needed to answer the questions. The *Q/A RAG* baseline is the most comparable to PLUM, as it has access to the same data. It achieves 83.5% accuracy for $k = 3$. The performance of PLUM is just shy of this, highlighting injecting conversation history into LLMs as a promising avenue for future research on personalization.

6 Discussion

6.1 Take-aways

We show that PLUM allows LLMs to efficiently remember user conversations. We can identify three key takeaways from our experiments that allow this success. First, the number of times the training samples for a conversation are shown to the model is highly important. We found that $e = 10$ provides the best trade-off between the accuracy of ‘yes’ and ‘no’ questions. A lower or higher number of epochs generally leads the model to err towards ‘yes’. More importantly, we found that the weighted CE loss is necessary for the successful recall of the question-answer pairs. In our experiments without the weighted loss or different configurations of weighting tokens, we observed

Method	Accuracy (%)		
	Yes	No	Overall
0-shot Performance	0.0	100.0	50.0
Perfect Conv. Recall	100.0	79.0	89.5
Perfect Sum. Recall	83.0	78.0	80.5
Conv. RAG ($k=1$)	86.0	80.0	83.0
Conv. RAG ($k=2$)	84.0	84.0	<u>84.0</u>
Conv. RAG ($k=3$)	84.0	89.0	86.5
Sum. RAG ($k=1$)	56.0	85.0	70.5
Sum. RAG ($k=2$)	67.0	81.0	74.0
Sum. RAG ($k=3$)	68.0	74.0	71.0
Q/A RAG ($k=1$)	66.0	94.0	80.0
Q/A RAG ($k=2$)	71.0	94.0	82.5
Q/A RAG ($k=3$)	73.0	94.0	83.5
PLUM	73.0	77.0	75.0
PLUM (w/ sys.)	71.0	92.0	81.5

Table 4: Performance of RAG-based baselines versus PLUM, with the best and second best overall accuracy in **bold** and underlined.

significant model deterioration. We also found that a balance of positive and negative samples is important for reinforcement of the knowledge boundary. Providing more positive samples deteriorates the performance on negative ones and vice versa.

6.2 Future Work

Our key takeaways indicate potential for future research. An area for further study is tuning the number of epochs per conversation, as some conversations may be easier to learn than others. Other avenues include further experiments on the loss, batch size and data sampling strategies. In particular, experiments exploring whether the findings of Chang et al. (2024) for the memorization of LLMs in pretraining would be interesting. Another area of research is improving the model’s answer consistency, as we observe large variations between time steps. Studying memorization profiles as proposed by Lesci et al. (2024) may be interesting here. Lastly, PLUM enables future research in personalization. For example, the injected knowledge of conversations can be used to reduce redundancy in responses or refine knowledge transfer between the LLM and user by building on past interactions.

7 Conclusion

In this work, we explore injecting knowledge of prior user conversations into LLMs. We propose PLUM, a pipeline for finetuning a LoRA adapter on question-answer pairs about prior conversa-

tions, maintaining conversation order. Moreover, we propose a custom loss for improved results. We achieve competitive results to RAG baselines, while not requiring to store each individual conversation. Our results indicate directly injecting personalization data into LLMs as an interesting avenue for future research, such as reducing redundancy in subsequent conversations or extending reasoning capabilities.

8 Limitations

Despite the contributions of this work, our work must be viewed in the context of a few limitations. It should be noted that all data used is in English and that we only verified PLUM on Llama 3 8B Instruct. Moreover, we limited our study to only 100 conversations, which can be seen as a reasonable but small subset of conversations over time. Lastly, we did not explore model performance on remembering conversations on the same versus clashing topics.

9 Ethical Considerations

We must carefully examine the ethical implications of our work. Remembering user conversations may lead to personal information being stored in the parametric knowledge of a LLM, which an adversary may extract (Mattern et al., 2023). This applies to personalization in general. Moreover, personalization may provide identifying information about the user leading to biases in the generated text (Wang et al., 2023; He et al., 2024; Weissburg et al., 2024). Furthermore, future work should consider how a model personalized with prior user conversations should talk to the user, as in Liao and Sundar (2021). We urge the reader to carefully consider the aforementioned points in their work extending or using PLUM to handle the user’s privacy with care.

References

Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. 2020. Piqa: Reasoning about physical commonsense in natural language. In *Thirty-fourth AAAI Conference on Artificial Intelligence*.

