

Comparing Retrieval-Augmentation and Parameter-Efficient Fine-Tuning for Privacy-Preserving Personalization of Large Language Models

Alireza Salemi

University of Massachusetts Amherst
Amherst, MA, United States
asalemi@cs.umass.edu

Hamed Zamani

University of Massachusetts Amherst
Amherst, MA, United States
zamani@cs.umass.edu

Abstract

Despite its substantial impact on various search, recommendation, and question answering tasks, privacy-preserving methods for personalizing large language models (LLMs) have received relatively limited exploration. There is one primary approach in this area through retrieval-augmented generation (RAG), which generates personalized outputs by enriching the input prompt with information retrieved from the user’s personal data. This paper studies an orthogonal approach to RAG that involves learning user-dependent LLM parameters through parameter-efficient fine-tuning (PEFT). This paper presents the first systematic study for exploration of PEFT for LLM personalization and provides an extensive comparison between RAG- and PEFT-based solutions, across a broad set of seven diverse datasets from the LaMP benchmark. Our results demonstrate that, on average, both RAG- and PEFT-based personalization methods yield 14.92% and 1.07% improvements over non-personalized LLMs, respectively. When combining RAG with PEFT, we observe a further improvement of 15.98%, highlighting the effectiveness of their integration in enhancing personalized text generation. Additionally, we identify a positive correlation between the amount of user data available and the effectiveness of PEFT. This finding suggests that RAG is particularly beneficial for cold-start users—users with limited personal data—while PEFT performs better when more user-specific data is available.

CCS Concepts

• Information systems → Personalization; • Computing methodologies → Natural language generation.

Keywords

Retrieval-augmented generation; personalization; text generation; text classification; parameter-efficient fine-tuning

ACM Reference Format:

Alireza Salemi and Hamed Zamani. 2025. Comparing Retrieval-Augmentation and Parameter-Efficient Fine-Tuning for Privacy-Preserving Personalization of Large Language Models. In *Proceedings of the 2025 International ACM SIGIR Conference on Innovative Concepts and Theories in Information Retrieval (ICTIR '25)*, July 18, 2025, Padua, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3731120.3744595>



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICTIR '25, Padua, Italy

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1861-8/2025/07

<https://doi.org/10.1145/3731120.3744595>

1 Introduction

Personalization has been a longstanding area of interest in the information retrieval community and has been extensively studied for tasks such as search [10], recommendation [22, 43, 48], question answering [18], and conversational search [5, 50]. With the rise of LLMs, personalizing them has recently emerged as a critical topic [21, 40, 41] due to its applications in various real-world systems, such as personalized recommender systems [3, 14], virtual assistants [20, 26], and content generation [1]. These systems benefit from tailoring responses and actions based on individual user preferences, leading to enhanced user experiences, greater user satisfaction, and more effective interactions [51]. Personalization allows LLMs to better understand and predict user needs, offering more relevant and contextually appropriate content or responses [41]. This ability to adjust to the unique characteristics of individual users, such as their interests, behaviors, and prior interactions, is what sets personalized systems apart from general-purpose models.

Retrieval-augmented generation (RAG) [23, 54] has emerged as an effective approach for enhancing various aspects of machine learning tasks, particularly in text generation. RAG [16, 23] integrates information retrieval with natural language generation, thereby improving the relevance and factual accuracy of the generated outputs. Unlike traditional large language models (LLMs), which rely solely on the static knowledge obtained during pre-training, RAG systems utilize a retriever to access external information at inference time, enabling contextually grounded and factually consistent generation [2, 45]. This capability allows RAG to produce more informed and relevant outputs by leveraging up-to-date or personal data. Currently, the most reliable approach to personalizing LLMs is through RAG personalization, where personalized information is retrieved from the user’s profile to construct a personalized input prompt for the LLM [40, 41]. This personalized prompt provides the LLM with relevant user context, enabling it to generate more tailored and contextually appropriate responses. This approach leverages the personalization capabilities of retrieval models to enhance the LLM’s output relevance.

Another potential approach to personalizing LLMs involves fine-tuning the LLM on user-specific data, such as documents authored by the user or other relevant information about the user. However, training a separate LLM for each user is computationally expensive and requires substantial storage resources, preventing scalability of this approach. To solve this, parameter-efficient fine-tuning (PEFT) techniques, such as low-rank adaptation (LoRA) [13], can be used to adjust a subset of the model’s parameters based on each user’s

data. LoRA has demonstrated significant reductions in memory consumption and storage requirements while increasing performance for tasks such as language modeling, reasoning, text generation, and question answering [36, 56, 57]. Thus, using LoRA to train the LLM on individual user data personalizes the LLM while avoiding the need to store a full set of parameters for each user, which would be computationally expensive in large-scale applications. Given the success story of PEFT in the mentioned applications, this paper presents a study of a established PEFT method for LLM personalization, by selectively tuning the parameters with user-specific data. This way PEFT can potentially align the model's knowledge and behavior with the user's preferences and requirements. Furthermore, RAG- and PEFT-based approaches can be integrated to enhance personalization, leveraging the strengths of both RAG and fine-tuning techniques for more effective user adaptation.

A key concern in personalization of LLMs is that using private user data to train LLMs may compromise user privacy. This risk is heightened when a shared model is used across multiple users and learns from all of their data. To address this, we focus on methods that use and train LLMs exclusively on user-specific data, ensuring that no data is shared between models for different users. Under this framework, both methods for personalizing LLMs—RAG and PEFT—maintain the user's privacy, as they neither update model nor generate input prompts using data from other users. Even though the LLM leverages private personal information to generate more tailored responses for the user, this data is only accessible to the user themselves. Since the information is not shared with other users or exposed to the model in a way that would affect other users' data, this approach does not constitute a privacy violation.

To systematically study these two schools of thoughts for personalizing LLMs, this paper studies a PEFT method for LLM personalization and compares it against more established RAG-based solutions in a privacy-preserving setting. Specifically, we focus on utilizing PEFT methods for personalizing LLMs, as this is an underexplored avenue in this area. To achieve this, we conduct an extensive set of experiments using seven diverse datasets from the Language Model Personalization (LaMP) benchmark [41]. LaMP comprises three classification tasks and four text generation tasks, each designed to evaluate different aspects of LLM personalization. In this benchmark, each input is treated as a separate user, with its own distinct input, expected output, and user profile. This structure makes LaMP an ideal testbed for evaluating the effectiveness of the personalization methods explored in this study. We address the following research questions to compare these two methods:

- **RQ1: How do PEFT- and RAG-based approaches perform for LLM personalization?** Our experiments reveal that personalizing LLMs using retrieval-augmented generation results in an average improvement of 14.92% over the non-personalized baseline, while parameter-efficient fine-tuning based personalization leads to only a modest 1.07% improvement.
- **RQ2: How does the combination of PEFT and RAG impact the personalization performance?** Our results show that combining both RAG and PEFT yields the best overall performance on personalized tasks, achieving a 15.98% improvement over the non-personalized baseline. This suggests that integrating both approaches is the most effective strategy for personalizing LLMs.

- **RQ3: How does profile size affect performance?** To deepen our understanding, we provide an analysis to explain why PEFT does not perform as well for personalizing LLMs. Our findings indicate a positive correlation between the size of the user profile and performance improvement, suggesting that the limited data per user is a key factor contributing to PEFT's underperformance.

- **RQ4: How does data presence in training corpus affect performance?** We observe that improvements from PEFT are most significant when the dataset used for training is not publicly accessible and is unlikely to be used in the pre-training data. Notably, PEFT achieves the highest gains on the Personalized Email Subject Generation task, which is based on a private dataset.

