

Exploring Persona Sentiment Sensitivity in Personalized Dialogue Generation

Yonghyun Jun and Hwanhee Lee*

Department of Artificial Intelligence, Chung-Ang University
{zgold5670, hwanheelee}@cau.ac.kr

Abstract

Personalized dialogue systems have advanced considerably with the integration of user-specific personas into large language models (LLMs). However, while LLMs can effectively generate personalized responses, the influence of persona sentiment on dialogue quality remains underexplored. In this work, we conduct a large-scale analysis of dialogues generated using a range of polarized user profiles. Our experiments reveal that dialogues involving negatively polarized users tend to overemphasize persona attributes. In contrast, positively polarized profiles yield dialogues that selectively incorporate persona information, resulting in smoother interactions. Furthermore, we find that personas with weak or neutral sentiment generally produce lower-quality dialogues. Motivated by these findings, we propose a dialogue generation approach that explicitly accounts for persona polarity by combining a turn-based generation strategy with a profile ordering mechanism and sentiment-aware prompting. Our study provides new insights into the sensitivity of LLMs to persona sentiment and offers guidance for developing more robust and nuanced personalized dialogue systems.*

1 Introduction

Personalized dialogue systems have evolved through advances in dataset construction (Li et al., 2016; Zhang, 2018), training methodologies (Kingma, 2013; Bowman et al., 2015), and studies on integrating user-specific preferences to enhance conversational abilities (Zhang, 2019; Roller, 2020). With the advent of Large Language Models (LLMs), embedding personas via system prompts (Yang et al., 2023; Wang et al., 2023b) has become effective. This progress has enabled applications ranging from personalized chatbots (Shus-

*Corresponding author.

*Our implementation is publicly available at <https://github.com/imsongpasimin/PesonaSensitivity>

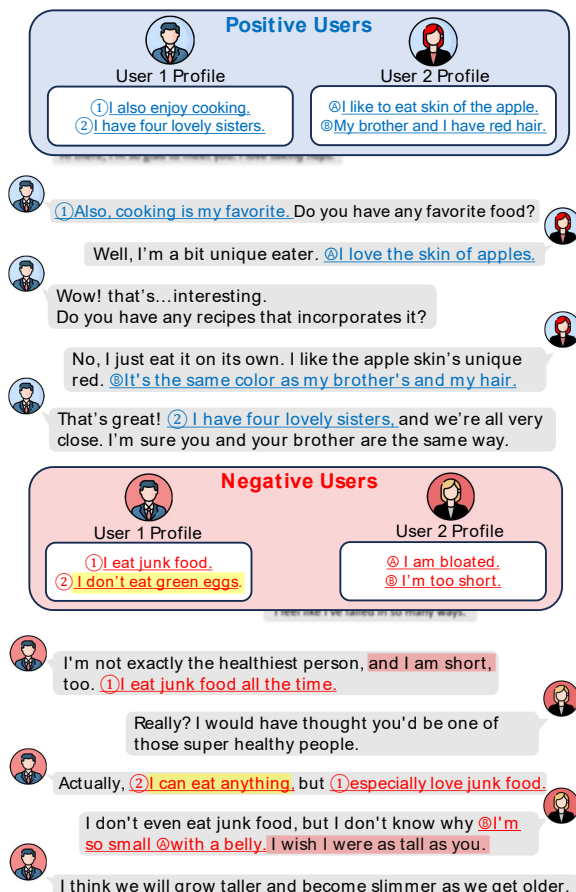


Figure 1: Examples of dialogues generated by polarized users. We indicate **Positive personas** in blue text and **Negative personas** in red text. Matching bullet points denote where user persona attributes appear in the utterance, and we highlight them in yellow if they contradict. Incoherent segments of the dialogue are marked with a red background.

ter et al., 2022; Lee et al., 2023) to enhanced performance in downstream tasks through inter-LLM conversations (Park et al., 2023) and synthetic dataset generation (Jandaghi et al., 2023). These developments underscore the growing importance of personalization in dialogue systems. However, despite these advances, recent studies have shown that LLMs are highly sensitive to contextual senti-

ment polarity (Liu et al., 2024; Wu et al., 2024b). In contrast, the impact of assigned persona sentiment on dialogue quality remains less explored. For example, Figure 1 illustrates dialogues generated by users with distinctly positive and negative personas. In dialogues driven by positive personas, the integration of persona attributes occurs in a manner that supports coherence and aligns naturally with the flow of conversation. Conversely, dialogues featuring negative personas often display contradictions and inconsistencies, as the persona attributes may be overemphasized or misaligned with the dialogue context. These findings highlight the necessity of systematically examining how variations in persona polarity influence LLM behavior.

In this paper, we construct and analyze personalized dialogues generated across a spectrum of polarized profiles. (§3) Our extensive analysis of personalized dialogues generated across a diverse spectrum of polarized profiles reveals several key insights: *1) Dialogues between negatively polarized users tend to overemphasize persona attributes*, which leads to increased contradictions and reduced overall coherence. *2) Dialogues between positively polarized users selectively incorporate persona information*, resulting in interactions that are both coherent and contextually appropriate. *3) Profiles with weak or neutral sentiment generally yield lower-quality dialogues*, as evidenced by significant differences in various metrics.

These insights underscore the need for methods that explicitly account for persona polarity in the design of personalized dialogue systems. To address this challenge, we propose a dialogue generation framework that combines a turn-based strategy, a profile ordering mechanism, and a sentiment-aware prompting. Specifically, personas exhibiting neutral or weak sentiment are placed earlier in the conversation, whereas more positive personas are positioned later. A concise prompt further instructs the model to attend carefully to negative or neutral sentiments when generating responses. These approaches promote more consistent and coherent interactions (§4). Our findings shed light on LLM sensitivity to persona sentiment and offer valuable guidance for developing more nuanced and effective personalized dialogue systems.

2 Study Design and Experimental Setup

This section explains the overall basic setup for our study. We first introduce the dataset and describe

how we construct polarized user profiles. Next, we present our research questions, detail the dialogue generation strategies, and describe the metrics used to assess the dialogues. We use four RTX 4090 GPUs to generate and evaluate all the dialogues.

2.1 Dataset

Early research on personalized dialogue systems typically used two-party conversations paired with user profiles. These profiles are commonly represented in three ways (Chen et al., 2024):

- *Descriptive sentences*: Each persona is described by a few sentences (often five) that together form a user profile (Zhang, 2018; Dinan et al., 2020).
- *Sparse key-value attributes*: Personas are specified using concise attributes (Qian et al., 2017; Gao et al., 2023).
- *User history*: Each speaker is identified by an ID that can be used to retrieve external persona information (Salemi et al., 2023).

Although several datasets exist, most trace back to PersonaChat (Zhang, 2018) and ConvAI2 (Dinan et al., 2020). In our experiments, we adopt the ConvAI2 dataset and define user personas primarily through *descriptive sentences*.

2.2 Polarized User Profile Construction

We extract polarized user personas from the ConvAI2 dataset using a sentiment classifier and then build coherent profiles by combining multiple personas while avoiding contradictions.

Polarity-Aware Persona Sampling Each persona in the ConvAI2 spans various content domains and carries a distinct sentiment polarity. We use the `distilbert-base-uncased-finetuned-sst-2-english` model (HF Canonical Model Maintainers, 2022) to classify the sentiment of each sentence, retaining only those with a confidence score above 0.99 as *polarized personas*. Consequently, as shown in Figure 2, 2,691 are classified as positive, 1,006 as negative, and 2,429 as neutral (weak), indicating that polarized personas constitute the majority (around 60% of all personas). Notably, approximately 17% of the dataset’s personas are negatively polarized—a higher proportion than one might intuitively expect. These figures suggest that polarized personas are not merely reflect a contrived scenario fabricated for the experiment. For detailed statistics, see Appendix F.




Figure 2: Distribution of persona sentences by polarity in the ConvAI2 dataset. We only map predictions with a confidence score of at least 0.99 to their respective polarity, classifying all remaining cases as neutral (weak).

Synthesizing User Profiles To create a user profile, we combine K polarized personas, where K is the number of personas we include in each profile. To avoid internal contradictions, we follow the procedure in (Jandaghi et al., 2023): starting with an initial persona, we iteratively add additional personas only if the `nli-deberta-v3-large` (He et al., 2021) detects no conflicts with the existing set. This approach yields three profile types:

- **Negative Profile:** only negative personas.
- **Positive Profile:** only positive personas.
- **Mixed Profile:** a random mix of positive and negative personas.

We generate 10K unique profiles for each type, which serve as the foundation for our dialogue generation experiments.

2.3 Research Questions

Building on the established user-polarized profiles, we structure our study around several inquiries. Considering Large Language Models (LLMs) can generate dialogues that reflect a given persona without additional tuning (Tseng et al., 2024), our study focuses on the following research questions:

RQ1. *Are LLMs Sensitive to Users’ Polarity?*

RQ2. *If so, How to Make LLMs Robust to Polarity?*

To address these questions, we evaluate a range of dialogue generation methods that capture the nuanced influences of polarized personas, using multiple evaluation metrics to comprehensively assess performance.