Hoyeon Chang, Jinho Park, Seonghyeon Ye, Sohee Yang, Youngkyung Seo, Du-Seong Chang, and Minjoon Seo. 2024. How do large language models acquire factual knowledge during pretraining? In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, Xingmei Wang, and 1 others. 2024. When large language models meet personalization: Perspectives of challenges and opportunities. *World Wide Web*, 27(4):42.

Xiang Chen, Ningyu Zhang, Xin Xie, Shumin Deng, Yunzhi Yao, Chuanqi Tan, Fei Huang, Luo Si, and Huajun Chen. 2022. Knowprompt: Knowledge-aware prompt-tuning with synergistic optimization for relation extraction. In *Proceedings of the ACM Web conference 2022*, pages 2778–2788.

Pengyu Cheng, Jiawen Xie, Ke Bai, Yong Dai, and Nan Du. 2023. Everyone deserves a reward: Learning customized human preferences. *arXiv preprint arXiv:2309.03126*.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *ArXiv*, abs/1803.05457.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.

Zhihao Fan, Yeyun Gong, Zhongyu Wei, Siyuan Wang, Yameng Huang, Jian Jiao, Xuan-Jing Huang, Nan Duan, and Ruofei Zhang. 2020. An enhanced knowledge injection model for commonsense generation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2014–2025.

Peng Fu, Yiming Zhang, Haobo Wang, Weikang Qiu, and Junbo Zhao. 2023. Revisiting the knowledge injection frameworks. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10983–10997.

Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, and 5 others. 2024. [A framework for few-shot language model evaluation](#).

Zorik Gekhman, Gal Yona, Roei Aharoni, Matan Eyal, Amir Feder, Roi Reichart, and Jonathan Herzig. 2024. [Does fine-tuning LLMs on new knowledge encourage hallucinations?](#) In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 7765–7784, Miami, Florida, USA. Association for Computational Linguistics.

Jerry Zhi-Yang He, Sashrika Pandey, Mariah L Schrum, and Anca Dragan. 2024. Cos: Enhancing personalization and mitigating bias with context steering. *arXiv preprint arXiv:2405.01768*.

- Liam Hebert, Krishna Sayana, Ambarish Jash, Alexandros Karatzoglou, Sukhdeep Sodhi, Sumanth Dodapaneni, Yanli Cai, and Dima Kuzmin. 2024. Persoma: Personalized soft prompt adapter architecture for personalized language prompting. *arXiv preprint arXiv:2408.00960*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Geoffrey E Hinton, Simon Osindero, and Yee-Whye Teh. 2006. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7):1527–1554.
- Zhuoqiao Hong, Haocheng Wang, Zaid Zada, Harshvardhan Gazula, David Turner, Bobbi Aubrey, Leonard Niekerken, Werner Doyle, Sasha Devore, Patricia Dugan, and 1 others. 2024. Scale matters: Large language models with billions (rather than millions) of parameters better match neural representations of natural language. *bioRxiv*, pages 2024–06.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Qiushi Huang, Shuai Fu, Xubo Liu, Wenwu Wang, Tom Ko, Yu Zhang, and Lilian Tang. 2023. [Learning retrieval augmentation for personalized dialogue generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2523–2540, Singapore. Association for Computational Linguistics.
- Qiushi Huang, Xubo Liu, Tom Ko, Bo Wu, Wenwu Wang, Yu Zhang, and Lilian Tang. 2024. [Selective prompting tuning for personalized conversations with LLMs](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 16212–16226, Bangkok, Thailand. Association for Computational Linguistics.
- Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised dense information retrieval with contrastive learning. *Transactions on Machine Learning Research*.
- Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. 2023. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. *arXiv preprint arXiv:2310.11564*.
- Diederik P Kingma. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi Rui Tam, Keith Stevens, Abdullah Barhoum, Duc Nguyen, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. 2023. [Openassistant conversations - democratizing large language model alignment](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 47669–47681. Curran Associates, Inc.
- Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, and 1 others. 2024. Longlamp: A benchmark for personalized long-form text generation. *arXiv preprint arXiv:2407.11016*.
- Pietro Lesci, Clara Meister, Thomas Hofmann, Andreas Vlachos, and Tiago Pimentel. 2024. [Causal estimation of memorisation profiles](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15616–15635, Bangkok, Thailand. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, and 1 others. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, and Michael Bendersky. 2024a. Learning to rewrite prompts for personalized text generation. In *Proceedings of the ACM on Web Conference 2024*, pages 3367–3378.
- Cheng Li, Mingyang Zhang, Qiaozhu Mei, Yaqing Wang, Spurthi Amba Hombaiah, Yi Liang, and Michael Bendersky. 2023. Teach llms to personalize—an approach inspired by writing education. *arXiv preprint arXiv:2308.07968*.
- Xinyu Li, Zachary C Lipton, and Liu Leqi. 2024b. Personalized language modeling from personalized human feedback. *arXiv preprint arXiv:2402.05133*.
- Mengqi Liao and S Shyam Sundar. 2021. How should ai systems talk to users when collecting their personal information? effects of role framing and self-referencing on human-ai interaction. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–14.
- Junru Lu, Siyu An, Mingbao Lin, Gabriele Pergola, Yulan He, Di Yin, Xing Sun, and Yunsheng Wu. 2023. Memochat: Tuning llms to use memos for consistent long-range open-domain conversation. *arXiv preprint arXiv:2308.08239*.
- Yun Luo, Zhen Yang, Fandong Meng, Yafu Li, Jie Zhou, and Yue Zhang. 2023. An empirical study of catastrophic forgetting in large language models during continual fine-tuning. *arXiv preprint arXiv:2308.08747*.

