

# Persona-DB: Efficient Large Language Model Personalization for Response Prediction with Collaborative Data Refinement

Chenkai Sun<sup>1</sup>, Ke Yang<sup>1</sup>, Revanth Gangi Reddy<sup>1</sup>, Yi R. Fung<sup>1</sup>  
Hou Pong Chan<sup>1</sup>, Kevin Small<sup>2</sup>, ChengXiang Zhai<sup>1</sup>, Heng Ji<sup>1</sup>

<sup>1</sup>University of Illinois Urbana-Champaign, <sup>2</sup>Amazon  
{chenkai5, czhai, hengji}@illinois.edu

## Abstract

The increasing demand for personalized interactions with large language models (LLMs) calls for methodologies capable of accurately and efficiently identifying user opinions and preferences. Retrieval augmentation emerges as an effective strategy, as it can accommodate a vast number of users without the costs from fine-tuning. Existing research, however, has largely focused on enhancing the retrieval stage and devoted limited exploration toward optimizing the representation of the database, a crucial aspect for tasks such as personalization. In this work, we examine the problem from a novel angle, focusing on how data can be better represented for more data-efficient retrieval in the context of LLM customization. To tackle this challenge, we introduce Persona-DB, a simple yet effective framework consisting of a hierarchical construction process to improve generalization across task contexts and collaborative refinement to effectively bridge knowledge gaps among users. In the evaluation of response prediction, Persona-DB demonstrates superior context efficiency in maintaining accuracy with a significantly reduced retrieval size, a critical advantage in scenarios with extensive histories or limited context windows. Our experiments also indicate a marked improvement of over 10% under cold-start scenarios, when users have extremely sparse data. Furthermore, our analysis reveals the increasing importance of collaborative knowledge as the retrieval capacity expands.

## 1 Introduction

The increasing demand for artificial intelligence-based services that can accurately predict and adapt to individual user preferences underscores the importance of personalization in today’s digital landscape. Personalization enhances user experience,

fostering trust and engagement beyond mere convenience. This need is particularly pronounced in applications involving large language models (LLMs), which serve as the foundation for a wide range of services, from personalized content moderation and medical assistance to educational platforms (Chen et al., 2023a; Baek et al., 2023).

In this work, we aim to improve LLM personalization, the model’s capability of customizing responses based on user background. For instance, understanding a user’s penchant for concise versus detailed explanations can tailor the verbosity of responses, while knowledge of their professional domain can refine the relevance of examples provided. Specifically, we explore from the perspective of retrieval-augmentation (Lewis et al., 2020; Zhang et al., 2023), which has the distinct advantage of mitigating the need for the computational burdens associated with training personalized models for a multitude of users and the magnitude of LLMs. In the retrieval-augmentation approach, relevant user-related data logs, such as historical social media posts and user descriptions, are fetched by a retriever. This information is then integrated into the instruction prompts for downstream LLMs to personalize responses.

While resource-efficient, the efficacy of the method hinges on the quality of the underlying database. Current solutions for retrieval augmentation, however, have limited exploration in the facet and often depend largely on expansive yet shallow logs of user interactions and summaries, which can prevent the model from achieving optimal performance. Firstly, such an approach would make it difficult to capture the depth of user information, a vital component in the model performance in answering queries. Secondly, the approach discourages retrieval efficiency, since it needs to gather scattered information pieces to provide sufficient pertinent data for retrieval augmentation. For instance, a higher-level opinion, such as an attitude

---

The code is available at <https://github.com/chenkaisun/Persona-DB>

toward a political party, can often be a more broadly applicable and retrieval-efficient insight compared to an opinion about a specific entity within the party, where the former requires analyzing historical behaviors. Moreover, existing methods have yet to explore the potential synergies that could arise from intelligently leveraging the interconnectedness of different users’ data to fill knowledge gaps.

Our paper delves into the critical research question: *How do we address context efficiency in retrieval-augmented personalization?* In other words, considering the cost of model inference, how do we maintain the accuracy of the generated response by retrieving less? To address previous gaps, we propose a solution based on the premise that the construction of more structured, self-improved, and interrelated databases can enhance the LLM’s ability to retrieve relevant information and to generalize. This leads us to the introduction of our framework Persona-DB, designed to encapsulate extensive user contexts and histories, fostering more accurate and individualized interactions.

The framework is underpinned by two integral components designed to enhance user databases for more generalizable inference. The first component involves a simple yet effective stage that involves distillation and induction on the user database to encapsulate and extend user personas such as opinions and preferences. The hierarchical construction process enables the application of learned insights across diverse task contexts. The second component involves interconnecting relevant databases among users, drawing inspiration from the principles of collaborative filtering in recommendation systems. Recognizing that an individual’s data may sometimes be insufficient to address the query (e.g., due to sparse data or domain-specific gaps), we facilitate the amalgamation of other users’ data that may contain relevant interactions, based on the previous studies that users with analogous mindsets tend to exhibit similar behaviors and preferences (Reimer et al., 2022; Torelli and Kaikati, 2009). For instance, a user who participates in outdoor activities and values environmental sustainability, yet has only a few interactions related to renewable energy projects, can be matched with users of a similar environmental ethos who have engaged in discussions about renewable energy. This assists us in inferring the original user’s viewpoints on solar energy initiatives, leveraging shared values

to supplement the absence of interactions in that domain. In our work, we specifically achieve this by matching users based on shared persona keys identified in the first component and composing information pieces from the individual’s database and the collaborative databases. Our approach aims to make databases both more precise and context-efficient for downstream LLM usage, reducing the need to retrieve extensive data.

We evaluate the method on *Response Forecasting for News Media*, where the downstream LLM is asked to predict users’ responses to news messages, and on *OpinionQA*, where the LLM predicts individual survey responses (Sun et al., 2023; Santurkar et al., 2023). Our experimental evaluation reveals that Persona-DB consistently outperforms the baselines across various retrieval sizes, notably achieving better results when the maximal retrieval size is 10 times smaller than that of the baseline. The results indicate that our framework is advantageous when dealing with extensive user histories or limited context windows in LLMs. In scenarios where the user has starkly scarce data, our framework outperforms the baseline by a substantial margin (e.g., 11% in the correlation metric), demonstrating the effectiveness of our approach to generalize under the case of new or infrequent users. In our analysis, we further explore the impact of knowledge contributed by collaborator databases, finding that its importance to overall performance escalates as the retrieval capacity grows.