2.4 Dialogue Generation Strategy

Traditional personalized dialogue systems use a *turn-based* approach to generate responses sequen-

tially based on a given persona and dialogue context (Zhou et al., 2023; Tseng et al., 2024). However, recent advances in LLMs have enabled a *dual-persona joint generation* method that encodes both personas simultaneously to generate the entire dialogue in one pass (Liu et al., 2020; Xu et al., 2022; Jandaghi et al., 2023). For resource efficiency, we primarily use *dual-persona joint generation* for large-scale dialogue production and analysis, while also experimenting with *turn-based generation* to identify optimal strategies (§4. Both methods use greedy decoding (zero-temperature), with a maximum length of 4096 tokens for dual-persona joint generation and 128 tokens for turn-based generation. Prompt templates for both strategies are provided in Appendix E. After generation, we filter out outliers, such as refusal-to-answer responses or repeated utterances inside a dialogue.

2.5 Evaluation Metrics

Open-domain dialogue is inherently one-to-many, meaning a single utterance can have multiple valid responses. Moreover, ambiguous conversation topics require multiple evaluative perspectives (Wang et al., 2024; Samuel et al., 2024). In our work, we focus on two key dimensions: *consistency*—how well a dialogue reflects a user’s persona—and *coherence*—the logical, smooth flow of conversation.

Consistency We adopt the **C score** (Madotto et al., 2019), a common metric in personalized dialogue systems that calculates entailment scores for user persona-utterance pairs using an NLI model. Detailed adaptations of the **C score** for our task are provided in Appendix D.1. Since summing raw scores may lead to Simpson’s Paradox (Simpson, 1951) due to varying entailment and contradiction counts, we introduce the *Contradiction Ratio* (**Contd.**), which measures the proportion of contradictions among all cases. Additionally, we propose a novel metric called the *perplexity gap* (**P gap**). Using GPT2-large (Radford et al., 2019), we compute the difference between the perplexity of a dialogue D and the perplexity of D conditioned on two user profiles U_1 and U_2 , quantifying how much the user profiles influence dialogue generation.

Coherence We assess dialogue coherence using the perplexity (**Perp.**) of the dialogue (Song et al., 2021; Zhou et al., 2023), estimated with GPT2-large (Radford et al., 2019). Additionally, we employ two automated coherence evaluation frameworks: **Q-DCE** (Ye et al., 2021), which uses a

BERT-based encoder to predict a coherence score, and **PairEval** (Park et al., 2024), which utilizes a fine-tuned LLaMA-2 model to compare generated dialogues with reference responses. Details for these metrics are provided in Appendix D.2.

GPT Score (G-Eval) Recent studies have shown that LLM-based evaluations can serve as an effective alternative to human annotation (Chiang and Lee, 2023). Inspired by CoT evaluation framework for summarization tasks (Liu et al., 2023), we adapt a similar approach with GPT-4o (Hurst et al., 2024) to evaluate dialogue quality, following (Jandaghi et al., 2023). Our prompt format, along with explicit guidelines for each evaluation dimension, allows the model to provide consistent judgments. Full prompt details are available in Appendix E.

Human Evaluation Beyond the automated metrics, we also conduct a human evaluation on a subset of samples. For each configuration under evaluation, we randomly sample 40 persona–dialogue pairs generated by Qwen. Three fluent English speakers then rate these samples using the same criteria as the LLM-based automatic evaluation, but on a condensed 1–3 scale. We set the final score for each configuration to the average of the three evaluators’ ratings.

3 Are LLMs Sensitive to Users’ Polarity?

3.1 Does Dialogue Quality Diverge According to Polarized User-Pairing?

Inspired by psychological research indicating that the quality of conversations varies with participants’ traits (Hassan et al., 2019), this section explores how polarized profiles and pairing types affect dialogue outcomes under various configurations. We conduct a large-scale, detailed study employing four backbone LLMs and eight metrics to verify the robustness of the results.

3.1.1 Experiment Setup

Pairing User Profiles We leverage the polarized user profiles described in §2.2, along with the original ConvAI2 (Dinan et al., 2020) user profiles, to construct a range of dialogue scenarios—including both homogeneous (e.g., two **negative profiles**) and heterogeneous (e.g., one **negative profile** and one **positive profile** configurations.) Specifically, we create five pairing types, each with 3K user pairs:

Original Pairing: Two original ConvAI2 profiles, serving as a baseline.

Negative Pairing: Two **Negative Profiles** are paired.

Mixed Pairing: Two **Mixed Profiles** are paired.

Positive Pairing: Two **Positive Profiles** are paired.

Opposite Pairing: One **Negative** and one **Positive Profiles** are paired. To reduce bias caused by prompt ordering (Lu et al., 2021), half the pairs begin with the negative, and half begin with the positive.

Experiment Pipeline For each pairing configuration, we feed all user pairs into four modern LLMs—LLaMa-3.1-8B-Instruct (Dubey et al., 2024), Qwen-2.5-7B-Instruct (Yang et al., 2024), Ministral-8B-Instruct (Jiang et al., 2024), and Gemma-2-9B-Instruct (Team et al., 2024)—which collectively generate 60K dialogues and filter out outliers, as mentioned in §2.4. We ensure that every model processes the same user pairs in the identical order. Dialogues where the model refuses to respond or produces repetitive utterances are discarded, resulting in approximately 58K valid dialogues. We provide a detailed distribution of dialogues by pairing scenario and model in Appendix F. Finally, we apply four consistency metrics and four coherence metrics to all dialogues, aggregating the results by pairing type and LLM. For utilizing GPT-4o as backbone LLM, we report the results in Appendix C.

3.1.2 Results & Analysis

LLM Generation Diverges with Polarity Configurations As shown in Table 3, dialogues produced under the *Negative pairing* configuration yield higher **C scores** and **Contd.** values than those under the *Positive pairing*. Specifically, although the gap between entailment and contradiction cases is greater for *Negative pairings*, the proportion of contradictions is also higher. Moreover, coherence is generally lower, implying that the model overemphasizes persona-related content. In such scenarios, frequent entailment goes hand in hand with frequent contradiction, reducing overall coherence, even consistency (i.e. **P gap**, **GPT score**)

Conversely, *Positive pairing* results in fewer entailment cases and fewer contradictions, leading to very low **Contd.** scores and overall high coherence. This suggests that the model tends to forgo incorporating certain persona elements if they undermine dialogue flow. Meanwhile, the *Mixed* configuration combines both tendencies, placing most metrics

Model	Pairing	Consistency				Coherence			
		C score ↑	Contd. ↓	P Gap ↓	G-eval ↑	Perp. ↓	Q-DCE ↑	PairEval ↑	G-eval ↑
LLaMa-3.1-8B	<i>Original</i>	0.391	14.33	-0.43	4.28	5.31	3.14	2.79	<u>4.39</u>
	<i>Negative</i>	0.444	14.71	-0.27	4.27	5.33	3.01	2.74	4.57
	<i>Positive</i>	0.428	9.83	-0.46	4.44	3.40	3.17	2.84	4.65
	<i>Mixed</i>	0.396	13.95	-0.34	4.32	<u>5.37</u>	3.09	2.77	4.51
	<i>Opposite</i>	<u>0.352</u>	13.62	-0.32	<u>4.20</u>	5.30	3.09	2.77	4.47
Qwen-2.5-7B	<i>Original</i>	0.392	<u>15.33</u>	<u>-0.73</u>	4.50	7.05	3.06	2.69	4.34
	<i>Negative</i>	0.520	13.48	-0.80	4.55	7.36	3.07	2.67	4.27
	<i>Positive</i>	0.452	8.84	-0.96	4.67	7.04	3.14	2.75	4.38
	<i>Mixed</i>	<u>0.404</u>	12.99	-0.82	4.45	7.09	3.03	2.70	4.43
	<i>Opposite</i>	0.409	12.58	-0.77	<u>4.33</u>	7.13	<u>3.02</u>	<u>2.67</u>	<u>4.24</u>
Ministral-8B	<i>Original</i>	0.555	<u>10.61</u>	-0.95	4.38	5.98	3.11	2.66	4.11
	<i>Negative</i>	0.778	9.93	-0.97	4.36	<u>7.27</u>	3.11	<u>2.61</u>	3.95
	<i>Positive</i>	0.595	5.78	-1.15	4.51	5.80	3.16	2.67	4.21
	<i>Mixed</i>	0.651	9.65	<u>-0.80</u>	4.43	6.06	3.10	2.62	4.01
	<i>Opposite</i>	<u>0.540</u>	10.48	-0.81	<u>4.27</u>	5.88	<u>3.08</u>	2.62	<u>3.92</u>
Gemma-2-9B	<i>Original</i>	0.391	16.10	-0.69	4.33	6.47	3.09	2.52	3.91
	<i>Negative</i>	0.423	13.57	-0.80	4.35	6.06	3.08	2.39	3.77
	<i>Positive</i>	0.465	7.58	-0.90	4.45	5.83	3.16	2.56	4.07
	<i>Mixed</i>	0.383	12.86	-0.77	4.39	5.62	3.08	2.44	3.85
	<i>Opposite</i>	<u>0.322</u>	13.41	<u>-0.64</u>	4.19	6.31	3.12	2.42	3.81

Table 1: Combined consistency and coherence results. Metrics with (↑) indicate higher is better, while those with (↓) indicate lower is better. Best scores per model and pairing are bolded and worst scores underlined. Highest values across models are in blue, lowest in red, with a gradient highlighting favorable trends as *Negative*, *Positive* Pairing, and *Mixed* scores increase with polarity.

somewhere between those of the *Negative* and *Positive* configurations as illustrated in the color gradients of Table 3. In summary, dialogues generated between *Negative users* strongly reflect each user’s persona—often at the expense of coherence and even consistency—whereas those between *Positive users* selectively apply persona elements, thus preserving both consistency and coherence.