- Jessica Maghakian, Paul Mineiro, Kishan Panaganti, Mark Rucker, Akanksha Saran, and Cheng Tan. 2022. Personalized reward learning with interaction-grounded learning (igl). *arXiv preprint arXiv:2211.15823*.
- Wenyu Mao, Jiancan Wu, Weijian Chen, Chongming Gao, Xiang Wang, and Xiangnan He. 2024. Reinforced prompt personalization for recommendation with large language models. *arXiv preprint arXiv:2407.17115*.
- Ariana Martino, Michael Iannelli, and Coleen Truong. 2023. Knowledge injection to counter large language model (llm) hallucination. In *European Semantic Web Conference*, pages 182–185.
- Justus Mattern, Fatemehsadat Miresghallah, Zhijing Jin, Bernhard Schoelkopf, Mrinmaya Sachan, and Taylor Berg-Kirkpatrick. 2023. [Membership inference attacks against language models via neighbourhood comparison](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 11330–11343, Toronto, Canada. Association for Computational Linguistics.
- Nick Mecklenburg, Yiyu Lin, Xiaoxiao Li, Daniel Holstein, Leonardo Nunes, Sara Malvar, Bruno Silva, Ranveer Chandra, Vijay Aski, Pavan Kumar Reddy Yannam, and 1 others. 2024. Injecting new knowledge into large language models via supervised fine-tuning. *arXiv preprint arXiv:2404.00213*.
- Sheshera Mysore, Zhuoran Lu, Mengting Wan, Longqi Yang, Steve Menezes, Tina Baghaee, Emmanuel Barajas Gonzalez, Jennifer Neville, and Tara Safavi. 2023. Pearl: Personalizing large language model writing assistants with generation-calibrated retrievers. *arXiv preprint arXiv:2311.09180*.
- Abhiman Neelakanteswara, Shreyas Chaudhari, and Hamed Zamani. 2024. [RAGs to style: Personalizing LLMs with style embeddings](#). In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, pages 119–123, St. Julians, Malta. Association for Computational Linguistics.
- Oded Ovadia, Menachem Brief, Moshik Mishaeli, and Oren Elisha. 2023. Fine-tuning or retrieval? comparing knowledge injection in llms. *arXiv preprint arXiv:2312.05934*.
- Chanwoo Park, Mingyang Liu, Kaiqing Zhang, and Asuman Ozdaglar. 2024. Principled rlhf from heterogeneous feedback via personalization and preference aggregation. *arXiv preprint arXiv:2405.00254*.
- Sriyash Poddar, Yanming Wan, Hamish Ivison, Abhishek Gupta, and Natasha Jaques. 2024. Personalizing reinforcement learning from human feedback with variational preference learning. *arXiv preprint arXiv:2408.10075*.
- Chris Richardson, Yao Zhang, Kellen Gillespie, Sudipta Kar, Arshdeep Singh, Zeynab Raeesy, Omar Zia Khan, and Abhinav Sethy. 2023. Integrating summarization and retrieval for enhanced personalization via large language models. *arXiv preprint arXiv:2310.20081*.
- Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, and 1 others. 1995. Okapi at trec-3. *Nist Special Publication Sp*, 109:109.
- Alireza Salemi, Surya Kallumadi, and Hamed Zamani. 2024a. Optimization methods for personalizing large language models through retrieval augmentation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 752–762.
- Alireza Salemi, Julian Killingback, and Hamed Zamani. 2025a. Expert: Effective and explainable evaluation of personalized long-form text generation. *arXiv preprint arXiv:2501.14956*.
- Alireza Salemi, Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, Tao Chen, Zhuowan Li, Michael Bendersky, and Hamed Zamani. 2025b. Reasoning-enhanced self-training for long-form personalized text generation. *arXiv preprint arXiv:2501.04167*.
- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024b. [LaMP: When large language models meet personalization](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7370–7392, Bangkok, Thailand. Association for Computational Linguistics.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. [Social IQa: Commonsense reasoning about social interactions](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4463–4473, Hong Kong, China. Association for Computational Linguistics.
- Xiaoteng Shen, Rui Zhang, Xiaoyan Zhao, Jieming Zhu, and Xi Xiao. 2024. Pmg: Personalized multimodal generation with large language models. In *Proceedings of the ACM on Web Conference 2024*, pages 3833–3843.
- Yiping Song, Rui Yan, Xiang Li, Dongyan Zhao, and Ming Zhang. 2016. Two are better than one: An ensemble of retrieval and generation-based dialog systems. *arXiv preprint arXiv:1610.07149*.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 3008–3021. Curran Associates, Inc.

