Style Amnesia: Investigating Speaking Style Degradation and Mitigation in Multi-Turn Spoken Language Models

Anonymous ACL submission

Abstract

In this paper, we show that when spoken language models (SLMs) are instructed to speak in a specific speaking style at the beginning of a *multi-turn* conversation, they cannot maintain the required speaking styles after several turns of interaction; we refer to this as the **style amnesia** of SLMs. We focus on paralinguistic speaking styles, including emotion, accent, volume, and speaking speed. We evaluate three proprietary and two open-source SLMs, demonstrating that none of these models can maintain a consistent speaking style when instructed to do so. We further show that when SLMs are asked to recall the style instruction in later turns, they can recall the style instruction, but they fail to express it throughout the conversation. We also show that explicitly asking the model to recall the style instruction can partially mitigate style amnesia. In addition, we examine various prompting strategies and find that SLMs struggle to follow the required style when the instruction is placed in system messages rather than user messages, which contradicts the intended function of system messages.

1 Introduction

Spoken language models (SLMs) can take speech input and generate speech responses. Unlike text-only large language models (LLMs), SLMs integrate audio encoders and vocoders (Kong et al., 2020) to support end-to-end (E2E) speech understanding and generation (Défossez et al., 2024; Zeng et al., 2024; Fang et al., 2025; Wu et al., 2025; Huang et al., 2025). While this E2E design introduces additional challenges beyond text processing, such as handling paralinguistic features and non-textual information, current SLMs have demonstrated the capability to detect user attributes like emotion, accent, and gender. By leveraging these acoustic cues, SLMs can adjust their outputs to produce more contextually appropriate responses (Wu et al., 2025; Google, 2025b). Beyond passive perception, SLMs are also capable of actively following the user-specified speaking style in single-turn interactions (Zeng et al., 2024; Wu et al., 2025; Chiang et al., 2025; Zhan et al., 2025).

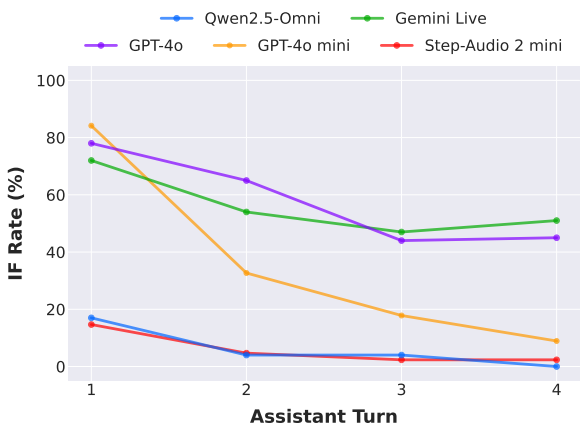

Figure 1: When instructed to consistently speak sadly throughout the conversation, SLMs try their best to follow the instruction in the first turn, but the instruction-following rate degrades rapidly in subsequent turns.

Despite these advances, most prior works focus on *single-turn* evaluation when evaluating the expressive speech generation of SLMs. It is thus unclear whether SLMs can consistently follow a user-specified speaking style in a *multi-turn* spoken interaction. To address this gap, we investigate the speaking style consistency of SLMs in multi-turn spoken conversations. Precisely, we instruct an SLM we want to evaluate to "*follow a specific* ***speaking style*** *throughout the conversation*" at the beginning of the dialogue and then interact with the SLM in a multi-turn conversation. The speaking style can be emotion, accent, volume, or speaking speed. After collecting model responses at each turn, we assess the style instruction-following (IF) rate using style-specific automatic judges.

Our results show that the speaking style instruction-following rate degrades over interaction turns, a phenomenon we refer to as *style amne-*

sia. Figure 1 presents an example when the SLM is instructed to consistently speak sadly, but no SLM can consistently do this. Importantly, we find that most models exhibit a higher style IF rate in the first turn, while the IF rate gradually declines over subsequent turns. These findings indicate that maintaining consistent speaking styles over multiple turns remains challenging for current SLMs. We conduct comprehensive experiments to understand style amnesia and explore potential methods to resolve this issue.

Our contribution can be summarized as follows:

- We identify that in multi-turn conversations, SLMs fail to consistently follow the given style instruction given in the first turn.
- We show that explicitly prompting SLMs to recall the initial style instruction can alleviate style amnesia, but it does not entirely resolve the issue.
- We conduct a comprehensive analysis to examine the impact of prompt techniques on speaking style consistency.

2 Related Works

2.1 Context Loss in Multi-Turn Dialogue

Although LLMs support multi-turn interactions, recent studies have shown that their performance often degrades as dialogues become longer (Kwan et al., 2024; Han et al., 2025; Qamar et al., 2025; Li et al., 2025; Laban et al., 2025). These studies indicate the shortcomings of current LLMs, such as difficulties with multi-turn instruction following (Han et al., 2025; Li et al., 2025), loss of information in long dialogue (Kwan et al., 2024; Qamar et al., 2025), and errors caused by fragmented information across turns (Laban et al., 2025).

In terms of SLMs evaluation, SpokenWOZ (Si et al., 2023) highlights the difficulty of aggregating fragmented information across turns. C^3 (Ma et al., 2025) focuses on bilingual spoken dialogue and evaluates models' capabilities in handling complex multi-turn conversational phenomena, including omission and coreference. ContextDialog (Kim et al., 2025), on the other hand, demonstrates that SLMs exhibit significantly lower memorization performance in long dialogue compared to their text-based counterparts. Overall, these works assess the ability of SLMs to retain and utilize information in long spoken dialogues but leave the expressiveness of generated speech unevaluated.

	Multi-Turn	Style Eval.	Interactive	Turn Analysis
SpokenWOZ (Si et al., 2023)	✓	✗	✗	✗
C^3 (Ma et al., 2025)	✓	✗	✗	✗
ContextDialog (Kim et al., 2025)	✓	✗	✗	✗
Vstyle (Zhan et al., 2025)	✗	✓	✗	✗
Game-Time (Chang et al., 2025)	✗	✓	✗	✗
StyleSet (Chiang et al., 2025)	✗	✓	✗	✗
URO-Bench (Yan et al., 2025)	✓	✓	✗	✗
VocalBench (Liu et al., 2025)	✓	✓	✗	✗
VoxDialogue (Cheng et al., 2025)	✓	✓	✗	✗
Multi-Bench (Deng et al., 2025)	✓	✓	✓	✗
Ours	✓	✓	✓	✓

Table 1: Comparison with related works. *Style Eval.* refers to the assessment of speaking style.

2.2 Speaking Style Assessment

Beyond basic intelligibility and audio quality evaluation (Saeki et al., 2022; Fang et al., 2025), a growing body of work has started to focus on the expressiveness of SLMs (Zeng et al., 2024; Wu et al., 2025). For example, Vstyle (Zhan et al., 2025) evaluates whether SLMs can adapt their voice style according to the given instruction. Game-Time (Chang et al., 2025) focuses on the temporal control of SLMs. Besides, Chiang et al. (2025) demonstrate that Large Audio-Language Models (LALMs) are capable of serving as an automatic judge to evaluate the voice style. Their experiment shows high agreement between human evaluation results and those from Gemini-2.5 Pro (Comanici et al., 2025).

Although previous works have evaluated the expressiveness of SLMs, voice style consistency in multi-turn interactions remains largely underexplored. URO-bench (Yan et al., 2025), VocalBench (Liu et al., 2025), and VoxDialogue (Cheng et al., 2025) use predefined dialogues as the first $k-1$ turns and ask SLMs to generate the k-th response, which does not support turn-level analysis. In contrast to approaches based on predefined multi-turn dialogues, we adopt a model-based user simulator to interact with the evaluated SLMs. We argue that this setting more closely reflects real-world conversational scenarios.

A concurrent work, Multi-Bench (Deng et al., 2025), also evaluates SLMs in multi-turn interactive dialogues. However, their evaluation aggregates performance into a single global score to assess overall emotional intelligence. In contrast, our turn-level analysis provides fine-grained and informative insights into how inconsistencies emerge across a wide range of speaking styles, including emotion, accent, speed, and volume. A comparison with prior work is summarized in Table 1.

3 Evaluating Multi-Turn Speaking Style Following of SLMs

We aim to evaluate whether SLMs can follow a speaking style instruction given at the beginning of a dialogue throughout the whole conversation. In this section, we first introduce our motivation in Section 3.1, and then describe the evaluation framework, dataset, and evaluation metrics in the following subsections.

