Comparative Personalization for Multi-document Summarization

Haoyuan Li, Snigdha Chaturvedi
University of North Carolina at Chapel Hill,
{haoyuanl, snigdha}@cs.unc.edu

Abstract

Personalized multi-document summarization (MDS) is essential for meeting individual user preferences of writing style and content focus for summaries. In this paper, we propose that for effective personalization, it is important to identify fine-grained differences between users' preferences by comparing preferences of different users. Motivated by this, we propose ComPSum, a personalized MDS framework. It first generates a structured analysis of a user by comparison between users. The generated structured analysis is then used to guide the generation of personalized summaries. To evaluate ComPSum, we propose AuthorMap, a fine-grained reference-free evaluation framework for personalized MDS. It evaluates the personalization of a system based on the authorship attribution between two personalized summaries generated for different users. For robust evaluation of personalized MDS, we construct PerMSum, a personalized MDS dataset in the review and news domain. We evaluate the performance of ComPSum, showing that it outperforms strong baselines. The code is available at https://github.com/leehaoyuan/comparative_personalize_sum.

1 Introduction

Multi-document summarization (MDS) aims to generate a summary with the salient information from multiple documents on a certain topic, such as news articles about an event (Fabbri et al., 2019) or reviews of a product (Bražinskas et al., 2020). However, previous works show that a one-size-fit-all approach is insufficient for summarization (Jones et al., 1999; Ter Hoeve et al., 2022). Summarization models should account for the audience and purpose since different users can have different *personal preferences* on *writing style* and *content focus* (Jang et al., 2023; Salemi et al., 2024). While writing style refers to the manner or tone in which the summaries are written, content focus refers to which aspects are emphasized about a certain topic. Users can have different preferences for *writing style* (Huang et al., 2024). For example, for reviews, some users may prefer a formal tone, while others may prefer a conversational tone. User preferences can also differ on *content focus* (Daumé III and Marcu, 2006). Some users may focus on price and utility while others might focus on quality and durability. To meet these user preferences, personalized MDS is essential.

Personalized MDS is related to personalized text generation. Recent works on personalized text generation are based on Large Language Models (LLMs) and use the *profile* of individual users–set of documents previously authored by the user. They either retrieve related documents from a user's profile (Salemi et al., 2024; Li et al., 2023a), include a summary of the user's profile (Richardson et al., 2023), or tune different models for different users based on their profiles (Tan et al., 2024). However, most of these works only include general features of the user and ignore how they are different from other users. To identify the finer differences, it is important to compare a user's profile documents with comparable profile documents written by other users (Qiu et al., 2025). Ideally, the profile documents of two users can be comparable if they are on the same topic but differ on personal preferences. In general personalized text generation, identifying such comparable profile documents of different users can be difficult since the differences between profile documents of different users can stem from either personal preferences or topic differences. Contrarily, for MDS, since all input documents are about the same topic (e.g. reviews about the same product), their differences are more likely to stem from differences in personal preferences.

Motivated by this, we propose ComPSum (**Com**parative **P**ersonalization for Multi-Document **Sum**marization), a personalized MDS framework. ComPSum considers two key

Proceedings of the 15th International Joint Conference on Natural Language Processing and the 5th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics, pages 966–985
November 6-10, 2026 ©2026 Association for Computational Linguistics

Style analysis: User X's writing style is characterized by a formal, journalistic tone, often incorporating quotes and direct references to sources. The language is precise and concise, with a focus on conveying information clearly and accurately. Unlike other users, who may use more casual or sensational language, User X's writing is more polished and professional. **Content analysis**: User X's profile texts are notable for their in-depth analysis of celebrity news, focusing on the deeper implications and societal context of the events. Unlike other users, who often stick to a surface-level summary of the news, User X delves into the complexities of the issue, often incorporating historical context, social commentary, and personal insights.

Table 1: Structured analysis on writing style and content focus generated by ComPSum.

preference dimensions: writing style and content focus (Zhang et al., 2024). It first generates a structured analysis of these two dimensions for a user by comparing profile documents with other documents authored by different users on the same topic as shown in Tab. 1. The generated structured analysis is then used to guide the generation of personalized summaries that capture the user's individual preference.

Apart from generating personalized summaries, their evaluation is also a major challenge, as there are no reference summaries and conventional metrics like ROUGE (Lin, 2004) cannot be applied. To address this, we propose AuthorMap, a fine-grained reference-free evaluation framework for personalized MDS that independently assesses writing style and content focus since they are independent of each other (Salemi et al., 2025). To control the inherent differences between systems, AuthorMap evaluates the personalization of a system based on the authorship attribution between two personalized summaries generated for different users by the same system. We perform automatic and human evaluation of AuthorMap and show that it achieves reasonable accuracy.

For robust evaluation of personalized MDS, we construct PerMSum, an MDS dataset in reviews and news domains. For evaluation, we need documents labeled with their authors (users). While there are MDS datasets with user labels in the review domain (Ni et al., 2019), the news domain still lacks such datasets. To construct such a dataset, we use news articles from the All The News dataset [1]. We collect and process 14K document sets and 1.4K users. These document sets are then combined with the document sets sampled from the Amazon dataset (Ni et al., 2019) to form PerMSum, which results in 45K document sets and 5.3K users.

[1]https://components.one/datasets/all-the-news-2-news-articles-dataset/

We evaluate the performance of ComPSum using four LLMs and two domains, showing that it consistently outperforms strong baselines across LLMs and domains. Our contributions are three-fold:

- ComPSum, a personalized MDS framework with comparative personalization;
- AuthorMap, a reference-free evaluation framework for personalized MDS;
- PerMSum, a personalized MDS dataset in the news and review domain;

2 Related Work

Personalized text generation aims to generate a personalized text for a user using his profile, making it closely related to personalized MDS. Salemi et al. (2024) retrieve related documents from a user's profile and Li et al. (2023a) train models to summarize and synthesize the retrieved documents. To mitigate the information loss in retrieval, Richardson et al. (2023) additionally include a summary of user profile. However, these works model user individually. Recently, Sun et al. (2025) use other similar users' profile to infer a user's profile when existing data about the user is sparse. Qiu et al. (2025) improve personalized review generation by comparing a user's review with other user's review. However, their design is specific to the review domain. We perform experiments to show that ComPSum outperforms their method in Sec. 7.4. Another line of works (Daumé III and Marcu, 2006; Padmakumar et al., 2025) generates personalized summaries using user-provided queries, which differs from personalized MDS that uses user profiles.

Most previous works on personalized text generation (Salemi et al., 2024; Ao et al., 2021) use reference-based metrics that evaluate the similarity between generated texts and human-written references, such as ROUGE (Lin, 2004) and BERTScore (Zhang et al., 2019). Recently, Salemi et al. (2025) propose to evaluate personalized text generation from writing style and content, which are similar to the dimensions used by AuthorMap, but it assesses how well the generated text aligns with the reference text, which differs from the reference-free AuthorMap. For reference-free evaluation, Wang et al. (2023) propose AuPEL, which evaluates personalization based on pairwise authorship attribution (Bozkurt et al., 2007) between two personalized texts from different systems for the same user. However, AuPEL overlooks various dimensions of personalization and does not con-

Algorithm 1 Pseudo Code for ComPSum

Require: Document set D, user profile P_u, retrieval model $\mathcal{R}$, LLM
$P_u^{1:k} = \mathcal{R}(D, P_u, k)$ ▷ Retrieve top-k relevant profile documents
for all $p_u^i \in P_u^{1:k}$ **do**
  $p_{\neg u}^i = \mathcal{R}(p_u^i, C_{\neg u}^{p_i}, 1)$ ▷ Retrieve most dissimilar comparative document
end for
$a_u = \text{LLM}(p_u^1, p_{\neg u}^1, \ldots, p_u^k, p_{\neg u}^k)$ ▷ Generate structured analysis
$s_u = \text{LLM}(P_u^{1:k}, a_u, D)$ ▷ Generate personalized summary
return s_u

trol the inherent differences in writing styles or content focuses between different systems. Zhang et al. (2025) uses aspect and sentiment similarity between personalized summaries and user profiles to evaluate personalization. However, their design is specific to the review domain.

3 Problem Statement

The input of personalized MDS is a document set D containing multiple documents on the same topic to be summarized. For personalization for user u, the input also contains a profile P_u containing multiple profile documents $p_u^i \in P_u$ authored by the user u. Given these inputs, the output of personalized MDS is a personalized summary s_u of the document set D that capture the individual preference of user u as expressed in their profile P_u.

4 ComPSum

In this section, we present ComPSum, our framework for personalized MDS. We describe how it generates a structured analysis of a user and then how it generates a personalized summary using the analysis. Its pseudo code is in Algo. 1.

Generating structured analysis: ComPSum uses an LLM to generate a structured analysis a_u of user u along two dimensions: writing style and content focus. To generate the structured analysis, following Salemi et al. (2024), ComPSum first retrieves the top k documents from user u's profiles P_u using a retrieval model $\mathcal{R}$. For the retrieval query, ComPSum uses the concatenation of all documents of document set D as the query and retrieves k profile documents most similar to it: $\mathcal{R}(D, P_u, k)$.

However, generating the structured analysis a_u only based on the retrieved profile documents can make the analysis focus on general features of the user u but ignore finer differences compared to other users. To address this issue, comparing documents on the same topic written by other users can be useful. Therefore, for each retrieved profile document $p_u^i \in \mathcal{R}(D, P_u, k)$, ComPSum identifies a set of documents, $C_{\neg u}^{p_i}$, that belong to the same document set as p_u^i (and hence are on the same topic) but are written by different users. ComPSum then retrieves one comparative document $p_{\neg u}^i \in C_{\neg u}^{p_i}$ that is most dissimilar to p_u^i using the retrieval model $\mathcal{R}$. By comparing every pair of profile document p_i and its comparative document $p_{\neg u}^i$, ComPSum can focus on distinguishing features that set the user apart from others. So, it instructs an LLM to compare the pair and generate the structured analysis a_u consisting of two distinct components, each highlighting the differences in writing style and content focus that distinguish user u from others:

$$a_u = LLM(p_u^1, p_{\neg u}^1, ..., p_u^k, p_{\neg u}^k) \quad (1)$$

The prompt for generating structured analysis is shown in Tab. 11. An example structured analysis of is in Tab. 1. The structured analysis captures finer differences between users.