Finally, we discuss the limitations of each approach to provide a clearer comparison. For instance, while PEFT is more resource-intensive than RAG due to the need for fine-tuning and storing user-specific parameters, it is generally faster at inference time since it does not require retrieving external information. In contrast, RAG is more efficient in terms of storage and adaptation to new information but incurs additional latency during retrieval. By examining these trade-offs, we provide a comprehensive evaluation of each method's advantages and limitations. To promote further exploration of this important area, we will open-source our implementation and integrate it with the LaMP codebase, ensuring that researchers can build upon and validate our findings.¹

2 Related Work

2.1 Retrieval-Augmented Generation

RAG [16, 23] is a framework that enhances natural language generation by integrating information retrieval with the generation process, improving the relevance and factual accuracy of the generated content. Unlike traditional large language models (LLMs), which rely solely on the knowledge acquired during pre-training, RAG systems retrieve external information via a retriever to incorporate contextually grounded and factually consistent knowledge during inference [2, 45]. This integration allows RAG to generate outputs that are more informed and relevant by leveraging real-time external data. The versatility of RAG spans several domains, including knowledge-grounded text generation [16, 23, 37], multi-modal reasoning [4, 39, 42], personalized generation [40, 41], and mitigating hallucinations in model outputs [44]. This adaptability makes RAG a powerful tool for applications that require up-to-date or user-specific information. In this paper, we employ RAG as one of the methods for personalizing LLMs, following the approach presented by Salemi et al. [41], which is one of the leading methods in the field of personalized LLMs.

2.2 Parameter Efficient Fine-Tuning

PEFT enables the adaptation of large language models (LLMs) to specific tasks without the need for full model retraining, significantly reducing computational costs while maintaining strong performance [11, 12, 29]. Among various PEFT methods, Low-Rank Adaptation (LoRA) is particularly effective for efficiently fine-tuning

¹The code and data used in our experiments in this paper are available at: <https://github.com/LaMP-Benchmark/LaMP>

LLMs. It introduces low-rank decomposition into weight matrices and injects trainable low-rank matrices into otherwise frozen model weights [13, 27, 28]. This technique minimizes the number of trainable parameters while preserving the model’s expressiveness, making it a scalable and computationally efficient alternative to traditional full fine-tuning. Due to its ability to significantly reduce the cost of model training, PEFT has gained attention for personalizing LLMs, allowing models to adapt to diverse user needs with relatively low resource consumption. However, prior work has typically maintained a shared pool of adapters across users, rather than training a separate adapter for each individual user [47]. This shared-adapter approach, while efficient, introduces potential risks of data leakage and privacy concerns, as information from different users may inadvertently influence the model’s behavior. This paper investigates the use of LoRA in a privacy-preserving manner by training adapters exclusively on each user’s data. Additionally, we explore the effectiveness of retrieval-augmented generation for LLM personalization. By comparing the two methods, this study examines their individual contributions to personalizing LLMs and explores potential synergies between PEFT and RAG, offering insights into optimizing LLM personalization through both parameter-efficient fine-tuning and retrieval-based augmentation while ensuring privacy protection.

2.3 Personalizing LLMs

Personalizing LLMs is a critical research area with applications in search, recommendation, and text generation [8, 34, 41, 53]. Salemi et al. [41] introduced a retrieval-augmented generation based approach for personalizing LLMs and proposed the LaMP benchmark as a framework for evaluating personalized text generation. Another line of research has focused on developing personalized writing assistants [25, 31, 33] and autonomous agents [55]. Various approaches for LLM personalization have been proposed, including training retrieval models based on user feedback for text generation [40], optimizing LLMs with personalized feedback [17], and generating personalized prompts automatically [24]. In parallel, recent studies have explored parameter-efficient fine-tuning [47], where a shared pool of adapters is trained on data from multiple users to personalize LLMs. However, this approach raises significant privacy concerns, as it may inadvertently lead to data leakage between users, compromising sensitive information. To address these concerns, this paper compares two prominent methods for personalizing LLMs—RAG and PEFT—in a privacy-preserving manner. Specifically, we explore training the LLM separately on each user’s data to ensure privacy, while maintaining the model’s ability to personalize responses effectively. This study seeks to enhance LLM personalization techniques while safeguarding user privacy by limiting model updates to user-specific data.

3 Problem Formulation

A language model M takes a prompt x as input and generates an output \hat{y} . However, the generated output is typically general, relying on the knowledge embedded in the model. To generate a personalized response, it is possible to incorporate a set of personalized information P_i , also known as the user profile, which provides tailored details about the user to the LLM. This allows the model to

produce responses that are more specific to the user’s context and preferences. This paper focuses on personalized text generation, aiming to produce outputs that are tailored to the preferences of a user. We assume access to a dataset $T = \{(x_i, y_i, P_i)\}_{i=1}^{|T|}$, where x_i is the input prompt from user u_i , y_i is the expected output for user u_i , and P_i is the user profile. Here, a user profile P_i consists of a set of structured or unstructured text documents for the user u_i , denoted as $P_i = \{d_{(i,j)}\}_{j=1}^{|P_i|}$. Note that $d_{(i,j)}$ can be either structured or unstructured documents, including input/output examples for the user or textual documents written by the user.

In this paper, we aim to leverage the user-specific information—provided by the user or collected through interactions with the user—available in the profile P_i to construct a personalized LLM $M_i = \text{PERSONALIZE}(M, P_i)$, by applying a transformation function PERSONALIZE to the original LLM M . This function can either modify the parameters of M to create M_i or alter the input to the LLM based on the profile P_i . We focus on comparing different methods for designing the transformation PERSONALIZE while maintaining privacy. Specifically, privacy is preserved by ensuring that no information from other users is used to personalize the LLM for a given user. A privacy violation occurs if the LLM reveals confidential information from one user to another. In this paper, we apply methods that work solely on each user’s data, ensuring that the information is only revealed to the corresponding user, thus preventing any information leakage or privacy violations.

4 Methods for Personalizing LLMs

There are two main approaches to personalizing LLMs. The first approach involves personalizing the input prompt provided to the model, which encourages the model to generate responses that are more tailored to the specific user. The second approach focuses on changing the parameters of the LLM through training the LLM using user-specific data to help the model learn the user’s preferences, writing style, and background knowledge, enabling it to generate responses that are better aligned with the individual user’s needs and expectations. A major concern in personalizing LLMs is the potential for privacy violations. If the language model utilizes personal information from a user, there is a risk that sensitive data could be inadvertently revealed, compromising the user’s privacy. This concern becomes particularly significant when information from one user is used to train a shared model, which may be accessed by multiple users. In such cases, there is a possibility that private information from one user could be leaked to others, violating privacy and undermining trust in the system.

One straightforward approach to mitigate privacy risks in personalizing LLMs is to ensure that a personalized LLM for a specific user is exclusively used by that user. In this setup, the model can only access that user’s information during both the training and inference phases. By restricting access to the user’s personal data, this approach effectively eliminates the possibility of revealing sensitive information to other users, thus maintaining privacy and minimizing the risk of data leakage. With this in mind, this paper focuses on studying and investigating how each of the two introduced approaches contributes to personalizing LLMs in a fully privacy-preserving context. Specifically, we examine scenarios where the data of one user cannot be used to optimize the LLM for other users,

Table 1: Prompts template used to augment the input of the LM with the user profile. We follow previous work by Salemi et al. [41] to implement these functions. `concat` is a function that concatenates the strings in its first argument by placing the string in the second argument between them. `add_to_paper_title` is a function designed to add the string in its first argument to the paper’s title in the Personalized Citation Identification task. `PPEP` is a function that create the prompt for each entry in the retrieved profile entries. `[INPUT]` is the task’s input.

