

Evaluating the Expressive Appropriateness of Speech in Rich Contexts

Tianrui Wang^{1,2}, Ziyang Ma^{2,3}, Yizhou Peng², Haoyu Wang¹, Zhikang Niu³,
Zikang Huang¹, Yihao Wu², Yi-Wen Chao², Yu Jiang¹, Yuheng Lu¹,
Guanrou Yang³, Xuanchen Li¹, Hexin Liu², Chunyu Qiang^{1,4}, Cheng Gong⁵,
Yifan Yang³, Tianchi Liu⁶, Junyu Wang¹, Nana Hou², Meng Ge¹,
Fuming You⁷, Wei Yang⁷, Zhongqian Sun⁷, Haifeng Hu⁷, Xiaobao Wang^{1*},
Eng Siong Chng², Xie Chen³, Longbiao Wang^{1*}, Jianwu Dang¹

¹Tianjin Key Laboratory of Cognitive Computing and Application, School of Artificial Intelligence, Tianjin University, ²Nanyang Technological University,
³Shanghai Jiao Tong University, ⁴Kuaishou Technology, ⁵TeleAI, China Telecom,
⁶National University of Singapore, ⁷Tencent

Correspondence: {longbiao_wang, wangxiaobao}@tju.edu.cn

Abstract

Evaluating expressive speech remains challenging, as existing methods mainly assess emotional intensity and overlook whether a speech sample is expressively appropriate for its contextual setting. This limitation hinders reliable evaluation of speech systems used in narrative-driven and interactive applications, such as audiobooks and conversational agents. We introduce **CEAEval**, a Context-rich framework for **Evaluating Expressive Appropriateness** in speech, which assesses whether a speech sample expressively aligns with the underlying communicative intent implied by its discourse-level narrative context. To support this task, we construct CEAEval-D, the first context-rich speech dataset with real human performances in Mandarin conversational speech, providing narrative descriptions together with fifteen dimensions of human annotations covering expressive attributes and expressive appropriateness. We further develop CEAEval-M, a model that integrates knowledge distillation, planner-based multi-model collaboration, adaptive audio attention bias, and reinforcement learning to perform context-rich expressive appropriateness evaluation. Experiments on a human-annotated test set demonstrate that CEAEval-M substantially outperforms existing speech evaluation and analysis systems.

1 Introduction

Automatic speech evaluation has long supported tasks such as data filtering and model optimization (Ribeiro et al., 2011). With the rapid deployment of speech dialogue systems (Ji et al., 2024) and audiobook-generation models (Park et al., 2025), the expressive quality has become a critical factor shaping user experience (Cong et al., 2021).

*Longbiao Wang is the primary corresponding author and Xiaobao Wang is the co-corresponding author.

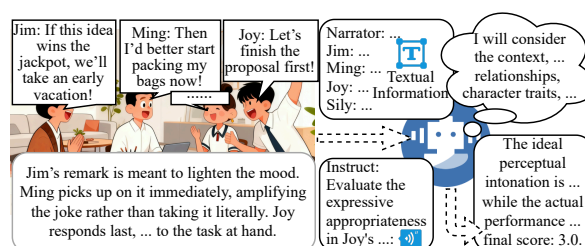


Figure 1: Overview of the proposed context-rich expressive appropriateness evaluation task. The dialogue example is shown in English for illustrative purposes.

Traditional speech evaluation methods primarily focus on word accuracy (Anastassiou et al., 2024), naturalness (Fu et al., 2018), signal quality (Reddy et al., 2021), or emotional intensity (Im et al., 2022) at the utterance level. However, these metrics are insufficient for determining whether speech expressiveness aligns with contextual intent. Expressive appropriateness can only be meaningfully assessed once conversational context and discourse progression are made explicit, as these factors strongly constrain the range of appropriate expressive realizations (Tawari and Trivedi, 2010). As illustrated in Figure 1, Joy’s utterance could be perceived as reproachful or even angry when considered in isolation; however, within the given conversational context, an expressive realization conveying restrained amusement is the most appropriate, which cannot be captured by existing evaluation methods based on emotion classification or intensity prediction (Ma et al., 2024; Zhou et al., 2022).

Despite increasing interest in expressive speech evaluation, existing resources remain insufficient for studying expressive appropriateness under rich contextual settings. As summarized in Table 1, prior benchmarks primarily target speech naturalness, dialogue quality, or role-play interaction, and vary substantially in their use of generated speech

Benchmark / Work	Real Speech	Real Context	Long-range Context	Multiple Turn	CoT-based Reasoning	Number of Annotation Dimensions	Task Focus
WavReward (Ji et al., 2025)	✗	✗	✗	✓	✓	1	Spoken Dialogue Quality
SpeechJudge (Zhang et al., 2025)	✗	✗	✗	✗	✓	2	Speech Naturalness
Speech-DRAME (Shi et al., 2025)	✓	✗	✗	✓	✗	13	Role-play Interaction
SpeechRole (Jiang et al., 2025)	✓	✗	✗	✓	✓	0	Role-play Interaction
CEAEval (ours)	✓	✓	✓	✓	✓	15	Contextual Expressive Appropriateness

Table 1: Comparison of CEAEval with existing expressive speech evaluation benchmarks. Long-range context refers to conversational contexts exceeding 10 dialogue turns.

or context, long-range discourse, and annotation granularity.