Generating personalized summary: Using the structured analysis a_u and retrieved profile documents $p_u^{1:k}$, ComPSum generates a personalized summary s_u for the document set D:

$$s_u = LLM(p_u^1, ..., p_u^k, a_u, D) \quad (2)$$

Specifically, ComPSum instructs the LLM to generate a summary s_u that cover the main information from the document set D while mimics the writing style and content focus based on retrieved profile documents $p_u^{1:k}$ and structured analysis a_u. To prevent information leakage, ComPSum explicitly instructs the LLM to only include contents from the document set D but not from the profile texts. The prompt for generating summaries is in Tab. 12.

5 AuthorMap

In this section, we describe a fine-grained reference-free evaluation framework, AuthorMap. It evaluates personalization along two key dimensions: writing style and content focus. The underlying idea behind AuthorMap is that if the generated summary is well-personalized, it will be possible to infer the user's preferences from the summary and use them for the task of authorship attribution. AuthorMap considers two profiles, P_{u_1} and P_{u_2}, and two personalized summaries, s_{u_1} and s_{u_2}, of the same input document set, D, generated for user u_1 and u_2 respectively. AuthorMap then evaluates whether each profile P_{u_*} can be correctly attributed to its user based on personalized summaries.

Algorithm 2 Pseudo Code for AuthorMap

1: **Input**: personalized summaries s_{u_1}, s_{u_2}, user profile P_{u_1}, P_{u_2}, number of retrieved profiles n
2: $P^s_{u_1} = \mathcal{R}(s_{u_1} \circ s_{u_2}, P_{u_1}, n)$
3: $P^s_{u_2} = \mathcal{R}(s_{u_1} \circ s_{u_2}, P_{u_2}, n)$
4: count = 0 ▷ count number of correct judgements
5: **for** $(s_a, s_b) \in \{(s_{u_1}, s_{u_2}), (s_{u_2}, s_{u_1})\}$ **do**
6: count += $(u_1 == LLM_{judge}(P^s_{u_1}, s_a, s_b))$
7: count += $(u_2 == LLM_{judge}(P^s_{u_2}, s_a, s_b))$
8: **end for** ▷ Mitigate position bias of LLM judge
9: **Return** (count>=3) ▷ Aggregate judgments

AuthorMap evaluate writing style and content focus separately as they are independent dimensions (Jafaritazehjani et al., 2020). However, user preferences for these dimensions may vary across topics. To account for this, AuthorMap retrieves a subset of profile documents for each user that are topically related to the summaries, $P^s_{u_*}$. Specifically, AuthorMap retrieves the top n profile documents from each user's profile P_{u_*} that are most similar to the concatenation s_{u_1} and s_{u_2} using the retrieval model $\mathcal{R}$: $\mathcal{R}(s_{u_1} \circ s_{u_2}, P_{u_*}, n)$, where $\circ$ denotes concatenation (lines 2 and 3 of Algo. 2). Using the concatenated summaries avoids bias toward either summary. For each retrieved profile $P^s_{u_*}$, AuthorMap instructs a judge LLM to predict which user, u_1 or u_2, is more likely to be the author of the retrieved profile $P^s_{u_*}$ given their personalized summaries, s_{u_1} and s_{u_2} for the given dimension:

$$\hat{u}_1 = LLM_{judge}(P^s_{u_1}, s_{u_1}, s_{u_2}) \quad (3)$$

$$\hat{u}_2 = LLM_{judge}(P^s_{u_2}, s_{u_1}, s_{u_2}) \quad (4)$$

where $\hat{u}_*$, the predicted author of the profile $P^s_{u_*}$, can be u_1, u_2 or tie. If the summary s_{u_1} is well personalized for user u_1, the LLM judge will be able to attribute the profile $P^s_{u_1}$ to user u_1. The same should also apply to user u_2.

To mitigate positional bias (Huang et al., 2023), AuthorMap performs such prediction twice with different orders of s_{u_1} and s_{u_2}, which results in four predictions in total (line 5 of Algo. 2). AuthorMap evaluates the personalization capability of a personalized MDS system as the percentage of samples where the judge LLM correctly predicts the author of the retrieved profile in the majority of four predictions (line 9 of Algo. 2). A larger percentage value indicates better personalization capability on the corresponding dimension of the personalized MDS system. The prompts for AuthorMap are shown in Tab. 13 and 14.

6 PerMSum

In this section, we describe the construction of the PerMSum dataset. We first describe how PerMSum obtains document sets labeled with a user (its author). We then describe how to select samples from PerMSum for the evaluation using AuthorMap.

Obtaining Document Set with User Label: For the news domain, PerMSum uses the All the News dataset[2], which includes metadata of author names, publishing media, and publishing dates. However, directly using news articles for personalized MDS has two issues. First, news articles with explicit mentions of authors or outlets (e.g., "X reports in New York") can be undesirable shortcuts for personalization. To address this, PerMSum removes all sentences with author names or publishing media. Second, news with more than three authors or media organizations might not reflect the preference of their individual authors. Therefore, PerMSum excludes them from being used as profile documents. After that, PerMSum clusters documents into document sets based on token overlap, named entities, and publishing dates as Liu et al. (2022) so that each document set is about the same event.

For the review domain, PerMSum uses Amazon book reviews (Ni et al., 2019) following Wang et al. (2023). PerMSum obtains document sets from reviews following Bražinskas et al. (2019) by only keeping English reviews of 50 to 150 words. PerMSum also filters out reviews written by users who write more than 200 reviews to avoid overrepresentation of certain authors.

For both domains, PerMSum only considers users that write at least 10 documents. When generating personalized summaries of a document set for a user u, PerMSum uses all documents written by the same user u that are not part of the document set as the user profile P_u. PerMSum splits the users into training, validation, and test sets using the user split (Salemi et al., 2024). Document sets are also split to prevent information leakage, ensuring no overlap across splits. More details of data curation process are in App. A.3.

Sample selection for AuthorMap: Each sample for AuthorMap consists of a document set D and two personalized summaries for users u_1 and u_2. However, randomly sampling two users for evaluation can face sparsity in personalization (Dong et al., 2024). For example, for a user whose profile

[2]https://components.one/datasets/all-the-news-2-news-articles-dataset/

documents are all entertainment news, it can be difficult to get enough information from his profile for generating a personalized summary of international news. To address this issue, for each document set D, PerMSum only selects user pairs u_1 and u_2 who write documents belonging to the document set D. To prevent the system from copying user-authored contents, we remove all documents written by the users u_1 or u_2 from the input document set. The generated pairs of personalized summaries are then used for evaluation using AuthorMap. To prevent users from dominating the evaluation, PerMSum limits each user to appear in at most 100 samples. Statistics of PerMSum are in Tab. 2.

7 Experiments

In this section, we describe experiments on AuthorMap and ComPSum.

7.1 Implementation Details

For AuthorMap, we use Llama3.3-70b-Instruct (AI@Meta, 2024) as the LLM judge. AuthorMap can also use other LLMs as the LLM judge. We perform evaluation using Gemma-3-27b-it (Team, 2025a) in App. A.7 and find that AuthorMap is independent of the choice of the LLM judge. For evaluation, AuthorMap retrieves $n = 5$ profile documents using BM25 (Robertson et al., 1995). Motivated by Huang et al. (2024), the LLM is instructed to additionally focus on linguistic features when evaluating writing style.

For ComPSum, we experiment with Llama3.1-8b-Instruct (AI@Meta, 2024), Qwen2.5-14B-Instruct (Yang et al., 2024), Llama3.3-70b-Instruct, and Qwen3-30B-A3B-Instruct-2507 (Team, 2025b). For personalization, ComPSum retrieves $k = 5$ profile documents and comparative documents using BM25. The token limit for personalized summaries is 100 words. All LLMs used in experiments use the default sampling parameter. Hyperparameters and prompts are tuned on the validation set. More implementation details are in App. A.1.

7.2 Automatic Evaluation of AuthorMap

In this section, we perform large-scale automatic evaluation to evaluate the accuracy of AuthorMap. Since there is no existing reference personalized summaries, we use documents from the same document set but written by two users of interest for this evaluation. To mimic the setting of AuthorMap, the documents are truncated to 100 words, matching the length of a typical generated personalized summary. We report the accuracy of AuthorMap on human-written documents in the test sets of PerMSum. For comparison, we also report the accuracy of AuthorMap when profile documents P_{u_*} are not retrieved but randomly sampled from user profiles, following the setting of Wang et al. (2023). The results are in Tab. 3.

As in Tab. 3, AuthorMap shows reasonable accuracy on the documents, showing that it can reliably evaluate personalization. Besides, AuthorMap outperforms its variant without retrieval, showing that retrieval is useful to capture the varying preferences of users on different topics.

Writing style and content focus are independent dimensions (Salemi et al., 2025) and AuthorMap should evaluate them independently. To evaluate this, we use accuracy of AuthorMap on pairs of paraphrased documents with altered writing styles. Specifically, we evaluate two types of pairs : d_{u_1} vs $para(d_{u_2})$, and d_{u_1} vs $para(d_{u_1})$, where d_{u_*} denotes the original document written by user u_*, $para(d_{u_*})$ denotes a paraphrased document following the writing style of the other user. The first type of pairs isolates content focus under similar writing styles, while the second isolates writing style under similar content focus. We use Llama3.3-70b-Instruct for paraphrasing (prompt in App. A.4). Even though the paraphrasing may not be perfect as it may not completely mimic the given writing style, it is sufficient to test the independence based on changes in AuthorMap's accuracies. If AuthorMap evaluates writing style and content focus independently, AuthorMap should show higher accuracy on d_{u_1} vs $para(d_{u_2})$ when evaluating content focus since the document pair has differ in content focus but not much in style. Conversely, when evaluating writing style, AuthorMap should show higher accuracy on d_{u_1} vs $para(d_{u_1})$ since the document pairs differ in style but not much in content focus. The results are shown in Tab. 4.

From the table, we observe the desirable pattern: AuthorMap shows higher accuracy on d_{u_1} vs $para(d_{u_2})$ when evaluating content focus and higher accuracy on d_{u_1} vs $para(d_{u_1})$ when evaluating writing style. The result shows that AuthorMap can evaluate personalization based on either writing style or content focus independently.

7.3 Human Evaluation of AuthorMap

We perform a human evaluation to evaluate the correlation between AuthorMap and human judg-

	#User	#Doc. Set	#Sample	Prof. Size	Doc. Set. Size	Doc. Len.
News	828/293/296	10730/1393/1463	-/2085/2360	39.72	3-10	216.72
Review	2400/763/766	27725/1878/1795	-/2774/2757	19.77	8	86.40

Table 2: Statistics of PerMSum. We report numbers of users (#User), document sets (#Doc. Set), and evaluation samples (#Sample) for the training, validation, and test sets, respectively, separated by '/', average profile size per user, size of document sets, and document lengths.