3.1 Motivation

In the real world, each user may have different preferences for their SLM's speaking style. Consequently, it is important that an SLM can follow the speaking style instruction specified by the user. The user may specify some style instructions for the SLM and expect it to follow them throughout the dialogue. While SLMs are capable of following style instructions in *single-turn* interaction, it is unrealistic to expect the user to restate the same style instruction in each round of the interaction. Thus, whether an SLM can consistently follow a speaking style given at the beginning of the conversation is an important ability of SLMs. Next, we introduce our evaluation framework to evaluate this ability.

3.2 Evaluation Framework

Given an SLM, our goal is to provide it with **speaking style instructions** and ask it to adhere to this style throughout the **multi-turn conversation** based on a specified topic. We then evaluate each turn of the SLM-generated response. An illustration of the overall framework is shown in Figure 2. In this section, we introduce three important components in our framework: (1) the speaking style instructions given to the SLM, (2) the conversation topic, and (3) how to interact with the SLM to form a multi-turn dialogue with a user simulator.

3.2.1 Speaking Style Instructions

At the beginning of the dialogue, we will instruct the SLM to follow a specific speaking style. Unless otherwise specified, we give the instruction in the first user turn, as shown in Figure 2. In our paper, we focus on paralinguistic speaking styles, as this is what makes SLMs different from text-only LLMs. We include four types of paralinguistic attributes, each with multiple possible values. The included speaking styles are listed as follows:

- **Emotion**: Sad, happy, angry, or neutral tone.
- **Accent**: North American or Indian accents.

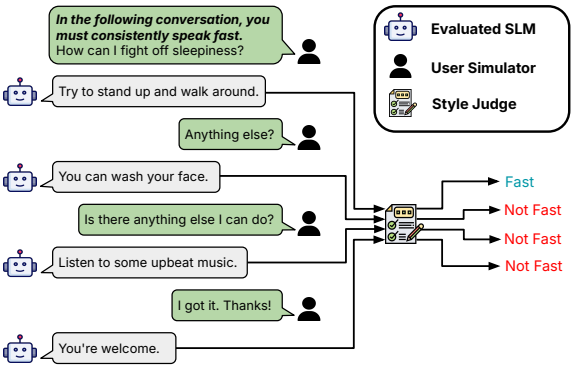

Figure 2: The overview of the evaluation framework. The SLM can speak fast in the first turn as instructed, but it fails to maintain this style in later dialogue turns.

- **Volume**: A higher or lower volume.
- **Speed**: A faster or slower pace.

Among many possible paralinguistic speaking styles, we select the above attributes and values since they can be automatically evaluated, as detailed in Section 3.3.

3.2.2 Dialogue Topics

In the first user turn of the dialogue, we provide a conversation opener for each dialogue. The conversation opener is used to control the topic that the SLM should talk about during each evaluation, ensuring that different SLMs are evaluated under a common conversational setting. We select the topic from Soda (Kim et al., 2023), a large-scale dataset of social interaction dialogues. The details about how we select the topics from Soda are shown in Appendix A. Eventually, we collect 100 diverse conversation openers to initiate the conversations. We manually inspect the selected topics and filter out improper topics, such as queries about personal preferences that are unsuitable for machine conversations. An example conversation opener is shown in Figure 2 "How can I fight off sleepiness?"

3.2.3 Multi-Turn Interaction Using a User Simulator

At the core of our evaluation framework is the multi-turn interaction between the evaluated SLM and a user. To enable back-and-forth interactions with an evaluated SLM, we build a *user simulator* with a cascade SLM. The cascade SLM is a speech-in-speech-out system composed of an ASR, a text-only LLM, and a TTS model. The ASR model takes the speech input from the evaluated

SLM and transcribes it into text.[1] A text-only LLM, powered by GPT-5 mini (OpenAI, 2025), takes the text input and generates a text response based on the transcription produced by the evaluated SLM. Last, the TTS model, GPT-4o mini TTS (OpenAI, 2024), converts the text output from the LLM into speech and sends it back to the evaluated SLM to continue the conversation.

3.2.4 Summary

By combining 10 speaking styles and 100 topics, each SLM is evaluated with **1,000** dialogues. Currently, we do not compose multiple types of speaking styles in a single instruction; this is mainly because SLMs cannot even properly follow a single speaking style instruction consistently, as later shown in Section 4. However, our evaluation framework can be easily extended to the composition of multiple types of speaking styles, and we leave this as future work.

3.3 Evaluation Metrics

To quantify how well SLMs adhere to the target speaking style, we adopt the **instruction-following (IF) rate**, denoted as IF, as our primary evaluation metric. Given a target speaking style s from Section 3.2.1, e.g., *speak sadly*, or *speak fast*, etc., and dialogue topic i, the evaluated SLM generates output $o_{i,j}$ at turn j, the IF rate for style s at turn j denoted as $IF_j(s)$ and defined by:

$$IF_j(s) = \frac{\sum_{i=1}^{N} \mathbb{1}(o_{i,j}, s)}{N} \times 100\%. \quad (1)$$

Here, $N = 100$ is the total number of topics, and $\mathbb{1}(\cdot)$ is a binary indicator determined by a style-specific judge (described later in Section 3.3.1), which reflects whether the generated output $o_{i,j}$ follows the given style instruction s. This metric indicates the percentage of generated speech that correctly follows the given style instruction.

Our goal is to quantify how well the style IF ability changes over the conversation turns. To do so, we define two metrics to measure how good the style IF is across multiple interaction turns.

(1) First-turn IF rate IF_1. This metric serves as a reference to evaluate the capability of the SLM to follow the specific style in the first turn. If an SLM fails to achieve a reasonable IF_1, we may not expect it to perform better in later turns.

[1]Most existing SLMs output both speech and its corresponding transcription at the same time. In this case, we directly use the text output from the evaluated SLM and skip the ASR module in the user simulator.

(2) Degradation rate D. This metric quantifies how the IF rate decays over subsequent turns relative to the initial performance IF_1. It is defined as the average absolute difference between the instruction-following rates of subsequent SLM turns and that of the first SLM turn.

$$D = \sum_{j=2}^{K} \frac{\max(IF_1(s) - IF_j(s), 0)}{K - 1}. \quad (2)$$

Here K denotes the total number of assistant turns of the evaluated SLM. We apply the $\max(\cdot, 0)$ operator to focus solely on style degradation, thereby avoiding the cancellation of positive and negative differences during averaging. For IF_1, larger values indicate better first-turn control over the requested voice style. A smaller D indicates less style degradation. In this paper, we set $K = 4$.

Aside from measuring the speaking style, we also measure the semantic coherence of the dialogue to ensure that the content is valid using LLM-as-a-judge (Chiang and Lee, 2023). Since the semantic coherence of the dialogue is not our primary focus, we leave the details in Appendix C.3. Overall, all SLMs we evaluate maintain reasonable semantic consistency across turns.

3.3.1 Automatic Judges for Styles

An automatic style judge takes the target speaking style s and a speech input $o_{i,j}$ and returns 0 or 1 to indicate whether the $o_{i,j}$ aligns with the target style s. Different automatic judges are adopted for different types of speaking styles.

Emotion We use Emotion2vec-Large (Ma et al., 2024) to predict the emotion of $o_{i,j}$. Although it is a 9-class model, we focus on the probabilities for happiness, sadness, anger, and neutral. We take the maximum probability among these four categories as the final prediction and verify whether it aligns with s. If the prediction is aligned, the indicator function $\mathbb{1}(\cdot)$ is 1; otherwise, it is 0.

Accent We use Voxlect-English-Dialect-Whisper-Large-v3 (Feng et al., 2025) to predict the accent of $o_{i,j}$. From this 16-class model, we extract the probabilities for North American and Indian English, taking the maximum as the predicted accent of $o_{i,j}$. If it is aligned with s, $\mathbb{1}(\cdot)$ is 1; otherwise $\mathbb{1}(\cdot)$ is 0.

Volume We measure volume using Loudness Units Full Scale (*LUFS*) with PyLoudnorm (Steinmetz and Reiss, 2021). Since subjective terms

like "*loud*" or "*quiet*" lack universal definitions, we adopt a relative evaluation standard. We first establish a baseline by querying the SLM with a neutral instruction: "*You are a text-to-speech model. Please read the given text at a normal volume without adding or omitting anything*," and generate another speech $o'_{i,j}$. $\mathbb{1}(\cdot)$ is 1 if the *LUFS* of $o_{i,j}$ is higher (for *loud* instruction) or lower (for *quiet* instruction) than $o'_{i,j}$; otherwise $\mathbb{1}(\cdot)$ is 0.

Speed We measure speaking rate using Words Per Minute (*WPM*) with Parakeet TDT v2 (Nvidia, 2025). Similar to volume, we compare the SLM generated speech $o_{i,j}$ with the TTS-generated speech $o'_{i,j}$ using the same model to determine whether $o_{i,j}$ is faster (or slower) than $o'_{i,j}$.