- Jingwei Sun, Zhixu Du, and Yiran Chen. 2024. Knowledge graph tuning: Real-time large language model personalization based on human feedback. *arXiv preprint arXiv:2405.19686*.
- Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin, and Meng Jiang. 2024. Democratizing large language models via personalized parameter-efficient fine-tuning. *arXiv preprint arXiv:2402.04401*.
- Kiran Vodrahalli, Santiago Ontanon, Nilesh Tripuraneni, Kelvin Xu, Sanil Jain, Rakesh Shivanna, Jeffrey Hui, Nishanth Dikkala, Mehran Kazemi, Bahare Fatemi, and 1 others. 2024. Michelangelo: Long context evaluations beyond haystacks via latent structure queries. *arXiv preprint arXiv:2409.12640*.
- Hongru Wang, Wenyu Huang, Yang Deng, Rui Wang, Zezhong Wang, Yufei Wang, Fei Mi, Jeff Z Pan, and Kam-Fai Wong. 2024. Unims-rag: A unified multi-source retrieval-augmented generation for personalized dialogue systems. *arXiv preprint arXiv:2401.13256*.
- Nan Wang, Qifan Wang, Yi-Chia Wang, Maziar Sanjabi, Jingzhou Liu, Hamed Firooz, Hongning Wang, and Shaoliang Nie. 2023. Coffee: Counterfactual fairness for personalized text generation in explainable recommendation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 13258–13275.
- Ruize Wang, Duyu Tang, Nan Duan, Zhongyu Wei, Xuanjing Huang, Jianshu Ji, Guihong Cao, Daxin Jiang, and Ming Zhou. 2021. *K-Adapter: Infusing Knowledge into Pre-Trained Models with Adapters*. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1405–1418, Online. Association for Computational Linguistics.
- Iain Weissburg, Sathvika Anand, Sharon Levy, and Hae-won Jeong. 2024. LLMs are biased teachers: Evaluating LLM bias in personalized education. *arXiv preprint arXiv:2410.14012*.
- Ronald J Williams and David Zipser. 1989. A learning algorithm for continually running fully recurrent neural networks. *Neural computation*, 1(2):270–280.
- John T Wixted and Ebbe B Ebbesen. 1991. On the form of forgetting. *Psychological science*, 2(6):409–415.
- Stanisław Woźniak, Bartłomiej Koptyra, Arkadiusz Janz, Przemysław Kazienko, and Jan Kocoń. 2024. Personalized large language models. *arXiv preprint arXiv:2402.09269*.
- Yuwei Wu, Xuezhe Ma, and Diyi Yang. 2021. *Personalized response generation via generative split memory network*. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1956–1970, Online. Association for Computational Linguistics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. *HellaSwag: Can a machine really finish your sentence?* In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Saber Zerhouni and Michael Granitzer. 2024. Personarag: Enhancing retrieval-augmented generation systems with user-centric agents. *arXiv preprint arXiv:2407.09394*.
- Kai Zhang, Yangyang Kang, Fubang Zhao, and Xiaozhong Liu. 2024a. *LLM-based medical assistant personalization with short- and long-term memory coordination*. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2386–2398, Mexico City, Mexico. Association for Computational Linguistics.
- Xinrong Zhang, Yingfa Chen, Shengding Hu, Zihang Xu, Junhao Chen, Moo Hao, Xu Han, Zhen Thai, Shuo Wang, Zhiyuan Liu, and Maosong Sun. 2024b. *∞Bench: Extending long context evaluation beyond 100K tokens*. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15262–15277, Bangkok, Thailand. Association for Computational Linguistics.
- Yujia Zhou, Qiannan Zhu, Jiajie Jin, and Zhicheng Dou. 2024. Cognitive personalized search integrating large language models with an efficient memory mechanism. In *Proceedings of the ACM on Web Conference 2024*, pages 1464–1473.
- Yuchen Zhuang, Haotian Sun, Yue Yu, Qifan Wang, Chao Zhang, and Bo Dai. 2024. Hydra: Model factorization framework for black-box LLM personalization. *arXiv preprint arXiv:2406.02888*.