	News		Review	
	style	content	style	content
AuthorMap	76.65	71.64	89.00	82.69
w/o retrieval	75.10	68.71	87.88	81.27

Table 3: Automatic Evaluation of AuthorMap. AuthorMap shows high accuracy on human-written documents and outperforms its variant without retrieval.

	News		Review	
	style	content	style	content
d_{u_1} vs $para(d_{u_2})$	61.34	58.76	70.34	70.58
d_{u_1} vs $para(d_{u_1})$	68.77	31.68	77.93	54.10

Table 4: Accuracy of AuthorMap on paraphrased human written documents. It shows that AuthorMap can independently evaluate writing style and content focus.

ment. Following Wang et al. (2023), each human evaluation sample contain two personalized summaries generated for different users and a set of retrieved profile documents written by one of the users. For each sample, human annotators are instructed to select which summary better aligns with the writing style or content focus expressed by the profile documents. We then calculate the accuracy of AuthorMap by measuring proportion of samples where AuthorMap and human annotation based on majority select the same summary.

We randomly sample 50 samples from each of the news and reviews domains, resulting in 100 samples in total. We use personalized summaries generated by ComPSum with Llama3.1-8b, which shows the medium-level performance in Tab. 5. Each sample is annotated by 3 annotators from Amazon MTurk. Evaluation details are in App. A.2.

Among 100 samples, the Randolph's Kappa (Randolph, 2005) between three annotators is 0.42, showing a moderate inter-annotator agreement. For the news, AuthorMap achieves an accuracy of 78% for writing styles and 76% for content focus. For the reviews, AuthorMap achieves 72% for writing style and 86% for content focus. Using bootstrap resampling, the 95% confidence intervals for the estimated accuracy are [66%, 88%] for both style and content on the news subset, and [64%, 86%] for style and [76%, 94%] for content on the review subset, which is statistically significant from random guessing. The results show that AuthorMap achieves high accuracy when compared to human annotated labels.

7.4 Evaluation of ComPSum

In this section, we evaluate the qualities of personalized summaries generated by ComPSumconsidering personalization and general qualities. For personalization, we consider writing style and content focus using AuthorMap. For general qualities, we consider factuality, which measures whether summaries only contain information supported by the input, and relevance, which measures whether summaries only include important information from document sets (Fabbri et al., 2021). We use FactScore (Min et al., 2023) for factuality and G-Eval (Liu et al., 2023) for relevance. To match the scales of other measures, we map G-Eval's score to 1-100. We report the geometirc average of the four measures as an overall score to measure the overall quality, where higher values indicate better quality.

We compare ComPSum with the following baselines: (i) General, which generates summaries without user profiles; (ii) RAG (Salemi et al., 2024), which generates a personalized summary based on retrieved profile documents; (iii) CICL (Gao and Das, 2024), which extends RAG by additionally retrieving comparative documents from other users; (iv) RAG+Summ (Li et al., 2023a), which summarizes the retrieved profile documents to guide generation of personalized summaries; (v) DPL (Qiu et al., 2025), which first generates an analysis for each retrieved profile document by comparing its comparative documents and then aggregate these analyses into a profile summary to generate the personalized summary; (vi) Rehearsal (Zhang et al., 2025), which iteratively refines a general summary using user and supervisor agents. Implementation details and computation cost analysis of these baselines are shown in App. A.5. The results on the test set of PerMSum are in Tab. 5.

	News					Review				
	style	content	fact.	rele.	overall	style	content	fact.	rele.	overall
Llama3.1-8b-Instruct										
General	43.60	36.48	97.86	98.51	62.58	42.79	38.69	98.58	93.94	62.58
RAG	53.35	49.61	98.00	96.05	70.65	52.52	54.12	98.16	92.71	71.32
CICL	53.22	48.89	97.34	95.47	70.12	54.56	55.83	96.87	88.91	71.57
RAG+Summary	54.58	50.89	98.03	97.17	71.72	54.56	57.10	97.67	92.02	72.74
DPL	53.94	47.90	97.91	96.30	70.26	54.70	59.13	97.05	87.68	72.43
Rehearsal	99.36	99.49	23.16	28.75	50.65	97.97	98.40	57.28	37.62	67.51
ComPSum	59.75	53.94	98.01	95.32	**74.07**	59.13	57.89	98.03	91.99	**74.54**
Qwen2.5-14B-Instruct										
General	44.92	40.42	98.21	98.71	64.77	42.03	42.61	97.89	93.55	63.64
RAG	52.88	48.26	98.07	96.77	70.15	48.78	51.18	97.76	91.76	68.79
CICL	51.69	48.64	97.11	96.68	69.71	54.66	55.90	96.95	88.91	71.64
RAG+Summary	52.75	49.32	98.17	97.93	70.72	54.88	59.75	97.40	90.97	73.42
DPL	50.97	46.31	98.04	97.37	68.90	50.71	57.53	96.83	89.42	70.89
Rehearsal	98.69	99.41	22.86	26.94	49.58	95.79	98.69	53.20	33.63	64.13
ComPSum	57.92	57.08	97.96	96.58	**74.78**	60.36	63.92	96.51	89.98	**76.08**
Llama3.3-70b-Instruct										
General	41.10	32.03	98.75	99.03	59.90	45.15	38.73	98.36	94.17	63.44
RAG	48.64	39.96	98.59	97.98	65.83	52.85	45.01	98.82	93.53	68.48
CICL	50.04	44.41	98.55	98.03	68.07	57.42	50.53	98.62	93.70	71.96
RAG+Summary	49.96	41.57	98.70	98.08	66.96	57.13	51.22	98.68	93.43	72.07
DPL	49.32	41.40	98.81	98.17	66.71	55.68	47.91	98.63	93.70	70.47
Rehearsal	99.49	100.00	21.74	24.90	48.17	97.53	98.55	62.03	33.81	67.00
ComPSum	53.90	45.38	98.64	98.03	**69.74**	62.11	53.18	98.62	93.76	**74.34**
Qwen3-30B-A3B-Instruct-2507										
General	43.35	33.73	97.89	99.51	61.43	46.24	35.28	98.05	97.88	62.90
RAG	53.52	47.08	97.58	99.29	70.29	58.22	47.48	98.04	96.79	71.57
CICL	53.52	46.23	97.52	99.31	69.96	58.22	49.55	97.84	96.85	72.31
RAG+Summary	56.82	50.72	97.49	99.27	72.67	66.50	57.50	97.43	96.31	77.39
DPL	55.38	50.97	97.52	99.32	72.31	61.34	55.83	97.47	95.28	75.09
Rehearsal	98.81	99.49	29.39	28.81	53.72	94.04	93.90	68.17	47.66	73.19
ComPSum	65.08	62.71	96.89	99.07	**79.11**	71.47	66.35	96.01	95.16	**81.13**

Table 5: Evaluation of ComPSum. A higher value indicates better performance. The best-performing method based on overall score is **bolded**. ComPSum shows the best performance.

From the table, we observe that ComPSum consistently outperforms all baselines in personalization quality and overall score. Additionally, ComPSum shows comparable performance on general quality and still outperforms other baselines on personalization when controlling the general quality (App. A.6). This suggests that ComPSum improves personalization without sacrificing general quality compared to other baselines. All improvements on overall scores over the second-best methods are statistically significant under paired bootstrap resampling ($p < 0.05$) (Koehn, 2004). We also evaluate ComPSum on Llama3.1-8b for additional two runs with different random seeds and show results in Tab. 10. The results show that the performance improvement of ComPSum is stable.

Although Rehearsal performs well on personalization, its performance on general qualities is pretty low. Many summaries generated by Rehearsal resemble user profiles rather than faithful summaries of input document sets. Analysis of its 'modification suggestions' shows that they often suggest adding information that only exists in user profiles but not input document sets. This can be because that neither the user agent nor the supervisor agent has access to input document sets when generating the suggestions. The finding also shows the importance of considering both personalization and general qualities during evaluation.

ComPSum also outperforms DPL, particularly in the news domain, despite both leveraging comparative documents for user analysis. This can be caused by two reasons. First, DPL is designed specifically for reviews and focus on aspects like emotional style that do not generalize to domains beyond reviews. Second, DPL generates initial analyses from single profile documents, which are insufficient to capture user preferences and subsequent aggregation cannot fully recover the lost information. Contrarily, ComPSum generates the structured analysis of a user directly conditioned on profile documents. A comparison with DPL using the same aspect as ComPSum is in Sec. 7.5.

All other methods outperform General, which does not use user profiles. The result shows that user profiles in PerMSum capture user preferences useful for personalized summarization, even though the profile texts (news or reviews) differ slightly in genre from summaries. This result is consistent with works on cross-genre authorship (Ma et al., 2025) that authors show consistent stylistic choices or content focus across genres.

We also evaluate the performance of ComPSum when using Gemma3-27b as the LLM judge. We

	Llama3.1-8b		Qwen2.5-14B		Llama3.3-70b		Qwen3-30b-a3b	
	News	Review	News	Review	News	Review	News	Review
General	67.19	67.72	67.94	67.65	61.99	64.62	65.20	66.72
RAG	71.52	72.88	71.54	72.04	69.58	72.48	71.15	73.94
CICL	71.94	74.05	71.74	72.93	70.28	72.36	71.04	74.45
RAG+Summary	73.38	75.31	73.15	75.81	70.77	74.28	72.65	78.57
DPL	72.19	74.27	71.80	73.12	71.23	72.22	73.18	76.22
Rehearsal	50.69	67.43	49.54	63.68	48.13	66.44	53.55	72.98
ComPSum	**74.90**	**76.57**	**75.50**	**77.01**	**71.95**	**77.93**	**77.92**	**81.16**

Table 6: Evaluation of ComPSum using Gemma3-27b as the LLM judge. ComPSum consistently outperforms other baselines.

report the overall scores across LLMs and benchmarks in Tab. 6. We find that ComPSum still consistently outperforms other baselines. The full results are shown in Tab. 17.