4 Main Experiments

4.1 Experimental Setup

We evaluate three proprietary SLMs: GPT-4o[2] (OpenAI, 2024), GPT-4o mini[3] (OpenAI, 2024), Gemini Live[4] (Google, 2025b), and two open-source end-to-end SLMs: Qwen2.5-Omni (Xu et al., 2025) and Step-Audio 2 mini (Wu et al., 2025). We select these two representative open-source SLMs because they provide official vLLM implementations (Kwon et al., 2023) to speed up the inference process.

Apart from these E2E SLMs, we introduce a cascaded SLM baseline comprised of GPT-5 mini and Gemini-TTS (Google, 2025a). The TTS in the baseline receives the style instruction in each turn. The performance of the cascaded baseline serves as the performance upper-bound. The hyperparameter choices of each model are provided in Appendix B.1.

4.2 Results

We report the first-turn IF rate IF_1 and the regradation rate D in Figure 3.[5] We also provide the illustration of style amnesia in Figures 1 and 4, and IF rate for each turn in Appendix C.1.

For the **Emotion** speaking style, the Cascaded Baseline demonstrates consistent IF rate across turns, with degradation within 3.0%. In contrast, Gemini Live and GPT-4o show a degradation rate ranging from 13.7% to 26.7% for Anger and Sadness. GPT-4o mini exhibits even larger degradation, reaching 34.7% and 65.3% for Anger and Sadness, respectively. Qwen2.5-Omni and Step-Audio 2 mini also show around 14.0% degradation for Sadness. However, their degradation rates for Anger are only 3.7% and 1.0%, because both models already struggle to perform this style effectively in the first turn.

For the **Accent** speaking style, Gemini Live shows stable performance when maintaining an Indian English accent, suggesting strong control over this attribute. In contrast, GPT-4o mini still exhibits nearly 50% degradation for the Indian English accent. Step-Audio 2 mini achieves a 32.0% IF_1 but a 5.0% degradation rate. Qwen2.5-Omni performs well with the North American accent but fails to produce the Indian accent.

Interestingly, we observe that nearly all SLMs perform better when generating happiness, neutral tone, and North American English than other style attributes. We hypothesize that this discrepancy arises because these attributes correspond to the default speaking styles of the evaluated SLMs. To validate this assumption, we analyze the emotion and accent distributions of samples generated during the speed and volume evaluations, as these prosodic features can coexist with emotion and accent. The distributions shown in Appendix C.2 support our hypothesis: most SLMs tend to default to happy or neutral tones and North American accents, thereby leading to a low degradation rate.

We also observe style amnesia when SLMs are instructed to maintain **Volume** and **Speed**. Regarding **Volume**, results indicate that speaking loudly is generally more challenging for SLMs than speaking quietly, even in their first turn. Despite this, models capable of volume adjustment still exhibit some degradation. In terms of **Speed**, most SLMs demonstrate reasonable control in the first turn but show significant degradation over time. An exception is Qwen2.5-Omni, which fails to control speed even in the first turn, as evidenced by an IF_1 score below 50%, a value comparable to a random baseline in our pairwise comparison setting.

Based on these experiments, we demonstrate that SLMs are prone to losing control of speaking style in multi-turn conversations.

[2] gpt-4o-audio-preview-2025-06-03

[3] gpt-4o mini-audio-preview-2024-12-17

[4] gemini-2.5-flash-native-audio-preview-09-2025

[5] GPT-4o and GPT-4o mini sometimes return only text transcription without speech. When this occurs, we query the models up to three times, but a few samples still fail. We discuss this issue in Appendix C.5. In the following section, all metrics are computed only on successful cases.

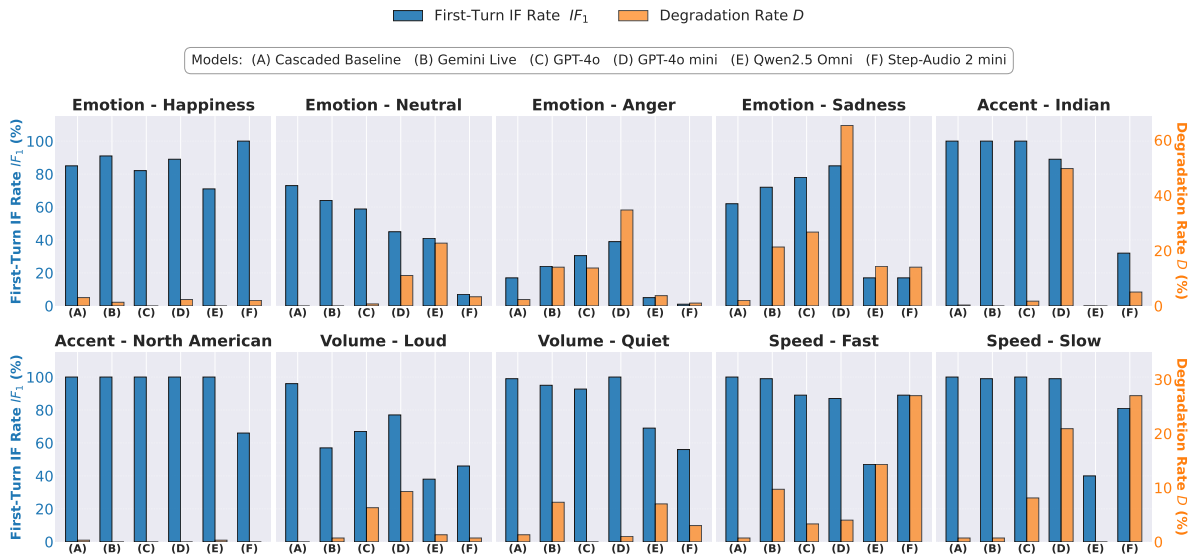

Figure 3: The first-turn IF rate IF_1 and degradation rate D across different speaking styles.

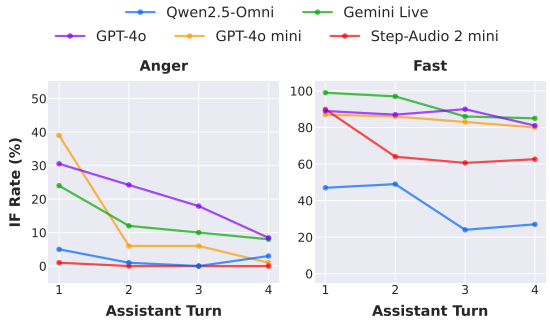

Figure 4: Visualization of style amnesia.

5 Analysis

5.1 The Effect of Prompt Position

In this section, we conduct experiments under different prompt positions to investigate their effect on style amnesia. In instruction-guided language models, system messages are designed with higher priority than user messages to establish global behaviors and safety constraints (Touvron et al., 2023; Wallace et al., 2024). A prior study also demonstrates that system prompts have a more profound impact on model behavior than user prompts do. (Neumann et al., 2025).

Although system messages are important for controlling LLMs, it remains unclear how the placement of speaking style instructions affects SLMs. To investigate this, we conduct an experiment comparing the performance of SLMs when instructions are placed in system messages versus user messages. We select five speaking instructions that most SLMs can follow for this experiment. The IF rate in the first turn is illustrated in Figure 5, and the IF rate for each turn is shown in Appendix C.4.

Surprisingly, we find that most SLMs cannot follow the instructions placed in system messages. When asking SLMs to speak sadly through system messages, GPT-4o, GPT-4o mini, and Step-Audio 2 mini show performance drops of approximately 30%, 50%, and 20%, respectively. Furthermore, GPT-4o mini displays nearly an 80% drop when asked to perform the Indian English accent. A similar issue occurs with instructions related to speaking rate. SLMs nearly ignore the speed instruction in system messages, as their performance is only comparable to the random baseline.

The results presented here focus on the IF rate in the first turn. However, the IF rate for each turn reported in Appendix C.4 show that style amnesia occurs in both settings. Through the experiments above, we identify another crucial issue prevalent in current SLMs. These findings drive us to place the instruction in user messages when conducting other experiments, ensuring that SLMs can follow the instructions more effectively.