Positive Configurations Yield Better Dialogues

As shown in Table 1, dialogues generated under the *Positive pairing* configuration exhibit consistently high consistency and coherence across nearly all tested models. Overall, except for a few exceptions, it achieves the highest scores across all pairing settings. This result indicates that providing an LLM with a **Positive profile** can produce high-quality dialogues—surpassing even those derived from *Original pairing* configuration—and reinforces the idea that LLMs are sensitive to user polarity.

However, in the *Opposite pairing* configuration, where a **Positive Profile** is paired with a **Negative Profile**, the resulting dialogues degrade to levels comparable to (or even worse than) those seen with *Negative pairing*. In other words, the advantage of the **Positive Profile** vanishes, further confirming

that LLMs respond sensitively to user polarity.

Impact of Backbone Model Selection Our analysis reveals that the choice of backbone model significantly influences dialogue characteristics beyond the pairing method. As shown in Table 1, both Ministral-8B and Qwen-2.5-7B tend to produce dialogues with higher consistency. In particular, Ministral-8B outperforms other models on three of the four consistency metrics, while Qwen-2.5-7B achieves the highest score on the G-eval metric. In contrast, LLaMa-3.1-8B excels in coherence, leading across all four coherence metrics. Meanwhile, Ministral-8B and Gemma-2-9B show lower coherence scores, suggesting that even though these models maintain consistency, they have difficulty generating fully coherent interactions. This observation is further confirmed by the variations seen in the G-eval outcomes.

Human Evaluation We verify whether the trends identified by automatic metrics persist under human judgment by evaluating dialogues generated with the *Original pairing*, *Negative pairing*, and *Positive pairing* configurations.

As shown in Table 2, human raters corroborate the automatic findings: dialogues generated

Pairing	Consistency	Coherence
<i>Original</i>	2.36	2.01
<i>Negative</i>	2.40	2.12
<i>Positive</i>	2.51	2.30

Table 2: Human evaluation results on consistency and coherence among the sentiment of user profiles.

from polarized users—especially *Positive pairings*—exhibit the highest quality. Notably, the decline in scores for the *Original pairing* is far more pronounced in the human evaluation than in the automatic assessment, further confirming the validity of our observations.


3.2 Does Dialogue Quality Diverge According to Users’ Polarity Level?

In the previous section §3.1, we analyzed the impact of pairings between *polarized users*. However, a more detailed quantitative interpretation of user personas’ sentimental polarity is necessary to clarify whether the observed effects arise from user polarity or the pairing strategy. In this section, we address this uncertainty by examining dialogue outcomes across systematically varied polarity levels. We construct user profiles that span multiple degrees of polarity and investigate how dialogue consistency and coherence evolve in each scenario.


3.2.1 Experimental Setup

Defining Polarity Level We aim to determine whether dialogues between users with strongly polarized personas (e.g., clearly positive or negative) differ from those with weak or neutral emotional content. We define each persona’s polarity level based on the confidence score produced by the classification model used in §2.2: values closer to 0 indicate a highly negative polarity, those closer to 1 indicate a highly positive polarity, and scores near 0.5 represent a neutral persona. Because half of all personas lie at sentimental extremes, uniformly dividing confidence scores would yield imbalanced distributions across levels. Instead, we partition the score range into nine intervals that ensure a more balanced representation of persona instances. Further details on the partition criteria and the distribution across levels are provided in Appendix F.

Experiment Pipeline We fix the number of personas K to 1 in each user profile to minimize external factors such as the effect of combining or ordering multiple personas. The profile’s polarity level follows directly from that single persona, and



(a) C score of per Confidence Score Interval



(b) Paireval Score of per Confidence Score Interval

Figure 3: Consistency and coherence across nine polarity levels. Dot x -coordinates show the median confidence score (0: most negative, 1: most positive).

we pair profiles of the same level together. For each level, we create 500 distinct user pairs and generate dialogues using LLaMa-3.1-8B and Qwen-2.5-7B models. We then compute the average consistency and coherence scores—measured respectively by **C score** and **PairEval**, for each level.

3.2.2 Results & Analysis

Stronger Sentimental Extremity in Persona Yields Higher-Quality Dialogue As illustrated in Figure 3, both consistency and coherence scores exhibit a U-shaped trend with respect to polarity level. In other words, dialogues derived from highly polarized personas, whether strongly positive or strongly negative, tend to outperform those derived from more neutral personas. This pattern is consistent across all tested models, suggesting that personalizing a model with a strongly polarized persona enhances dialogue quality.

Positive Polarity Results in Particularly High-Quality Dialogue Furthermore, dialogues generated from positively polarized profiles demonstrate superior consistency and coherence compared to those from negatively polarized profiles. Notably, the consistency score for dialogues produced by the LLaMa-3.1-8B model at the highest polarity level

is up to seven times greater than that at the lowest level, while the Qwen-2.5-7B model shows an average improvement of approximately 0.3 points in coherence. These findings indicate that positive polarity yielding especially pronounced gains.

We further analyze the underlying causes of the performance gap discussed in this section in Appendix A.

4 How to Make LLMs Robust to Polarity?

Our previous findings indicate that persona polarity influences dialogue quality in LLMs. This observation underscores the need for methods that leverage polarity traits to build a more robust personalized dialogue system. To that end, we conduct experiments using real-world profiles in an original pairing configuration to determine how to make LLMs robust to polarity.

4.1 Adjusting Generation Strategy

In addition to the *dual-persona joint generation* method, we adopt a *turn-based agent conversation* approach (§2). In this strategy, each turn is generated using a single user’s profile, thereby avoiding the complexities associated with mixed polarity inputs (§3.1.2) and allows for better handling of persona-specific characteristics. Although this approach requires a separate model call for each turn, our experiments show that models with around 3B parameters can generate high-quality dialogues without requiring extensive computational resources. Consequently, we employ two LLMs, LLaMa-3.2-3B and Qwen-2.5-3B, in this configuration to reduce latency and computational load while maintaining performance.

When applying the turn-based approach, it is necessary to specify a dialogue endpoint. For a fair comparison, we align the number of turns with the distribution observed in the joint-generation method. For instance, if joint generation produces 10 dialogues—6 of which have 8-turns and 4 of which have 10—then we assign an 8-turn endpoint with a 60% probability and a 10-turn endpoint with a 40% probability before initiating turn-based dialogue generation.

Additionally, generating utterances on a turn-by-turn basis sometimes results in excessively long responses that cluttered the overall dialogue. To address this, we add an instruction in the prompt to produce short and concise utterances. The detailed

prompt is provided in Appendix E.

4.2 User Profile Ordering & Sentiment-aware Prompting

Analysis in §3 revealed that profiles containing negative or neutral personas produce lower-quality dialogues compared to those with exclusively positive personas. Since modifying user personas is undesirable, we propose a simple but effective polarity-based ordering technique to enhance dialogue performance without altering the original content. Inspired by (Wu et al., 2024a) showing that LLMs perform better when desired persona content appears early in the prompt, we design three ordering strategies based on the confidence score from §2.2:

- *Score Ascending*: Arrange personas in order of increasing confidence (more negative first).
- *Score Descending*: Arrange personas in order of decreasing confidence (more positive first).
- *Center-Out Score Ascending*: Sort by $|(confidence\ score - 0.5)|$, adding a bias factor of $a = 0.05$ for positive personas to prioritize negatives.

We apply this ordering only to the turn-based approach, as the joint generation method combines both profiles into a single prompt.

Plus, we concatenate a lightweight prompt template: *"Please ensure that each user’s persona, especially negative or neutral personas, is well integrated into the dialogue and that the overall dialogue remains coherent"*. It selectively conditions the model on the sentiment of personas present in a user profile at the end of the generation prompt. Because large language models already handle positive personas reliably, the template explicitly instructs the model to be aware when responding to negative or neutral personas.

4.3 Results & Analysis

Turn-based Approach Improves Overall Quality.

For the LLaMa-3.2-3B model, the *turn-based* approach significantly enhances dialogue consistency compared to *dual-persona joint generation*. This improvement stems from separating users’ polarized personas into distinct inputs, thereby reducing internal complexity. In contrast, the Qwen-2.5-3B model demonstrates lower **G-eval** scores primarily because it blends the personas of both users into a single dialogue. When **User 1** initiates the conversation, its persona tends to dominate, leading **User 2** to adopt a similar framing. For example, if

Model	Strategy	Consistency			Coherence	
		C score	Contd.	G-eval	PairEval	G-eval
LLaMa-3.2	Joint	0.371	15.01	4.15	2.70	4.50
	Turn-based	0.609	7.92	4.14	2.79	4.56
	+ asc.	0.610	7.97	4.13	2.78	4.65
	+ dsc.	0.597	8.09	4.13	2.77	4.63
	+ c-asc.	0.617	7.39	4.21	2.79	4.67
	+ sap.	0.688	6.56	4.18	2.78	4.59
	+ c-sap.	0.717	6.07	4.25	2.84	4.68
Qwen-2.5	Joint	0.470	11.68	4.32	2.62	4.36
	Turn-based	0.557	10.45	4.02	2.65	4.60
	+ asc.	0.557	10.45	3.99	2.69	4.69
	+ dsc.	0.535	10.99	4.01	2.67	4.69
	+ c-asc.	0.570	10.07	4.08	2.69	4.71
	+ sap.	0.777	8.27	4.58	2.61	4.63
	+ c-sap.	0.774	7.49	4.59	2.69	4.77

Table 3: Dialogue consistency and coherence across generation strategies using LLaMa-3.2-3B and Qwen-2.5-7B with the original pairing. Abbreviations: *asc.* (score ascending), *dsc.* (score descending), *c-asc.* (center-out score ascending), *sap.* (sentiment-aware prompting), *c-sap.* (center-out score ascending + sentiment-aware prompting). Best results are in **bold**.