A Generating the Original Conversation

A.1 Conversation Prompt

We use the OpenAssistant Dataset (Köpf et al., 2023), a crowd-sourced dataset with content moderation, to sample 100 initial conversation prompts. These conversation prompts are human generated and in English. As we are looking for knowledge-based prompts, we hand-select 100 of these. The selected prompts do not contain personally identifiable information or offensive content.

A.2 Conversation Response Generation

We use the following system prompt to generate a response to the initial starting prompt from the OpenAssistant Dataset (Köpf et al., 2023):

You are a helpful LLM answering questions concisely. Do not ramble or generate repetitive output.

We then simply add the original conversation prompt. Examining the dataset, we found that some of the prompts are not formatted correctly so we apply simple corrections such as capitalizing the first word of a sentence and adding a question mark at the end of sentences.

B Question-Answer Pair Generation

We generate question-answer pairs via a two step process. First we infer a set of questions about the conversation. Then we simply prompt the model to answer the question about the conversation, providing the conversation as context. In total, we generate 4 types of question-answer pairs:

1. Positive Open-ended Questions
2. Negative Open-ended Questions
3. Positive Closed-ended Questions (‘Yes’ Questions)
4. Negative Closed-ended Questions (‘No’ Questions)

We generate open-ended questions about the conversation to elicit a summary style answer. We generate positive and negative questions of this type, where positive means that the topic was indeed discussed in the conversation, while negative means that the topic was not discussed and the model should politely express this. Similarly, we also generate closed-ended questions of this type. These are questions that should be answerable with a ‘yes’ and ‘no’.

B.1 4-shot Sample Conversations

In order to elicit open and closed-ended questions from the model, we perform 4-shot prompting. We write four knowledge-based single-turn conversations for this. We phrase some initial starting prompts as questions and others as instructions (e.g. ‘Tell me about ...’). We ensure that the responses to the questions have enough content to write a varied set of positive and negative sample questions.

B.2 4-shot Sample Questions

For each of the four few-shot sample conversations, we then write 12 positive and negative, open-ended and closed questions for each of the conversations. Recall that closed questions are of the format ‘Did we discuss...?’ to elicit a ‘yes’ and ‘no’ response. The open-ended questions begin with ‘What did we

discuss about ...?’ to elicit a longer, summary-style responses. The 4-shot prompting with example questions allows us to demonstrate to the model that we are looking for positive questions, questions that can directly be answered via the conversation, and negative questions, questions about information adjacent to the topic of discussion. Note that we define 12 questions of each style of question for each few-shot example to encourage the model to write a substantial number of questions. However, we do not directly prompt the model to generate this many questions later on, to make sure that the questions are of high quality. We discuss this further in Section B.3.

B.3 Question Generation

B.3.1 Positive Questions

In order to elicit the model to generate similar positive questions for our sample conversations generated on the base of the OpenAssistant dataset (Köpf et al., 2023), we prompt it in the following way with the 4-shot examples prepended. We define the system prompt x_{sys} as:

You are a helpful LLM specialized in inferring the topic of a conversation and writing questions about what was discussed about this topic in the conversation. Do not deviate from the topics and contents of the conversation. All questions must be answerable with ‘yes’. Only speak from the ‘we’ perspective of ‘you and the user’. You must start your sentence with ‘Did we discuss...’. Always specify the topic you are referring to, avoiding ambiguity, so that the question makes sense without the conversation. Never say ‘the conversation’.