Task	Per Profile Entry Prompt (PPEP)	Aggregated Input Prompt(AIP)
1: Citation Identification	" P_i [title]"	<code>add_to_paper_title(concat([PPEP(P_1), ..., PPEP(P_n)], ", and "), [INPUT])</code>
2: Movie Tagging	the tag for the movie: " P_i [description]" is " P_i [tag]"	<code>concat([PPEP(P_1), ..., PPEP(P_n)], ", and "). [INPUT]</code>
3: Product Rating	P_i [score] is the score for " P_i [text]"	<code>concat([PPEP(P_1), ..., PPEP(P_n)], ", and "). [INPUT]</code>
4: News Headline Generation	" P_i [title]" is the title for " P_i [text]"	<code>concat([PPEP(P_1), ..., PPEP(P_n)], ", and "). [INPUT]</code>
5: Scholarly Title Generation	" P_i [title]" is the title for " P_i [abstract]"	<code>concat([PPEP(P_1), ..., PPEP(P_n)], ", and "). Following the given patterns [INPUT]</code>
6: Email Subject Generation	" P_i [title]" is the title for " P_i [text]"	<code>concat([PPEP(P_1), ..., PPEP(P_n)], ", and "). [INPUT]</code>
7: Tweet Paraphrasing	" P_i [text]"	<code>concat([PPEP(P_1), ..., PPEP(P_n)], ", and ") are written by a person. Following the given patterns [INPUT]</code>

ensuring that each user’s personal information remains isolated and protected throughout the personalization process. The following sections provide a detailed explanation of these approaches.

4.1 RAG for Personalizing LLMs

In this approach, we employ a retrieval model to retrieve a set of personalized information from the user profile and use this information to construct a personalized prompt. The personalized prompt, which incorporates the user’s data, is then fed to the LLM to generate a tailored response. This method ensures that the output is specific to the individual user’s context and preferences.

Formally, given an input prompt x from the user, we use the query generation function ϕ_q to create a query. This query is then passed through the retriever R , which retrieves k documents from the user’s profile P_u . Finally, the prompt generation function ϕ_p combines the retrieved documents and the input prompt to generate a personalized prompt, which is used as the input to the LLM M to generate a more tailored response, formally, defined as:

$$\hat{y} = M(\phi_p(x, R(\phi_q(x), k))) \quad (1)$$

where to implement the query generation function ϕ_q , following Salemi et al. [41], we extract and use the non-template portions of the user’s input prompt as the query. For further details on the template used for generating inputs in the LaMP benchmark, we refer the reader to Salemi et al. [41]. Furthermore, we use the same function (ϕ_p) as Salemi et al. [41] to generate personalized prompts for the LLM, as detailed in Table 1.

Note that this approach does not modify the LLM itself. Instead, it adjusts its input, using a tailored prompt to the user based on the retrieved documents from the user profile. This allows us to personalize the LLM’s response without altering its underlying structure and parameters, which works on any black-box LLM. In our experiments, we used a wide range of retrieval models: BM25 [38] as a lexical-matching retrieval model, Contriever [15] as a semantic matching retrieval model, Recency [41] as a time-aware retrieval model, and RSPG [40] as an ensemble model that chooses an appropriate retrieval model per input. This model ensures that each

input is directed to the most suitable retrieval model for retrieving relevant documents that align with the given input. By selecting the optimal retrieval model, the system improves the relevance and accuracy of the retrieved information, thereby enhancing the overall response generation.

4.2 PEFT for Personalizing LLMs

Maintaining a separate LLM for each user is impractical for systems with large user bases. For instance, storing a model like LLaMA-3.1 [6] requires 45 GB per user. With 1 million users, this amounts to 45,000 TB of storage, which is nearly infeasible for most systems to support. In addition to storage challenges, serving the full set of LLM parameters for each user requires substantial computational resources that cannot be efficiently shared across users. In contrast, a LoRA adapter with $r = 8$ for the same model only requires 55 MB of storage. For 1 million users, this would total 55 TB, which is far more manageable for real-world applications. Moreover, by using the same LLM backbone for all users and loading individual LoRA adapters per user, the system can operate more efficiently from a computational standpoint. Therefore, using Parameter-Efficient Fine-Tuning provides a more cost-effective solution compared to training an entire LLM for each user.

This approach uses a user profile P_u to learn user-specific parameters, resulting in a personalized LLM M_u . There are various ways to achieve this personalization, and in our method, we apply LoRA (Low-Rank Adaptation) to the LLM M and train the model using the documents from P_u . LoRA fine-tunes LLMs by injecting trainable low-rank matrices into the model’s weight matrices. Instead of updating all the model weights during training, LoRA decomposes the weight update into two smaller, low-rank matrices $A \in \mathbb{R}^{d \times r}$ and $B \in \mathbb{R}^{r \times k}$, where d is the dimension of the input, k is the dimension of the output, and r is the rank parameter that controls the capacity of the low-rank approximation. The original weight matrix $W_0 \in \mathbb{R}^{d \times k}$ is kept frozen, and only the matrices A and B are trained. The low-rank matrices A and B are optimized to approximate the update to the original weights, such that the updated weights are $W = W_0 + AB$. This allows LoRA to efficiently

Table 2: Implementation of the input-output generation function CONVERT for PEFT personalization. The profiles in LaMP-2, LaMP-3, LaMP-4, LaMP-5, and LaMP-7 consist of input-output pairs, which are directly used as training pairs. However, for the LaMP-1 and LaMP-7, such pairs do not exist. For LaMP-1, we provide the model with a title and ask it to generate the abstract. For LaMP-7, we randomly divide a tweet into two parts and ask the model to generate the second part based on the prefix.

Dataset	Profile Format	Generated Input (x_i)	Generated Output (y_i)
LaMP-1: Personalized Citation Identification	title: [title] abstract: [abstract]	Write an abstract for this title: [title]	[abstract]
LaMP-2: Personalized Movie Tagging	description: [description] tag: [tag]	Which tag does this movie relate to among the following tags? Just answer with the tag name without further explanation. tags: [sci-fi, based on a book, comedy, action, twist ending, dystopia, dark comedy, classic, psychology, fantasy, romance, thought-provoking, social commentary, violence, true story] description: [description]	[tag]
LaMP-3: Personalized Product Rating	review: [review] score: [score]	What is the score of the following review on a scale of 1 to 5? just answer with 1, 2, 3, 4, or 5 without further explanation. review: [review]	[score]
LaMP-4: Personalized News Headline Generation	article: [article] title: [title]	Generate a headline for the following article: [article]	[title]
LaMP-5: Personalized Scholarly Title Generation	abstract: [abstract] title: [title]	Generate a title for the following abstract of a paper: [abstract]	[title]
LaMP-6: Personalized Email Subject Generation	email: [email] title: [title]	Generate a subject for the following email: [email]	[title]
LaMP-7: Personalized Tweet Paraphrasing	tweet: [tweet]	Complete the following tweet: [first part of the tweet]	[second part of the tweet]

adapt the model by learning a compact set of parameters, without requiring the full weight matrix to be updated, thereby reducing computational and memory costs. The parameter r , also known as the rank, is a crucial aspect of LoRA’s capacity. A larger r allows for more complex adaptations, while a smaller r ensures a more compact and efficient adaptation, but with potentially less expressiveness. In practice, r is chosen based on a trade-off between the computational efficiency (smaller r) and the model’s capacity to learn detailed user-specific adjustments (larger r). This low-rank adaptation approach not only ensures that the model can be personalized effectively for each user but also makes the process more computationally efficient, as it avoids the need to retrain or store a full set of model parameters for each user. Instead, only the smaller adaptation matrices A and B are learned and stored, making it feasible to personalize the model for a large number of users without excessive resource consumption.