User 1’s persona is *"I enjoy sewing"*, the dialogue might begin with, *"I heard you’re good at sewing"*, prompting **User 2** to continue discussing sewing. As a result, **User 2**’s responses inadvertently mix in **User 1**’s persona, while **User 1** seldom reaffirms the persona initially designated for **User 2**. This behavior ultimately lowers the overall consistency score. Nonetheless, contradictions between personas and utterances actually decrease, boosting contradiction metrics (**C score** and **Contd.**). Also, we can resolve this issue by prompt engineering, which can be further explored.

Regarding coherence, the overall score improves for both models, with a notable enhancement in the Qwen-2.5-3B model. This outcome suggests that sequentially processing an updated dialogue context supports better integration of previously generated content.

Additional Profile Ordering is Helpful. As shown in Table 3, we show that the center-out ascending method yields modest yet consistent improvements across all five evaluation metrics compared to the other ordering strategies. In this approach, we arrange profiles so that those with moderate polarity appear first, while profiles with extreme negative traits are prioritized over clearly positive ones. This ordering leverages the model’s sensitivity to the arrangement of persona information (Wu et al., 2024a), enhancing personalization for ambiguous personas that are less effectively integrated (§3.2) while enabling well-incorporated

positive personas (§3) to more robustly support the overall context.

Synergy between Profile Ordering and Sentiment-Aware Prompting Furthermore, as shown in Table 3, adding the sentiment-aware prompt alone produces some improvements for both models. More importantly, jointly applying the center-out score-ascending strategy and sentiment-aware prompting yields a strong synergy that markedly elevates dialogue quality. In particular, the coherence metric increases by around 0.1 points for both models, demonstrating substantial gains in conversational coherence. We attribute this improvement to positioning the personas that require heightened attention earlier in the prompt, thereby capitalizing on the uni-directional decoding of large language models in concert with the sentiment-aware instruction.

Human Evaluation We further manually compare turn-based generation with and without our profile ordering and sentiment-aware prompting techniques.

Strategy	Consistency	Coherence
Turn-based	1.80	2.40
+ c-sap.	2.27 (+0.47)	2.43 (+0.03)

Table 4: Human evaluation results for consistency and coherence of dialogues generated by the turn-based approach, with and without profile ordering and sentiment-aware prompting.

As shown in Table 4, under the original pairing condition, applying our profile ordering and sentiment-aware prompting techniques to the turn-based generation improves consistency by 0.47 points and coherence by 0.03 points compared with the baseline. These observations further reinforce the robustness and efficacy of our proposed approach.

5 Related Works

5.1 Personalized Dialogue Generation

Personalized dialogue systems have increasingly focused on leveraging user-specific information for more contextually aligned interactions (Li et al., 2016; Zhang, 2018, 2019; Roller, 2020). Early approaches typically involved training generative models with VAE (Lee et al., 2022) to ensure dialogue coherence or NLI components (Song et al., 2021; Chen et al., 2023b; Zhou et al., 2023) to

capture persona representations. These persona representations often defined as descriptive sentences (Zhang, 2018; Dinan et al., 2020) or key-value attributes (Qian et al., 2017; Gao et al., 2023). With the advent of large language models and subsequent instruction tuning methods (Ouyang et al., 2022), persona information can now be more flexibly embedded directly into system prompts (Yang et al., 2023; Wang et al., 2023b).

Recent research has also advanced multi-agent (Park et al., 2023) or dual-persona (Xu et al., 2022; Jandaghi et al., 2023) strategies, enabling two distinct personas to converse within a single session. This approach enhances personalized-chatbot capabilities (Shuster et al., 2022; Lee et al., 2023) and supports large-scale synthetic data generation (Jandaghi et al., 2023) for further personalization.

5.2 Dialogue Evaluation

Evaluating open-domain dialogue is inherently multifaceted, reflecting diverse aspects such as coherence, fluency, and persona consistency (Wang et al., 2024; Samuel et al., 2024). Traditional metrics like ROUGE (Lin, 2004) and BLEU (Papineni et al., 2002) often fail to capture higher-level qualities. Consequently, specialized metrics leveraging pretrained models—including C score (Madotto et al., 2019) (for consistency), QuantiDCE (Ye et al., 2021), and PairEval (Park et al., 2024) (for coherence)—have gained traction (Ghazarian et al., 2022; Li et al., 2024). We adopt these automated metrics for quantifiable evaluation.

In parallel, LLM-based evaluation strategies have rapidly emerged as a cost-effective alternative to human annotation (Chiang and Lee, 2023). Leveraging Chain-of-Thought prompting (Wei et al., 2022) further enhances evaluative transparency, allowing models to articulate their reasoning (Liu et al., 2023). In our work, we integrate both traditional and LLM-based metrics to comprehensively assess persona-driven dialogue.

5.3 Sentimental Sensitivity in LLMs

Large Language Models (LLMs) are known to be highly sensitive to a variety of factors, including prompt order, language, cultural context, and sentiment (Lu et al., 2021; Dang et al., 2024; Shen et al., 2024; Kwok et al., 2024). As a difference perspective from previous works, we focus on sentimental sensitivity in LLMs in our work. Although instruction tuning (Ouyang et al., 2022) and RLHF (Dang

et al., 2024) can mitigate these effects, recent studies still show that contextual sentiment can strongly influence model outputs (Liu et al., 2024; Wu et al., 2024b). However, while much of the existing research focuses on sentiment that arises naturally in generated text, less work has considered sentimental polarity embedded naturally in explicit persona definitions. This underexplored avenue is central to our investigation, as it can profoundly affect both the coherence and consistency of persona-driven dialogues.

6 Conclusion

In this work, we demonstrate that LLMs are sensitive to the sentimental polarity of the persona during generating personalized dialogues. While positively polarized users yield smooth interactions, negatively polarized users overemphasize persona attributes and lead to lower coherence and even consistency. Furthermore, we find that personas with weak or neutral sentiment generally produce lower-quality dialogues. This findings suggest that the quality of the dialogue decreases if the user profile contains negative, or even neutral personas. Consequently, we introduce a dialogue generation approach that leverages persona polarity through a turn-based strategy with profile ordering, achieving more consistent and coherent outputs. Our findings underscore the importance of incorporating persona sentiment into personalized dialogue systems.

Limitation

Despite our extensive configurations designed to show that our observations are not confined to specific conditions, there remain areas for improvement. First, **the source dataset used in our study is limited**. To control for variations in LLM sensitivity that might arise from different domains or features of personas, we employed only a single dataset, focusing on sensitivity purely driven by sentimental polarity. Future work could expand upon our findings by employing alternative persona representations, including *Sparse key-value attributes* or *User history* in §2.1. Nonetheless, the ConvAI2 dataset we used contains a substantial volume of data, offering sufficient permutations of persona configurations to validate our research.

Second, **the inherent bias of the backbone model**. Automated evaluation metrics may not align with human perception if the underlying models themselves are biased. Similarly, if the polarity

classification model is biased, it could potentially misclassify a persona’s polarity based on certain keywords (e.g., those referencing race or culture). To mitigate such issues, we employ a broad set of diverse metrics to offset individual-model biases, and adopted confidence thresholds to classify polarity only when the model surpasses a pre-defined threshold (0.99). Furthermore, prior work indicates that BERT-based models exhibit fewer biases (Wang et al., 2023a), making complete inversions of our intuitive judgments relatively rare.

Finally, **context length varies along the configurations**. Especially, negative polarity often involves negation words (e.g., “not,” “n’t”), potentially yielding longer sentences. However, Table 7 indicates that the overall difference in length is minimal. Additionally, dialogues generated by LLaMa-3.2-3B using a turn-based approach are notably longer, complicating direct comparisons. Nonetheless, it is noteworthy that these more extended utterances maintain high coherence, suggesting a promising direction for dialogue quality and consistency.

Ethics Statement

This research utilizes only fictional persona profiles from the ConvAI2 dataset, ensuring that the handling or storage of any real personal data is completely avoided. All dialogues generated by multiple large language models are entirely simulated; they represent virtual interactions created through algorithmic synthesis and not real human conversations. These simulated dialogues do not raise any privacy concerns, as they are devoid of actual personal data.

Acknowledgement

This research was supported by the Chung-Ang University Graduate Research Scholarship in 2025. This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [RS-2021-II211341, Artificial Intelligence Graduate School Program (Chung-Ang University)].

References

Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation*

measures for machine translation and/or summarization, pages 65–72.

Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*.

Nuo Chen, Ning Wu, Shining Liang, Ming Gong, Linjun Shou, Dongmei Zhang, and Jia Li. 2023a. Is bigger and deeper always better? probing llama across scales and layers. *CoRR*.