We then add the following instruction prompt x_{ins} :

From the user perspective, write as many questions as sensible about what was discussed in the following conversation.
`<START_CONVERSATION> USER: {user_prompt}
 YOU: {model_response} <END_CONVERSATION>`

As indicated above, we inject the user prompt and model response in place of {user_prompt} and {model_response}, respectively. The system prompt and instruction prompt for the other types of questions follow a similar structure. Note that we only ask the model to generate as many ques-

tions about the conversation as sensible. This is because some conversations may be more data-rich than others, which makes it easier to ask varied questions. We observed this behavior during prompt tuning. Please also note that we have removed line breaks for formatting purposes here.

B.3.2 Negative Questions

In order to elicit negative questions, we replace the system prompt x_{sys} with the following:

You are a helpful LLM specialized in inferring the topic of a conversation and writing 12 questions about closely related topics that were not covered in the conversation. Do not deviate from the overarching topic of the conversation. All questions must be answerable with ‘We did not discuss ...’. Only speak from the ‘we’ perspective of ‘you and the user’. Start your sentence with ‘What did we discuss about...’. Always specify the topic you are referring to, avoiding ambiguity, so that the question makes sense without the conversation. Never say ‘the conversation’.

Similarly, we replace the instruction prompt x_{ins} with the following:

From the user perspective, write 12 questions about something that was not discussed in the following conversation.
`<START_CONVERSATION> USER: {user_prompt}
 YOU: {model_response} <END_CONVERSATION>`

Please note that we include line breaks were appropriate, but we have removed these for the formatting of the paper. Please also note that we specify 12 questions here, rather than just asking the model to generate as many questions as sensible. This is because negative questions, questions about topics adjacent to the conversation, are easier to generate. However, please also notice that we then balance the number of positive and negative samples per conversation by taking the smaller number of the two.

B.4 Answer Generation

After splitting the questions generated, we prompt the model to answer each question individually. We also provide the user conversation as reference for the positive questions. We do not provide the conversation for the negative ones, as the information is not needed and we can rely on the performance

of the base model, which is to politely decline having knowledge of prior conversations.

We use the following positive system prompt x_{sys} to answer the questions about the conversation:

You are a helpful LLM trained to answer questions about prior conversations you had. You are very detailed and summarize the whole conversation to answer the question. You do not include details that are not in the conversation.

We then provide the conversation and the instruction prompt x_{ins} :

–Conversation– USER: {user_prompt}
 AGENT: {model_response} –Task– Please answer the following question about whether you have discussed the indicated topic with the user. Question: {conv_question} Answer:

The model then begins completing the prompt. As before, we have marked the corresponding sections to add in the user prompt, model response and generated conversation question. The negative instruction follows a similar wording, however, we do not provide the conversation as context there. This is because the question cannot be answered from the conversation and the model should politely decline. Please note that we have also removed line breaks here for formatting purposes.

B.5 Data Filtering

We perform some basic checks for filtering. Firstly, we check that the questions generated follow the expected structure, e.g. ‘Did we discuss ...?’. We also check that the generated answer yields the expected response, e.g. it contains ‘yes’. We also remove all duplicate questions.

C Prompts for Model Finetuning

We use the following system prompt x_{sys} for model finetuning:

You are a helpful LLM trained to answer questions about prior conversations you had. If you do not remember having discussed the topic, you state to the user that you do not remember having had this conversation.

We then follow-up with the following instruction prompt x_{ins} :

Please answer the following question about whether you have discussed the

indicated topic.

We then concatenate the question q_i and answer a_i in the following format:

Question {question} Answer: {answer}

We insert the question-answer pair samples for a conversation in place of {question} and {answer}, respectively. All prompts presented were selected by writing variants and manually examining the quality of the performance for a handful of conversations. Again, please note that we have removed line breaks in our prompts for formatting purposes.

D Finetuning

We perform teacher-forcing for finetuning, using the Adam optimizer (Kingma, 2014) with a learning rate of 0.001. We set the random seed to $s = 42$. All other hyperparameters are described in the main text. We train our models in a distributed manner, with a maximum number of 8 NVIDIA A100 80 GB being used for an experiment.