7.5 Ablation Study of ComPSum

In this section, we validate the design of ComPSum via five ablated variants: (i) **w/o comp. doc.**, which generates structured analysis of a user without comparative documents; (ii) **w/o structure**, which instructs LLM to generate a profile summary of a user instead of separate analysis of writing style and content focus; (iii) **w/ sim. comp.**, which retrieves comparative documents as documents that are most similar to profile documents; (iv) **w/ mid. comp.**, which retrieves comparative documents as documents that have medium similarity to profile documents; (v) **w/ multi. stage**, which generates the structured analysis in multiple stages, similar to DPL, but focusing on writing style and content focus instead of dimensions used by DPL. For these ablated variants, we report their overall scores in Tab. 7 which evaluates personalization and general qualities as described in Sec. 7.4 on the test set of PerMSum. The detailed scores for each dimension and implementation details are shown in App. A.8.

From the table, ComPSum outperforms w/o comp. doc. and w/o structure, showing the effectiveness of comparative documents and structured analysis. Besides, ComPSum outperforms w/ multi. stage, which shows that generating analysis based on multiple profile documents is more effective than summarizing analyses based on one profile document. Besides, the results also show that retrieving the most dissimilar documents is more effective than retrieving the most similar documents (w/ sim. comp.), documents with medium similarity (w/ mid. comp.), or retrieving from precomputed clusters (w/ multi. stage). Overall, ComPSum outperforms all of its ablated version based on average performance across LLMs and datasets. Example summaries generated by ComPSum are in App. A.10. We perform ablation study on number of retrieved documents in App. A.11 and find that retrieve 5 profile documents performs the best.

7.6 Analysis Generated by ComPSum

In this section, we examine whether comparative documents increase the diversity of structured analysis for different users. For each sample in the test set, we measure the average similarity between the structured analysis generated for two users given the same input document set. We then compare the similarity produced by ComPSum with those from its ablated variant, **w/o comp. doc.**, which generates structured analysis without using comparative documents. We report cosine similarity of structured analysis's embedding generated by gte-Qwen2-1.5B-instruct (Li et al., 2023b) in Tab. 8.

From the table, ComPSum has lower similarity score than w/o comp. doc., showing that comparative documents leads to more diverse structured analysis for different users. We also show example structured analysis for writing style generated by ComPSum and w/o comp. doc. in Tab. 9. The example shows that the structured analysis generated by ComPSum additionally includes comparison with other users, which helps the MDS system to better differentiate different users. Example analysis and retrieved comparative documents are in App. A.9.

8 Conclusion

We propose ComPSum, a personalized MDS framework. It captures the finer differences between users by comparing profile documents with documents authored by different users. We also propose AuthorMap, a reference-free fine-grained evaluation framework. We show that AuthorMapachieves reasonable accuracy using human and automatic evaluation. We also construct PerMSum, a personalized MDS dataset in the news and review domain. We evaluate the performance of ComPSum, showing

	Llama3.1-8b		Qwen2.5-14b		Llama3.3-70b		Qwen3-30b-a3b		
	News	Review	News	Review	News	Review	News	Review	Avg.
ComPSum	**74.07**	74.54	**74.78**	76.08	**69.74**	74.34	**79.11**	81.13	**75.47**
w/o comp. doc.	72.49	73.22	71.08	75.91	67.40	73.39	74.68	79.52	73.46
w/o structure	71.43	**78.43**	69.34	75.88	69.69	**76.78**	75.37	**82.10**	74.88
w/ sim. comp.	71.93	74.02	73.12	**77.04**	69.22	74.40	77.95	80.88	74.82
w/ mid.comp.	73.36	74.50	74.60	76.54	69.69	73.63	77.16	80.06	74.94
w/ multi. stage	69.87	75.15	68.62	74.77	67.14	75.33	75.28	80.43	73.33

Table 7: Ablation study of ComPSum. The best-performing method is **bolded**. All differences between the best and the second-best performing methods are statistically significant using paired bootstrap resampling ($p < 0.05$). ComPSum outperforms its ablated version.

	Llama3.1-8b		Qwen2.5-14b		
	News	Review	News	Review	Avg.
ComPSum	80.93	82.63	78.81	81.18	80.89
w/o comp. doc.	83.18	82.74	81.73	82.42	82.52

Table 8: Average similarity between structured analysis of different users. Comparative documents can make structured analysis more diverse for different users.

ComPSum: User X's writing style is characterized by a clear and concise narrative voice, often incorporating direct quotes and specific details... **Unlike other users, who may rely on sensational language or emotional appeals**, User X's tone is measured and informative, making their content feel more authoritative and trustworthy.

w/o comp. doc.: User X's writing style is characterized by a conversational tone and a focus on storytelling. They often use anecdotes and quotes from celebrities to illustrate their points, making their content feel more relatable and engaging. The text is also well-structured and easy to follow, with a clear and concise writing style...

Table 9: Structured analysis for writing style generated by ComPSum and w/o comp. doc. The structured analysis generated by ComPSum includes comparison with other users (in **bold**), which helps in better personalization.

that it outperforms strong baselines.

9 Limitation

One limitation of ComPSum is its reliance on comparable documents that share the same topic as the profile documents. Otherwise, the differences between profile documents and comparable can stem from topic difference but not individual preference differences. Therefore, ComPSum cannot be directly applied to general personalized text generation as identifying such comparable documents can be difficult for other tasks. Future work could explore methods for automatically identifying or generating comparable documents to broaden the applicability of ComPSum to more general personalized text generation.

When constructing the PerMSum dataset, we define the profile documents for a user as documents written by the user. However, for the news domain, defining the profile documents as documents clicked or liked by the user seems to be more similar to the real-world application. Unfortunately, existing publicly available news datasets with user interaction data are not suitable for personalized MDS. For example, the PENS dataset (Ao et al., 2021) lacks publishing dates of news articles, which makes it difficult to efficiently create input document sets where all documents are about the same event. MIND dataset (Wu et al., 2020) lacks both publishing dates and full texts of news articles. Therefore, we construct PerMSum using articles from All the News dataset, which is the only large-scale dataset that includes publishing dates, author information, and full article texts.

10 Ethical Consideration

The datasets we use are all publicly available. All the models used in this paper are publicly accessible. The inference and finetuning of models are performed on four Nvidia A6000 or Nvidia L40 GPUs. We do not annotate any data on our own.

We perform human evaluation experiments on Amazon Mechanical Turk. The annotators were compensated at a rate of $20 per hour. During the evaluation, human annotators were not exposed to any sensitive or explicit content.

We use LLMs to polish the writing of the paper.

11 Acknowledgment

This work was supported by the NSF grant IIS-2047232.

References

AI@Meta. 2024. Llama 3 model card.

Xiang Ao, Xiting Wang, Ling Luo, Ying Qiao, Qing He, and Xing Xie. 2021. Pens: A dataset and generic framework for personalized news headline generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 82–92.

Ilker Nadi Bozkurt, Ozgur Baghoglu, and Erkan Uyar. 2007. Authorship attribution. In *2007 22nd international symposium on computer and information sciences*, pages 1–5. IEEE.

Arthur Bražinskas, Mirella Lapata, and Ivan Titov. 2019. Unsupervised opinion summarization as copycat-review generation. *arXiv preprint arXiv:1911.02247*.

Arthur Bražinskas, Mirella Lapata, and Ivan Titov. 2020. Unsupervised opinion summarization as copycat-review generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5151–5169.

Hal Daumé III and Daniel Marcu. 2006. Bayesian query-focused summarization. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 305–312.

Yijiang River Dong, Tiancheng Hu, and Nigel Collier. 2024. Can LLM be a personalized judge? In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10126–10141, Miami, Florida, USA. Association for Computational Linguistics.

Alexander R Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409.

Alexander Richard Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir Radev. 2019. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084.

Xiang Gao and Kamalika Das. 2024. Customizing language model responses with contrastive in-context learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18039–18046.

Baixiang Huang, Canyu Chen, and Kai Shu. 2024. Can large language models identify authorship? In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 445–460.

Kung-Hsiang Huang, Philippe Laban, Alexander R Fabbri, Prafulla Kumar Choubey, Shafiq Joty, Caiming Xiong, and Chien-Sheng Wu. 2023. Embrace divergence for richer insights: A multi-document summarization benchmark and a case study on summarizing diverse information from news articles. *arXiv preprint arXiv:2309.09369*.

Somayeh Jafaritazehjani, Gwénolé Lecorvé, Damien Lolive, and John Kelleher. 2020. Style versus content: A distinction without a (learnable) difference? In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2169–2180, Barcelona, Spain (Online). International Committee on Computational Linguistics.

Joel Jang, Seungone Kim, Bill Yuchen Lin, Yizhong Wang, Jack Hessel, Luke Zettlemoyer, Hannaneh Hajishirzi, Yejin Choi, and Prithviraj Ammanabrolu. 2023. Personalized soups: Personalized large language model alignment via post-hoc parameter merging. *arXiv preprint arXiv:2310.11564*.

K Sparck Jones et al. 1999. Automatic summarizing: factors and directions. In *Advances in automatic text summarization*, 1, pages 1–12. MIT press Cambridge, MA.

Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain. Association for Computational Linguistics.

Cheng Li, Mingyang Zhang, Qiaozhu Mei, Yaqing Wang, Spurthi Amba Hombaiah, Yi Liang, and Michael Bendersky. 2023a. Teach llms to personalize–an approach inspired by writing education. *arXiv preprint arXiv:2308.07968*.

Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023b. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281*.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Y Liu, X Zhang, D Wegsman, N Beauchamp, and L Wang. 2022. Politics: Pretraining with same-story article comparison for ideology prediction and stance detection. *Findings of the Association for Computational Linguistics: NAACL 2022*.

Yang Liu, Dan Iter, Yichong Xu, Shuohang Wang, Ruochen Xu, and Chenguang Zhu. 2023. G-eval: Nlg evaluation using gpt-4 with better human alignment. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522.

Marcus Ma, Duong Minh Le, Junmo Kang, Yao Dou, John Cadigan, Dayne Freitag, Alan Ritter, and Wei Xu. 2025. Crossnews: A cross-genre authorship verification and attribution benchmark. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 24777–24785.

Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Wei Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In *EMNLP*.

Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in*

Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pages 188–197, Hong Kong, China. Association for Computational Linguistics.

Vishakh Padmakumar, Zichao Wang, David Arbour, and Jennifer Healey. 2025. Principled content selection to generate diverse and personalized multi-document summaries. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 29884–29899, Vienna, Austria. Association for Computational Linguistics.

Yilun Qiu, Xiaoyan Zhao, Yang Zhang, Yimeng Bai, Wenjie Wang, Hong Cheng, Fuli Feng, and Tat-Seng Chua. 2025. Measuring what makes you unique: Difference-aware user modeling for enhancing llm personalization. pages 21258–21277.