Ruijun Chen, Jin Wang, Liang-Chih Yu, and Xuejie Zhang. 2023b. Learning to memorize entailment and discourse relations for persona-consistent dialogues. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 12653–12661.

Yi-Pei Chen, Noriki Nishida, Hideki Nakayama, and Yuji Matsumoto. 2024. Recent trends in personalized dialogue generation: A review of datasets, methodologies, and evaluations. *arXiv preprint arXiv:2405.17974*.

Cheng-Han Chiang and Hung-yi Lee. 2023. Can large language models be an alternative to human evaluations? *arXiv preprint arXiv:2305.01937*.

John Dang, Arash Ahmadian, Kelly Marchisio, Julia Kreutzer, Ahmet Üstün, and Sara Hooker. 2024. Rllhf can speak many languages: Unlocking multilingual preference optimization for llms. *arXiv preprint arXiv:2407.02552*.

Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Emily Dinan, Varvara Logacheva, Valentin Malykh, Alexander Miller, Kurt Shuster, Jack Urbanek, Douwe Kiela, Arthur Szlam, Iulian Serban, Ryan Lowe, et al. 2020. The second conversational intelligence challenge (convai2). In *The NeurIPS’18 Competition: From Machine Learning to Intelligent Conversations*, pages 187–208. Springer.

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

Jingsheng Gao, Yixin Lian, Ziyi Zhou, Yuzhuo Fu, and Baoyuan Wang. 2023. Livechat: A large-scale personalized dialogue dataset automatically constructed from live streaming. *arXiv preprint arXiv:2306.08401*.

Sarik Ghazarian, Nuan Wen, Aram Galstyan, and Nanyun Peng. 2022. Deam: Dialogue coherence evaluation using amr-based semantic manipulations. *arXiv preprint arXiv:2203.09711*.

- Narehan Hassan, Nur Athirah, and Rozilah Abdul Aziz. 2019. [The influence of personality traits on communication competence](#). *International Journal of Academic Research in Business and Social Sciences*, 9.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. Debertav3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*.
- HF Canonical Model Maintainers. 2022. [distilbert-base-uncased-finetuned-sst-2-english \(revision bfdd146\)](#).
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Pegah Jandaghi, XiangHai Sheng, Xinyi Bai, Jay Pujara, and Hakim Sidahmed. 2023. Faithful persona-based conversational dataset generation with large language models. *arXiv preprint arXiv:2312.10007*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Diederik P Kingma. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Louis Kwok, Michal Bravansky, and Lewis D Griffin. 2024. Evaluating cultural adaptability of a large language model via simulation of synthetic personas. *arXiv preprint arXiv:2408.06929*.
- Jing Yang Lee, Kong Aik Lee, and Woon Seng Gan. 2022. Improving contextual coherence in variational personalized and empathetic dialogue agents. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7052–7056. IEEE.
- Joosung Lee, Minsik Oh, and Donghun Lee. 2023. P5: Plug-and-play persona prompting for personalized response selection. *arXiv preprint arXiv:2310.06390*.
- Jiwei Li, Michel Galley, Chris Brockett, Georgios P Spithourakis, Jianfeng Gao, and Bill Dolan. 2016. A persona-based neural conversation model. *arXiv preprint arXiv:1603.06155*.
- Xue Li, Jia Su, Yang Yang, Zipeng Gao, Xinyu Duan, and Yi Guan. 2024. Dialogues are not just text: Modeling cognition for dialogue coherence evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 18573–18581.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Qian Liu, Yihong Chen, Bei Chen, Jian-Guang Lou, Zixuan Chen, Bin Zhou, and Dongmei Zhang. 2020. You impress me: Dialogue generation via mutual persona perception. *arXiv preprint arXiv:2004.05388*.
- 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. *arXiv preprint arXiv:2303.16634*.
- Yang Liu, Xichou Zhu, Zhou Shen, Yi Liu, Min Li, Yujun Chen, Benzi John, Zhenzhen Ma, Tao Hu, Zhi Li, et al. 2024. Do large language models possess sensitive to sentiment? *arXiv preprint arXiv:2409.02370*.
- Adian Liusie, Potsawee Manakul, and Mark JF Gales. 2023. Zero-shot nlg evaluation through pairwise comparisons with llms. *arXiv preprint arXiv:2307.07889*.
- I Loshchilov. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2021. Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity. *arXiv preprint arXiv:2104.08786*.
- Andrea Madotto, Zhaojiang Lin, Chien-Sheng Wu, and Pascale Fung. 2019. Personalizing dialogue agents via meta-learning. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 5454–5459.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- ChaeHun Park, Minseok Choi, Dohyun Lee, and Jaegul Choo. 2024. Paireval: Open-domain dialogue evaluation with pairwise comparison. *arXiv preprint arXiv:2404.01015*.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.
- Qiao Qian, Minlie Huang, Haizhou Zhao, Jingfang Xu, and Xiaoyan Zhu. 2017. Assigning personality/identity to a chatting machine for coherent conversation generation. *arXiv preprint arXiv:1706.02861*.

- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- S Roller. 2020. Recipes for building an open-domain chatbot. *arXiv preprint arXiv:2004.13637*.
- Alireza Salemi, Sheshera Mysore, Michael Bendersky, and Hamed Zamani. 2023. Lamp: When large language models meet personalization. *arXiv preprint arXiv:2304.11406*.
- Vinay Samuel, Henry Peng Zou, Yue Zhou, Shreyas Chaudhari, Ashwin Kalyan, Tanmay Rajpurohit, Ameet Deshpande, Karthik Narasimhan, and Vishvak Murahari. 2024. Personagym: Evaluating persona agents and llms. *arXiv preprint arXiv:2407.18416*.
- Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria, and Rada Mihalcea. 2024. Understanding the capabilities and limitations of large language models for cultural commonsense. *arXiv preprint arXiv:2405.04655*.
- Kurt Shuster, Jing Xu, Mojtaba Komeili, Da Ju, Eric Michael Smith, Stephen Roller, Megan Ung, Moya Chen, Kushal Arora, Joshua Lane, et al. 2022. Blenderbot 3: a deployed conversational agent that continually learns to responsibly engage. *arXiv preprint arXiv:2208.03188*.
- Edward H Simpson. 1951. The interpretation of interaction in contingency tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 13(2):238–241.
- Haoyu Song, Yan Wang, Kaiyan Zhang, Wei-Nan Zhang, and Ting Liu. 2021. Bob: Bert over bert for training persona-based dialogue models from limited personalized data. *arXiv preprint arXiv:2106.06169*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Yu-Ching Hsu, Jia-Yin Foo, Chao-Wei Huang, and Yun-Nung Chen. 2024. Two tales of persona in llms: A survey of role-playing and personalization. *arXiv preprint arXiv:2406.01171*.
- Danqing Wang, Kevin Yang, Hanlin Zhu, Xiaomeng Yang, Andrew Cohen, Lei Li, and Yuandong Tian. 2024. Learning personalized alignment for evaluating open-ended text generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13274–13292.
- Yiwei Wang, Yujun Cai, Muhao Chen, Yuxuan Liang, and Bryan Hooi. 2023a. Primacy effect of chatgpt. *arXiv preprint arXiv:2310.13206*.
- Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, et al. 2023b. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. *arXiv preprint arXiv:2310.00746*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Sean Welleck, Jason Weston, Arthur Szlam, and Kyunghyun Cho. 2018. Dialogue natural language inference. *arXiv preprint arXiv:1811.00671*.
- Bin Wu, Zhengyan Shi, Hossein A Rahmani, Varsha Ramineni, and Emine Yilmaz. 2024a. Understanding the role of user profile in the personalization of large language models. *arXiv preprint arXiv:2406.17803*.
- Zhen Wu, Ritam Dutt, and Carolyn Rose. 2024b. Evaluating large language models on social signal sensitivity: An appraisal theory approach. In *Proceedings of the 1st Human-Centered Large Language Modeling Workshop*, pages 67–80.
- Chen Xu, Piji Li, Wei Wang, Haoran Yang, Siyun Wang, and Chuangbai Xiao. 2022. Cosplay: Concept set guided personalized dialogue generation across both party personas. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 201–211.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.
- Fan Yang, Zheng Chen, Ziyang Jiang, Eunah Cho, Xiaojian Huang, and Yanbin Lu. 2023. Palr: Personalization aware llms for recommendation. *arXiv preprint arXiv:2305.07622*.
- Zheng Ye, Liucun Lu, Lishan Huang, Liang Lin, and Xiaodan Liang. 2021. Towards quantifiable dialogue coherence evaluation. *arXiv preprint arXiv:2106.00507*.
- Saizheng Zhang. 2018. Personalizing dialogue agents: I have a dog, do you have pets too. *arXiv preprint arXiv:1801.07243*.
- Y Zhang. 2019. Dialogpt: Large-scale generative pre-training for conversational response generation. *arXiv preprint arXiv:1911.00536*.
- Junkai Zhou, Liang Pang, Huawei Shen, and Xueqi Cheng. 2023. Simoap: Improve coherence and consistency in persona-based dialogue generation via over-sampling and post-evaluation. *arXiv preprint arXiv:2305.11130*.

A Why Are Large Language Models Sensitive to Persona Sentiment?