## 2 Persona-DB

The primary challenge in enhancing retrieval augmentation for personalized LLM interactions is the effective and accurate representation of user context (i.e., the interaction data that inform the LLM’s personalization process) within the constraints of LLM context windows and amidst the vast diversity of user behaviors. To effectively represent personas, it is crucial to balance comprehensiveness and data efficiency. We aim to encapsulate detailed user information, such as experiences and opinions, while ensuring that the data is efficiently retrievable. In our context, we define personalization efficiency in terms of information density and predictive power. For instance, a personalization-efficient database will be able to maintain a downstream model with fewer retrieval items. Despite advancements in retrievers and encoders, existing methods, which predominantly rely on user logs and summaries,

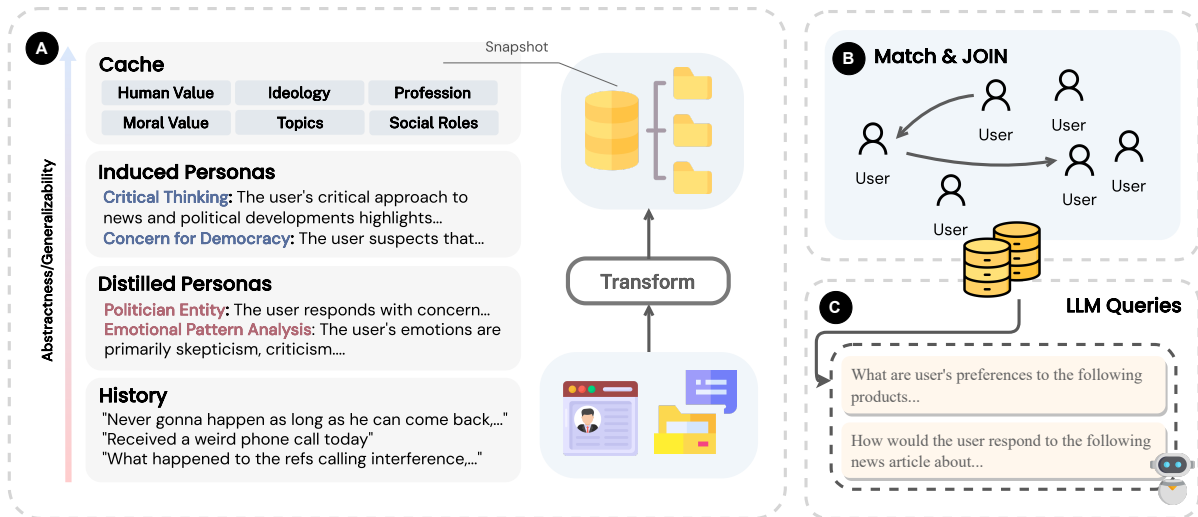


Figure 1: The image outlines the Persona-DB workflow, which starts by distilling and inducing abstract personas from users’ interaction histories. It then leverages the cache layer to facilitate the joining of relevant user databases, effectively borrowing knowledge to fill contextual gaps in the primary user’s data. This enriched data pool is subsequently used by the retrieval for personalized inference.

do not fully leverage the predictive capabilities of user personas or facilitate the exploration of relations between users. Recognizing the limitations, we introduce Persona-DB, a novel framework designed to address key challenges in efficient user context representation. Our approach is guided by two key intuitions. Firstly, high-level (or abstract) user personas, which have been demonstrated by previous social psychology works (Reimer et al., 2022; Torelli and Kaikati, 2009) to effectively predict human actions across different social contexts, can provide a more generalizable basis for persona representation. Secondly, inspired by collaborative filtering prevalent in recommendation systems, we explore the potential of applying similar principles in the context of retrieval-augmented personalization. Specifically, our framework consists of two components, (1) the creation of a database by distilling and extending user data into more generalizable persona constructs and (2) a JOIN operation to bridge knowledge gaps between users by using keys identified in the previous stage. The workflow is depicted in Figure 1.

## 2.1 Hierarchical Refinement

User histories, which capture a user’s past behaviors, are often noisy, with key information dispersed and sometimes only inferable through sequential analysis. Given the scenario where an LLM (or Agent) can only use a limited amount of information due to context limitations, it becomes crucial to appropriately transform the database for

more efficient downstream usage. One natural method for such transformation is creating summaries of events, yet it discourages generalization at some level. For instance, a person’s nature and beliefs (e.g., being supportive of an initiative) can typically be more generalizable across contexts than a reaction under a specific scenario (e.g., positive attitude toward a particular entity within the initiative). From previous works in social psychology, it was suggested that abstract personas such as values are generalizable to predicting future user actions under different social contexts (Reimer et al., 2022; Ajzen, 1991). This indicates that extracting higher-level variables from user data may introduce additional benefits to personalization accuracy while requiring lighter retrieval. To achieve this step, we choose to use an instruction-tuned LLM (Wei et al., 2021; Chung et al., 2022) to structure and infer user facts and opinions to automatically enhance the database. In our experiment, we use the same LLM for both hierarchical refinement and downstream prediction.

Our design schema is a hierarchy consisting of History, Distilled Persona, Induced Persona, and Cache, where an example snapshot is shown in Figure 1. The History layer contains the original user records. On the other hand, the Distilled Persona (abbreviated DP) is an induced dictionary of persona-related facts from History, such as superficial opinions and pattern analysis. The Induced Persona (abbreviated IP)

layer comprises inferred higher-level information from the previous two components, offering a more abstract view of the user. For instance, DP can include specific observations, such as the user’s criticism of a governmental decision to reduce public spending on community services, while IP can include conjectures such as that the user has a broad concern regarding social justice, derived from observations in DP. Lastly, Cache contains consistent human-defined high-level persona categories to assist relevancy matching in the collaboration stage of the framework. We utilize an LLM as an analyzer to infer and enrich the database with both user facts and opinions. In particular, each layer is populated by feeding the information from the lower layers to the LLM for processing. The structure aims to create more packed user-level features to assist data efficiency. By being able to utilize high-level, generalizable personas, the downstream model can apply these insights across a broad range of contexts, enhancing the model’s ability to personalize interactions.

## 2.2 Collaborative Refinement

While using a user’s data to predict intention is straightforward, it can introduce significant challenges when the user’s interactions are irrelevant to a domain presented in the query or when the user simply does not have many interactions. In fact, it was discovered previously that only 25% of highly active users generate 97% of the content on Twitter (McClain et al., 2021), indicating that such a case is dominant in the social media context. To address the issues, we introduce a stage to interconnect knowledge between users, under the assumption that users with similar mindsets would make similar decisions (Reimer et al., 2022). By enabling a user to retrieve and integrate information from relevant users, we establish a collaborative database. This enriches the user context by the inclusion of potentially more domain-matching content from collaborators, thereby enhancing the generalizability and relevance of the user context representation. As demonstrated in our experiment results, such an approach indeed brings enormous benefits in cold-start situations when the user has scarce data on the platform.