To train the LLM using LoRA on a user profile P_u , each document $d \in P_u$ is first transformed into a pair of input-output sequences, where the LLM receives the input and generates the corresponding output. We define this transformation process using the function $(x_i, y_i) = \text{CONVERT}(d_i)$, where x_i represents the input sequence derived from d_i and y_i is the target output. For each document d_i in P_u , we apply the conversion function to obtain the input-output pair (x_i, y_i) . These input-output pairs form a training dataset for each user, enabling the LLM to be trained specifically for that user’s profile. By training the model on these personalized pairs, the LLM learns to generate tailored outputs that align with the user-specific information encoded in their profile. The LLM is trained on these pairs by minimizing the sequence-to-sequence cross-entropy loss [46]. This loss function encourages the model to generate the target output y_i given the input x_i .

There are various ways to implement the CONVERT function depending on the structure of the user profile. If the user profile consists of input-output pairs, such as previous inputs from the user and the corresponding preferred outputs (e.g., rated by thumbs up or directly written by the user), these pairs can be directly used to train the model. In this case, the function $\text{CONVERT}(d_i)$ would simply map each document d_i to the corresponding input-output pair (x_i, y_i) , which can then be used for training the LLM. Alternatively, when the user profile does not consist of explicit input-output pairs, such as when the profile contains previous comments or documents written by the user, these pairs can be automatically generated from each document in the profile. In such cases, text completion can be used as a task to teach the LLM the token distribution preferred by the user. The CONVERT function can then generate input-output pairs by selecting parts of the user’s previous text as the input and the next segment of text as the target output. This allows the model to learn patterns and preferences based on the user’s writing style, tone, and content. Whenever the user profile contains explicit input-output pairs, we use those directly for training. If the profile does not contain such pairs, we define text completion as the training task. Specifically, we randomly select 10-20% of the text as the input and use the remaining portion as the output. The conversion function for these tasks is summarized in Table 2.

4.3 PEFT-RAG for Personalizing LLMs

This approach integrates PEFT and RAG to enhance LLM personalization by leveraging their complementary strengths. PEFT facilitates efficient adaptation by fine-tuning a small number of additional parameters, reducing computational and memory overhead while allowing the model to capture user-specific preferences. However, PEFT alone has limitations, as it relies on static updates and

Table 3: Statistics of the datasets within the LaMP benchmark [41] with time-based data separation configuration used in our experiments in this paper. We use the time-based setting to study the effect of recency on the retrieval model’s performance. Notably, this is the only setting where users are shared between the training and test sets, allowing us to evaluate PEFT methods.

Task	#train	#dev	#test	Input Length	Output Length	#Profile Size	#classes
LaMP-1: Personalized Citation Identification	6542	1500	1500	51.43 ± 5.70	-	84.15 ± 47.54	2
LaMP-2: Personalized Movie Tagging	5073	1410	1557	92.39 ± 21.95	-	86.76 ± 189.52	15
LaMP-3: Personalized Product Rating	20000	2500	2500	128.18 ± 146.25	-	185.40 ± 129.30	5
LaMP-4: Personalized News Headline Generation	12500	1500	1800	29.97 ± 12.09	10.07 ± 3.10	204.59 ± 250.75	-
LaMP-5: Personalized Scholarly Title Generation	14682	1500	1500	162.34 ± 65.63	9.71 ± 3.21	87.88 ± 53.63	-
LaMP-6: Personalized Email Subject Generation	4821	1250	1250	454.87 ± 889.41	7.37 ± 2.78	55.67 ± 36.32	-
LaMP-7: Personalized Tweet Paraphrasing	13437	1498	1500	29.72 ± 7.01	16.96 ± 5.67	15.71 ± 14.86	-

may struggle to learn less common user preferences effectively [32]. RAG mitigates this issue by retrieving user-specific information at inference time, ensuring that personalization remains dynamic and contextually relevant without requiring extensive fine-tuning. By combining both approaches, the model benefits from long-term adaptation through fine-tuned parameters and real-time personalization through retrieval.

To do this, first, we train the LLM M on a user profile P_u using the method described in Section 4.2, which results in the personalized LLM M_u . Next, we apply the RAG personalization approach outlined in Section 4.1, denoted as:

$$\hat{y} = M_u(\phi_p(x, R(\phi_q(x), k))) \quad (2)$$

where M_u is a personalized model with PEFT as described in Section 4.2, R is a retriever, ϕ_q and ϕ_p are the query and prompt generation functions as shown in Table 1, and k is the number of retrieved documents. The key difference from Equation 1 is that here we first train the LLM on the user-specific profile to get a personalized train LLM M_u and learn the user preferences, the utilizing RAG as explained in Section 4.1 to further personalize the LLM.

5 Experiments

This section details our findings about comparing RAG and PEFT for privacy-preserving personalization of LLMs.

5.1 Experimental Setup

Datasets & Tasks. Our experiments utilize the LaMP benchmark [41], which is specifically designed to evaluate LLM personalization. Each instance in this benchmark represents a user and includes an input prompt, an expected output, and a user profile comprising structured or unstructured documents. LaMP covers seven personalized tasks, including three text classification tasks and four text generation tasks. Table 3 provides detailed dataset statistics. We adopt the time-based configuration of LaMP, as it ensures that the same users appear in both the training and test sets, enabling effective model training on user profiles for the PEFT approach. For evaluation, following previous work [40, 41], we use accuracy for binary classification (LaMP-1), accuracy and F1 score for categorical

classification (LaMP-2), and MAE and RMSE for ordinal classification (LaMP-3). For text generation tasks (LaMP-4 to LaMP-7), we assess performance using ROUGE-1 and ROUGE-L [30].

RAG Pipeline Configuration. For personalizing LLMs using retrieval-augmented generation, we adopt the experimental setup from Salemi et al. [41] and Salemi et al. [40]. Specifically, we utilize the BM25 [38] retrieval model, implemented in the *rank_bm25* library², along with Contriever³ [15], Recency-based retrieval, and RSPG [40]. In all experiments, we retrieve $k = 4$ documents from the user profile to personalize the LLM. Following Salemi et al. [41], we use FlanT5-XXL⁴ [6], a model with 11 billion parameters, as our base LLM. The model is configured with an input length of 512 tokens and an output length of 128 tokens. For text generation, we employ beam search [9] with a beam size of 4. All experiments are conducted using the Hugging Face library⁵ [52].

PEFT Pipeline Configuration. To train the LLMs for each user, we utilize the PEFT library⁶. Each model is trained for 50 epochs on the corresponding user profile with a learning rate of 5×10^{-4} , applying a linear scheduler with a warm-up phase covering 5% of the total training steps. We use the Adam optimizer [19] with a weight decay of 10^{-4} , and a batch size of 16 is achieved through gradient accumulation. LoRA is employed with a dropout rate of 0.1 and a scaling factor $\alpha = 32$, applied to all key, query, and value projections in the transformer [49]. Following Salemi et al. [41], we use FlanT5-XXL⁷ [6], a model with 11 billion parameters, as our base LLM. The model is configured with an input length of 512 tokens and an output length of 128 tokens. For text generation, we employ beam search [9] with a beam size of 4. The experiments are conducted on up to 32 Nvidia A100 GPUs (80GB VRAM) with 128GB RAM over a period of up to 7 days. In total, our experiments consumed over 10,000 GPU hours. To mitigate computational costs, we train an LLM only for users present in the test set rather than for all users in the benchmark. As a result, 37,560 LoRA adapters were trained, occupying approximately 18 TB of disk space.