We have observed the performance gap across different persona polarities in §3. As we utilize an LLM without any additional fine-tuning, we could infer that LLMs inherently exhibit sensitivity to persona sentiment. Consequently, we explain that this sensitivity is mainly due to the **scarcity of negative expressions in the pretraining data**, which may have been intentionally limited to prevent toxicity. And Figure 2 clearly illustrates a lower proportion of negative personas compared to other polarities. Such data imbalance is evident even in widely-used foundational datasets like ConvAI2, suggesting similar limitations in collecting sufficient negative persona data in other datasets or real-world contexts. Additionally, we also find that LLM’s sensitivity can stem from post-training techniques that steer models toward generating responses preferred by humans. For the experiment, we sample 1,000 dialogues per model under both the negative and positive configurations described in §3.1.1. For the neutral configuration, we follow a similar approach by grouping five intermediate-stage personas from Figure 3 into pairs, generating 1,000 dialogues for each model accordingly.

Pairing	Persona	Dialogue
Negative	1.71	2.67 (+0.96)
Positive	4.25	4.52 (+0.27)
Neutral	3.06	4.09 (+1.03)

Table 5: Average sentiment scores (1–5 Likert) for dialogues and their personas, rated by GPT-4o. Lower scores denote negativity, higher positivity; red values mark the dialogue’s sentiment gain over its persona.

As shown in Table 5, regardless of model type, we observe a consistent tendency to produce more positive dialogues than the assigned personas. Models consistently generate more positive dialogues, with the gap widening when dialogue quality decreases. This trend suggests that **post-training practices bias models toward positive outputs, leading to confusion** when persona sentiment diverges significantly, lowering dialogue quality.

B Which Factors Influence the Polarity-Based Trend?

In §3, we observe that dialogues between negative users tend to yield higher **C score** and **Contd.** values while remaining relatively incoherent, whereas

dialogues from positive users generally exhibit lower **Contd.** values and improved coherence. This section investigates whether these trends persist under variations in key experimental factors or if new patterns emerge. Specifically, we examine the influence of: (1) the number of personas K in a user profile, (2) the polarity ratio within a user profile, and (3) the model size.

B.1 Varying Persona Count

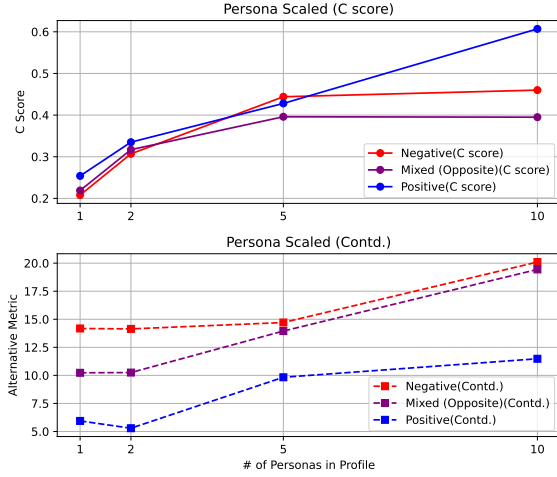
Experiment Pipeline We vary the value of number of personas K in each user profile to 1, 2, and 10, then examine how these changes affect consistency and coherence in dialogues. Adopting the method from §2.2, we generate a total of 1K user profiles per polarity and create 1K user pairs according to the *Negative*, *Positive*, and *Mixed Pairing* strategies described in §3. When each profile contains only one persona, *Mixed Pairing* is not feasible; we therefore use *Opposite Pairing* in that case. We measure consistency with the **C score** and **Contd.** and coherence with the **PairEval** score.

Positive Pairing Exhibits a Direct Positive Relationship between K and Dialogue Quality


As shown in Figure 4, dialogues generated under *Positive pairing* achieve the highest consistency and coherence scores except for C score in $K=5$ setting. Specifically, Figure 4a reveals a positive correlation between K and the **C score** for *Positive Pairing*, despite of growing **Contd.** values. It suggests that as the number of positive personas grows, LLMs can still effectively incorporate them while suppressing the increase in the number of contradictions. Interestingly, as illustrated in Figure 4b, coherence also improves with more personas in *Positive Pairing*, contrary to our initial assumption that additional personas might artificially distort the conversation context.

Negative and Mixed Pairing Show No Clear Improvement

In contrast, *Negative* and *Mixed Pairings* exhibit only marginal gains in consistency and coherence as K increases. Notably, when K is either small or very large, contradictions accumulate faster than entailments, causing the C score to drop below that of positive pairings. In particular, *Mixed Pairing* displays a declining trend in both metrics, presumably because the model struggles to integrate multiple personas polarized at opposite extremes.



(a) C score and Contd. for Scaled K Values. The solid lines represent the trend in **C score**, while the dotted lines indicate the trend in **Contd.**



(b) Paireval Score for Scaled K values.

Figure 4: Consistency and Coherence Comparison Between Pairing Recipe. K is set to 1, 2, 5, and 10. We use Llama-3.1-8B for backbone LLM.

B.2 Impact of Polarity Ratio

Experiment Pipeline Next, we fix K to 5 and vary only the proportion of positive personas to assess its influence on dialogue generation. Building on the *Mixed Pairing* setup in §3, we categorize user profiles by their positive-persona ratios and measure the consistency of the resulting dialogues with the **C score** and **Contd.** Since each paired user may have a different proportion of positive personas, we calculate consistency of the dialogue separately for each user, then average the scores across each ratio category. We use two backbone LLMs, Llama-3.1-8B and Ministral-8B, which yield highly consistent dialogue in §3.1.

A Notable Anomaly in Mixed Ratios As shown in Figure 5, the **C score** decreases as the proportion of positive personas grows, in line with our initial expectations. However, **Contd.** does not steadily decline; it first rises and then starts to drop once the positive ratio exceeds 50%. This pattern consis-

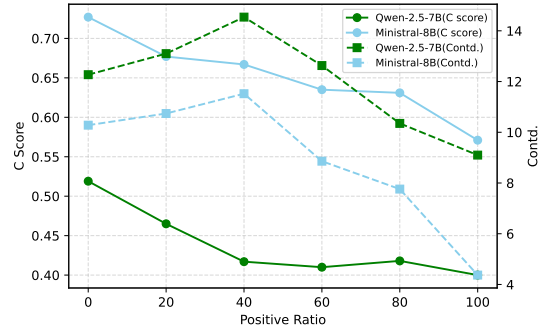



Figure 5: C Score by Positive Persona Ratio. Using LLaMa-3.1-8B and Ministral-8B, higher positive ratios correlate with improved dialogue consistency.

tently appears in both models, implying that when a profile contains negative personas mixed with a slightly smaller share of positive personas, the model becomes more overemphasized to incorporate the user persona. Nonetheless, because the distribution of data across different positive ratios is uneven and our sample size is relatively small, these findings serve only as preliminary observations rather than definitive conclusions.

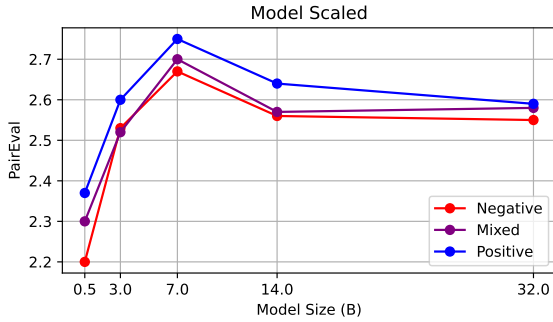
B.3 Effect of Model Size

Experiment Pipeline To determine whether polarity-based trends remain consistent across different model sizes, we employ various Qwen-2.5 models ranging from 0.5B to 3B, 7B, 14B, and 32B, alongside the default 7B model. For practical purposes, we sample 1K user pairs from the pool using the *Negative Pairing*, *Positive Pairing*, and *Mixed Pairing* strategies in §3, ensuring identical inputs for each model. The dialogue metrics follow those in Appendix B.1.

Polarity-Based Trends Fluctuate with Model Size In §3.1, we discussed why dialogues between negative users tend to yield higher **C score** and **Contd.** values yet remain relatively incoherent. By contrast, positive-user dialogues generally show lower **Contd.** values and higher coherence. As illustrated in Figure 6a, these tendencies shift with model size. When the model is 3B or smaller (i.e., 0.5B or 3B), negative-user dialogues exhibit markedly low **C score** and high **Contd.**, indicating that smaller-scale models handle negative personas less effectively than positive personas. However, once the model size surpasses 7B (e.g., 7B, 14B, 32B), the original pattern reemerges, suggesting that while larger models gain the capacity to process negative personas more accurately, they continue to overemphasize them in the dialogue.



(a) C score and Contd. for Different Model Sizes. The solid lines represent the trend in C score, while the dotted lines indicate the trend in Contd.



(b) Paireval Score for Scaling the Model Size.

Figure 6: Consistency and Coherence Comparison Between Model Size. We use Qwen-2.5-0.5B, 3B, 7B, 14B and 32B for various scaled models.

Dialogue Quality Does Not Always Increase with Model Size We initially hypothesized that a larger model would generate higher consistency and coherence; however, the experimental results do not fully support this assumption. For consistency, the 3B model actually achieves the highest average **C score** and **Contd.**, irrespective of user polarity. Although performance improves somewhat as model size scales from 7B to 32B, the gains are not substantial relative to the increase in size. In terms of coherence, the performance peaks at 7B and then declines at larger scales. These findings align with (Chen et al., 2023a)’s conclusion that model size does not necessarily translate into higher downstream task performance.