E Further Ablations

E.1 LoRA Architecture

To gain a better understanding of how to best configure the LoRA adapter, we vary the layers the LoRA adapter attaches to and the adapter size. We found that attaching to all linear layers in the model yields the best results, as 81.5% accuracy, compared to only attaching to attention layers, yielding an accuracy of only 63.0%. We also vary the adapter size r , with $r = 8$ and $r = 32$ only achieving an accuracy of 76.5%. Overall, we find that a LoRA adapter attaching to all linear layers, with a size of $r = 16$ and $a = 64$, performs best.

E.2 Batch Size

To gain a better understanding of the effect of the batch size we replicate our setup with different batch sizes where $b = \{1, 16, 32\}$. Note that the batch size we use for PLUM is $b = 8$. We observe that a batch size of $b = 16$ yields slightly improved results at 78.5% compared to a $b = 8$ at 75.0% accuracy. However, the oscillations between ‘yes’ and ‘no’ are smaller at $b = 8$, indicated by the decreased gap between the ‘yes’ and ‘no’ accuracy, as well as the standard deviation for the accuracy across time.

E.3 Reproducibility

To ensure that our results are reproducible, we run PLUM (without a system prompt) on two further

seeds ($s = 7$ and $s = 73$). We observe similar performance at an overall accuracy of 71.5% and 78.0%.

F Performance on Llama 3 70B Instruct Generated Split

Figure 6 summarizes the results of PLUM in comparison to RAG on the dataset generated with Llama 3 70B Instruct model. We observe that RAG significantly benefits from the data generated by the larger model. For example, Q/A RAG with $k = 1$ improves from an accuracy of 80.0% to 89.0%. In contrast, PLUM does not benefit as much from the data generated by the larger model, with the accuracy only increasing from 75.0% to 78.0% in the version without the system prompt. In the case of PLUM with a system prompt, accuracy even deteriorates. This could be due to the model having a harder time remembering user conversations outside of its own model distribution. Using a different model to generate conversations may cause the model to have to remember new knowledge as well as whether a topic has been discussed or not, which is a more challenging task. Therefore, these results should be seen in context of PLUM requiring finetuning versus RAG simply performing retrieval.

G Licenses

We use the OpenAssistant Dataset (Köpf et al., 2023), which is available under the Apache license 2.0. We also rely on a number of benchmarking datasets and tools, such as the MMLU dataset (Hendrycks et al., 2021), HellaSwag (Zellers et al., 2019) and the Language Model Evaluation Harness framework (Gao et al., 2024), which are available under the MIT license. ARC (Clark et al., 2018) is available under the Creative Commons Attribution license, while PIQA (Bisk et al., 2020) is available under the Apache 2.0 License. SiQA (Sap et al., 2019) is available under the CC0 1.0 Universal license. We use these datasets and tools in accordance with their licenses for research only.