Justus J Randolph. 2005. Free-marginal multirater kappa (multirater k [free]): An alternative to fleiss' fixed-marginal multirater kappa. *Online submission*.

Chris Richardson, Yao Zhang, Kellen Gillespie, Sudipta Kar, Arshdeep Singh, Zeynab Raeesy, Omar Zia Khan, and Abhinav Sethy. 2023. Integrating summarization and retrieval for enhanced personalization via large language models. *arXiv preprint arXiv:2310.20081*.

Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. 1995. Okapi at trec-3. *Nist Special Publication Sp*, 109:109.

Alireza Salemi, Julian Killingback, and Hamed Zamani. 2025. Expert: Effective and explainable evaluation of personalized long-form text generation. *Preprint*, arXiv:2501.14956.

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)*, pages 7370–7392.

Chenkai Sun, Ke Yang, Revanth Gangi Reddy, Yi Fung, Hou Pong Chan, Kevin Small, ChengXiang Zhai, and Heng Ji. 2025. Persona-db: Efficient large language model personalization for response prediction with collaborative data refinement. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 281–296.

Zhaoxuan Tan, Qingkai Zeng, Yijun Tian, Zheyuan Liu, Bing Yin, and Meng Jiang. 2024. Democratizing large language models via personalized parameter-efficient fine-tuning. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6476–6491, Miami, Florida, USA. Association for Computational Linguistics.

Gemma Team. 2025a. Gemma 3.

Qwen Team. 2025b. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

Maartje Ter Hoeve, Julia Kiseleva, and Maarten de Rijke. 2022. What makes a good and useful summary? incorporating users in automatic summarization research. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 46–75.

Yaqing Wang, Jiepu Jiang, Mingyang Zhang, Cheng Li, Yi Liang, Qiaozhu Mei, and Michael Bendersky. 2023. Automated evaluation of personalized text generation using large language models. *arXiv preprint arXiv:2310.11593*.

Fangzhao Wu, Ying Qiao, Jiun-Hung Chen, Chuhan Wu, Tao Qi, Jianxun Lian, Danyang Liu, Xing Xie, Jianfeng Gao, Winnie Wu, et al. 2020. Mind: A large-scale dataset for news recommendation. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 3597–3606.

An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, and Zhihao Fan. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Yanyue Zhang, Yulan He, and Deyu Zhou. 2025. Rehearse with user: Personalized opinion summarization via role-playing based on large language models. *arXiv preprint arXiv:2503.00449*.

Zhehao Zhang, Ryan A Rossi, Branislav Kveton, Yijia Shao, Diyi Yang, Hamed Zamani, Franck Dernoncourt, Joe Barrow, Tong Yu, Sungchul Kim, et al. 2024. Personalization of large language models: A survey. *arXiv preprint arXiv:2411.00027*.

A Appendix

A.1 Implementation Details

We obtain the above prompts by first tuning the prompt and number of retrieved documents of `AuthorMap` to optimize its accuracy in automatic

	News					Review				
	style	content	fact.	rele.	overall	style	content	fact.	rele.	overall
RAG	53.08	47.34	98.14	97.11	69.95 (0.62)	54.42	53.14	98.18	92.43	71.56 (0.31)
CICL	54.04	49.37	97.27	96.08	70.66 (0.54)	57.25	56.76	97.17	89.34	72.87 (1.13)
RAG+Summary	56.51	51.64	98.08	97.24	72.63 (1.00)	56.64	57.63	97.87	92.21	73.66 (0.95)
DPL	54.77	48.18	98.01	97.2	70.81 (0.50)	57.02	57.53	97.21	89.83	73.15 (0.76)
Rehearsal	99.42	99.46	23.58	28.75	50.88 (0.47)	97.69	98.73	57.16	37.81	67.57 (0.08)
ComPSum	59.22	53.56	98.07	96.51	74.02 (0.05)	60.98	58.06	97.95	92.09	75.17 (0.55)

Table 10: Average performance of ComPSum and baselines on Llama3.1-8b-Instruction for three runs. The variance of overall score is in brackets.

You are a helpful assistant. Respond only with a JSON object including two key elements:
{
"content_analysis": a single-paragraph analysis of unique aspect and content preferences,
"style_analysis": a single-paragraph analysis of unique writing styles,
}
You are asked to analyze the distinctive features of User X's profile texts in comparison to those written by other users on similar topics. Specifically, you should first generate an analysis of unique aspect and content preferences that set User X apart. It should be written in a consecutive paragraph with less than 150 words (denoted as "content_analysis"). You should then generate an analysis of unique writing styles that set User X apart. It should be written in a consecutive paragraph with less than 150 words (denoted as "style_analysis").
****
Below are pairs of profile texts. Each pair describes the same product—one version is written by the User X, and the other by a different user.
Pair 1(written by User X): profile document 1
Pair 1(written by a different user): comparative document 1
****
Pair 2(written by User X): profile document 2
Pair 2(written by a different user): comparative document 2
****
...

Table 11: Prompt used by ComPSum to generate structured analysis.

You are requested to generate a personalized summary with less than 100 words for multiple query texts about the same product for the User X based on the User X's profile texts and the anlysis of User X's aspect and content preference as well as writing style. The personalized summary should cover the main information of the query texts while mimic the aspect and content preference and writing style of the User X.
****
A list of profile texts for the User X are shown below:
{profile_document}
****
The analysis of the User X are shown below, where "content_analysis" denotes the analysis of aspect and content preference, "style_analysis" denotes the analysis of writing style:
{structured_analysis}
****
A list of query texts to be summarized are shown below:
{input_document_set}
****
Please write a single personalized summary with less than 100 words for the query texts. The summary should only include contents from the query texts but not from the profile texts. Do not list sources of contents in the summary. Please directly output the personalized summary without any explanation. The summary should not be first person.

Table 12: Prompt used by ComPSum to generate personalized summaries.

You are a helpful assistant. Respond only with a JSON object including two key elements:
{
"analysis": Reasoning behind your answer,
"answer": query texts more likely to be written by the author of the profile texts (Query Text 1 or Query Text 2 or Tie)
}
You are given a set of profile texts with a certain author and two query texts (Query Text 1 and Query Text 2) on the same topic.
Your task is to determine which query text is more likely to be written by the author of the profile texts solely based on writing style. Specifically, first identify the differences in writing style between Query Text 1 and Query Text 2. Focus on linguistic features such as phrasal verbs, modal verbs, punctuation, rare words, affixes, quantities, humor, sarcasm, typographical errors, and misspellings. Then, detemine which query text's writing style is more closely aligned with the writing style presented in the profile texts. Please disregard the differences in contents and aspects during the comparison. If you cannot determine which query text is more likely to be written by the author of the profile texts solely based on writing style, output Tie.
The profile texts written by a certain author are shown below. The profile texts are delimited with two vertical bars: ||.
The query texts are shown below.
Query Text 1:
****
Query Text 2:

Table 13: Prompt used by AuthorMap to evaluate writing style.

You are a helpful assistant. Respond only with a JSON object including two key elements:
{
"analysis": Reasoning behind your answer,
"answer": query texts more likely to be written by the author of the profile texts (Query Text 1 or Query Text 2 or Tie)
}
You are given a set of profile texts with a certain author and two query texts (Query Text 1 and Query Text 2) on the same topic.
Your task is to determine which query text is more likely to be written by the author of the profile texts solely based on content and aspect preference.
Specifically, first identify the differences in content and aspect focus between Query Text 1 and Query Text 2. Then, detemine which query text's content and aspect focus is more closely aligned with the content and aspect presented in the profile texts. Please disregard the differences in writing style during the comparison. If you cannot determine which query text is more likely to be written by the author of the profile texts solely based on content and aspect preference, output Tie.
The profile texts written by a certain author are shown below. The profile texts are delimited with two vertical bars: ||.
The query texts are shown below.
Query Text 1:
****
Query Text 2:

Table 14: Prompt used by AuthorMap to evaluate content focus.

evaluation on the validation set as described in Sec. 7.2. We also adjust the prompts to ensure the generated judgments follow the given output format. We then keep the prompt for AuthorMap fixed and tune the prompt and number of retrieved documents of ComPSum to optimize its performance on the validation set measured by AuthorMap. We also adjust the prompts to ensure the generated structured analysis follows the given output format and the final generated summary does not have direct comparison with other users.

To match the token limit of generated summaries, all retrieved profile documents of AuthorMap and ComPSum are truncated to 100 words. This truncation not only matches the token limit of generated summaries but also prevents models from exploiting document length as a shortcut for personalization, as profile documents written by different users can vary considerably in length. For instance, in the news domain, the standard deviation of average profile lengths across users is 60 words without truncation, a large variation relative to the average length of 216 words (Table 2).

During inference, we set the temperature as 0.6 and top-p as 0.9 for Llama3.1-8b-Instruct and Llama3.3-70b-Instruct. For Qwen2.5-14B-Insturct, we set the temperature as 0.7, top-p as 0.8, top-k as 20, and repetition penalty as 1.05. Additionally, for Llama3.3-70B-Instruct, we apply 8-bit quantization to accelerate inference.

A.2 Human Evaluation of AuthorMap

Each human evaluation sample contains two personalized summaries generated for different users and a set of profile documents written by one of the users. We retrieve 5 profile documents using the concatenation of two personalized summaries as the query with BM25, just as AuthorMap. The annotators for human evaluation are recruited from Amazon Mechanical Turk. The annotators should be from English-speaking countries and have HIT Approval Rates greater than 98%. The interface of human evaluation is shown in Fig. 1.

A.3 Preprocessing detail of PerMSum

To obtain user labels, we use metadata of the All the News dataset and the Amazon dataset. For the All the News dataset, around 60 percent of documents have non-empty author information. For these documents that have valid author information and have fewer than or equal to three authors, we label the first author as the author of the document. For the Amazon dataset, since all documents have exactly one author in their metadata, we just use the author in the metadata as the author of the document.

For the news domain, PerMSum clusters the news articles into document sets based on token overlap, named entities, and publishing dates, following Liu et al. (2022). Specifically, each news article is treated as a node in a graph. If two news articles are published within two days, share at least one named entity in their titles or first three sentences, have cosine similarities based on TF-IDF embedding over 0.30, there will be a line between these articles. The news articles are then clustered based on the maximum cliques of the graph. We filter out clusters where more than three articles are written by the same author to prevent the author from having too much impact on that cluster. To control the context length, we further divide clusters that contain more than 10 news articles into smaller clusters and truncate all news articles to 300 words.