C GPT-4o Results

We extended our analysis by incorporating GPT-4o—the leading closed-source LLM in many chatbot benchmarks—into our backbone set to test whether the phenomenon identified in **RQ1** persists. The experimental setup mirrors that described in §3.1.1.

As shown in Table 6, the same pattern emerges with GPT-4o: the *positive-pairing* configuration

Pairing	Consistency			Coherence	
	C score ↑	Contd. ↓	G-eval ↑	PairEval ↑	G-eval ↑
<i>Original</i>	0.385	7.92	4.14	2.79	4.56
<i>Negative</i>	0.626	11.11	4.27	2.75	4.44
<i>Positive</i>	0.651	4.12	4.42	2.83	4.69
<i>Mixed</i>	0.638	8.60	4.32	2.76	4.54
<i>Opposite</i>	0.589	8.64	4.23	2.75	4.44

Table 6: Combined consistency and coherence results of the dialogues generated by GPT-4o.

yields the highest dialogue quality across all metrics, whereas configurations containing negative pairings or mixed sentiment polarities perform noticeably worse. In particular, the *original-pairing* configuration—which reflects real-world distributions—suffers a larger quality drop than we observed with open-source LLMs. This result reinforces the robustness of our findings and warns of potential, unforeseen performance degradation in downstream applications.

D Evaluation

D.1 Consistency

C score **C score** (Madotto et al., 2019) is a widely used metric for assessing the consistency of personalized utterances. Let the user profile for user n be $U^{(n)}$ and its constituent personas be $p_i^{(n)}$. Similarly, denote the dialogue by D , composed of utterances $u_i^{(n)}$ from user n . We feed these persona-utterance pairs into a fine-tuned NLI model to obtain entailment scores, then sum these values for each utterance across all personas:

$$\begin{aligned}
 U^{(n)} &= \{p_1^{(n)}, p_2^{(n)}, \dots, p_k^{(n)}\} \quad , \\
 D &= \{u_1^{(1)}, u_1^{(2)}, u_2^{(1)}, u_2^{(2)}, \dots\} \\
 \text{NLI}(u_i^{(n)}, p_j^{(n)}) &= \begin{cases} 1 & \text{entailment,} \\ 0 & \text{neutral,} \\ -1 & \text{contradiction.} \end{cases} \\
 C(u_i^{(n)}) &= \sum_{j=1}^k \text{NLI}(u_i^{(n)}, p_j^{(n)})
 \end{aligned}$$

We fine-tune a BERT-large (Devlin, 2018) model on the DialogNLI dataset (Welleck et al., 2018) by utilizing AdamW optimizer (Loshchilov, 2017) for NLI, achieving near-state-of-the-art accuracy (88.9%) on the test set, using a learning rate of 1×10^{-5} and a batch size of 16.

Originally, the C score was designed to measure the consistency of a single response given a dialogue context. To adapt it for entire dialogues, we

split the dialogue into individual utterances, match each utterance of speaker n to its user profile $U^{(n)}$, compute all persona-utterance entailment scores, and sum them. We then average these sums across all utterances to obtain the dialogue-level C score:

$$C(D) = \frac{1}{|D|} \sum_i \left(C(u_i^{(1)}) + C(u_i^{(2)}) \right),$$

where $|D|$ is the total number of turns in the dialogue.

Contd. Because the C score alone may be susceptible to Simpson’s Paradox (Simpson, 1951), we propose the Contradiction Ratio (**Contd.**) to provide a more detailed distribution of entailments versus contradictions. For example, suppose the entailment scores for a particular utterance across personas are $[0, 0, 0, 1, 0]$ and $[1, 1, 1, -1, 0]$. These yield C scores of 1 and 2, respectively, even though the latter contains a contradiction. Consequently, Contd. measures the percentage of contradictions out of all entailments or contradictions. Formally:

$$\begin{aligned} \# \text{ of Contradictions (C\#)} &= \sum_{n=1}^2 \sum_i \sum_{j=1}^k \mathbf{1}(\text{NLI}(u_i^{(n)}, p_j^{(n)}) = -1), \\ \# \text{ of Entailments (E\#)} &= \sum_{n=1}^2 \sum_i \sum_{j=1}^k \mathbf{1}(\text{NLI}(u_i^{(n)}, p_j^{(n)}) = 1). \\ \text{Contd.} &= \left(\frac{\text{C\#}}{\text{C\#} + \text{E\#}} \right) \times 100, \end{aligned}$$

where $\mathbf{1}(\cdot)$ is the indicator function.

D.2 Coherence

Q-DCE & PairEval Traditional reference-based metrics (e.g., BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005)) often fail to correlate well with human judgments of dialogue coherence (Ghazarian et al., 2022; Li et al., 2024), inspiring the emergence of reference-free automatic evaluation. Among these, Q-DCE uses a ranking-based loss function to pre-train an encoder model (Devlin, 2018), then applies knowledge distillation to produce quantitative coherence scores on a scale of 1 to 5.

With the rise of powerful LLMs, reference-free evaluation has increasingly shifted toward large language models (Liusie et al., 2023). We adopt *PairEval* (Park et al., 2024), wherein a fine-tuned

LLM compares a target response against alternative responses given the same dialogue context. However, these studies often focus on single-turn evaluations, so we treat every utterance from the second to the last as a candidate response. All preceding utterances are considered the dialogue context for that response. Formally, let $D = \{u_1, u_2, \dots, u_N\}$ represent the dialogue, and let $M(c, u) = s$ denote an automated evaluation model that produces a coherence score s for a response u given context c . We define the coherence of the dialogue $\text{Coh}(D)$ as:

$$\text{Coh}(D) = \frac{1}{N-1} \sum_{i=2}^N M(u_{1:i-1}, u_i).$$

E Prompt Templates

In Figure 7, we present the prompt templates for the two generation strategies introduced in §2.4 and employed in §4. Left part shows the LLM prompt used in §3 and §4 whereas right part illustrates the LLM prompt applied in §4. In the latter, two LLMs are personalized with distinct user profiles, and additional instructions are included to limit the length of the responses.

Furthermore, Figure 8 depicts the GPT evaluation prompt template introduced in §2.5. Both templates provide detailed scoring criteria to facilitate more accurate evaluations by the models. Since it was common for dialogues to receive a perfect score of 5, we added the instruction “BE STRICT TO YOUR EVALUATION” to promote a wider distribution of scores. Additionally, to distinguish our approach from conventional metrics, a stronger penalty for contradictions was imposed.

F Dialogue Statistics

Table 7 presents statistics for the ConvAI2 source dataset, the user profiles constructed in §2.2, and the datasets used in experiments addressing each research question. The symbol “” indicates that the omitted value deviates only marginally from the reported statistic, while “-” is used when a statistic is either difficult to compute or not applicable. For §3.1, only statistics for Llama-3.1-8B are provided for convenience, as its trends are consistent with those observed for §3.2.

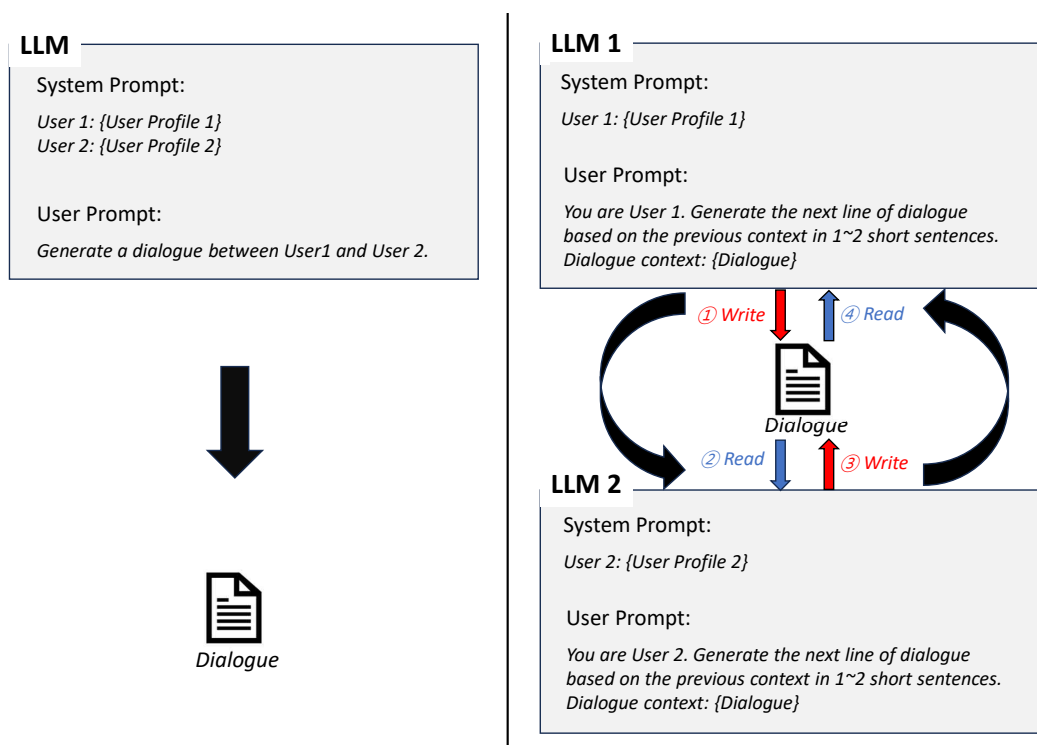


Figure 7: Prompt Template with Simple Framework Illustration for Two Generation Strategies. First one indicates *dual-persona joint generation*, and second one indicates *turn-based approach*.