In our design, we name this procedure a JOIN stage as an analogy to the JOIN clause in SQL. Formally, let  $DB_i$  be the persona database for user  $i$ , let  $DB_c$  be the database for the current user, and let  $P_i$  indicate retrieval-based embedding prompt

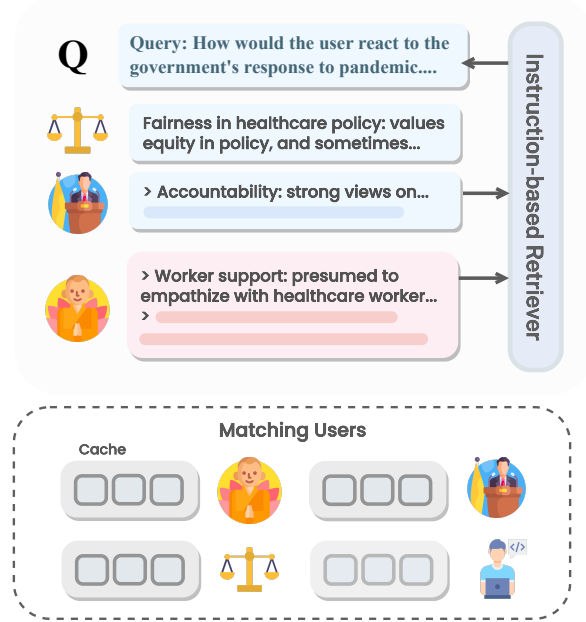


Figure 2: During the retrieval-augmentation stage, the retriever selects data from the user’s and the collaborative databases and composes them at a set ratio to inform the LLM. This strategy aims to enable the model to address challenges like sparse user interactions (e.g., cold-start) and domain irrelevance, offering an effective approach to LLM personalization in environments lacking user graphs.

for user  $i$ . We utilize instruction-based embedding for matching in order to help the model understand the intents behind the query instead of merely comparing semantic similarity. Specifically, the process involves first encoding the Cache of each user’s persona database using a general domain instruction-tuned  $LLM_{\text{embed}}$  that has been trained on large-scale text similarity data (Su et al., 2022; Asai et al., 2022),

$$\mathbf{X}_i = LLM_{\text{embed}}(P_i, DB_i[\text{Cache}]) \quad \forall i \in |U| \quad (1)$$

We then use the embeddings as the key to compute similarities between the current user’s database and others, followed by retrieval and joining of the most relevant databases. Formally,

$$\psi(\mathbf{X}_i, \mathbf{X}_c) = \frac{\mathbf{X}_c \cdot \mathbf{X}_i}{\|\mathbf{X}_c\| \|\mathbf{X}_i\|} \quad (2)$$

$$DB_{\text{Join}}^c = \bigg\|_{i \in \text{TopK}(\psi)} DB_i \quad (3)$$

where  $DB_{\text{Join}}^c$  is the resulting collaborative database. During the retrieval stage, we use  $LLM_{\text{embed}}$  (with a different prompt) to retrieve task-relevant information from both the database from collaborating users and the current user. We show

Method	$\phi_{int}$ (%)		$\phi_p$ (%)	
	$r_s$	$r$	MiF1	MaF1
Majority	-	-	43.41	20.18
Random	0.62	0.41	35.51	30.55
H-Retrieval	40.96	41.17	58.52	47.36
H-Recency	40.7	41.49	58.71	47.49
History (Full)	42.8	43.09	59.58	48.8
IntSum	44.89	45.05	59.96	47.32
Persona-DB	<b>47.67</b>	<b>47.88</b>	<b>62.66</b>	<b>50.59</b>
w/o JOIN	44.66	45.0	61.79	50.39
w/o IP	43.21	43.81	60.73	47.79
w/o DP	44.42	44.51	60.44	48.72

Table 1: The table presents the results from the RFPN dataset, with the best overall performance highlighted in bold. Our framework (with top-40 retrieval) consistently surpasses baseline methods, notably achieving higher accuracy than the method using full user history databases. Additionally, the results include an ablation study demonstrating the standalone effectiveness of Persona-DB without the individual components.

an example of this procedure in Figure 2. In the process of amalgamating information from the two distinct databases, we define a composition ratio denoted by  $x$ . This ratio indicates the allocation of the downstream retrieval capacity,  $r$ , such that  $\lceil x \cdot r \rceil$  represents the proportion of data selected from the collaborative database,  $DB_{Join}^c$ , while the remaining capacity,  $r - \lceil x \cdot r \rceil$ , is sourced from the individual user database,  $DB_c$ . Subsequently, aggregated items are integrated into a prompt, which is then fed to a downstream LLM to produce personalized responses relevant to the user’s query.

### 3 Experiment

#### 3.1 Dataset Overview

To achieve personalization, a model needs to be capable of predicting users’ opinions and preferences regarding a wide range of potential responses or products. In this work, we utilize two datasets, Response Forecasting for Personas in News Media (RFPN) (Sun et al., 2023) and OpinionQA (Santurkar et al., 2023), to evaluate the model’s performance in personalized response prediction. Detailed statistics are provided in the Appendix.

The RFPN dataset comprises 13.3k responses from 8.4k users to 3.8k news headlines sourced from Twitter. We utilize the dataset’s test set, in which the history length can exceed 300. We evaluate the model performance on predicting the sentiment polarity  $\phi_p$  (categorized as *Positive*, *Negative*, or *Neutral*) and the ordinal intensity  $\phi_{int}$  (on a scale from 0 to 3), based on a given persona and news

media message. Here, a persona is constructed from user attributes, including profiles and historical posts. The prediction task allows for reliable evaluation, simplifying the challenge of assessing text generation in personalized conversation with lengthy user contexts, which can be challenging for both LLM and Human. The data also naturally allows the developed model to perform personalized content moderation and recommendation. The OpinionQA dataset includes survey data from Pew Research’s American Trends Panels, and each sample consists of a single-answer multiple-choice question and the user’s selection. In OpinionQA, a persona is defined by both the respondent’s demographic information and historical responses. We specifically use *Biomedical and Food Issues*, *Global Attitudes*, and *America in 2050* topics from the dataset; this allows us to assess the method’s performance across diverse scenarios.

#### 3.2 Experimental Setup

We follow the evaluation from Sun et al. (2023) for the first task. Specifically, we assess intensity predictions using Spearman and Pearson correlation coefficients, denoted as  $r_s$  and  $r$ , respectively. These metrics help quantify the association between the ordinal predicted and actual intensity scales, allowing for credits even when exact matches are not achieved. The sentiment polarity multi-class classification is evaluated using the traditional Micro-F1 and Macro-F1 scores, denoted **MiF1** and **MaF1**, respectively. For the multiple-choice-based OpinionQA task, we measure performance using accuracy. In addition to the main results, we also perform analyses of our method on the RFPN dataset.

To test the efficacy of the Persona-DB components, we compare our framework against different variants that use the original user historical data as the database, including *H-Retrieval*, where the most relevant entries are used, and *H-Recency*, which uses the most recent entries. We also compare with Richardson et al. (2023) (abbreviated *IntSum*), which uses task-aware user summaries generated by LLMs. Additionally, we evaluate the performance of Persona-DB without collaborative refinement in a variant named Persona-DB w/o JOIN (abbreviated P-DB w/o J). For the experiments in 3.3 and 3.4, we use 40 as the retrieval capacity. We also include naive strategies including *Random* and *Majority*, where the former makes predictions randomly, and the latter follows the

majority label. A reference to using full history, *History (Full)*, offers insight into the performance of complete history utilization. In the main and case studies, we set 25% as the composition ratio. To accommodate the context window limitation caused by an excessively large number of interactions, we limit the number of items to within the range that LLM can accept. The retrieval process utilizes an instruction-tuned LLM encoder, `text-embedding-ada-002`<sup>1</sup>, to facilitate both the collaborative user retrieval and the downstream retrieval augmentation. Inference is performed using `gpt-3.5-turbo-0613`<sup>2</sup> via the official OpenAPI, with a fixed seed across runs. The same model is also utilized in the first stage of our framework to process user histories.