A.4 Prompt for Paraphrasing

To test whether AuthorMap can independently measure writing style and content focus, we instruct Llama3.3-70b-Instruct to generate paraphrased doc-

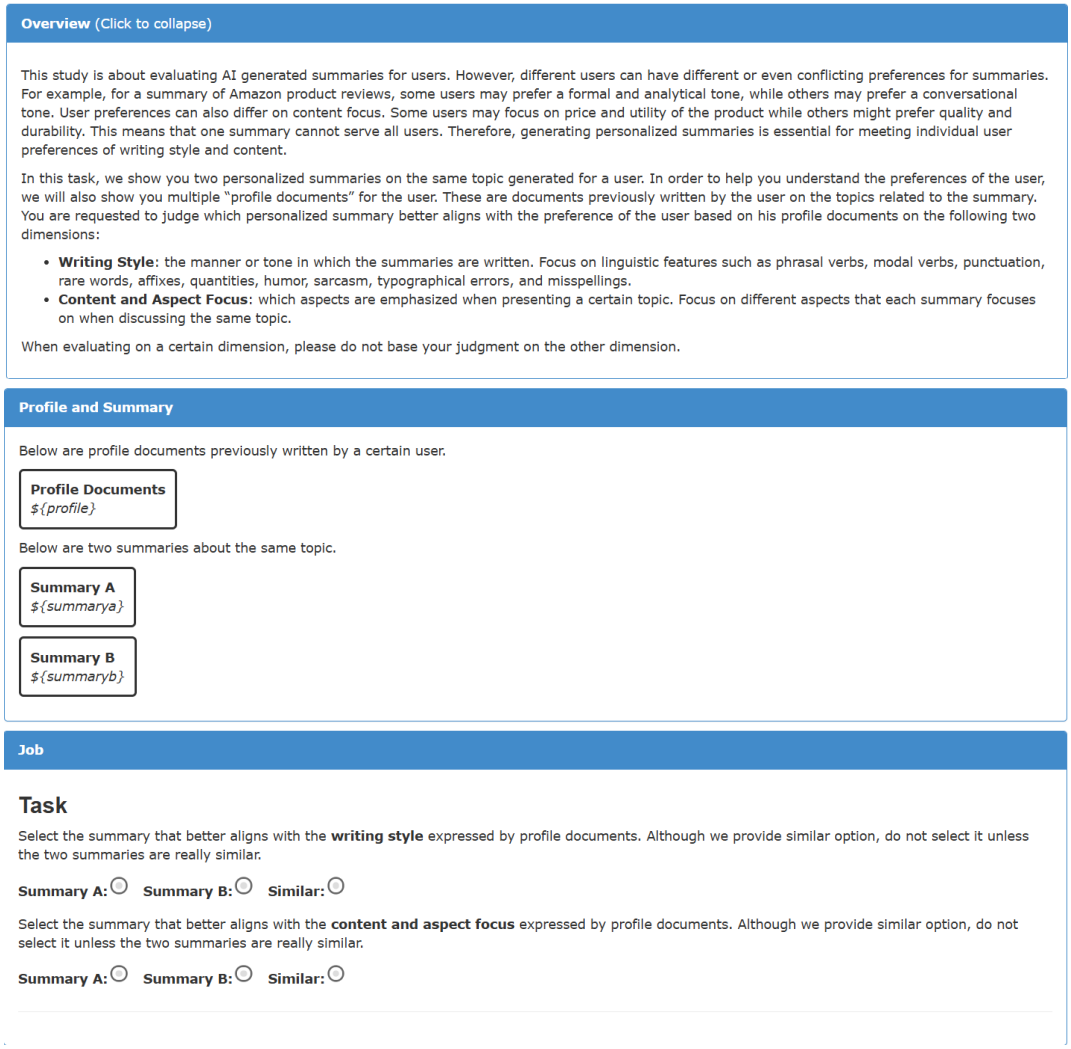


Figure 1: Interface for Human Evaluation

Given a set of profile texts with a certain writing style and a query document, please paraphrase the query document following the writing style of the profile texts. The paraphrased query document should have exatcly the same information as the original query document. The paraphrased query document should not include any additional information from the profile texts that are not mentioned by the original query document. The paraphrased query document should have the same length as the original query document.
Profile texts are shown below. The profile texts are delimited with two vertical bars: ||.
{profile_document}
Query Document:
Please directly output the paraphrased query document

Table 15: Prompt for paraphrasing.

uments that are originally written by certain users to follow the writing style of other users. We show the prompt for paraphrasing in Tab. 15

A.5 Experiment Details for Evalation of ComPSum

To evaluate factuality, we use FactScore (Min et al., 2023), which measures the proportion of atomic content units that are supported by input document sets. To evaluate relevance, we use G-Eval (Liu et al., 2023), which rates summaries based on their relevance from 1 to 5. For FactScore, we use Llama3.1-8b-Instruct to extract ACUs and judge whether ACUs are supported by input document sets. For G-Eval, we use Llama3.3-70b-Instruct to rate the relevance of summaries. To reduce the computation cost, we report the results for FactScore and G-EvaL on a subset of test set of PerMSum, which contains 25 percent of samples.

For the General baseline, we randomly sample two general summaries for each input document set without using any user profiles and use these sampled summaries for evaluation. For all other baselines, we retrieve 5 profile documents using BM25 similarly as ComPSum for a fair comparison. In the original implementation of DPL, comparable documents are retrieved based on the embeddings of user profiles. However, in PerMSum, not all documents have a valid user. Therefore, we instead retrieve comparable documents based on the embeddings of the documents themselves. For each retrieved profile document, we then select the four most dissimilar documents written by other users as the comparable documents.

We compare the computational cost of ComPSum and other baselines using Llama-3.1-8B-Instruct on a single NVIDIA RTX A6000 Ada GPU. On the news subset, the simplest RAG baseline requires approximately 30 seconds to retrieve profile documents and 15 minutes to generate summaries. Compared with RAG, CICL incurs an additional 30 seconds to retrieve comparable profile documents. RAG+Summary requires an additional 15 minutes to generate profile summaries. DPL generates an analysis for each retrieved profile document, resulting in an additional 75 minutes of computation. Rehearsal takes approximately 150 minutes to generate personalized summaries. ComPSum incurs an additional 30 seconds to retrieve comparable profile documents and 15 minutes to generate the structured user analysis. Overall, the computational cost of ComPSum is comparable to RAG+Summary, while being substantially lower than DPL and Rehearsal.

A.6 Evaluation of ComPSum Controling General Quality

To evaluate whether ComPSum improves personalization at the expense of general quality compared to other personalization baselines, we evaluate personalization of ComPSum and other baselines on subsets of test sets of PerMSumthat control for general quality. Specifically, each subset contains only samples on which ComPSum achieves perfect scores (100) for both relevance and factuality. Under this setting, competing personalization baselines can only perform worse in terms of general quality. Our hypothesis is that if ComPSum substantially sacrifices general quality to improve personalization, it would not outperform other baselines on these controlled subsets. We report the personalization performance measured by AuthorMap in Tab. 16, along with the percentage of samples included in each subset (shown in parentheses).

From the table, we observe that ComPSum still outperforms other baselines on personalization except for Rehearsal. The results show that ComPSum does not significantly sacrifices general quality to improve personalization compared to other personalization baselines.

A.7 Evaluation of ComPSum using Gemma3

In the main content, AuthorMap uses Llama3.3-70b-Instruct as the LLM judge to evaluate ComPSum. However, AuthorMap can also use other LLMs as the LLM judge. We show the evaluation results when using Gemma-3-27b-it (Team, 2025a) as the LLM judge in Tab. 17.

	News		Review		Avg.
	Style	Content	Style	Content	
			Llama3.1-8b-Instruct(0.57,0.36)		
General	44.75	36.30	43.39	39.38	40.95
RAG	52.02	46.88	52.20	52.91	51.00
CICL	53.05	48.79	55.21	52.91	52.49
RAG+Summary	54.81	49.52	56.11	57.21	54.42
DPL	52.31	47.32	54.91	57.82	53.09
Rehearsal	99.34	99.41	98.60	98.50	98.96
ComPSum	58.93	51.07	58.42	53.21	55.40
			Qwen2.5-14B-Instruct(0.60 0.25)		
General	43.26	39.94	42.17	40.90	41.57
RAG	51.52	47.28	48.66	42.17	47.41
CICL	51.66	49.19	55.01	54.30	52.54
RAG+Summary	52.79	48.91	53.88	54.87	52.61
DPL	49.82	45.73	50.35	53.88	49.95
Rehearsal	98.73	99.36	95.20	98.59	97.97
ComPSum	56.25	56.39	57.40	62.06	58.02
			Llama3.3-70b-Instruct(0.70,0.43)		
General	42.73	32.38	44.70	37.33	39.29
RAG	48.11	39.32	52.81	43.79	46.01
CICL	49.73	44.23	55.96	47.52	49.36
RAG+Summary	49.01	41.65	56.21	49.25	49.03
DPL	47.88	40.63	55.46	45.70	47.42
Rehearsal	99.32	100.00	97.21	98.04	98.64
ComPSum	52.78	43.99	59.60	49.34	51.43
			Qwen3-30B-A3B-Instruct-2507(0.44,0.19)		
General	41.18	34.01	45.14	36.19	39.13
RAG	51.45	45.16	56.81	43.19	49.15
CICL	54.26	46.51	56.23	44.16	50.29
RAG+Summary	56.78	50.29	66.34	52.92	56.58
DPL	57.29	49.20	58.17	52.72	54.35
Rehearsal	98.79	99.19	93.89	90.84	95.68
ComPSum	63.66	60.47	63.81	62.84	62.70

Table 16: Evaluation of personalization quality on the subset of PerMSumwhere general quality of ComPSum is controlled. ComPSum still outperforms other baselines except for Rehearsal.

From the table, we observe that ComPSum still performs the best, similar as the results when using Llama3.3-70b-Instruct as the LLM judge (Tab. 5). The results show that AuthorMap is independent of the choice of the LLM judge.

A.8 Experiment Details for Ablation Study

For ComPSum w/o structure, we instruct the model to generate a profile summary for profile documents and their corresponding comparative documents without specifying that model should focus on the writing style or content focus. For a fair comparison, the token limit of the profile summary is the same as the structured analysis of ComPSum. For ComPSum w/multi. stage, we first instruct LLMs to generate a structured analysis of writing style and content focus for each retrieved profile document by comparing it with four comparable documents retrieved similarly as described in App. A.5. We then instruct LLMs to synthesize all structured analysis for each profile document into a final structured analysis. For a fair comparison, the token limit of the profile summary is the same as the structured analysis of ComPSum. For ComPSum w/ sim. comp. retrieves the The full results of ablation study are reported in Tab. 18. To reduce the computation cost, we report the results for FactScore and G-Eval on subset of test set of PerMSum, which contains 25 percent of samples.