Model Setup		Accuracy (%)			Accuracy over Time (%)		
		Yes	No	Overall	Yes	No	Overall
	PLUM	73.0	77.0	75.0	68.2 ± 17.9	79.2 ± 17.3	73.7 ± 6.9
	PLUM (w/ sys.)	71.0	92.0	81.5	79.7 ± 18.3	59.2 ± 28.1	<u>69.5 ± 9.3</u>
Epochs	$e = 1$	82.0	30.0	56.0	58.7 ± 29.4	52.2 ± 29.9	55.4 ± 4.8
	$e = 1$ (w/ sys.)	70.0	46.0	58.0	45.7 ± 25.7	66.3 ± 22.4	56.0 ± 3.8
	$e = 5$	70.0	78.0	74.0	60.6 ± 23.4	70.4 ± 26.8	65.5 ± 9.1
	$e = 5$ (w/ sys.)	84.0	57.0	70.5	81.8 ± 23.4	42.2 ± 28.4	62.0 ± 7.8
	$e = 15$	39.0	91.0	65.0	55.6 ± 22.7	79.5 ± 15.2	67.5 ± 7.6
	$e = 15$ (w/ sys.)	91.0	52.0	71.5	82.9 ± 19.3	51.4 ± 25.0	67.2 ± 7.6
	$e = 20$	71.0	55.0	63.0	46.8 ± 20.4	79.4 ± 13.5	63.0 ± 6.2
	$e = 20$ (w/ sys.)	0.0	0.0	0.0	30.8 ± 39.0	30.5 ± 36.8	30.6 ± 33.6
Batch	$b = 1$	84.0	40.0	62.0	44.9 ± 30.3	58.3 ± 30.3	51.6 ± 15.4
	$b = 16$	83.0	74.0	<u>78.5</u>	67.8 ± 24.7	68.5 ± 26.4	68.2 ± 9.5
	$b = 32$	72.0	75.0	73.5	57.3 ± 24.9	72.8 ± 25.1	65.0 ± 9.1
CE	$e = 1$	5.0	98.0	51.5	1.9 ± 3.0	98.9 ± 1.7	50.4 ± 0.7
	$e = 10$	21.0	13.0	17.0	54.7 ± 35.0	41.9 ± 35.2	48.3 ± 16.7
	$e = 20$	0.0	0.0	0.0	26.6 ± 37.0	22.3 ± 32.4	24.5 ± 26.6
	w/ sys.	38.0	44.0	41.0	54.9 ± 34.1	53.5 ± 32.3	54.2 ± 10.1
Loss Var.	(q_i, a_i) -only CE	88.0	25.0	56.5	63.1 ± 33.4	55.1 ± 35.7	59.1 ± 13.5
	(q_i, a_i) -only	57.0	95.0	76.0	59.4 ± 17.7	85.3 ± 18.0	72.4 ± 6.9
	a_i -only	40.0	90.0	65.0	46.9 ± 21.6	75.5 ± 22.4	61.2 ± 7.1
	a_i -only (w/ sys.)	88.0	35.0	61.5	82.5 ± 21.6	35.0 ± 22.7	58.8 ± 5.2
Data	70B Model Gen.	30.0	80.0	55.0	39.1 ± 27.4	84.7 ± 16.6	61.9 ± 10.1
	Upsampled ‘Yes’	75.0	76.0	75.5	64.3 ± 21.3	70.9 ± 24.5	67.6 ± 7.1
	Upsampled ‘No’	69.0	83.0	76.0	67.3 ± 20.5	70.0 ± 23.2	68.7 ± 7.9
LoRA	Att.-only, $r = 16, \alpha = 64$	92.0	34.0	63.0	75.1 ± 26.1	42.9 ± 28.7	59.0 ± 7.2
	Lin., $r = 16, \alpha = 32$	85.0	65.0	75.0	69.5 ± 24.1	69.9 ± 23.1	69.7 ± 8.1
	Lin., $r = 8, \alpha = 64$	64.0	89.0	76.5	54.4 ± 22.5	85.3 ± 18.1	69.8 ± 8.1
	Lin., $r = 32, \alpha = 64$	63.0	90.0	76.5	66.3 ± 20.2	80.6 ± 19.5	73.5 ± 7.8
Seed	$seed = 7$	62.0	81.0	71.5	69.6 ± 19.0	68.5 ± 22.1	69.0 ± 6.6
	$seed = 73$	78.0	78.0	78.0	66.2 ± 22.2	78.0 ± 22.4	72.1 ± 9.8

Table 5: Model performance on various ablations. The best and second best overall accuracy and accuracy over time are in **bold** and underlined. For completeness, we have included all ablations run.

Method	Test Split Accuracy			Llama 70B Gen. Split Accuracy		
	Yes	No	Overall	Yes	No	Overall
0-shot Performance	0.0	100.0	50.0	0.0	100.0	50.0
Perfect Conv. Recall	100.0	79.0	89.5	100.0	97.0	98.5
Perfect Sum. Recall	83.0	78.0	80.5	100.0	95.0	97.5
Conv. RAG (k=1)	86.0	80.0	83.0	93.0	95.0	94.0
Conv. RAG (k=2)	84.0	84.0	<u>84.0</u>	92.0	89.0	90.5
Conv. RAG (k=3)	84.0	89.0	86.5	93.0	92.0	92.5
Sum. RAG (k=1)	56.0	85.0	70.5	86.0	96.0	91.0
Sum. RAG (k=2)	67.0	81.0	74.0	90.0	89.0	89.5
Sum. RAG (k=3)	68.0	74.0	71.0	95.0	84.0	89.5
Q/A RAG (k=1)	66.0	94.0	80.0	82.0	96.0	89.0
Q/A RAG (k=2)	71.0	94.0	82.5	93.0	96.0	<u>94.5</u>
Q/A RAG (k=3)	73.0	94.0	83.5	94.0	97.0	95.5
PLUM	73.0	77.0	75.0	83.0	73.0	78.0
PLUM (w/ sys.)	71.0	92.0	81.5	89.0	45.0	67.0

Table 6: Performance of RAG-based baselines versus PLUM on the Llama 8B and 70B model generated data splits. The best and second best overall accuracy are in **bold** and underlined.