### 3.3 General Case Results Discussion

We evaluate the effectiveness of the Persona-DB framework and the baselines, setting the top-k as 40 for both our variants and the baselines in the retrieval-augmentation step. From the evaluation on RFPN (Table 1), it is evident that H-Recency and H-Retrieval underperform compared to History (Full), as expected, due to their less comprehensive view of the user’s information. On the other hand, Persona-DB consistently outperforms these retrieval-based methods across all evaluation metrics with a large margin in correlation, demonstrating that our framework can achieve superior performance under the same retrieval capacity, compared to using historical posts only. The table also indicates that while leveraging a user’s complete history offers a broad view of user data during inference, our approach using only the top-k personas, in fact, exceeds the performance of History. This suggests that our strategy of self-improving and composing interrelated personas achieves better personalization efficiency as it uses much less data while achieving better results.

To underscore the significance of the collaborative stage within our framework further, we examine their absence’s effect on performance. Omitting JOIN operations results in a marked decrease in correlation metrics and F1 scores, affirming its essential role in addressing knowledge gaps among users. Interestingly, even without implementing JOIN, the hierarchical database exhibits greater predictive ca-

<sup>1</sup><https://platform.openai.com/docs/models/embeddings>

<sup>2</sup><https://platform.openai.com/docs/models/gpt-3-5-turbo>

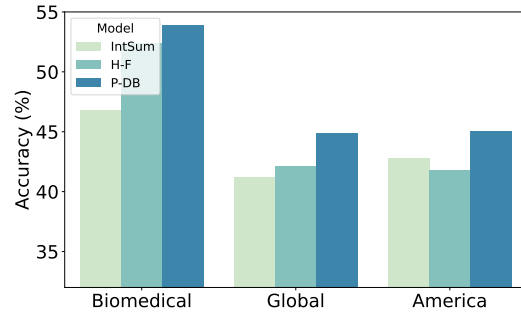


Figure 3: Performance comparison on the OpinionQA task. The plot shows that Persona-DB outperforms the baselines consistently.

pability than that of the baselines using a log-only database. This observation reiterates the value of high-level insights in improving generalizability in the downstream models. We also observe that removing the intermediate persona layers (i.e. IP or DP), leads to a decrease in performance, underscoring their importance in enhancing the model’s capability. Additionally, we conduct a validation study to assess the quality of the LLM analyzer in stage 1 (detailed in the Appendix). Although the extraction is not perfect due to the use of LLM, it still allows both components of our framework to work. This indicates that the method is robust against error propagation at some level, and further enhancements to the LLM to generate higher-quality insights may yield improved performance for Persona-DB.

We also present results on OpinionQA in Figure 3. The plot reveals a similar pattern where Persona-DB surpasses the baseline models across topics. Notably, we see that compared with the baselines, it performs particularly well in the *Global Altitudes* category, consisting of questions on international affairs. This enhanced performance can be attributed to the importance of opinions from users within similar demographic groups in predicting current users’ opinions on this topic.

### 3.4 Analysis on Lurkers and Frequent Users

In our evaluation, we delve into the evaluation of Persona-DB across two types of common yet challenging cases: Lurkers and Frequent Users. This segmentation allows us to understand the framework’s robustness and performance in scenarios representing the extremes of user interaction volumes. We perform the analysis on the RFPN task as it features high diversity in the history lengths.

**Lurkers** The term "Lurker" refers to users with minimal interaction history, a common scenario in

Method	$\phi_{int}$ (%)		$\phi_p$ (%)	
	$r_s$	$r$	MiF1	MaF1
Lurkers (Cold-Start)				
History (Full)	40.66	41.62	67.96	46.02
IntSum	43.15	45.41	69.9	46.05
Persona-DB	<b>55.05</b>	<b>56.75</b>	<b>77.67</b>	<b>57.29</b>
w/o JOIN	47.07	53.34	73.79	49.79
Frequent Users				
H-Retrieval	30.63	33.34	56.08	42.02
H-Recency	38.74	40.95	58.75	46.21
IntSum	44.65	46.58	60.83	44.67
History (Full)	43.29	45.07	61.72	48.63
Persona-DB	<b>49.03</b>	<b>49.71</b>	<b>65.58</b>	<b>52.72</b>
w/o JOIN	41.41	44.44	63.5	49.67

Table 2: The two case studies, focusing on users with extremely scarce (Lurkers) and long (Frequent Users) data histories, demonstrate that our framework markedly improves performance for both ends of the spectrum. These results underscore the effectiveness of our methods in cold-start scenarios, and the ability to process and personalize content for users with extensive histories.

recommendation systems known as the cold-start problem. Being capable of personalizing responses for lurkers indicates a framework’s ability to offer tailored interactions to new users without the prerequisite of extensive historical accumulation. We select 100 users with the sparsest (nonempty) interaction records, averaging only 13.81 interactions per user. The results, as illustrated in Table 2, highlight Persona-DB’s superior performance across all metrics compared to the baselines. Specifically, our framework shows a 11% improvement in Pearson correlation over the baseline method, showcasing its exceptional ability to maintain accuracy with minimal user data. This efficiency can be attributed to the framework’s capability to discern user characteristics from limited interactions, leveraging reliable, frequent collaborators’ insights to enhance personalization accuracy. Remarkably, even before applying the JOIN operation, the transformed database exhibits noticeably improved performance, underscoring the predictive power of LLMs in identifying generalizable personas.

**Frequent Users** Furthermore, we assess the framework’s performance with frequent users, characterized by voluminous interaction histories, potentially introducing noise to the retrieval process. For this analysis, we focused on the top 300 users with the longest histories. The evaluation is detailed in Table 2. Interestingly, the performance of H-Retrieval is notably lower than that of the

Top-K	Method	Pearson	Accuracy
40	Persona-DB	<b>47.88</b>	<b>62.66</b>
40	IntSum	45.05	59.96
40	Hist-Recency	41.49	58.71
30	Persona-DB	<b>47.45</b>	<b>61.89</b>
30	IntSum	41.93	59.38
30	Hist-Recency	42.08	59.58
10	Persona-DB	<b>45.46</b>	<b>60.35</b>
10	IntSum	45.33	60.15
10	Hist-Recency	42.21	59.19

Table 3: Our framework is capable of maintaining performance above that of baselines, even when the retrieval size is significantly reduced.

recency-based method. One hypothesis for this observation is that when the user history becomes frequent, the retrieval of relevant information in a small capacity becomes harder as there are more semantically similar yet non-relevant items. The table also demonstrates Persona-DB’s consistent out-performance over alternative methods. This finding emphasizes the framework’s adeptness at tailoring content for users with extensive interaction histories. The success can be largely attributable to the hierarchical representation’s ability to synthesize and distill relevant information from a broad dataset, effectively mitigating noise and enhancing personalization accuracy for both frequent users and their collaborators.