A.9 Example Analysis Generated by ComPSum

In this section, we show additional examples of structured analysis generated by ComPSum and its ablation variant without comparative documents for writing style and content. We show the example analysis for writing style in Fig. 2. We show the example analysis for content in Fig. 3. We show the example comparative documents retrieved for generating structured analysis in Tab. 21 and 20.

A.10 Example Summary Generated by ComPSum

In this section, we show example summaries generated by ComPSumand its ablation variant without comparative documents. We compare two summaries personalized for different users for the same input document set and show them in Fig. 4. We observe that summaries generated by ComPSum show more diverse styles and content focuses.

A.11 Ablation Study on Number of Retrieved Docuemnts

In the main content, ComPSum retrieves $m = 5$ profile documents for generating personalized summaries following Salemi et al. (2024). To examine the impact of the number of retrieved documents, we perform an ablation study on the number of retrieved documents and report the results in Tab. 19. From the table, we can observe that ComPSum performs the best when 5 documents are retrieved.

A.12 Failure Analysis of ComPSum

The most common failures of ComPSum are that generated personalized summaries can contain contents only presented in user profiles, but not in input documents. For example, for some input documents about Hulu ending its free service, a personalized summary generated by ComPSum can be 'Roku's public offering raised $252 million, valuing the company at $2 billion, while it competes against tech giants like Apple, Amazon, and Google in the streaming market. Meanwhile, YouTube announced a streaming TV service, offering live-streams and online video content for $35 a month, covering six accounts. Netflix expanded to 130 new countries, tripling its distribution. In

ComPSum	
User 1: User X's writing style is characterized by a conversational tone, often incorporating direct quotes from celebrities and using a more informal language. This approach makes the content feel more approachable and relatable. In contrast, other users tend to write in a more formal and objective tone, focusing on delivering factual information without adding a personal touch. User X also tends to use a more narrative-driven approach, weaving together multiple sources and ideas to create a cohesive story, **which sets them apart from other writers who may rely on a more straightforward reporting style.**	**User 2:** User X's writing style is characterized by a clear and concise narrative voice, often incorporating direct quotes and specific details to support their points. Their texts tend to be well-structured and easy to follow, with a focus on conveying complex information in an accessible way. **Unlike other users, who may rely on sensational language or emotional appeals**, User X's tone is measured and informative, making their content feel more authoritative and trustworthy. This style suggests that User X prioritizes clarity and accuracy in their writing, setting **them apart from users who may prioritize other factors like entertainment value or clickbait headlines**.
w\o comp. doc.	
User 1: The writing style of User X appears to be conversational and engaging, with a focus on storytelling and anecdotal evidence. The use of quotes from celebrities and behind-the-scenes details adds a sense of intimacy and immediacy to the articles. The tone is generally lighthearted and celebratory, particularly in the pieces about celebrity collaborations and magazine covers. However, when discussing more serious topics like gun violence and aging, the tone shifts to one of gravity and concern. Overall, the writing style is approachable and accessible, suggesting that User X aims to inform and entertain their audience	**User 2:** User X's writing style is characterized by a conversational tone and a focus on storytelling. They often use anecdotes and quotes from celebrities to illustrate their points, making their content feel more relatable and engaging. The text is also well-structured and easy to follow, with a clear and concise writing style. User X tends to use a more informal tone, often incorporating colloquial expressions and contractions, which creates a sense of familiarity with the reader. Overall, the writing style is approachable and accessible, making it easy for readers to stay engaged and interested in the content.

Figure 2: Example analysis of writing style generated by ComPSum and its ablation variant without comparative documents. The structured analysis generated by ComPSum additionally includes comparison with other users (in **bold**), which helps in better personalization.

ComPSum	
User 1: User X tends to focus on the positive and empowering aspects of celebrity news, often highlighting the personal growth and self-acceptance of public figures, such as Miley Cyrus embracing her past, Kim Kardashian West's loving family moments, and Cameron Diaz's openness about aging. **In contrast, other users often focus on controversy, drama, or sensationalized aspects of celebrity news.** User X also tends to emphasize the intersection of celebrity culture and social issues, such as gun violence and body positivity, in a way that is uplifting and inclusive.	**User 2:** User X's profile texts stand out for their in-depth analysis of current events, particularly in the realms of feminism and social justice. They delve into the nuances of the topics, often referencing specific quotes and details from the individuals involved. This level of depth and specificity suggests that User X values accuracy and context when discussing complex issues. Additionally, their texts often highlight the personal stories and experiences of individuals, such as Miley Cyrus and Bella Thorne, to illustrate the human impact of these issues. This approach creates a more empathetic and relatable tone**, setting User X apart from other users who may focus on more sensational or surface-level aspects of the news**.
w\o comp. doc.	
User 1: The content preferences of User X seem to revolve around entertainment, celebrity culture, and social issues. The topics of the provided texts include Miley Cyrus's past controversy, the Kardashian-Jenner family, and celebrities speaking out on social issues. The content also touches on fashion, with articles about celebrity appearances in Calvin Klein campaigns and magazine covers. Additionally, there is a focus on female empowerment and self-expression, as seen in the discussion of Miley Cyrus's unapologetic attitude and Cameron Diaz's openness about aging. This suggests that User X is interested in celebrity news, social justice, and feminist themes.	**User 2:** The content preferences of User X appear to be centered around social justice, feminism, and celebrity culture. The user frequently discusses topics such as body positivity, sexual harassment, and assault, often using high-profile figures like Miley Cyrus and Bella Thorne as examples. User X also seems to be interested in news and updates on social movements like #TimesUp and the prosecution of Harvey Weinstein. Additionally, they have a strong interest in celebrity relationships and personal stories, as evident from their coverage of Miley Cyrus's marriage and new music releases. Overall, User X's content preferences suggest a focus on current events, social issues, and celebrity culture.

Figure 3: Example analysis of content focus generated by ComPSum and its ablation variant without comparative documents. The structured analysis generated by ComPSum additionally includes comparison with other users (in **bold**), which helps in better personalization.

	News					Review				
	style	content	fact.	rele.	overall	style	content	fact.	rele.	overall
Llama3.1-8b-Instruct										
General	49.36	42.84	97.86	98.51	67.19	49.98	45.43	98.58	93.94	67.72
RAG	54.45	51.06	98.00	96.05	71.52	56.12	55.24	98.16	92.71	72.88
CICL	55.13	52.29	97.34	95.47	71.94	59.33	58.83	96.87	88.91	74.05
RAG+Summary	56.91	53.50	98.03	97.17	73.38	60.07	59.58	97.67	92.02	75.31
DPL	54.87	52.50	97.91	96.30	72.19	59.69	59.91	97.05	87.68	74.27
Rehearsal	99.66	99.49	23.16	28.75	50.69	97.97	97.96	57.28	37.62	67.43
ComPSum	58.81	57.29	98.01	95.32	**74.90**	62.01	61.46	98.03	91.99	**76.57**
Qwen2.5-14B-Instruct										
General	49.60	44.32	98.21	98.71	67.94	50.26	45.52	97.89	93.55	67.65
RAG	53.05	52.03	98.07	96.77	71.54	54.54	55.04	97.76	91.76	72.04
CICL	55.38	50.95	97.11	96.68	71.74	57.40	57.17	96.95	88.91	72.93
RAG+Summary	56.91	52.33	98.17	97.93	73.15	60.71	61.42	97.40	90.97	75.81
DPL	54.87	50.74	98.04	97.37	71.80	57.46	57.46	96.83	89.42	73.12
Rehearsal	98.81	98.98	22.86	26.94	49.54	94.34	97.45	53.20	33.63	63.68
ComPSum	59.07	58.14	97.96	96.58	**75.50**	63.21	64.07	96.51	89.98	**77.01**
Llama3.3-70b-Instruct										
General	45.34	33.31	98.75	99.03	61.99	47.67	39.48	98.36	94.17	64.62
RAG	53.82	45.08	98.59	97.98	69.58	56.30	53.26	98.76	93.16	72.48
CICL	52.47	48.14	98.55	98.03	70.28	56.26	53.19	98.70	92.79	72.36
RAG+Summary	54.41	47.63	98.70	98.08	70.77	57.64	57.27	98.68	93.43	74.28
DPL	55.52	47.80	98.81	98.17	71.23	55.81	53.15	98.50	93.10	72.22
Rehearsal	99.32	99.83	21.74	24.90	48.13	95.67	97.11	62.03	33.81	66.44
ComPSum	56.78	48.81	98.64	98.03	**71.95**	67.69	59.14	98.70	93.35	**77.93**
Qwen3-30b-A3b-Instruct-2507										
General	48.39	38.35	97.89	99.51	65.20	49.25	41.93	98.05	97.88	66.72
RAG	53.43	49.49	97.58	99.29	71.15	55.55	56.69	98.04	96.79	73.94
CICL	52.20	50.38	97.52	99.31	71.04	59.23	54.73	97.84	96.85	74.45
RAG+Summary	55.04	52.29	97.49	99.27	72.65	63.68	63.78	97.43	96.31	78.57
DPL	57.25	51.74	97.52	99.32	73.18	59.58	61.02	97.47	95.28	76.22
Rehearsal	98.47	98.64	29.39	28.81	53.55	91.07	95.85	68.17	47.66	72.98
ComPSum	63.43	60.55	96.89	99.07	**77.92**	68.87	68.97	96.01	95.16	**81.16**

Table 17: Evaluation results of ComPSum when using Gemma-3-27b-it as the LLM judge. A higher value indicates better performance. The best-performing method based on overall score is **bolded**. The evaluation results using Gemma-3-27b-it is consistent with evaluation results using Llama3.3-70b-Instruct in Tab. 5.

a related development, Hulu will end its free, ad-supported service and shift its content to Yahoo View, a new platform offering free TV episodes with ads, while Hulu focuses on subscription tiers and a live TV service.' Only the last sentence in the summary is related to the event. We checked the unrelated sentences and found that they are actually from the user profile. To mitigate this issue, we instruct models to only include contents from input documents but not from the profile texts, as shown in Fig. 12.