5.2 Validation of Emotion and Accent Judges

As described in Section 3.3.1, we evaluate **Speed** and **Volume** using deterministic, signal-level metrics. In contrast, **Emotion** and **Accent** are evaluated using learned models designed to approximate human perception. To ensure the reliability of the automatic judges for **Emotion** and **Accent**, we conduct a human validation study to assess the correlation between our automatic judges and human annotators. In addition, we compare our automatic

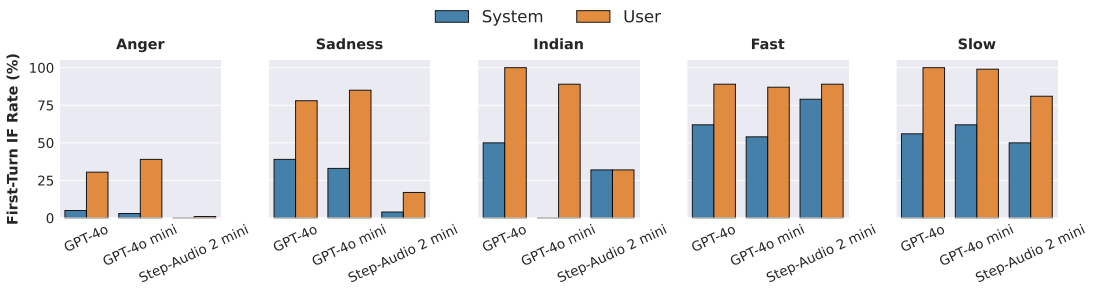

Figure 5: The difference of first-turn IF IF_1 rate when instructions are placed in system and user messages.

Task	Model	Cohen's Kappa	MCC
Accent	Gemini-2.5 Pro	0.741	0.747
	Voxlect	**0.809**	**0.811**
Emotion	Gemini-2.5 Pro	0.464	0.487
	Emotion2vec	**0.476**	**0.511**

Table 2: The correlation between human annotators and automatic judge models.

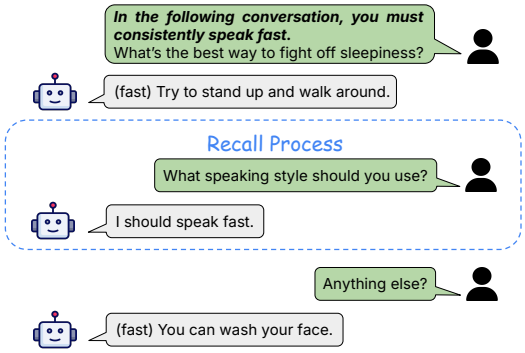

Figure 6: The illustration of the recall process.

judges with Gemini-2.5 Pro (Comanici et al., 2025), a robust LALM capable of general-purpose speech understanding. This model has been shown to exhibit high correlation with human annotators when used as an automatic judge in a prior LALM-as-a-judge study (Chiang et al., 2025).

We employ annotators via Amazon Mechanical Turk to evaluate a randomly sampled subset of 720 speech clips. This subset is constructed by selecting five samples per turn across four conversational turns, six speaking styles, and six evaluated models. The six evaluated styles include Happiness, Neutral, Anger, Sadness, Indian English accent, and North American English accent. Each sample is evaluated by three annotators, who are asked whether the generated speech matches the required style. Details of the human evaluation setup are provided in Appendix D.1.

After collecting the annotations, we derive the final label using majority voting, compute Cohen's Kappa (Cohen, 1960) for inter-annotator agreement, and use Matthews Correlation Coefficient (MCC) (Matthews, 1975) to evaluate judge reliability against the final labels. While Cohen's Kappa evaluates the degree of consensus among raters, MCC provides a balanced measure of classification quality even with imbalanced datasets. Therefore, we report both metrics for reference.

The results, shown in Table 2, indicate that our selected judge achieves the highest reliability. For the accent classification task, Voxlect outperforms Gemini-2.5 Pro, achieving the highest agreement with human annotations, with a Cohen's Kappa of 0.809 and an MCC of 0.811. In the emotion classification task, Emotion2vec-Large demonstrates the strongest reliability among three emotion judges, obtaining a Cohen's Kappa of 0.476 and an MCC of 0.511. These results are similar to the crowdsourced agreement levels found in common speech emotion datasets, as detailed in Appendix D.2.

6 SLM Recall Process

Building on the experiments in Section 4.2, we observe that SLMs suffer from style amnesia: their speaking style IF rate begins to degrade after only one turn and often worsens as the conversation progresses. This raises a key research question: *Do SLMs forget the initial instruction, or just fail to follow the specified style?* To address this, we introduce a *recall process* for every turn after the first, in which the SLM is prompted to restate the initial speaking style before processing the following user input. The recall process is illustrated in Figure 6. Using this method, we first utilize the recall probe to measure whether the SLM retains the original instruction in Section 6.1. Subsequently, we evaluate the effectiveness of integrating this recall process to mitigate style amnesia in Section 6.2.

6.1 Do SLMs Forget the Instruction?

To quantify whether the model remembers the initial style instruction s, we define the recall rate R as follows. Given a style s and dialogue topic i, the recall rate at assistant turn j, denoted $R_j(s)$, is:

$$R_j(s) = \frac{\sum_{i=1}^{N} \mathbb{1}_{\text{recall}}(r_{i,j}, s)}{N} \times 100\%, \quad (3)$$

where $r_{i,j}$ denotes the response generated by the SLM to the recall query q_{recall} before generating its response at the user turn j, and N is the number of dialogue topics. The binary indicator function $\mathbb{1}_{\text{recall}}(\cdot)$ is 1 if the recalled instruction matches the original style instruction, and 0 otherwise. We use GPT-5 mini to judge the correctness of the recalled instruction. The evaluation prompt is shown in Appendix B.3.

We conduct the following experiments on models capable of producing a wide range of speaking styles, as well as on styles that most models can reliably generate. The results are summarized in Table 3. Interestingly, most SLMs remember the initial instruction quite well. The three proprietary models show a near-perfect recall rate. Step-Audio 2 mini shows weaker performance with declining recall rates across turns, yet it still achieves a recall rate of 55.0% to 89.0%. These results reveal a clear gap between proprietary and open-source SLMs in memorization ability. From our manual inspection, when Step-Audio 2 mini fails to recall, some responses contain an incorrect style instruction, while others simply ignore the question and produce irrelevant or meaningless outputs.

Notably, although GPT-4o mini exhibits a 65.3% degradation in performance when producing a sad speaking style and a 20% degradation for speaking slowly, it still maintains a high recall rate. A similar pattern is observed with Gemini Live and GPT-4o. These observations indicate that the models do not forget the instructions, suggesting that style amnesia is not caused by memory loss. Instead, the models retain the instructions but fail to follow the requested style effectively.

6.2 Recall Process Mitigates Style Amnesia

In this section, we investigate whether explicitly asking SLMs to recall the instruction can mitigate style amnesia. The results, presented in Table 3, show that SLMs equipped with the recall process notably reduce degradation. Even for models with relatively low degradation, such as Gemini Live and GPT-4o, the recall process still provides measurable improvements. GPT-4o mini, which suffers from substantial degradation, achieves roughly a 25% reduction in the average degradation rate, demonstrating the effectiveness of the recall process. Step-Audio 2 mini shows slightly improved but largely comparable degradation across the four tasks, likely due to its lower recall rate relative to other SLMs.

Model	Style	R_2	R_3	R_4	Degradation Rate D		
					Base	+Recall	Improv.
Gemini Live	Indian	100.0	100.0	100.0	**0.0**	**0.0**	0.0
	Sadness	100.0	99.0	97.0	21.3	**17.3**	+4.0
	Fast	100.0	100.0	99.0	9.7	**6.3**	+3.4
	Slow	100.0	99.0	100.0	**0.7**	**0.7**	0.0
GPT-4o	Indian	100.0	100.0	100.0	1.7	**0.0**	+1.7
	Sadness	100.0	100.0	100.0	26.7	**14.9**	+11.8
	Fast	100.0	100.0	100.0	3.3	**0.3**	+3.0
	Slow	100.0	100.0	100.0	8.1	**4.4**	+3.7
GPT-4o mini	Indian	93.0	93.0	93.8	49.7	**14.9**	+34.8
	Sadness	100.0	100.0	100.0	65.3	**30.3**	+35.0
	Fast	99.0	100.0	99.0	4.0	**0.0**	+4.0
	Slow	100.0	100.0	100.0	20.9	**1.5**	+19.4
Step-Audio 2 mini	Indian	89.0	72.0	70.0	**5.0**	5.3	-0.3
	Sadness	85.0	56.0	55.0	14.0	**11.0**	+3.0
	Fast	83.0	62.0	57.0	**27.0**	29.0	-2.0
	Slow	83.0	64.0	66.0	27.0	**24.7**	+2.3

Table 3: The recall rate R of SLMs and the impact of the recall process on the degradation rate D across different styles. *Base* represents direct inference without the recall process, while + *Recall* indicates performance using the SLM integrated with the recall process.

Although the recall process alleviates style amnesia, some degree of decay remains unavoidable. This suggests that the inconsistency originates from fundamental limitations of current SLMs and must be addressed in future model designs.