### 3.5 Analysis on Varying Retrieval Sizes

To minimize context window cost in the inference stage, it is crucial to improve data efficiency in retrieval (i.e., the retrieved set remains generalizable while being reduced in size). This section explores how varying retrieval sizes affect our framework’s performance.

Table 3 demonstrates the robustness of our method, illustrating that *accuracy is maintained above the baselines even as the retrieval cap is significantly reduced*. This underscores the method’s resilience to variations in context window size for LLMs, highlighting its capability to deliver precise predictions with minimal data retrieval. The compelling performance, even with a constrained database, supports our hypothesis that the personas derived from our hierarchical construction possess enhanced predictive power. Moreover, we can see that as the retrieval size decreases, the advantage of using collaborative entries shrinks. This trend likely stems from the increased significance of a user’s own data in situations where retrieval capacity is limited.

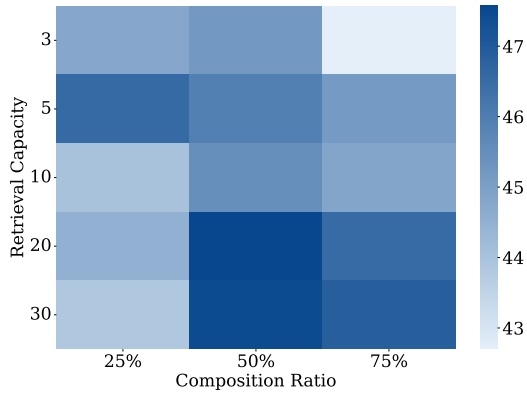


Figure 4: The figure illustrates the shift of correlation performance metric as capacity and proportion of collaborative content changes. The trends show that collaborative retrieval becomes more important as the retrieval size grows.

### 3.6 The Impact of Collaborative Composition

We further analyze how the proportion of collaborative items influences Pearson correlation performance across different retrieval capacities. The heatmap depicted in Figure 4 reveals an interesting trend: for larger retrieval sizes, incorporating a higher percentage of data from collaborative databases results in improved performance. Conversely, for smaller retrieval sizes, an increased reliance on collaborative data sometimes detrimentally affects performance. This pattern can be attributed to the reason that typically, a user’s database contains only a limited number of entries directly relevant to addressing a query effectively. Once the retrieval size surpasses this small pool of pertinent data, the inclusion of novel, domain-relevant insights from collaborative databases becomes increasingly beneficial. However, in situations with extremely limited retrieval capacity, the critical, predictive nature of the user’s own data takes precedence. In these instances, over-reliance on external data sources can, therefore, undermine performance. The pattern underscores the importance of maintaining a strategic balance in the composition of the retrieval set to optimize LLM personalization.

## 4 Related Work

Personalization plays a crucial role in the effectiveness of large language models (LLMs), particularly in ensuring their alignment with individual preferences and interests (Hwang et al., 2023; Kirk et al., 2023; Mao et al., 2023; Han et al., 2023; Chan et al., 2021; Ge et al., 2024; Salemi et al., 2023). With the advent of LLMs, the land-

scape of personalization has also been significantly transformed, enabling enhanced user engagement and the delivery of services tailored to individual needs (Chen et al., 2023b) and cultural preferences (Li et al., 2023b; Fung et al., 2023, 2024; Zhang et al., 2024). Traditional personalization approaches for language models often use auxiliary encoders to learn user representation from user profile data (King and Cook, 2020; Qiu et al., 2021; Wang et al., 2021; Li et al., 2023a). For instance, to personalize pre-trained language models, UserIdentifier (Miresghallah et al., 2022) prepends a user identifier prefix at the beginning of the input text during the fine-tuning process. Yet such works would require training resources and can be challenging in situations of extensive user data and large models. Recently, there have been works focusing on utilizing and improving retrieval augmentation due to its capability to efficiently borrow external knowledge lacking in the pre-trained data (Wang and McAllester, 2020; Lewis et al., 2020) without the cost of fine-tuning downstream LLMs. In Zhang et al. (2023) and Asai et al. (2022), the authors introduced new embedding models to better support the diverse retrieval augmentation needs. Furthermore, Srinivasan et al. (2022) introduced a practical approach for distilling retrieval-augmented LLMs, and Guu et al. (2020) demonstrated the effectiveness of retrieval-augmented language model pre-training for open-domain question answering. Yet limited exploration has been done on optimizing the data to be retrieved to improve personalization.

## 5 Conclusions and Future Work

In conclusion, we introduce Persona-DB, a framework designed to enhance the accuracy and context efficiency of retrieval-based LLM personalization through collaborative data refinement. The framework leverages a hierarchical approach to user data representation, enabling the construction of insights that are generalizable across task contexts, and introduces a collaborative mechanism to effectively bridge communal knowledge for improved query response when the user lacks sufficient relevant data. Our experimental results indicate a marked performance improvement, particularly in scenarios characterized by sparse user data, thus addressing a critical challenge in LLM personalization. Furthermore, we perform analysis to study the impact of data composition on retrieval



effectiveness. Future work will focus on developing methods that refine the framework’s capability to dynamically adapt the matching process based on user feedback and interaction patterns.

## Limitations

We employ LLMs for abstracting personas in our work. Due to inherent bias and imperfection in the LLM, the quality of the extraction in hierarchical refinement can also affect the generalizability on diverse personas. We believe that this shortfall can be addressed through continuous improvement of the LLM’s capabilities. Moreover, in this study, our focus primarily lies on improving retrieval efficiency in downstream task scenarios, where we evaluate the model’s capability to preserve accuracy while reducing retrieval capacity, which in turn reduces the cost associated with online inference. In the hierarchical refinement stage, there is an existence of a one-time inference cost associated with pre-processing user data, a trade-off for enhancing data efficiency in downstream usage. We believe that using a smaller distilled model will alleviate the cost for this stage.

## Ethics Statements

Our personalization approach involves storing and retrieving information about user preferences and personalities, which requires careful safety protocol and ethical consideration when deployed. While the extraction of generalized user personas and collaborative data composition offers powerful personalization capabilities, the sharing and combining of information between users requires consideration of privacy implications. The process should be implemented with appropriate privacy protections such as data anonymization, secure storage, and user consent.

Additionally, we recognize important considerations for responsible development: First, personas can risk oversimplifying and stereotyping users (as they are caricatures of people), potentially leading to biased or unrealistic representations that is unable to capture the true complexity and diversity of human behavior. Second, personalization technology enables detailed user profiling that could be misused for surveillance or targeting of individuals and groups. To mitigate these risks, the developers need to emphasize privacy-preserving techniques and aim to maintain user agency by being transparent about data usage.

Moreover, the datasets we used in this work are from publicly accessible repositories from existing publications. We only provide pointers to these existing data repositories and the data loading script that contains no user information. We will not share the datasets themselves. We will also not release any intermediate representations generated by models which may contain user information.