²Available at https://github.com/dorianbrown/rank_bm25

³Available at <https://huggingface.co/facebook/contriever>

⁴Available at <https://huggingface.co/google/flan-t5-xxl>

⁵Available at <https://huggingface.co/>

⁶Available at <https://huggingface.co/docs/peft/en/index>

⁷Available at <https://huggingface.co/google/flan-t5-xxl>

Table 4: Performance of different methods (RAG-based and PEFT-based) used for LLM personalization on the datasets in the LaMP benchmark. The results indicate that RAG methods are more effective than PEFT for personalizing LLMs. Moreover, combining RAG and PEFT achieves the best performance in LLM personalization.

Dataset	Metric	No Personalization	PEFT Personalization				RAG Personalization				PEFT-RAG Personalization			
			$r = 8$	$r = 16$	$r = 32$	$r = 64$	BM25	Recency	Contriever	RSPG	$r = 8$	$r = 16$	$r = 32$	$r = 64$
LaMP-1: Personalized Citation Identification	Accuracy \uparrow	0.502	0.502	0.502	0.504	0.506	0.626	0.622	0.636	0.672	0.670	0.668	0.671	0.671
LaMP-2: Personalized Movie Tagging	Accuracy \uparrow	0.359	0.360	0.360	0.360	0.359	0.387	0.377	0.396	0.430	0.430	0.431	0.430	0.430
	F1 \uparrow	0.276	0.278	0.278	0.278	0.277	0.306	0.295	0.304	0.339	0.341	0.342	0.341	0.341
LaMP-3: Personalized Product Rating	MAE \downarrow	0.308	0.308	0.307	0.306	0.301	0.298	0.296	0.299	0.264	0.264	0.265	0.264	0.259
	RMSE \downarrow	0.611	0.607	0.607	0.602	0.600	0.611	0.605	0.616	0.568	0.568	0.570	0.564	0.562
LaMP-4: Personalized News Headline Generation	ROUGE-1 \uparrow	0.176	0.178	0.177	0.178	0.178	0.186	0.189	0.183	0.203	0.203	0.204	0.204	0.203
	ROUGE-L \uparrow	0.160	0.162	0.162	0.163	0.163	0.171	0.173	0.169	0.186	0.186	0.186	0.187	0.186
LaMP-5: Personalized Scholarly Title Generation	ROUGE-1 \uparrow	0.478	0.478	0.478	0.477	0.478	0.477	0.475	0.483	0.480	0.481	0.480	0.480	0.479
	ROUGE-L \uparrow	0.428	0.429	0.429	0.428	0.428	0.427	0.426	0.433	0.429	0.431	0.431	0.431	0.431
LaMP-6: Personalized Email Subject Generation	ROUGE-1 \uparrow	0.335	0.342	0.342	0.341	0.343	0.412	0.403	0.401	0.433	0.436	0.436	0.436	0.437
	ROUGE-L \uparrow	0.319	0.325	0.326	0.325	0.326	0.398	0.389	0.386	0.418	0.422	0.422	0.422	0.421
LaMP-7: Personalized Tweet Paraphrasing	ROUGE-1 \uparrow	0.449	0.449	0.449	0.449	0.449	0.446	0.444	0.440	0.461	0.460	0.460	0.460	0.460
	ROUGE-L \uparrow	0.396	0.397	0.397	0.396	0.396	0.394	0.393	0.390	0.409	0.409	0.409	0.408	0.409

5.2 Main Findings

RQ1: How do PEFT- and RAG-based approaches perform for LLM personalization? The results of both PEFT-based and RAG-based personalization methods, along with the non-personalized baseline, are reported in Table 4. Our findings about this question in Table 4 indicate that incorporating PEFT improves performance over non-personalized LLMs in 5 out of the 7 evaluated datasets, highlighting its effectiveness in leveraging parameter-efficient tuning for personalization. Similarly, the RAG approach demonstrates consistent performance improvements across all datasets, suggesting that retrieval-based personalization is a robust method for enhancing LLM outputs. When comparing PEFT with RAG, our analysis in Table 4 reveals that RAG is a more effective personalization strategy. As shown in Table 4, PEFT leads to a modest average performance improvement of 1.07% over non-personalized LLMs, whereas the RAG-based approach yields a significantly higher average improvement of 14.92%. This substantial performance gap underscores the advantage of retrieval-augmented generation in adapting LLMs to personalized contexts.

It is worth noting that different retrieval models in Table 4 exhibit varying levels of effectiveness in improving performance. Among them, RSPG [40], which dynamically selects the most suitable retrieval model for each instance, consistently outperforms all other retrieval-based methods, achieving the highest overall performance. This highlights the importance of adaptive retrieval mechanisms in optimizing personalization. Additionally, we observe that the rank parameter r in PEFT also influences model performance, though its impact varies across datasets. Specifically, on LaMP-1, LaMP-3, and LaMP-6, increasing r leads to noticeable performance improvements, suggesting that a higher rank enables better adaptation to personalized contexts in these tasks. However, for other datasets, the effect of adjusting r is less pronounced, indicating that the benefits of increasing parameter efficiency may be task-dependent.

RQ2: How does the combination of PEFT and RAG impact the personalization performance? To answer this research question, we integrate the best-performing retrieval model from Table 4 with each user’s personalized LLM, trained using PEFT, to perform RAG personalization with PEFT. This approach aims to leverage both retrieval-augmented generation and parameter-efficient fine-tuning to maximize personalization effectiveness. The results of this experiment are presented in Table 4. Our findings indicate that combining RAG with PEFT yields performance improvements over the standard RAG approach in 4 out of the 7 evaluated tasks. This suggests that fine-tuning a personalized LLM alongside retrieval-based augmentation allows the model to better adapt to user-specific preferences and contextual nuances. Moreover, this combined approach results in an overall 15.98% improvement over the non-personalized LLM, surpassing standard RAG personalization by an additional 0.44% in relative performance gain. Although the improvement over RAG alone is modest, it demonstrates that incorporating PEFT can provide additional benefits, particularly in scenarios where retrieval alone may not fully capture personalization needs. These results highlight the potential of combining PEFT with RAG as a viable strategy for further enhancing LLM personalization, offering a balanced approach that benefits from both retrieval-augmented generation and parameter-efficient adaptation.

RQ3: How does profile size affect performance? We create a regression plot⁸ that visualizes the relationship between the number of documents in a user’s profile and the relative improvement achieved by the best-performing personalized LLM using PEFT- and RAG-based personalization compared to the non-personalized LLM. Here, we define improvement as 1 when there is a performance gain over the non-personalized baseline and -1 when there is no gain. To ensure that our analysis focuses on meaningful variations, we exclude users who experience no change in performance. The resulting plot is shown in Figure 1. This figure reveals that in 5 out

⁸<https://seaborn.pydata.org/generated/seaborn.regplot.html>




Figure 1: Correlation (with 95% confidence interval) between profile item count and performance improvement of PEFT- and RAG-based personalization in comparison with no personalization on the tasks in the LaMP benchmark. The results indicate a positive correlation between the improvement of PEFT over non-personalized LLMs and the profile size, suggesting that the limited amount of personalized data per user contributes to PEFT’s underperformance. Conversely, RAG exhibits a negative correlation with profile size, indicating that as the profile grows, the retrieval model faces challenges in retrieving relevant information for effective LLM personalization.

of the 7 datasets, there is a positive correlation between the number of items in a user’s profile and the performance improvement achieved through PEFT-based personalization. This suggests that as more user-specific data becomes available for fine-tuning, the model is better able to adapt to individual preferences and generate more personalized responses. However, in the case of the LaMP-5 task, we observe either a negative or zero correlation. One possible explanation for this anomaly is the nature of the user profiles, which consist of abstracts from papers authored by the user. Since many of these papers are collaborative efforts, users with larger profiles tend to be senior researchers who may not have been directly involved in the writing process. Our analysis further supports this hypothesis, as we found that in 94% (17 out of 18) of the cases where performance dropped, the user was not the primary author on most papers in their profile. This suggests that training on such collaborative data is less effective for personalizing the LLM to the user’s actual writing style and preferences.