7 Conclusion

In this paper, we identify that SLMs suffer from *style amnesia* in multi-turn conversations, which leads to inconsistent speaking styles across turns. We also demonstrate that SLMs can recall the instruction when asked, but fail to perform it explicitly. Additionally, we indicate that placing style instructions in system messages results in a significant performance drop.

SLMs exhibit a relatively strong ability to memorize specified speaking styles. However, the retention of stylistic information does not necessarily translate into stylistic expression at generation time. Closing this gap between style retention and stylistic control remains an important direction for future research. We hope these findings offer valuable insights for the community and contribute to the development of more reliable SLMs.

Limitations

Practically, role-playing is an important scenario that requires maintaining consistency across multiple turns. The current lack of reliable automatic judges for assessing speech role-playing behaviors limits the breadth of our experiments. Similarly, the current analysis in this paper is based on a single style. In practice, users may dynamically update the instructions during conversations, so SLMs may need to perform a composite style, such as "speak fast and sadly." However, based on our preliminary experiments, most current SLMs struggle to perform this kind of complex style. Additionally, the current judge shows a low correlation with this type of out-of-domain data.

Nonetheless, these limitations do not weaken the conclusions of our study. Even with the styles we can evaluate, SLMs already exhibit noticeable degradation in multi-turn dialogues, suggesting that similar or even greater challenges may arise in more complex settings. We therefore leave the analysis of these scenarios to future work, as it requires further advances in both SLMs and judge models.

We do not see specific harm in our paper.

References

Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeannette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42(4):335–359.

Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. 2014. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing*, 5(4):377–390.

Kai-Wei Chang, En-Pei Hu, Chun-Yi Kuan, Wenze Ren, Wei-Chih Chen, Guan-Ting Lin, Yu Tsao, Shao-Hua Sun, Hung-yi Lee, and James Glass. 2025. Gametime: Evaluating temporal dynamics in spoken language models. *arXiv preprint arXiv:2509.26388*.

Xize Cheng, Ruofan Hu, Xiaoda Yang, Jingyu Lu, Dongjie Fu, Zehan Wang, Shengpeng Ji, Rongjie Huang, Boyang Zhang, Tao Jin, and 1 others. 2025. Voxdialogue: Can spoken dialogue systems understand information beyond words? In *The Thirteenth International Conference on Learning Representations*.

Cheng-Han Chiang and Hung-yi Lee. 2023. Can large language models be an alternative to human evaluations? In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15607–15631, Toronto, Canada. Association for Computational Linguistics.

Cheng-Han Chiang, Xiaofei Wang, Chung-Ching Lin, Kevin Lin, Linjie Li, Radu Kopetz, Yao Qian, Zhendong Wang, Zhengyuan Yang, Hung-yi Lee, and Lijuan Wang. 2025. Audio-aware large language models as judges for speaking styles. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 467–480, Suzhou, China. Association for Computational Linguistics.

Jacob Cohen. 1960. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46.

Gheorghe Comanici and 1 others. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.

Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou, Edouard Grave, and Neil Zeghidour. 2024. Moshi: a speech-text foundation model for real-time dialogue. *arXiv preprint arXiv:2410.00037*.

Yayue Deng, Guoqiang Hu, Haiyang Sun, Xiangyu Zhang, Haoyang Zhang, Fei Tian, Xuerui Yang, Gang Yu, and Eng Siong Chng. 2025. Multi-bench: A multi-turn interactive benchmark for assessing emotional intelligence ability of spoken dialogue models. *arXiv preprint arXiv:2511.00850*.

Qingkai Fang, Shoutao Guo, Yan Zhou, Zhengrui Ma, Shaolei Zhang, and Yang Feng. 2025. Llama-omni: Seamless speech interaction with large language models. In *The Thirteenth International Conference on Learning Representations*.

Tiantian Feng, Kevin Huang, Anfeng Xu, Xuan Shi, Thanathai Lertpetchpun, Jihwan Lee, Yoonjeong Lee, Dani Byrd, and Shrikanth Narayanan. 2025. Voxlect: A speech foundation model benchmark for modeling dialects and regional languages around the globe. *arXiv preprint arXiv:2508.01691*.

Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378.

Google. 2025a. Advanced audio dialog and generation with Gemini 2.5. https://blog.google/technology/google-deepmind/gemini-2-5-native-audio/.

Google. 2025b. Gemini Live: A more helpful, natural and visual assistant. https://blog.google/products/gemini/gemini-live-updates-august-2025/.

Chi Han, Xin Liu, Haodong Wang, Shiyang Li, Jingfeng Yang, Haoming Jiang, Zhengyang Wang, Qingyu Yin, Liang Qiu, Changlong Yu, Yifan Gao, Zheng Li, Bing Yin, Jingbo Shang, and Heng Ji. 2025. Can

language models follow multiple turns of entangled instructions? In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 25445–25460, Suzhou, China. Association for Computational Linguistics.

Ailin Huang, Boyong Wu, Bruce Wang, Chao Yan, Chen Hu, Chengli Feng, Fei Tian, Feiyu Shen, Jingbei Li, Mingrui Chen, and 1 others. 2025. Step-audio: Unified understanding and generation in intelligent speech interaction. *arXiv preprint arXiv:2502.11946*.

Heeseung Kim, Che Hyun Lee, Sangkwon Park, Jiheum Yeom, Nohil Park, Sangwon Yu, and Sungroh Yoon. 2025. Does your voice assistant remember? analyzing conversational context recall and utilization in voice interaction models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 8984–9014, Vienna, Austria. Association for Computational Linguistics.

Hyunwoo Kim, Jack Hessel, Liwei Jiang, Peter West, Ximing Lu, Youngjae Yu, Pei Zhou, Ronan Bras, Malihe Alikhani, Gunhee Kim, and 1 others. 2023. Soda: Million-scale dialogue distillation with social commonsense contextualization. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12930–12949.

Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. In *Advances in neural information processing systems*, volume 33, pages 17022–17033.

Klaus Krippendorff. 1980. *Content analysis: An introduction to its methodology*. Sage publications.

Wai-Chung Kwan, Xingshan Zeng, Yuxin Jiang, Yufei Wang, Liangyou Li, Lifeng Shang, Xin Jiang, Qun Liu, and Kam-Fai Wong. 2024. MT-eval: A multi-turn capabilities evaluation benchmark for large language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 20153–20177, Miami, Florida, USA. Association for Computational Linguistics.

Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.

Philippe Laban, Hiroaki Hayashi, Yingbo Zhou, and Jennifer Neville. 2025. Llms get lost in multi-turn conversation. *arXiv preprint arXiv:2505.06120*.

Jinnan Li, Jinzhe Li, Yue Wang, Yi Chang, and Yuan Wu. 2025. StructFlowBench: A structured flow benchmark for multi-turn instruction following. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 9322–9341, Vienna, Austria. Association for Computational Linguistics.

Yu-Xiang Lin, Chih-Kai Yang, Wei-Chih Chen, Chen-An Li, Chien-yu Huang, Xuanjun Chen, and Hung-yi Lee. 2025. A preliminary exploration with gpt-4o voice mode. *arXiv preprint arXiv:2502.09940*.

Heyang Liu, Yuhao Wang, Ziyang Cheng, Ronghua Wu, Qunshan Gu, Yanfeng Wang, and Yu Wang. 2025. Vocalbench: Benchmarking the vocal conversational abilities for speech interaction models. *arXiv preprint arXiv:2505.15727*.

Chengqian Ma, Wei Tao, and Steven Y Guo. 2025. C3: A bilingual benchmark for spoken dialogue models exploring challenges in complex conversations. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 22789–22807.

Ziyang Ma, Zhisheng Zheng, Jiaxin Ye, Jinchao Li, Zhifu Gao, Shiliang Zhang, and Xie Chen. 2024. emotion2vec: Self-supervised pre-training for speech emotion representation. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 15747–15760.

Brian W Matthews. 1975. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure*, 405(2):442–451.

Anna Neumann, Elisabeth Kirsten, Muhammad Bilal Zafar, and Jatinder Singh. 2025. Position is power: System prompts as a mechanism of bias in large language models (llms). In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '25, page 573–598, New York, NY, USA. Association for Computing Machinery.

Nvidia. 2025. NVIDIA Speech AI Models Deliver Industry-Leading Accuracy and Performance. https://developer.nvidia.com/blog/nvidia-speech-ai-models-deliver-industry-\leading-accuracy-and-performance/.

OpenAI. 2024. GPT-4o mini: Advancing Cost-Efficient Intelligence. https://openai.com/index/gpt-4o-mini-advancing\-cost-efficient-intelligence/.

OpenAI. 2024. Gpt-4o system card. *Preprint*, arXiv:2410.21276.

OpenAI. 2025. Introducing GPT-5. https://openai.com/index/introducing-gpt-5/.

Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. Meld: A multimodal multi-party dataset for emotion recognition in conversations. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pages 527–536.

Ayesha Qamar, Jonathan Tong, and Ruihong Huang. 2025. Do LLMs understand dialogues? a case study on dialogue acts. In *Proceedings of the 63rd Annual Meeting of the Association for Computational*

Linguistics (Volume 1: Long Papers), pages 26219–26237, Vienna, Austria. Association for Computational Linguistics.

Takaaki Saeki, Detai Xin, Wataru Nakata, Tomoki Koriyama, Shinnosuke Takamichi, and Hiroshi Saruwatari. 2022. Utmos: Utokyo-sarulab system for voicemos challenge 2022. In *Proc. Interspeech 2022*, pages 4521–4525.

Shuzheng Si, Wentao Ma, Haoyu Gao, Yuchuan Wu, Ting-En Lin, Yinpei Dai, Hangyu Li, Rui Yan, Fei Huang, and Yongbin Li. 2023. Spokenwoz: A large-scale speech-text benchmark for spoken task-oriented dialogue agents. *Advances in Neural Information Processing Systems*, 36:39088–39118.

Christian J. Steinmetz and Joshua D. Reiss. 2021. pyloudnorm: A simple yet flexible loudness meter in python. In *150th AES Convention*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Eric Wallace, Kai Xiao, Reimar Leike, Lilian Weng, Johannes Heidecke, and Alex Beutel. 2024. The instruction hierarchy: Training llms to prioritize privileged instructions. *arXiv preprint arXiv:2404.13208*.

Boyong Wu, Chao Yan, Chen Hu, Cheng Yi, Chengli Feng, Fei Tian, Feiyu Shen, Gang Yu, Haoyang Zhang, Jingbei Li, and 1 others. 2025. Step-audio 2 technical report. *arXiv preprint arXiv:2507.16632*.

Jin Xu and 1 others. 2025. Qwen2. 5-omni technical report. *arXiv preprint arXiv:2503.20215*.

Ruiqi Yan, Xiquan Li, Wenxi Chen, Zhikang Niu, Chen Yang, Ziyang Ma, Kai Yu, and Xie Chen. 2025. URO-bench: Towards comprehensive evaluation for end-to-end spoken dialogue models. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 17211–17242, Suzhou, China. Association for Computational Linguistics.

Aohan Zeng, Zhengxiao Du, Mingdao Liu, Kedong Wang, Shengmin Jiang, Lei Zhao, Yuxiao Dong, and Jie Tang. 2024. Glm-4-voice: Towards intelligent and human-like end-to-end spoken chatbot. *arXiv preprint arXiv:2412.02612*.

Jun Zhan, Mingyang Han, Yuxuan Xie, Chen Wang, Dong Zhang, Kexin Huang, Haoxiang Shi, DongXiao Wang, Tengtao Song, Qinyuan Cheng, and 1 others. 2025. Vstyle: A benchmark for voice style adaptation with spoken instructions. *arXiv preprint arXiv:2509.09716*.

A Dataset Construction

To minimize topic-induced variance by averaging across diverse conversation contents, we select samples from the SODA dataset (Kim et al., 2023) to generate conversation openers. Soda is an English dialogue dataset covering a wide range of social interactions. Each dialogue in SODA includes a narrative context, speaker name, and a knowledge graph that defines the events and relationships within the dialogue. Soda is released under the Creative Commons Attribution 4.0 (CC BY 4.0) license. Accordingly, our use of the dataset is compliant with the license requirements, including proper attribution to the original authors.

To obtain the conversation openers, we prompt GPT-5 mini to generate a topic that encompasses the entire discussion, utilizing both the narrative and the first utterance of the dialogue. The prompt is shown in Figure 7. In this step, we ask the model to not only produce conversation openers but also perform filtering. If the model thinks that the dialogue is unsuitable for SLMs to discuss, it returns "no," allowing us to filter it out.

Finally, we manually inspect the generated topics and remove improper ones, such as queries regarding personal preferences or experiences.

B Implementation Details

B.1 Experimental Setup

We set the temperature of evaluated SLMs to 1, as we find that greedy decoding often caused some SLMs, such as GPT-4o, GPT-4o mini, and Step-Audio 2 mini, to produce audio with long silences at the end. For all other hyperparameters, we use the default values provided in the official examples.

For the user simulator, since text-only LLMs tend to generate verbose responses containing text unsuitable for speech, we provide the following instruction to align the outputs with the nature of spoken dialogue: "You are a chatbot. Please start a conversation by opening a new topic. Chat casually and feel free to role-play in different scenarios. If the conversation stalls, you can extend the topic. Keep each response under 20 English words. As this is a spoken dialogue, avoid using words or expressions that cannot be naturally spoken aloud."

B.2 Prompts for Evaluating Dialogue Coherence

The prompt is shown in Figure 8.

B.3 Prompts for Evaluating Recall Rate

The prompt is shown in Figure 9.

C Full Results

C.1 Full Main Results

The IF rate for each turn are shown in Table 5, showing that all evaluated SLMs exhibit style amnesia.

C.2 SLM Default Speaking Styles

We observe that most SLMs show higher consistency when generating Happiness, Neutral tone, and North American English than with other style attributes. This is likely because these attributes match the models' default speaking styles. To verify this, we examine the emotion and accent distributions of samples generated during the speed and volume evaluations. As shown in Table 4, most SLMs tend to perform well with happy or neutral tones and North American accents, which explains the less degradation for these styles.

Model	Emotion		Accent
	Happiness	Neutral	North American
Gemini Live	55.5	36.3	99.0
GPT-4o	52.3	41.4	100.0
GPT-4o mini	68.8	24.5	99.9
Qwen2.5-Omni	76.8	21.7	99.8
Step-Audio 2 mini	93.7	4.0	76.7

Table 4: The distribution of default speaking styles by model.

C.3 Dialogue Coherence Evaluation

We prompt GPT-5 mini to evaluate dialogue coherence, and the results are reported in Table 6. The results suggest that, in general, all SLMs are capable of participating in long conversations.

C.4 Prompt Position

After placing the style instruction at different positions, the IF rate for each turn is shown in Table 8. The results clearly indicate that placing the instruction in user messages yields significantly better performance than placing it in system messages.

C.5 GPT-4o Fails to Return Speech

In our experiments, we find that GPT-4o and GPT-4o mini sometimes return only a text transcription without any synthesized speech. When this occurs, we re-query the models up to three times with different random seeds. However, some samples still fail to produce speech. We report metrics computed only with samples that successfully return speech in Table 7.

Lin et al. (2025) indicate that GPT-4o exhibited a high refusal rate. We hypothesize that the issue stems from the model's internal safety guard mechanisms. Based on our observations, it is possible that these two models return no speech at each round. Since they are proprietary models to which we do not have access, it is difficult for us to definitively address this issue. As the percentage of failed cases is low and the observed style degradation is significant, we believe this issue will not change our conclusions.

D Automatic Judge Validation

D.1 Human Evaluation Setups

To validate the automatic judges for emotion and accent, we hire annotators on Amazon Mechanical Turk to conduct human evaluations. To ensure annotation quality, we set the following requirements for annotators:

- Annotators must be MTurk Masters, which ensures high-quality workers.
- Annotators must have an approval rate higher than 98%.
- Annotators must have more than 10,000 approved tasks.
- Annotators must be located in the United States to ensure familiarity with English.

In addition, we include an attention-check question to ensure that annotators actually listen to the audio before answering. The attention-check is a short audio clip that clearly states the correct answer, so annotators must listen to the audio to respond correctly. Each task contains four evaluation samples and one attention-check sample in random order. If an annotator fails the attention check, we reject the submission and reassign the task to a new annotator. The annotation interface is shown in Figure 10 and Figure 11.

We pay each annotator $0.15 per task to ensure that the payment meets minimum wage standards.

D.2 Inter-Annotator Agreement of Speech Emotion Recognition Datasets

IEMOCAP (Busso et al., 2008), CREMA-D (Cao et al., 2014), and MELD (Poria et al., 2019) are

three widely used speech emotion recognition datasets. These datasets provide large-scale speech emotion annotations through crowdsourcing, making significant contributions to the development of speech emotion recognition models. Emotion in speech is inherently subjective and may be perceived differently by different annotators. As a result, inter-annotator agreement (IAA) is often imperfect. This variability does not indicate annotation noise or low quality, but reflects the inherent subjectivity of emotion perception in speech.

Among them, IEMOCAP and MELD adopt Fleiss' Kappa (Fleiss, 1971) as the IAA indicator. Fleiss' Kappa is designed to assess the reliability of agreement among a fixed number of annotators assigning categorical labels to a set of items, measuring the extent to which the observed agreement exceeds what would be expected by chance. Each sample in the IEMOCAP dataset is annotated by three evaluators, yielding a Fleiss' Kappa of 0.27, whereas the MELD dataset demonstrates a higher agreement of 0.43 among three annotators.