## Acknowledgement

This research is based upon work supported in part by U.S. DARPA INCAS Program No. HR001121C0165. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

## References

- Icek Ajzen. 1991. The theory of planned behavior. *Organizational Behavior and Human Decision Processes*.
- Akari Asai, Timo Schick, Patrick Lewis, Xilun Chen, Gautier Izacard, Sebastian Riedel, Hannaneh Hajishirzi, and Wen-tau Yih. 2022. Task-aware retrieval with instructions. *arXiv preprint arXiv:2211.09260*.
- Jinheon Baek, Nirupama Chandrasekaran, Silviu Cucerzan, Sujay Kumar Jauhar, et al. 2023. Knowledge-augmented large language models for personalized contextual query suggestion. *arXiv preprint arXiv:2311.06318*.
- Hou Pong Chan, Lu Wang, and Irwin King. 2021. [Controllable summarization with constrained markov decision process](#). *Trans. Assoc. Comput. Linguistics*, 9:1213–1232.
- Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, Xingmei Wang, et al. 2023a. When large language models meet personalization: Perspectives of challenges and opportunities. *arXiv preprint arXiv:2307.16376*.
- Zheng Chen, Ziyang Jiang, Fan Yang, Zhankui He, Yupeng Hou, Eunah Cho, Julian McAuley, Aram Galstyan, Xiaohua Hu, and Jie Yang. 2023b. The first workshop on personalized generative ai@ cikum 2023: Personalization meets large language models. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 5267–5270.

- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Yi Fung, Tuhin Chakrabarty, Hao Guo, Owen Rambow, Smaranda Muresan, and Heng Ji. 2023. **NORM-SAGE: Multi-lingual multi-cultural norm discovery from conversations on-the-fly**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 15217–15230, Singapore. Association for Computational Linguistics.
- Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. 2024. **Massively multi-cultural knowledge acquisition & lm benchmarking**. *Preprint, arXiv:2402.09369*.
- Tao Ge, Xin Chan, Xiaoyang Wang, Dian Yu, Haitao Mi, and Dong Yu. 2024. Scaling synthetic data creation with 1,000,000,000 personas. *arXiv preprint arXiv:2406.20094*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *International conference on machine learning*, pages 3929–3938. PMLR.
- Chi Han, Jialiang Xu, Manling Li, Yi Fung, Chenkai Sun, Nan Jiang, Tarek Abdelzaher, and Heng Ji. 2023. Lm-switch: Lightweight language model conditioning in word embedding space. *arXiv preprint arXiv:2305.12798*.
- EunJeong Hwang, Bodhisattwa Prasad Majumder, and Niket Tandon. 2023. **Aligning language models to user opinions**. *CoRR*, abs/2305.14929.
- Milton King and Paul Cook. 2020. Evaluating approaches to personalizing language models. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 2461–2469.
- Hannah Rose Kirk, Bertie Vidgen, Paul Röttger, and Scott A Hale. 2023. Personalisation within bounds: A risk taxonomy and policy framework for the alignment of large language models with personalised feedback. *arXiv preprint arXiv:2303.05453*.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Junyi Li, Ninareh Mehrabi, Charith Peris, Palash Goyal, Kai-Wei Chang, Aram Galstyan, Richard Zemel, and Rahul Gupta. 2023a. On the steerability of large language models toward data-driven personas. *arXiv preprint arXiv:2311.04978*.
- Sha Li, Chi Han, Pengfei Yu, Carl Edwards, Manling Li, Xingyao Wang, Yi Fung, Charles Yu, Joel Tetreault, Eduard Hovy, and Heng Ji. 2023b. **Defining a new NLP playground**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 11932–11951, Singapore. Association for Computational Linguistics.
- Shengyu Mao, Ningyu Zhang, Xiaohan Wang, Mengru Wang, Yunzhi Yao, Yong Jiang, Pengjun Xie, Fei Huang, and Huajun Chen. 2023. Editing personality for llms. *arXiv preprint arXiv:2310.02168*.
- Colleen McClain, Regina Widjaya, Gonzalo Rivero, and Aaron Smith. 2021. The behaviors and attitudes of us adults on twitter.
- Fatemehsadat Miresghallah, Vaishnavi Shrivastava, Milad Shokouhi, Taylor Berg-Kirkpatrick, Robert Sim, and Dimitrios Dimitriadis. 2022. **UserIdentifier: Implicit user representations for simple and effective personalized sentiment analysis**. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL 2022, Seattle, WA, United States, July 10-15, 2022*, pages 3449–3456. Association for Computational Linguistics.
- Saif Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. Semeval-2018 task 1: Affect in tweets. In *Proceedings of the 12th international workshop on semantic evaluation*, pages 1–17.
- Zhuoshi Pan, Qianhui Wu, Huiqiang Jiang, Menglin Xia, Xufang Luo, Jue Zhang, Qingwei Lin, Victor Rühle, Yuqing Yang, Chin-Yew Lin, et al. 2024. LlmLingua-2: Data distillation for efficient and faithful task-agnostic prompt compression. *arXiv preprint arXiv:2403.12968*.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*.
- Zhaopeng Qiu, Xian Wu, Jingyue Gao, and Wei Fan. 2021. U-bert: Pre-training user representations for improved recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 4320–4327.
- Nils Karl Reimer, Mohammad Atari, Farzan Karimi-Malekabadi, Jackson Trager, Brendan Kennedy, Jesse Graham, and Morteza Dehghani. 2022. Moral values predict county-level covid-19 vaccination rates in the united states. *American Psychologist*, 77(6):743.
- 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*.

- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2023. Lamp: When large language models meet personalization. *arXiv preprint arXiv:2304.11406*.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? In *International Conference on Machine Learning*, pages 29971–30004. PMLR.
- Krishna Srinivasan, Karthik Raman, Anupam Samanta, Lingrui Liao, Luca Bertelli, and Mike Bendersky. 2022. Quill: Query intent with large language models using retrieval augmentation and multi-stage distillation. *arXiv preprint arXiv:2210.15718*.
- Hongjin Su, Weijia Shi, Jungo Kasai, Yizhong Wang, Yushi Hu, Mari Ostendorf, Wen-tau Yih, Noah A Smith, Luke Zettlemoyer, and Tao Yu. 2022. One embedder, any task: Instruction-finetuned text embeddings. *arXiv preprint arXiv:2212.09741*.
- Chenkai Sun, Jinning Li, Hou Pong Chan, Chengxiang Zhai, and Heng Ji. 2023. Measuring the effect of influential messages on varying personas. *arXiv preprint arXiv:2305.16470*.
- Carlos J Torelli and Andrew M Kaikati. 2009. Values as predictors of judgments and behaviors: The role of abstract and concrete mindsets. *Journal of personality and social psychology*, 96(1):231.
- Hai Wang and David McAllester. 2020. On-the-fly information retrieval augmentation for language models. *arXiv preprint arXiv:2007.01528*.
- Zhenghui Wang, Lingxiao Luo, and Diyi Yang. 2021. Personalized response generation with tensor factorization. In *Proceedings of the 1st Workshop on Natural Language Generation, Evaluation, and Metrics (GEM 2021)*, pages 47–57.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.
- 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*, pages 38–45, Online. Association for Computational Linguistics.
- Peitian Zhang, Shitao Xiao, Zheng Liu, Zhicheng Dou, and Jian-Yun Nie. 2023. Retrieve anything to augment large language models. *arXiv preprint arXiv:2310.07554*.
- Wenxuan Zhang, Hou Pong Chan, Yiran Zhao, Mahani Aljunied, Jianyu Wang, Chaoqun Liu, Yue Deng, Zhiqiang Hu, Weiwen Xu, Yew Ken Chia, Xin Li, and Lidong Bing. 2024. **Seallms 3: Open foundation and chat multilingual large language models for southeast asian languages**. *CoRR*, abs/2407.19672.