When aggregating results across all tasks, we observe an overall positive correlation between PEFT-based performance improvement and the number of items in a user’s profile. In contrast, we find a negative correlation between performance improvement using RAG-based personalization and the non-personalized baseline. This indicates that as the size of a user’s profile increases, retrieval models struggle to accurately identify and retrieve the most relevant documents for personalization, potentially leading to diminished gains. These findings provide insight into the strengths and limitations of PEFT- and RAG-based personalization. While PEFT benefits from having more user-specific training data, its effectiveness is

constrained when the available data is insufficient for robust adaptation. On the other hand, RAG’s performance appears to degrade as the profile size grows, likely due to challenges in retrieval quality. This suggests that one of the primary reasons PEFT does not consistently outperform RAG for personalizing LLMs is the limited amount of training data per user, which restricts the model’s ability to effectively learn individual preferences.

RQ4: How does data presence in training corpus affect performance? Another notable observation from Table 4 is that PEFT achieves the highest performance gains on the LaMP-6 task compared to other evaluated tasks. We hypothesize that this is due to the nature of the dataset and the pretraining corpus of FlanT5 [6]. Since FlanT5 is trained on publicly available datasets, it has not been exposed to private datasets such as LaMP-6, which contains personally identifiable information (PII) from the Avocado [35] corpus. As a result, the model encounters this corpus for the first time during fine-tuning, allowing it to learn new, task-specific patterns more effectively. In contrast, other datasets used in our experiments are primarily derived from public sources that were likely included in FlanT5’s pretraining corpus. This means that the model has already seen similar data during pretraining, leading to relatively smaller performance improvements when fine-tuned using PEFT. These findings suggest that applying PEFT to private user data can lead to considerable performance gains, particularly in scenarios where the LLM has not previously encountered the data. This underscores the potential of PEFT for enhancing personalization in cases where user-specific data is distinct from publicly available corpora.

6 Limitations of RAG and PEFT for LLM Personalization

This section discusses the limitations of personalizing LLMs using RAG and PEFT. In addition, it outlines future research directions that could help mitigate these challenges, paving the way for more effective personalization techniques that can be deployed in real-world systems that serve millions of users.

6.1 Resource Intensity

Personalizing large language models, particularly using PEFT with LoRA, can be computationally demanding. Training models with LoRA requires substantial resources, leading to increased costs and extended training times, which can hinder scalability in resource-constrained environments. Due to these limitations, our experiments in this paper were conducted on just one single large language model, *FlanT5-XXL* [6], which has 11 billion parameters, following Salemi et al. [41]. While investigating personalization across multiple LLM architectures would provide valuable insights, the computational costs are prohibitively high.

In this work, we utilized over 10,000 hours of A100 GPU computation for training and experimentation. Based on the average figures reported by Dodge et al. [7], running these experiments on cloud-based GPU providers would result in the generation of at least 400 kilograms of CO₂ emissions. Since our experiments were conducted locally, where energy efficiency may be lower, the actual carbon footprint could be even higher. This underscores the environmental impact of large-scale LLM experiments and the challenges of scaling personalized LLMs to millions of users. Expanding this study to include multiple LLMs would significantly increase the computational burden, making it infeasible at this time.

Beyond the costs associated with training, storing personalized large language models presents another challenge. If we assume on average each adapter requires 200 MB of disk space, a platform with 100 million users would require approximately 20 PB of storage just for storing adapters. Such large storage requirements pose a significant challenge for deploying personalized LLMs at scale. Addressing these limitations—both computational and storage-related—is essential for the practical deployment of personalized generative AI systems. Future research should explore solutions such as parameter-efficient methods, adaptive storage mechanisms, or dynamic retrieval approaches to mitigate these constraints and facilitate real-world applications of personalized large language models.

6.2 Adapter Loading and Retrieval Latency

In our comparison between RAG and PEFT methods, adapter loading and retrieval latency emerge as critical factors influencing overall system performance. Each approach presents unique challenges that impact deployment efficiency and real-time usability. For RAG-based models, retrieval latency is a significant concern. Querying external databases and fetching relevant information introduces time costs that can affect system responsiveness. High retrieval latency can be particularly problematic for real-time applications, where rapid responses are essential. Additionally, managing retrieval efficiency involves challenges such as effective indexing strategies, optimizing search algorithms, and handling large-scale

corpora, all of which can exacerbate latency issues. If retrieval is slow or inaccurate, it can degrade the model's effectiveness, as the retrieved information may not be timely or contextually relevant.

Conversely, PEFT-based approaches rely on adapting pre-trained models by loading user-specific adapters. While PEFT is more resource-efficient than full fine-tuning, adapter loading introduces its own overhead, especially in large-scale deployments. Initializing multiple adapters or switching between them dynamically can impose computational burdens that affect system throughput. This challenge is particularly pronounced in scenarios requiring frequent updates or real-time interactions, where loading delays may introduce unwanted lag. Additionally, when working with millions of personalized adapters, memory management becomes critical factors for maintaining scalability. Thus, while RAG-based personalization suffers from retrieval latency, PEFT-based personalization is constrained by adapter loading overhead. Addressing these bottlenecks—whether through optimized retrieval pipelines for RAG or efficient adapter-loading mechanisms for PEFT—is essential for enabling seamless personalization in LLM-powered applications. Future work should explore solutions such as cached retrieval strategies, lightweight indexing, model compression techniques, or dynamic adapter management to improve both efficiency and scalability in personalized LLM systems.

7 Conclusion

This paper presents the first systematic comparison of retrieval-augmentation and parameter-efficient fine-tuning for personalizing LLMs in a privacy-preserving setting. Additionally, we propose a hybrid approach that combines both methods to enhance LLMs' ability in personalized text generation. Our experiments on the LaMP benchmark demonstrate that both RAG and PEFT improve LLM performance on personalized tasks. However, RAG significantly outperforms PEFT in this setting. Moreover, we find that the best results are achieved when RAG and PEFT are combined, leading to a 15.98% improvement over the performance of a non-personalized LLM on the LaMP benchmark tasks. Further analysis reveals a positive correlation between the number of documents available for each user and the performance gain from PEFT-based personalization. This finding suggests that the primary limitation of PEFT is the lack of sufficient per-user training data, which restricts its effectiveness when used alone. Finally, we show that PEFT is particularly effective when trained on private user data, as the LLM has not encountered this data during pretraining. This highlights PEFT's potential for leveraging private information in a way that standard pretraining and retrieval-augmented approaches cannot.

Acknowledgment

The authors would like to thank Mohammad Aliannejadi for his feedback on a draft of this paper. This work was supported in part by the Center for Intelligent Information Retrieval, in part by Google, in part by Lowe's, and in part by Microsoft. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect those of the sponsor.