In comparison, the CREMA-D dataset uses Krippendorff's alpha (Krippendorff, 1980) to report the agreement. This metric is used to measure agreement across any number of observers and is particularly robust because it can handle incomplete data and various data scales while accounting for disagreements expected by chance. The authors collect samples with an average of 9.8 annotators per clip and exhibit a Krippendorff's alpha of 0.42. Furthermore, they report a self-consistency rate of approximately 70%, which represents the level of agreement when the same annotator evaluates the same sample at different times. In sum, these scores highlight the inherent subjectivity and complexity of emotional annotation in spoken dialogue.

Prompt for Dataset Construction

Task
You are given two pieces of information about a conversation:
(1) A short narrative context that describes the social situation.
(2) The original first utterance that started the dialogue.
Rewrite the first utterance into a stronger, more natural opening line that better fits the narrative context.

Guidelines
- This sentence will be used as the initial input to start a conversation with an AI assistant. If the given conversation is not appropriate for interacting with an AI. For example, if it's clearly directed toward a specific person, then respond only with "no."
- The utterance should be open-ended, encouraging multi-turn, in-depth discussions rather than prompting a single, definitive response.
- Keep the opener suitable for starting a conversation in this situation.
- Preserve the core intent/topic of the original first utterance when appropriate, but improve clarity, grounding, and engagement.
- You may slightly adjust the angle to better align with the narrative, but do NOT invent new facts beyond what the narrative implies.
- Do not mention "narrative" or "dialogue" or that you are rewriting; just produce the line.
- Output ONLY the rewritten opening line (no numbering, quotes, or extra text).

Narrative:
{narrative in SODA}

Original first utterance:
{first utterance in SODA}

Figure 7: Prompt for dataset construction.

Prompt for Dialogue Coherence Evaluation

Task
Your task is to evaluate the quality, coherence, and naturalness of a dialogue. The dialogue provided involves two participants: a “Referee” and a “Participant”.

Your job is to assess the **Participant’s responses** to the “Referee”. You must evaluate the naturalness, coherence, and overall reasonableness of the **Participant’s replies only**. Do not score the Referee’s sentence.
Focus on whether the Participant’s replies are logical, on-topic, and sound natural in the context of the conversation.

Evaluation Steps
1. ** Analyze the Dialogue Context **
Read the entire dialogue history to understand the conversational flow. Identify the turns belonging to the ’Referee’ and the ’Participant’.
2. ** Evaluate Participant’s Responses **
Review all responses made by the ’Participant’. Evaluate their overall quality based on the following criteria:
- **Coherence:**
Are the replies logically connected to the Referee’s statements? Do they make sense in context, or are they frequently off-topic?
- **Naturalness & Reasonableness:**
Do the replies sound like a real person would say them? Is the content reasonable and appropriate? Do the responses show appropriate depth, or are they overly simplistic/robotic?
3. ** Provide Analysis **
Summarize your findings. Justify your final score by highlighting specific examples of good (coherent, natural) or poor (incoherent, unnatural) responses from the Participant.
4. ** Report the Final Score **
Conclude your evaluation with the following format: Final score: [[score]]. Replace score with an integer in {{score_set}}. Keep the brackets as shown.

Scoring Rubric
- 1: **Completely Incoherent**: The Participant’s replies are semantically unrelated to the Referee’s statements. They are random, nonsensical, or completely off-topic.
- 2: **Mostly Incoherent**: The Participant’s replies are only vaguely related (e.g., catching a keyword but missing the point) or frequently introduce irrelevant topics, making the dialogue logically hard to follow.
- 3: **Partially Coherent**: The Participant’s replies are generally understandable and respond to the Referee, but contain clear logical leaps, topic drift, or semantic inconsistencies.
- 4: **Mostly Coherent**: The Participant’s replies are logical follow-ups and stay on-topic. The dialogue is semantically smooth, with only minor imprecision.
- 5: **Highly Coherent**: The Participant’s replies are semantically tightly-coupled to the Referee’s statements, logically sound, and accurately advance the conversation, making it very fluent."

Dialogue
{dialogue}

Figure 8: Prompt for dialogue coherence evaluation

Prompt for Recall Evaluation

User Instruction (Ground Truth):
{instruction}

Model Response:
{response}

Question:
Please evaluate the Model Response based on the User Instruction. Determine if the model correctly recalled the specific instruction given by the user.

Select one of the following categories:
(A) The response is not answering the Question, is unrelated, meaningless, or avoids the Question.
(B) The response gives an instruction but different from the User Instruction (Ground Truth).
(C) The response answers the question correctly but includes some meaningless sentences that are unrelated to the question.
(D) The response answers and is completely correct regarding the User Instruction.

Return only the single letter of the category (A, B, C, D).

Figure 9: Prompt for Recall Evaluation

Style	Model	Assistant Turn			
		1	2	3	4
Anger	Cascaded Baseline	17.0	14.0	17.0	13.0
	Gemini Live	24.0	12.0	10.0	8.0
	GPT-4o	30.5	24.2	17.9	8.4
	GPT-4o mini	39.0	6.0	6.0	1.0
	Step-Audio 2 mini	1.0	0.0	0.0	0.0
	Qwen2.5-Omni	5.0	1.0	0.0	3.0
Happiness	Cascaded Baseline	85.0	79.0	87.0	82.0
	Gemini Live	91.0	89.0	89.0	92.0
	GPT-4o	82.1	92.6	87.4	85.3
	GPT-4o mini	89.0	90.0	87.0	84.0
	Step-Audio 2 mini	100.0	99.0	98.0	97.0
	Qwen2.5-Omni	71.0	87.0	83.0	77.0
Neutral	Cascaded Baseline	73.0	79.0	75.0	74.0
	Gemini Live	64.0	69.0	75.0	69.0
	GPT-4o	58.8	62.9	59.8	56.7
	GPT-4o mini	45.0	35.0	40.0	27.0
	Step-Audio 2 mini	7.0	3.0	3.0	5.0
	Qwen2.5-Omni	41.0	20.0	17.0	18.0
Sadness	Cascaded Baseline	62.0	58.0	60.0	64.0
	Gemini Live	72.0	54.0	47.0	51.0
	GPT-4o	78.0	65.0	44.0	45.0
	GPT-4o mini	85.0	32.0	18.0	9.0
	Step-Audio 2 mini	17.0	4.0	4.0	1.0
	Qwen2.5-Omni	17.0	4.0	4.0	0.0

Style	Model	Assistant Turn			
		1	2	3	4
North American	Cascaded Baseline	100.0	99.0	100.0	100.0
	Gemini Live	100.0	100.0	100.0	100.0
	GPT-4o	100.0	100.0	100.0	100.0
	GPT-4o mini	100.0	100.0	100.0	100.0
	Step-Audio 2 mini	66.0	78.0	72.0	71.0
	Qwen2.5-Omni	100.0	99.0	100.0	100.0
Indian	Cascaded Baseline	100.0	100.0	99.0	100.0
	Gemini Live	100.0	100.0	100.0	100.0
	GPT-4o	100.0	100.0	98.0	97.0
	GPT-4o mini	89.0	57.0	35.0	26.0
	Step-Audio 2 mini	33.0	21.0	36.0	30.0
	Qwen2.5-Omni	0.0	0.0	0.0	0.0

Style	Model	Assistant Turn			
		1	2	3	4
Loud	Cascaded Baseline	96.0	97.0	98.0	99.0
	Gemini Live	57.0	55.0	60.0	73.0
	GPT-4o	67.0	68.0	59.0	56.0
	GPT-4o mini	77.0	74.0	68.0	61.0
	Step-Audio 2 mini	46.0	50.0	44.0	49.0
	Qwen2.5-Omni	38.0	36.0	36.0	41.0
Quiet	Cascaded Baseline	99.0	99.0	97.0	97.0
	Gemini Live	95.0	93.0	88.0	82.0
	GPT-4o	92.7	94.8	94.8	95.8
	GPT-4o mini	100.0	99.0	99.0	99.0
	Step-Audio 2 mini	56.0	66.0	52.0	51.0
	Qwen2.5-Omni	69.0	63.0	64.0	59.0

Style	Model	Assistant Turn			
		1	2	3	4
Fast	Cascaded Baseline	100.0	100.0	99.0	99.0
	Gemini Live	99.0	97.0	86.0	85.0
	GPT-4o	89.0	87.0	90.0	81.0
	GPT-4o mini	87.0	86.0	83.0	80.0
	Step-Audio 2 mini	89.0	64.0	60.0	62.0
	Qwen2.5-Omni	47.0	49.0	24.0	27.0
Slow	Cascaded Baseline	100.0	100.0	98.0	100.0
	Gemini Live	99.0	99.0	98.0	98.0
	GPT-4o	100.0	96.0	92.9	86.9
	GPT-4o mini	99.0	83.8	80.8	69.7
	Step-Audio 2 mini	81.0	67.0	53.0	42.0
	Qwen2.5-Omni	40.0	66.0	71.0	76.0

Table 5: IF rate across emotion, accent, volume, and speed styles.