## A Appendix

### A.1 Additional Details

Our experiments were conducted using the PyTorch framework (Paszke et al., 2019) and the Hugging-face Transformers library (Wolf et al., 2020). The sentiment intensity labels adhere to the definition in SemEval-2018 Task 1 (Mohammad et al., 2018), incorporating both magnitude and direction in the evaluation of sentiment intensity. In our work, we use concatenation of histories to implement the JOIN in Equation 3 after collaborative databases have been retrieved. Future work will explore condensing the collaborative databases using LLMs.

The data statistics for RFPN are shown in Table 4. For OpinionQA, each topic we used contains a test split of 1000 samples. In particular, *Biomedical and Food Issues* contains 67 questions in the survey and there are 2537 respondents in total, *Global Attitudes* contains 104 questions and 2596 respondents, and *America in 2050* contains 90 questions with 2524 respondents in total.

We show all prompts used in the work in the following figures. Figure 7 and Figure 8 represent the prompts for constructing the hierarchical database. Figure 9, Figure 10, and Figure 11 represent the inference prompt for baseline, Persona-DB w/o JOIN, and Persona-DB.

### A.2 Validation Study on LLM Extraction

As the extracted personas in stage 1 serve as an intermediate layer to the method, and the LLM analyzer itself can introduce errors, we conducted a validation study regarding the quality of the extraction. Specifically, we sampled 50 LLM extraction results from user histories and distributed them to three human raters (who are graduate students who passed an initial quiz of 8 samples) to verify the accuracy. On average, the raters assigned a score of 3.9/5, indicating a notable level of precision despite imperfections in the LLM’s output. While the extraction accuracy isn’t imperfect, our results show that the quality of the current annotation does have the merit of contributing to performance positively.

### A.3 Comparison with Prompt Compression

In our method, we focus on enhancing data efficiency through retrieval augmentation. Another line of approach involves compressing the prompt. Although prompt compression methods can be integrated with our approach, either before or after, it would be beneficial to conduct a comparison using

Split	Train	Dev.	Test
# Samples	10,977	1,341	1,039
# Headlines	3,561	1,065	843
# Users	7,243	1,206	961
Avg # Profile Tokens	10.75	11.02	10.50
Avg # Response Tokens	12.33	12.2	11.87
Avg # Headline Tokens	19.79	19.82	19.72

Table 4: Summary statistics for the RFPN dataset.

method	wasserstein	mse
H-Retrieval	72.77	4.45
H-Recency	72.71	4.44
IntSum	73.15	4.3
History (Full)	73.46	4.3
Persona-DB	<b>74.27</b>	<b>4.13</b>
w/o JOIN	73.13	4.32
w/o IP	72.88	4.49
w/o DP	72.98	4.4

Table 5: Results on additional metrics.

the datasets in our study. Specifically, we compare our method with LLMLingua-2 (Pan et al., 2024) with a 40% compression rate on the full history (e.g., the actual number of tokens produced can vary due to probabilistic inference). LLMLingua-2 employs a data distillation procedure to extract knowledge from an LLM for prompt compression. The comparative results are presented in Table 6 and Figure 6. Our initial analysis indicates that when comparing the two independent directions of methods separately, our method outperforms in the current setting; this may stem from our method’s integration of knowledge from collaborative users.

### A.4 Additional Evaluation Metrics and Qualitative Results

We show results including 1-Wasserstein distance-based alignment score and mean squared error metrics for sentiment intensity prediction in RFPN (Table 5). We see that the results match the trend in the main table.

We additionally show qualitative case studies in Figure 5 to demonstrate the benefits of collaborative strategy.

<p><b>News:</b> Justin Gatlin defeats world record holder Usain Bolt in the 100-meter final in</p>	<p><b>News:</b> "It is not up to women of color to save this country from itself. That's on all of us." Planned Parenthood President Cecile Richards urges women to "do better" in the fight for equality at the Women's March rally in Las Vegas.</p>
<p><b>User Original Historical Posts:</b> A mix of interactions with public figures and topics ranging from environmental concerns to personal travel and technology projects. There is little information related to athletics or opinions on specific athletes that would help predict a response. Example lines: "@LeoDiCaprio I agree totally!", "@BBCWorld good for you. Thanks for all the great races. I have technology that would use all waste in Jamaica and make electricity, pure water".</p> <p><b>Collaborative Refinement:</b> The collaborators' databases provide insights into general interests in significant events and achievements, which could suggest how a similar community might respond to major sports milestones. Example lines: "View toward NGSuperEagles: Best of luck, Shock the world", "View toward VictorMoses: Great efforts"</p>	<p><b>User Original Historical Posts:</b> Show a consistent pattern of political and social commentary. While these posts display a political stance, they don't specifically address the user's views on racial equality. Examples: "@danablankenhorn @ananavarro Me too Dana, me too!", "@TempleDrakeOO @ProgressivePgh @Doobsac @leftgear @robrousseau @NoFascistLies It's a simple case of 'Oldfartism'. The old fogeys in the House &amp; Senate are set in their ways."</p> <p><b>Collaborative Refinement:</b> The collaborators' opinions provide valuable insights into their views on social justice and equality, which are relevant to the post text. Their advocacy for fairness and equal rights for all suggests a broader and more inclusive understanding of social justice issues. Example lines: "View toward Equality: Advocate for fairness and equal rights for all.", "View toward Social Justice: Support for Democratic candidates and progressive policies.", "View toward Social equality: Supports Minority rights."</p>
<p><i>Baseline Prediction:</i> Neutral <i>Our Prediction:</i> Positive (✓)</p>	<p><i>Baseline Prediction:</i> Neutral <i>Our Prediction:</i> Positive (✓)</p>

Figure 5: Case Studies.

Method	$\phi_{int}$ (%)		$\phi_p$ (%)	
	$r_s$	$r$	MiF1	MaF1
LLMLingua-2	42.46	42.94	59.96	48.87
Persona-DB	<b>47.67</b>	<b>47.88</b>	<b>62.66</b>	<b>50.59</b>

Table 6: The table displays the RFPN task results when compared to prompt compression.

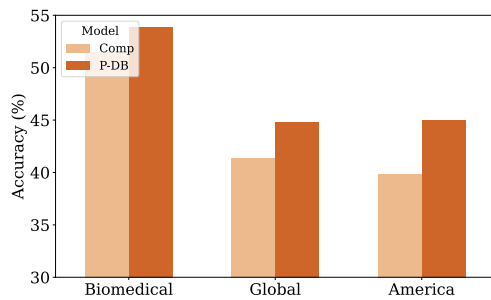


Figure 6: Comparison with prompt compression on OpinionQA task.

This is a sociology 101 mid-term exam.