	News					Review				
	style	content	fact.	rele.	overall	style	content	fact.	rele.	overall
Llama3.1-8b-Instruct										
ComPSum	59.75	53.94	98.01	95.32	**74.07**	59.13	57.89	98.03	91.99	74.54
w/o comp. doc.	57.03	50.69	98.16	97.29	72.49	56.08	57.21	97.07	92.28	73.22
w/o structure	58.64	47.33	98.08	95.62	71.43	68.35	67.26	95.58	86.11	**78.43**
w/ sim. comp.	59.75	47.29	97.95	96.71	71.93	58.51	57.17	97.74	91.80	74.02
w/ mid.comp.	57.68	53.86	97.30	95.80	73.36	60.09	56.93	98.05	91.83	74.50
w/ multi. stage	55.51	45.42	98.16	96.31	69.87	59.71	59.93	97.57	91.34	75.15
Qwen2.5-14B-Instruct										
ComPSum	57.92	57.08	97.96	96.58	**74.78**	60.36	63.92	96.51	89.98	76.08
w/o comp. doc.	55.51	48.73	98.03	96.25	71.08	60.22	61.74	97.98	91.17	75.91
w/o structure	53.56	45.42	97.84	97.12	69.34	60.76	63.19	96.59	89.39	75.88
w/ sim. comp.	57.92	51.44	97.83	98.05	73.12	63.23	64.50	96.15	89.83	**77.04**
w/ mid.comp.	58.53	54.97	98.01	98.24	74.60	63.50	61.43	96.92	90.80	76.54
w/ multi. stage	52.16	44.62	98.23	96.98	68.62	59.75	59.75	96.73	90.50	74.77
Llama3.3-70b-Instruct										
ComPSum	53.90	45.38	98.64	98.03	**69.74**	62.11	53.18	98.62	93.76	74.34
w/o comp. doc.	51.53	41.06	98.78	98.73	67.40	59.46	52.85	98.72	93.50	73.39
w/o structure	53.81	44.96	98.86	98.59	69.69	66.64	57.50	98.44	92.14	**76.78**
w/ sim. comp.	53.90	43.69	98.69	98.81	69.22	62.40	52.96	98.81	93.83	74.40
w/ mid.comp.	52.99	45.59	98.76	98.83	69.69	60.38	52.64	98.80	93.62	73.63
w/ multi. stage	50.21	41.53	98.72	98.75	67.14	63.48	55.03	98.45	93.63	75.33
Qwen3-30B-A3B-Instruct-2507										
ComPSum	65.08	62.71	96.89	99.07	**79.11**	71.47	66.35	96.01	95.16	81.13
w/o comp. doc.	61.61	52.88	96.24	99.21	74.68	69.58	62.69	95.99	95.49	79.52
w/o structure	59.92	56.36	96.33	99.22	75.37	72.05	69.44	95.74	94.85	**82.10**
w/ sim. comp.	63.22	61.36	95.91	99.25	77.95	70.82	66.64	95.72	94.71	80.88
w/ mid.comp.	63.90	57.88	96.91	98.92	77.16	69.60	65.28	95.68	94.51	80.06
w/ multi. stage	58.69	57.46	96.04	99.18	75.28	69.51	65.99	95.69	95.33	80.43

Table 18: Full ablation results of ComPSum. A higher value indicates better performance. The best-performing method based on overall score is **bolded**.

	News					Review				
	style	content	fact.	rele.	overall	style	content	fact.	rele.	overall
Llama3.1-8b-Instruct										
ComPSum	59.75	53.94	98.01	95.32	**74.07**	59.13	57.89	98.03	91.99	**74.54**
ComPSum(k=10)	55.17	44.83	98.18	96.71	70.12	55.17	51.07	98.02	92.06	72.36
ComPSum(k=2)	58.69	50.25	97.54	95.56	72.37	59.02	55.10	97.91	92.24	74.28

Table 19: Ablation results of ComPSum with different number of retrieved documents. A higher value indicates better performance. The best-performing method based on overall score is **bolded**.

Pair 1(written by User X): I've been a fan of this series since it began, especially due to the conflict between his adopted religion and his job. As the protagonist has drifted from this conflict to his estranged family and, now, his 34;partner 34; it's become more and more a formula spy thriller.
With this book I end my purchases, mostly because it's now a serial, if I can't read a complete story in one book I'm done. Especially as this reads like a hurry-up book, almost like the author, on Vince's passing, wanted to get a Mitch Rapp like book out there.
Sorry Alex, we had a fairly enjoyable run but it's now complete.
Pair 1(written by a different user): He has a good, writing style which captures the reader. His story was interesting, and somewhat different from other writers who have these type of story lines. David Baldacci gets into a pattern and the writer begins to figure out what will happen. Berenson still is a writer who is fresh enough for me to keep the book interesting.
Pair 2(written by User X): While I'm only half way thru this marvelous little book, I've enjoyed every page. Our protagonist is one of the most appealing characters ever.
I've been a reader of mostly SciFi/Fantasy and the occasional spy/adventure novels, but these genres have tended ever more into a dark/depressing/bleakness.
I may be, therefore, over-rating this book as it stands in total opposition to my regular reading.
Well, no, this is a wonderful book, reader satisfaction up there with "The Art of Racing In the Rain"! The highest praise I can offer.
Pair 2(written by a different user): Making his fiction debut at age seventy, Alan Bradley has crafted a charming novel that is sure to be one of the most loved mysteries of the year. The Sweetness at the Bottom of the Pie introduces a delightful, intrepid, acid-tongued new heroine to the genre.

Flavia de Luce, an eleven-year-old girl living in 1950s England, is an aspiring chemist, a lover of poisons and a terror to her older sisters. When she discovers a dead body in the garden of her family's estate, she adds another hat to her repertoire: detective.

The Sweetness at the Bottom of the Pie has a wonderfully entertaining plot, but it is Flavia that truly makes it shine. As this is only the first in what is projected to be a six-book series, readers are sure to enjoy the delights of Bradley and his creation for some time.

Table 20: Example profile documents and their comparative documents in reviews.

ComPSum	
User 1: It's a double-edged sword for Julian Assange - a UN ruling says he's been unlawfully detained, but the UK and Sweden aren't budging. The UN panel's decision is a "vindication" for Assange, but a British government spokesman called it "ridiculous." Assange has been living in the Ecuadorian embassy since 2012, avoiding arrest and extradition to Sweden over rape allegations. He's hailed the ruling as a "victory of historical importance," but the UK has made it clear they'll arrest him if he leaves the embassy. It's a stalemate that's left Assange stuck in limbo.	**User 2:** A United Nations panel has ruled that Julian Assange has been arbitrarily detained at the Ecuadorian embassy in London since his arrest in 2010 over allegations of sex crimes. The UN Working Group on Arbitrary Detention called on Sweden and the UK to let Assange move freely and ensure his safety, and to provide compensation. However, British authorities have stated they will arrest and extradite Assange to Sweden if he leaves the embassy. Assange has hailed the decision as a "vindication" and a "victory of historical importance."
w\o comp. doc.	
User 1: A United Nations panel has ruled that Julian Assange has been arbitrarily detained in the Ecuadorian embassy in London since 2010. The UN Working Group on Arbitrary Detention called on Sweden and the UK to respect Assange's rights, including his freedom of movement and physical integrity. The ruling is not legally binding, but Assange sees it as a "vindication" and a "victory of historical importance." The UK and Sweden have rejected the decision, with the UK's Foreign Secretary calling it "ridiculous." Assange remains in the embassy, where he has been living for over three years.	**User 2:** A United Nations panel has ruled that WikiLeaks founder Julian Assange has been arbitrarily detained by Sweden and the UK, with the decision calling for his release and compensation. Assange, who has been living in the Ecuadorian embassy in London since 2012, hailed the ruling as a "vindication" and a "victory of historical importance." However, the UK and Swedish governments have rejected the ruling, with the British Foreign Secretary describing it as "ridiculous" and stating that Assange will be arrested if he leaves the embassy.

Figure 4: Example summaries generated by ComPSum and its ablation variant without comparative documents.

Pair 1(written by User X): Update: Manhattan District Attorney Cyrus R. Vance, Jr. announced an additional set of charges against Harvey Weinstein on Monday, July 2. The Grand Jury has charged the disgraced producer with an additional count of Criminal Sexual Act in the First Degree for a forcible sexual act against a third woman in 2006 and two counts of Predatory Sexual Assault, a Class A-II felony. The felony "carries a minimum sentence of 10 years and a maximum sentence of life imprisonment." The original story continues below. On Tuesday, June 5, Harvey Weinstein pled not guilty to rape and criminal sex act charges in New York City, the Associated Press reports.
Pair 1(written by a different user): The Manhattan district attorney on Monday announced new criminal charges against film producer Harvey Weinstein, accusing him of sexually assaulting a third woman in addition to the two mentioned in earlier charges. The new indictment accuses Weinstein of forcibly performing oral sex on a woman in July 2006. The new charges, which include predatory sexual assault, carry a sentence of 10 years to life in prison, according to the office of District Attorney Cyrus Vance. Weinstein, 66, once one of Hollywood's most powerful filmmakers, has pleaded not guilty to the first set of charges and has denied ever having non-consensual sex.
Pair 2(written by User X): On Tuesday, Miley Cyrus took back the apology she issued after her June 2008 photoshoot with Vanity Fair, which showed the then-15-year-old as seemingly topless under a blanket. Ten years ago, the photo elicited a flurry of controversy that, in retrospect, has not aged well. It's hard to believe that had the photo been released now the public would have treated it with the same kind of disdain. "A lot of things have changed and I think the conversation has changed a lot," Cyrus told Jimmy Kimmel on Tuesday night. This change is good and necessary, of course, because women are no longer encouraged to be ashamed of their bodies, or immediately sexualized whenever they wear something that shows their midriff.
Pair 2(written by a different user): Miley Cyrus is taking back her apology for a 2008 Vanity Fair photo shoot that many deemed too scandalous for the then-15-year-old. A decade after the singer posed for the pictures by famed photographer Annie Leibovitz, Cyrus said she had no remorse about the shoot, which included an image in which she appeared to be topless, covered only by a sheet. "IM NOT SORRY," she captioned a tweet, Sunday, featuring the front page of the April 28, 2008, edition of the New York Post with the headline "Miley's Shame" and an article about her apology. "F— YOU #10yearsago," she wrote. Cyrus — then a Disney Channel star — issued a statement exclusively to it in 2008 apologizing for the pictures, saying she was "embarrassed." "My goal in my music and my acting is always to make people happy.

Table 21: Example profile documents and their comparative documents in news.