References

- [1] Bashar Alhafni, Vivek Kulkarni, Dhruv Kumar, and Vipul Raheja. 2024. Personalized Text Generation with Fine-Grained Linguistic Control. In *Proceedings of the 1st Workshop on Personalization of Generative AI Systems (PERSONALIZE 2024)*, Ameet Deshpande, Eunjeong Hwang, Vishvak Murahari, Joon Sung Park, Diyi Yang, Ashish Sabharwal, Karthik Narasimhan, and Ashwin Kalyan (Eds.). Association for Computational Linguistics, St. Julians, Malta, 88–101. <https://aclanthology.org/2024.personalize-1.8>
- [2] Akari Asai, Zequi Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=hSyW5go0v8>
- [3] Junyi Chen. 2023. A Survey on Large Language Models for Personalized and Explainable Recommendations. arXiv:2311.12338 [cs.IR] <https://arxiv.org/abs/2311.12338>
- [4] Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William Cohen. 2022. MuRAG: Multimodal Retrieval-Augmented Generator for Open Question Answering over Images and Text. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, Abu Dhabi, United Arab Emirates, 5558–5570. doi:10.18653/v1/2022.emnlp-main.375
- [5] Eunah Cho, Ziyang Jiang, Jie Hao, Zheng Chen, Saurabh Gupta, Xing Fan, and Chenlei Guo. 2021. Personalized Search-based Query Rewrite System for Conversational AI. In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, Alexandros Papangelis, Pawel Budzianowski, Bing Liu, Elnaz Nouri, Abhinav Rastogi, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Online, 179–188. doi:10.18653/v1/2021.nlp4convai-1.17
- [6] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2024. Scaling Instruction-Finetuned Language Models. *Journal of Machine Learning Research* 25, 70 (2024), 1–53. <http://jmlr.org/papers/v25/23-0870.html>
- [7] Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A. Smith, Nicole Decario, and Will Buchanan. 2022. Measuring the Carbon Intensity of AI in Cloud Instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency* (Seoul, Republic of Korea) (FAccT '22). Association for Computing Machinery, New York, NY, USA, 1877–1894. doi:10.1145/3531146.3533234
- [8] Andrew Fowler, Kurt Partridge, Ciprian Chelba, Xiaojun Bi, Tom Ouyang, and Shumin Zhai. 2015. Effects of Language Modeling and Its Personalization on Touchscreen Typing Performance. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems* (Seoul, Republic of Korea) (CHI '15). Association for Computing Machinery, New York, NY, USA, 649–658. doi:10.1145/2702123.2702503
- [9] Markus Freitag and Yaser Al-Onaizan. 2017. Beam Search Strategies for Neural Machine Translation. In *Proceedings of the First Workshop on Neural Machine Translation*, Thang Luong, Alexandra Birch, Graham Neubig, and Andrew Finch (Eds.). Association for Computational Linguistics, Vancouver, 56–60. doi:10.18653/v1/W17-3207
- [10] Qian Guo, Wei Chen, and Huaiyu Wan. 2021. AOL4PS: A Large-scale Data Set for Personalized Search. *Data Intelligence* 3, 4 (10 2021), 548–567. arXiv:https://direct.mit.edu/dint/article-pdf/3/4/548/1968580/dint_a_00104.pdf doi:10.1162/dint_a_00104
- [11] Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. 2024. Parameter-Efficient Fine-Tuning for Large Models: A Comprehensive Survey. arXiv:2403.14608
- [12] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-Efficient Transfer Learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 97)*, Kamalika Chaudhuri and Ruslan Salakhutdinov (Eds.). PMLR, 2790–2799. <https://proceedings.mlr.press/v97/houlsby19a.html>
- [13] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=nZeVkeFYf9>
- [14] Wenyue Hua, Lei Li, Shuyuan Xu, Li Chen, and Yongfeng Zhang. 2023. Tutorial on Large Language Models for Recommendation. In *Proceedings of the 17th ACM Conference on Recommender Systems*. 1281–1283.
- [15] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised Dense Information Retrieval with Contrastive Learning. arXiv:2112.09118
- [16] Gautier Izacard and Edouard Grave. 2021. Leveraging Passage Retrieval with Generative Models for Open Domain Question Answering. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, Paola Merlo, Jorg Tiedemann, and Reut Tsarfaty (Eds.). Association for Computational Linguistics, Online, 874–880. doi:10.18653/v1/2021.eacl-main.74
- [17] 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:2310.11564
- [18] Pranav Kasela, Marco Braga, Gabriella Pasi, and Raffaele Perego. 2024. SE-PQA: Personalized Community Question Answering. In *Companion Proceedings of the ACM Web Conference 2024* (Singapore, Singapore) (WWW '24). Association for Computing Machinery, New York, NY, USA, 1095–1098. doi:10.1145/3589335.3651445
- [19] Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*. San Diego, CA, USA.
- [20] Ahmet Baki Kocaballi, Shlomo Berkovsky, Juan C Quiroz, Liliana Laranjo, Huang Ly Tong, Dana Rezaadegan, Agustina Briatore, and Enrico Coiera. 2019. The Personalization of Conversational Agents in Health Care: Systematic Review. *J Med Internet Res* 21, 11 (7 Nov 2019), e15360. doi:10.2196/15360
- [21] Ishita Kumar, Snigdha Viswanathan, Sushrita Yerra, Alireza Salemi, Ryan A. Rossi, Franck Dernoncourt, Hanieh Deilamsalehy, Xiang Chen, Ruiyi Zhang, Shubham Agarwal, Nedim Lipka, and Hamed Zamani. 2024. LongLaMP: A Benchmark for Personalized Long-form Text Generation. arXiv:2407.11016 [cs.CL] <https://arxiv.org/abs/2407.11016>
- [22] Carson K. Leung, Evan W.R. Madill, and Qi Wen. 2023. A Privacy-Preserving Semi-Decentralized Personalized Recommendation System. In *2023 IEEE International Conference on Big Data (BigData)*. 1336–1343. doi:10.1109/BigData59044.2023.10386507
- [23] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (<conf-loc>, <city>Vancouver</city>, <state>BC</state>, <country>Canada</country>, <conf-loc>) (NIPS '20). Curran Associates Inc., Red Hook, NY, USA, Article 793, 16 pages.
- [24] Cheng Li, Mingyang Zhang, Qiaozhu Mei, Weize Kong, and Michael Bendersky. 2024. Learning to Rewrite Prompts for Personalized Text Generation. In *Proceedings of the ACM on Web Conference 2024 (WWW '24)*. ACM. doi:10.1145/3589334.3645408
- [25] 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:2308.07968
- [26] Hao Li, Chenghao Yang, An Zhang, Yang Deng, Xiang Wang, and Tat-Seng Chua. 2024. Hello Again! LLM-powered Personalized Agent for Long-term Dialogue. arXiv:2406.05925 [cs.CL] <https://arxiv.org/abs/2406.05925>
- [27] Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, and Tuo Zhao. 2024. LoftQ: LoRA-Fine-Tuning-aware Quantization for Large Language Models. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=LzPWwPAdY4>
- [28] Vladislav Lialin, Sherin Muckatira, Namrata Shivagunde, and Anna Rumshisky. 2024. ReLoRA: High-Rank Training Through Low-Rank Updates. In *The Twelfth International Conference on Learning Representations*. <https://openreview.net/forum?id=DLJznSp6X3>
- [29] Baohao Liao, Yan Meng, and Christof Monz. 2023. Parameter-Efficient Fine-Tuning without Introducing New Latency. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 4242–4260. doi:10.18653/v1/2023.acl-long.233
- [30] Chin-Yew Lin. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, Barcelona, Spain, 74–81. <https://aclanthology.org/W04-1013>
- [31] Zhuoran Lu, Sheshera Mysore, Tara Safavi, Jennifer Neville, Longqi Yang, and Mengting Wan. 2024. Corporate Communication Companion (CCC): An LLM-empowered Writing Assistant for Workplace Social Media. arXiv:2405.04656
- [32] Alex Mallen, Akari Asai, Victor Zhong, Rajarshi Das, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. When Not to Trust Language Models: Investigating Effectiveness of Parametric and Non-Parametric Memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (Eds.). Association for Computational Linguistics, Toronto, Canada, 9802–9822. doi:10.18653/v1/2023.acl-long.546
- [33] 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:2311.09180