Style	Gemini Live	GPT-4o	GPT-4o mini	Qwen2.5-Omni	Step-Audio 2 mini
Anger	3.80	4.19	4.13	4.09	3.53
Happiness	4.18	4.20	4.22	4.23	3.60
Sadness	4.31	4.28	4.20	4.17	3.49
Neutral	4.28	4.21	4.35	4.18	3.59
Fast	4.27	4.26	4.18	4.09	3.51
Slow	4.09	4.22	4.04	4.12	3.50
Loud	4.00	4.19	4.17	4.29	3.49
Quiet	4.18	4.23	4.19	4.05	3.59
Indian	4.17	4.22	4.25	4.36	3.54
North American	4.20	4.18	4.16	4.17	3.61

Table 6: Dialogue Coherence Evaluation Results.

Style	GPT-4o	GPT-4o mini	GPT-4o + recall	GPT-4o mini + recall
Anger	95	100	-	-
Happiness	95	100	-	-
Sadness	100	100	85	100
Neutral	97	100	-	-
Fast	100	99	99	99
Slow	99	89	99	89
Loud	100	100	-	-
Quiet	96	100	-	-
Indian	100	100	86	96
North American	100	100	-	-

Table 7: The number of successfully generated samples after up to three retries

Model	Style	Position	Turn			
			1	2	3	4
GPT-4o	Anger	System	5.0	4.0	1.0	0.0
		User	**30.5**	**24.2**	**17.9**	**8.4**
	Sadness	System	39.0	18.2	11.1	12.2
		User	**78.0**	**65.0**	**44.0**	**45.0**
	Indian	System	50.0	24.0	11.8	5.6
		User	**100.0**	**100.0**	**98.0**	**97.0**
	Fast	System	62.0	61.0	61.6	62.1
		User	**89.0**	**87.0**	**90.0**	**81.0**
	Slow	System	56.0	45.5	38.8	39.5
		User	**100.0**	**96.0**	**92.9**	**86.9**
GPT-4o mini	Anger	System	3.0	0.0	0.0	0.0
		User	**39.0**	**6.0**	**6.0**	**1.0**
	Sadness	System	33.0	8.0	4.0	2.0
		User	**85.0**	**32.0**	**18.0**	**9.0**
	Indian	System	0.0	1.0	0.0	2.0
		User	**89.0**	**57.0**	**35.0**	**26.0**
	Fast	System	54.0	53.0	63.0	62.0
		User	**87.0**	**86.0**	**83.0**	**80.0**
	Slow	System	62.0	55.0	44.0	47.0
		User	**99.0**	**83.8**	**80.8**	**69.7**
Step-Audio 2 mini	Anger	System	0.0	0.0	0.0	0.0
		User	**1.0**	0.0	0.0	0.0
	Sadness	System	4.0	3.0	0.0	**1.0**
		User	**17.0**	**4.0**	**4.0**	**1.0**
	Indian	System	**32.0**	**25.0**	28.0	**30.0**
		User	**32.0**	21.0	**33.0**	28.0
	Fast	System	79.0	**69.0**	**65.0**	**73.0**
		User	**89.0**	64.0	60.0	62.0
	Slow	System	50.0	55.0	51.0	**50.0**
		User	**81.0**	**67.0**	**53.0**	42.0

Table 8: IF rate comparison by prompt position. **Bold** indicates better performance between user and system messages.

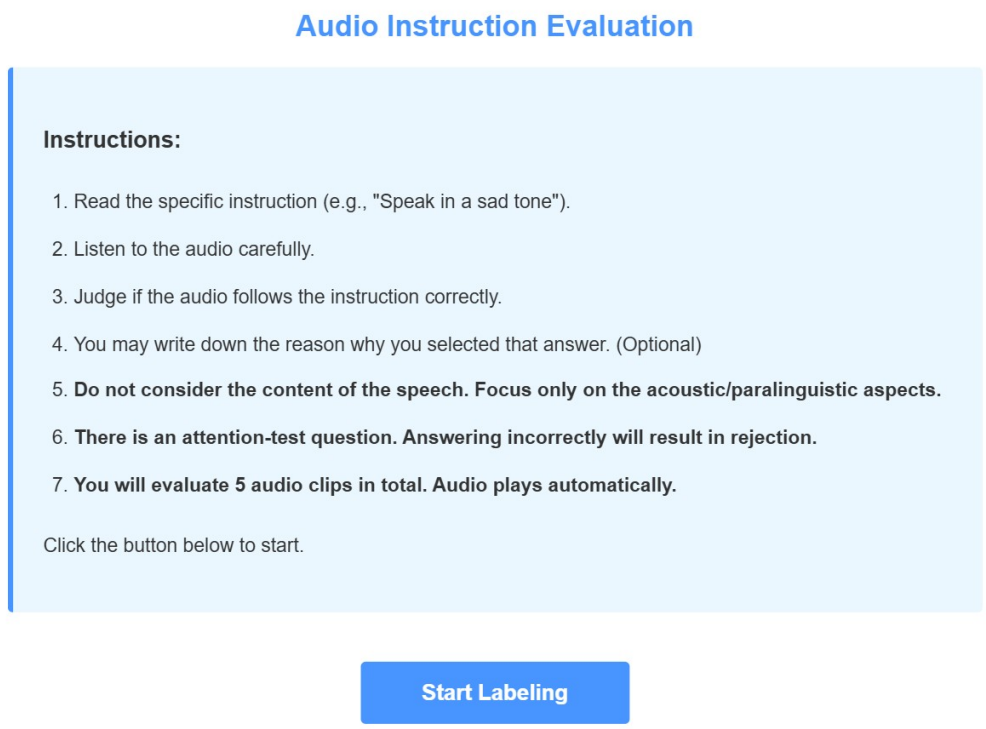

Figure 10: Instruction page

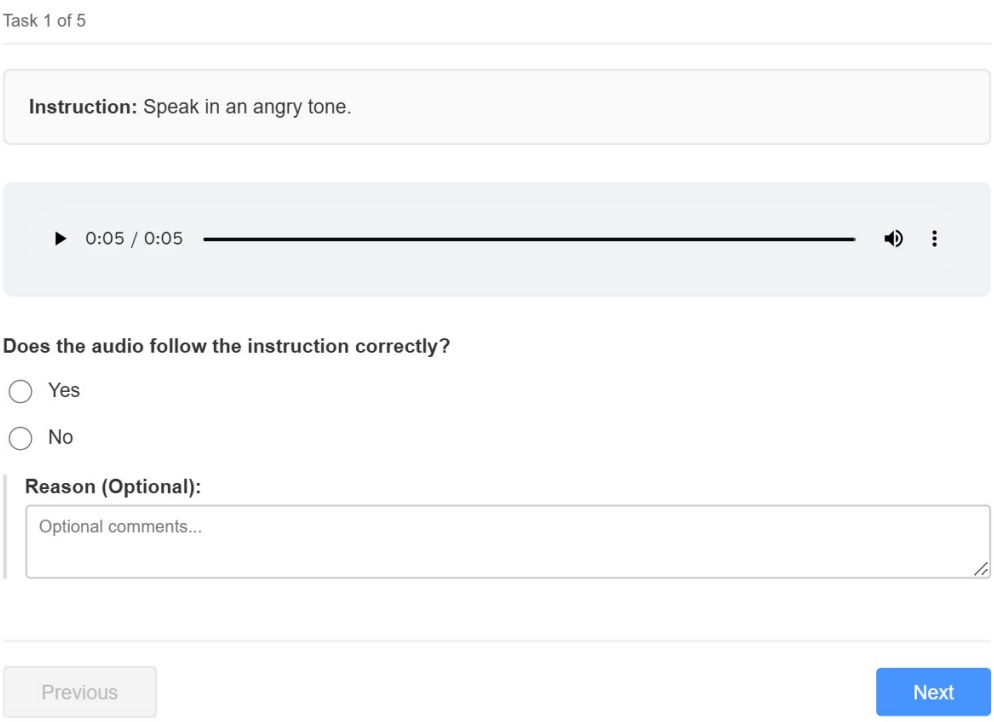

Figure 11: Annotation page