Describe the Twitter user (based on its profile and history) by filling in the categories below COMPLETELY and COMPREHENSIVELY, without losing information. Do not rename the categories. Separate by newline. Replace Twitter usernames with entity name when possible and remove all Twitter usernames when replacement is possible. You can write as long as you want and just make sure the answer is comprehensive.:

Human Values (i.e., what human values does the user likely have. Choose from "Conformity", "Tradition", "Security", "Power", "Achievement", "Hedonism", "Stimulation", "Self-Direction", "Universalism", "Benevolence". These are based on the basic theory of human values. Separate by commas. Do not justify),  
Moral Values (i.e., what moral values does the user likely have. choose from "authority", "betrayal", "care", "cheating", "degradation", "fairness", "harm", "loyalty", "purity", "subversion". Separate by commas),  
Ideologies (i.e., choose from political ideologies. Separate by commas),  
Possessions (i.e., what the user owns),  
Interested topics (i.e., what topics is the user interested in. This can be inferred from the profile and history. Separate by commas),  
Interested issues and the user's belief/view toward each issue (i.e., in the format of "<issue>:<belief/view>". Separate by commas.),  
Interested entities and the user's belief/view toward each entity (i.e., in the format of "<entity>:<belief/view>". Separate by commas.),  
Profession (e.g., jobs, specialty. Separate by commas),  
Social roles (e.g., mom, dad. Separate by commas),  
Behavior Pattern Analysis,  
Emotional Pattern Analysis,  
Engagement Pattern Analysis,  
Other Notes,.

Answer the categories concisely and comprehensively, don't put words like "Based on the user profile and historical posts, here is the user description" in your answer.  
If a category cannot be answered, fill in the exact word "None" for the corresponding category.

I will iteratively provide the user profile and user historical posts. Here is the information,

[user profile]: '{profile}',

[user historical posts](separated by semi-colons): '{history}'.

Figure 7: Prompt template used for extracting hierarchical user database.

This is a sociology 101 mid-term exam.

Given the most recent User Profile/History and [Extracted Personas] (derived from the full User Profile/History), reflect on them and generate higher-level "Interested Issues and Beliefs" and "Interested Entities and Beliefs" and "Analysis" that do not exist in [Extracted Personas]. Exclude the issues or entities existing in [Extracted Personas]. Higher-Level means more abstract. These are higher-level and generalizable beliefs and analysis that can be used to predict user's beliefs on many other aspects. Specifically, fill the categories below COMPREHENSIVELY, without losing information. Do not rename the categories. Do not put explanation or reasons. Think step by step. Take a deep breath.:

High-Level Interested issues and the user's belief/view toward each issue,  
High-Level Interested entities and the user's belief/view toward each entity,  
High-Level Analysis,  
Other Notes.

[User Profile]  
'{profile}'

[User History]  
'{history}'

[Extracted Personas]  
'{key}'.

Figure 8: Prompt template used for extracting hierarchical user database.

This is a sociology 101 mid-term exam.

Predict response from the user to the news headline in terms of exact comment words (i.e., what would user reply in comment), sentiment polarity (i.e., Positive, Neutral, Negative), and sentiment intensity (integer scaled between 0-3 inclusively where 0 means no intensity and 3 means the most intense). Note: when sentiment polarity is neutral, the sentiment intensity should be 0.

In other words, filling the <answer> in the categories below. Separate by newline:

Comment: <answer>

Sentiment Polarity: <answer>

Sentiment Intensity: <answer>

I will iteratively provide the news headline, user profile, and user historical posts. Here is the information,

[news headline]: '{post}'.

[user profile]: '{profile}'.

[user historical posts]: '{history}'.

Figure 9: Prompt template used for Baseline Inference.

This is a sociology 101 mid-term exam.

Predict response from the user to the news headline in terms of exact comment words (i.e., what would user reply in comment), sentiment polarity (i.e., Positive, Neutral, Negative), and sentiment intensity (integer scaled between 0-3 inclusively where 0 means no intensity and 3 means the most intense). Note: when sentiment polarity is neutral, the sentiment intensity should be 0.

In other words, filling the <answer> in the categories below. Separate by newline:

Comment: <answer>

Sentiment Polarity: <answer>

Sentiment Intensity: <answer>

I will iteratively provide the news headline, user profile, user structured profile (inferred by AI) and user history. The user history can contain analysis on the user or historical tweets from the user. Here is the information,

[news headline]: '{post}'.

[user profile]: '{profile}'.

[user structured profile]: '{key}'.

[user historical posts or analysis on the user]: '{history}'.

Figure 10: Prompt template used for Inference for Persona-DB without JOIN.

This is a sociology 101 mid-term exam.

Predict response from the user to the news headline in terms of exact comment words (i.e., what would user reply in comment), sentiment polarity (i.e., Positive, Neutral, Negative), and sentiment intensity (integer scaled between 0-3 inclusively where 0 means no intensity and 3 means the most intense). Note: when sentiment polarity is neutral, the sentiment intensity should be 0.

In other words, filling the <answer> in the categories below. Separate by newline:

Comment: <answer>  
Sentiment Polarity: <answer>  
Sentiment Intensity: <answer>

I will iteratively provide the news headline, user profile, user structured profile (inferred by AI), user/neighbors history (containing analysis on the user/neighbors or historical tweets from the user/neighbors). The neighbors are the selected users that could be similar to the current user, but they are not necessarily helpful, so treat it cautiously. Here is the information,

[news headline]: '{post}'.

[user profile]: '{profile}'.

[user structured profile]: '{key}'.

[user historical posts or analysis on the user]: '{history}'.

[neighbors historical posts or analysis on the neighbors]: '{neighbors}'.

Figure 11: Prompt template used for Inference for Persona-DB.

Query Prompt: "Encode this news headline to find the most relevant user information that might hint at their reaction to it: "

History Prompt: "Encode the user's historical tweet for retrieval. Even if not directly related to a news headline, this tweet may still hint at the user's opinions and personality. Consider the tweet carefully. Tweet: "

Induced Persona Prompt: "Encode the analysis on the user for retrieval. Even if not directly related to a news headline, this analysis on the user may still hint at the user's opinions and personality. Consider the analysis carefully. Analysis: "

Figure 12: Prompt template for retriever.

This is a sociology 101 mid-term exam.

Predict response from the user to the query from a Pew survey, ensuring your selection is limited to the provided choice list. Provide a brief explanation of your selection on the basis of the user profile.

In other words, filling the <answer> in the categories below. Separate by newline:

Answer: <answer>  
Explanation: <answer>

I will iteratively provide the query, choice list, user profile, user/neighbors history (containing analysis on the user/neighbors or historical question answer pairs from the user/neighbors). The neighbors are the selected users that could be similar to the current user, but they are not necessarily helpful, so treat it cautiously. Here is the information,

====query====  
{post}

====choice list====  
{choices}

====user profile====  
{profile}

====user historical posts or analysis on the user====  
{history}

====neighbors historical posts or analysis on the neighbors====  
{neighbors}

Figure 13: Prompt template used for Inference for Persona-DB in OpinionQA.