- [34] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Malleevich, Iliia Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy. 2019. Deep Learning Recommendation Model for Personalization and Recommendation Systems. arXiv:1906.00091 [cs.IR]
- [35] Oard, Douglas, Webber, William, Kirsch, David A., and Golitsynskiy, Sergey. 2015. Avocado Research Email Collection. doi:10.35111/WQT6-JG60
- [36] Rui Pan, Xiang Liu, Shizhe Diao, Renjie Pi, Jipeng Zhang, Chi Han, and Tong Zhang. 2024. LISA: Layerwise Importance Sampling for Memory-Efficient Large Language Model Fine-Tuning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=L8ifDX5XNq>
- [37] Fabio Petroni, Aleksandra Piktus, Angela Fan, Patrick Lewis, Majid Yazdani, Nicola De Cao, James Thorne, Yacine Jernite, Vladimir Karpukhin, Jean Mailard, Vassilis Plachouras, Tim Rocktäschel, and Sebastian Riedel. 2021. KILT: a Benchmark for Knowledge Intensive Language Tasks. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Kristina Toutanova, Anna Rumshisky, Luke Zettlemoyer, Dilek Hakkani-Tur, Iz Beltagy, Steven Bethard, Ryan Cotterell, Tanmoy Chakraborty, and Yichao Zhou (Eds.). Association for Computational Linguistics, Online, 2523–2544. doi:10.18653/v1/2021.naacl-main.200
- [38] Stephen E. Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike Gatford. 1994. Okapi at TREC-3. In *Text Retrieval Conference*. <https://api.semanticscholar.org/CorpusID:3946054>
- [39] Alireza Salemi, Juan Altmayer Pizzorno, and Hamed Zamani. 2023. A Symmetric Dual Encoding Dense Retrieval Framework for Knowledge-Intensive Visual Question Answering. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval* (<conf-loc>, <city>Taipei</city>, <country>Taiwan</country>, </conf-loc>) (SIGIR '23). Association for Computing Machinery, New York, NY, USA, 110–120. doi:10.1145/3539618.3591629
- [40] Alireza Salemi, Surya Kallumadi, and Hamed Zamani. 2024. Optimization Methods for Personalizing Large Language Models through Retrieval Augmentation. In *Proceedings of the 47th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval* (Washington, DC, USA) (SIGIR '24). (to appear).
- [41] Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2024. LaMP: When Large Language Models Meet Personalization. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 7370–7392. <https://aclanthology.org/2024.acl-long.399>
- [42] Alireza Salemi, Mahta Rafiee, and Hamed Zamani. 2023. Pre-Training Multimodal Dense Retrievers for Outside-Knowledge Visual Question Answering. In *Proceedings of the 2023 ACM SIGIR International Conference on Theory of Information Retrieval* (<conf-loc>, <city>Taipei</city>, <country>Taiwan</country>, </conf-loc>) (ICTIR '23). Association for Computing Machinery, New York, NY, USA, 169–176. doi:10.1145/3578337.3605137
- [43] Andriy Shepitsen, Jonathan Gemmell, Bamshad Mobasher, and Robin Burke. 2008. Personalized recommendation in social tagging systems using hierarchical clustering. In *Proceedings of the 2008 ACM Conference on Recommender Systems* (Lausanne, Switzerland) (RecSys '08). Association for Computing Machinery, New York, NY, USA, 259–266. doi:10.1145/1454008.1454048
- [44] Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. Retrieval Augmentation Reduces Hallucination in Conversation. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 3784–3803. doi:10.18653/v1/2021.findings-emnlp.320
- [45] Shamane Siriwardhana, Rivindu Weerasekera, Elliott Wen, Tharindu Kaluarachchi, Rajib Rana, and Suranga Nanayakkara. 2023. Improving the Domain Adaptation of Retrieval Augmented Generation (RAG) Models for Open Domain Question Answering. *Transactions of the Association for Computational Linguistics* 11 (2023), 1–17. doi:10.1162/tacl_a_00530
- [46] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to Sequence Learning with Neural Networks. In *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger (Eds.), Vol. 27. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2014/file/a14ac55a4f27472c5d894ec1c3c743d2-Paper.pdf
- [47] Zhaoxuan Tan, Zheyuan Liu, and Meng Jiang. 2024. Personalized Pieces: Efficient Personalized Large Language Models through Collaborative Efforts. arXiv:2406.10471
- [48] Liang Tang, Yexi Jiang, Lei Li, Chunqiu Zeng, and Tao Li. 2015. Personalized Recommendation via Parameter-Free Contextual Bandits. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Santiago, Chile) (SIGIR '15). Association for Computing Machinery, New York, NY, USA, 323–332. doi:10.1145/2766462.2767707
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- [50] Konstantinos N. Vavliakis, Maria Th. Kotouza, Andreas L. Symeonidis, and Pericles A. Mitkas. 2019. Personalization and the Conversational Web. In *Web Information Systems and Technologies*, Maria José Escalona, Francisco Dominguez Mayo, Tim A. Majchrzak, and Valérie Monfort (Eds.). Springer International Publishing, Cham, 56–77.
- [51] M.A. Walker, S.J. Whittaker, A. Stent, P. Maloor, J. Moore, M. Johnston, and G. Vasireddy. 2004. Generation and evaluation of user tailored responses in multimodal dialogue. *Cognitive Science* 28, 5 (2004), 811–840. doi:10.1016/j.cogsci.2004.06.002 2003 Rumelhart Prize Special Issue Honoring Aravind K. Joshi.
- [52] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Qun Liu and David Schlangen (Eds.). Association for Computational Linguistics, Online, 38–45. doi:10.18653/v1/2020.emnlp-demos.6
- [53] Gui-Rong Xue, Jie Han, Yong Yu, and Qiang Yang. 2009. User Language Model for Collaborative Personalized Search. *ACM Trans. Inf. Syst.* 27, 2, Article 11 (mar 2009), 28 pages. doi:10.1145/1462198.1462203
- [54] Hamed Zamani, Fernando Diaz, Mostafa Dehghani, Donald Metzler, and Michael Bendersky. 2022. Retrieval-Enhanced Machine Learning. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval* (<conf-loc>, <city>Madrid</city>, <country>Spain</country>, </conf-loc>) (SIGIR '22). Association for Computing Machinery, New York, NY, USA, 2875–2886. doi:10.1145/3477495.3531722
- [55] Kai Zhang, Yangyang Kang, Fubang Zhao, and Xiaozhong Liu. 2024. 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)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 2386–2398. <https://aclanthology.org/2024.naacl-long.132>
- [56] Longteng Zhang, Lin Zhang, Shaohuai Shi, Xiaowen Chu, and Bo Li. 2023. LoRA-FA: Memory-efficient Low-rank Adaptation for Large Language Models Fine-tuning. arXiv:2308.03303 [cs.CL] <https://arxiv.org/abs/2308.03303>
- [57] Justin Zhao, Timothy Wang, Wael Abid, Geoffrey Angus, Arnab Garg, Jeffery Kinnison, Alex Sherstinsky, Piero Molino, Travis Addair, and Devvret Rishi. 2024. LoRA Land: 310 Fine-tuned LLMs that Rival GPT-4, A Technical Report. arXiv:2405.00732 [cs.CL] <https://arxiv.org/abs/2405.00732>