You will be given two of user persona descriptions and a dialogue between these users.
 Your task is to rate the dialogue on one metric.
 Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.

Evaluation Criteria:
 Consistency (1-5) - the factual alignment between the Persona descriptions and the utterances in the dialogue. A factually consistent utterance contains only statements that are entailed by the source persona descriptions. Annotators were also asked to penalize dialogue that contained utterances contradicted by Persona descriptions.

- 1: Dialogue misses all personas in Persona descriptions, or contains many utterances contradicted by Persona descriptions.
- 2: Dialogue misses three or four personas in Persona descriptions, or contains several utterances contradicted by Persona descriptions.
- 3: Dialogue misses two personas in Persona descriptions, or contains few utterances contradicted by Persona description.
- 4: Dialogue misses one persona in Persona descriptions.
- 5: Dialogue reflects all of Persona descriptions perfectly.

Evaluation Steps:

1. Read the persona descriptions carefully and identify the main facts and details it presents.
2. Read the dialogue and compare it to the descriptions. Check if the dialogue contains any utterances contradicted by given descriptions.
3. Assign a score for consistency based on the Evaluation Criteria. BE STRICT IN YOUR EVALUATION.
4. Justify the rating by referring to specific aspects of the conversation that demonstrate its coherence or lack thereof.

Example:
 Persona descriptions:
 User 1 : {User Profile 1}
 User 2: {User Profile 2}

Dialogue:
 {dialogue}

Evaluation Form (scores ONLY):
 - Consistency:

(a) G-eval Prompt for Evaluating Consistency.

You will be given a pair of user personas. You will then be given one conversation between this persona pair.
 Your task is to rate the conversation on one metric. Please make sure you read and understand these instructions carefully. Please keep this document open while reviewing, and refer to it as needed.

Evaluation Criteria:
 Coherence (1-5) - the collective quality of all utterances. The conversation should be well-structured and well-organized. The conversation should not just be a heap of related information, but should build from utterance to a coherent body of conversation about a topic and previous context.

- 1: All utterances in dialogue are unrelated to each other, or the dialogue contains many utterances contradicted to the previous context.
- 2: Dialogue contains some utterances contradicted to the previous context.
- 3: Dialogue contains some utterances unrelated to the previous context.
- 4: Dialogue is somewhat fluent, but the flow of the topic is not smooth.
- 5: Dialogue is perfectly fluent and well-organized.

Evaluation Steps:

1. Read and understand the given conversation.
2. Evaluate the conversation based on the coherence of the utterances.
3. Rate the conversation on a scale of 1 to 5 based on Evaluation Criteria. BE STRICT IN YOUR EVALUATION.
4. Justify the rating by referring to specific aspects of the conversation that demonstrate its coherence or lack thereof.

Example:
 Conversation: {dialogue}

Evaluation Form (scores ONLY):
 - Coherence:

(b) G-eval Prompt for Evaluating Coherence

Figure 8: G-eval Prompts for Evaluating Dialogue Quality.

Source	Personas (profiles)	# of samples	# profile words	# dialogue turns	# dialogue words
ConvAI2	unique profiles	2310	63.34	-	-
	unique personas	6126	6.59	-	-
	negative personas (confidence score ≤ 0.01)	1006	7.08	-	-
	$0.01 < \text{confidence} \leq 0.1$	652	6.88	-	-
	$0.1 < \text{confidence} \leq 0.2$	158	7.02	-	-
	$0.2 < \text{confidence} \leq 0.4$	151	6.53	-	-
	$0.4 < \text{confidence} \leq 0.6$	150	6.63	-	-
	$0.6 < \text{confidence} \leq 0.8$	194	6.47	-	-
	$0.8 < \text{confidence} \leq 0.9$	191	6.76	-	-
	$0.9 < \text{confidence} \leq 0.99$	933	6.57	-	-
	positive personas (confidence score > 0.99)	2691	6.42	-	-
Synthesized user profiles	Negative Profiles	10000	35.68	-	-
	Mixed Profiles	10000	34.65	-	-
	Profiles Profiles	10000	35.07	-	-
§3.1	Original pairings (llama)	2296	63.27	16.01	414.09
	Negative pairings (llama)	2898	70.26	14.53	390.15
	Mixed pairings (llama)	2879	69.34	15.36	417.49
	Positive pairings (llama)	2864	67.29	14.86	423.67
	Opposite pairings (llama)	2826	69.47	15.82	411.73
	Original pairings (qwen)	2301	“	10.42	356.15
	Negative pairings (qwen)	2987	“	10.05	393.20
	Mixed pairings (qwen)	2916	“	10.43	390.12
	Positive pairings (qwen)	2851	“	10.46	374.32
	Opposite pairings (qwen)	2895	“	10.52	381.12
	Original pairings (minstral)	2288	“	10.73	287.17
	Negative pairings (minstral)	2940	“	9.88	289.77
	Mixed pairings (minstral)	2851	“	10.10	295.07
	Positive pairings (minstral)	2790	“	10.13	299.75
	Opposite pairings (minstral)	2801	“	10.52	299.20
	Original pairings (gamma)	2167	“	13.66	227.85
	Negative pairings (gamma)	2710	“	13.43	233.62
	Mixed pairings (gamma)	2795	“	13.97	239.88
	Positive pairings (gamma)	2863	“	12.72	230.19
	Opposite pairings (gamma)	2672	“	15.12	236.88
	Negative pairings (K=1) (llama)	500	13.61	10.29	281.76
§3.2	($0.01 < \text{conf} \leq 0.1$) pairings (llama)	500	13.64	11.46	289.12
	($0.1 < \text{conf} \leq 0.2$) pairings (llama)	500	14.36	11.40	275.91
	($0.2 < \text{conf} \leq 0.4$) pairings (llama)	500	13.41	11.27	228.41
	($0.4 < \text{conf} \leq 0.6$) pairings (llama)	500	12.84	10.95	263.11
	($0.6 < \text{conf} \leq 0.8$) pairings (llama)	500	13.49	11.45	231.23
	($0.8 < \text{conf} \leq 0.9$) pairings (llama)	500	13.02	11.27	253.59
	($0.9 < \text{conf} \leq 0.99$) pairings (llama)	500	12.96	12.64	241.78
	Positive pairings (K=1) (llama)	500	12.87	9.72	284.83
§4	Original pairing (llama-3B)	2203	63.26	10.33	231.17
	Original pairing (qwen-3B)	2254	“	11.55	246.13
	Original pairing (llama-turn-based)	2032	“	17.68	1020.30
	Original pairing (qwen-turn-based)	2190	“	12.03	382.63
	Original pairing (llama-turn-based + ascending)	2056	“	17.56	1048.87
	Original pairing (llama-turn-based + descending)	2060	“	17.86	1056.91
	Original pairing (llama-turn-based + center-out Ascending)	2060	“	17.43	1031.42
	Original pairing (qwen-turn-based + ascending)	2148	“	12.09	256.13
	Original pairing (qwen-turn-based + descending)	2181	“	12.09	266.45
	Original pairing (qwen-turn-based + center-out Ascending)	2155	“	12.09	262.42
Appendix B	Negative pairings (K=1)	500	13.61	10.29	281.76
	Opposite pairings (K=1)	500	13.34	10.04	280.69
	Positive pairings (K=1)	500	12.87	9.72	284.83
	Negative pairings (K=2)	992	29.34	11.67	324.25
	Mixed pairings (K=2)	994	27.44	11.89	333.66
	Positive pairings (K=2)	987	27.29	11.75	334.45
	Negative pairings (K=5)	2898	70.26	14.53	390.15
	Mixed pairings (K=5)	2879	69.34	15.36	417.49
	Positive pairings (K=5)	2864	67.29	14.86	423.67
	Negative pairings (K=10)	913	151.94	16.68	387.12
	Mixed pairings (K=10)	905	150.24	18.02	574.97
	Positive pairings (K=10)	890	146.09	16.56	475.81
	Positive Ratio 0%	125	“	“	“
	Positive Ratio 20%	479	“	“	“
	Positive Ratio 40%	885	“	“	“
	Positive Ratio 60%	869	“	“	“
	Positive Ratio 80%	415	“	“	“
	Positive Ratio 100%	106	“	“	“
	Negative pairings (0.5B)	974	70.19	10.33	231.17
	Mixed pairings (0.5B)	970	69.32	11.55	246.13
	Positive pairings (0.5B)	960	68.24	14.25	300.61
	Negative pairings (3B)	982	“	10.81	334.58
	Mixed pairings (3B)	984	“	11.24	328.64
	Positive pairings (3B)	985	“	11.54	327.08
	Negative pairings (7B)	2987	“	10.05	393.20
	Mixed pairings (7B)	2916	“	10.43	390.12
	Positive pairings (7B)	2851	“	10.46	374.32
	Negative pairings (14B)	995	“	12.51	333.07
	Mixed pairings (14B)	987	“	12.68	335.39
	Positive pairings (14B)	987	“	12.21	317.31
	Negative pairings (32B)	999	“	10.83	339.31
	Mixed pairings (32B)	980	“	10.64	328.76
	Positive pairings (32B)	959	“	10.15	310.81

Table 7: Personas and Dialogues Sttistics table combining ConvAI2, synthesized user profiles, and RQ1–RQ4 data.