From a data perspective, most existing datasets either focus on daily conversational speech with limited expressive range (Yan et al., 2025) or rely on synthesized speech (Ji et al., 2025) and generated contexts to approximate context-expression alignment (Zhan et al., 2025; Jiang et al., 2025). As a result, expressive behavior is often evaluated without grounding in authentic narrative structure or long-range discourse, which are crucial for determining whether a given expressive realization is contextually appropriate.

From a methodological perspective, most existing speech evaluation approaches operate at the single-utterance or short-context level, even when limited multi-turn dialogue is considered (Zhang et al., 2025). Contextual information is often summarized or truncated, which hinders modeling long-range dependencies between speech expressiveness and narrative progression (Shi et al., 2025). Recent work has begun to incorporate large language models (LLM) and chain-of-thought (CoT) reasoning to enable semantic- and discourse-level analysis (Ji et al., 2025; Zhang et al., 2025). However, particularly for base models with limited reasoning capacity, directly applying long textual reasoning to speech evaluation can cause attention to be dominated by text, suppressing speech perception under long-context multimodal inputs (Tian et al., 2025). This fundamentally constrains effective expressive appropriateness evaluation, which requires joint reasoning over both speech and long-range context.

To address these challenges, we introduce CEAEval, a unified framework for context-rich expressive appropriateness evaluation of Mandarin speech under rich contextual settings. CEAEval integrates long-range context modeling, fine-grained expressive perception, and stable reasoning. Our

main contributions are threefold:

- We formalize context-rich expressive appropriateness evaluation and introduce the first human-annotated dataset for this task, based on real Mandarin audiobook speech. The dataset contains long-range narrative context and 15 carefully designed annotation dimensions with high inter-annotator consistency.
- We propose a planner-judge decoupled evaluation framework for expressive appropriateness, which separates long-context textual reasoning from fine-grained perceptual scoring of speech. To alleviate text-dominated reasoning under long-context inputs, we further introduce an adaptive audio attention bias mechanism and reinforcement learning optimization.
- Experiments on the proposed task demonstrate that our method maintains strong agreement with human judgments as contextual length increases, achieving a linear correlation coefficient of 0.72 and an accuracy of 70.8%, while providing interpretable scoring rationales. Our demo, annotated data, model, and code will be released at <https://wangtianrui.github.io/ceaeval/>.

2 Related Work

2.1 General Speech Evaluation

A large body of prior work on speech evaluation focuses on intelligibility (Maniati et al., 2022), naturalness (Mittag et al., 2021), perceived quality (Wang et al., 2025b), and robustness to noise (Reddy et al., 2021), typically supported by human perceptual ratings. While effective for assessing acoustic quality and signal-level properties, these approaches are not designed to evaluate whether a speech sample is expressively appropriate for its contextual setting, as they largely operate

at the utterance level and do not incorporate rich discourse or narrative context.

2.2 Expressive Speech Evaluation and Data

Recent work has begun to explore expressive aspects of speech, including aesthetics and prosody (Yao et al., 2025), emotional cues in dialogue (Yan et al., 2025; Ji et al., 2025), and human preference judgments (Zhang et al., 2025). Some studies further incorporate role- or scene-related information to enrich expressive modeling (Jiang et al., 2025; Shi et al., 2025). Despite these advances, existing expressive speech datasets remain insufficient for context-rich expressive appropriateness evaluation. As summarized in Table 1, fine-grained human annotations under real speech, long-range contextual settings are largely absent.

2.3 Learning-based Speech Evaluation with Large Language Models

Recent speech evaluation methods increasingly adopt learning-based approaches to predict human perceptual judgments. Models such as QualityNet (Fu et al., 2018), DNSMOS (Reddy et al., 2021), and MOSNet (Lo et al., 2019) primarily focus on acoustic quality and signal-level attributes. To incorporate semantic- and discourse-level information, several recent works integrate large language models and CoT-style reasoning into speech evaluation frameworks (Ji et al., 2025; Zhang et al., 2025; Wang et al., 2025b; Jiang et al., 2025; Shi et al., 2025; Manku et al., 2025). However, these approaches typically rely on fine-tuning LLMs on speech modalities, which has been shown to degrade their original text reasoning capabilities (Tang et al., 2024). This limitation poses a fundamental challenge for expressive appropriateness evaluation, which requires robust reasoning over rich textual context and discourse structure.

3 Proposed Method

3.1 Task Definition and Problem Formulation

We study the task of context-rich expressive appropriateness evaluation for Mandarin conversational speech, which aims to assess whether the expressive realization of a spoken utterance aligns with the underlying content, discourse intent, and situational context given its rich narrative and conversational background. Following established principles in Chinese broadcast speech and reading aesthetics (Zhang, 2003), we assess expressive appro-

priateness by jointly considering emotional expression, prosodic realization (intonation and rhythm), recording conditions, and the appropriateness of paralinguistic vocalizations. These factors are evaluated in an integrated manner rather than in isolation, reflecting how expressiveness is perceived in natural speech. Detailed definitions and analyses of how each expressive attribute contributes to overall expressive appropriateness are provided in Appendix A. The evaluation output consists of a scalar appropriateness score ranging from 0 to 5, along with a structured reasoning process that analyzes relevant paralinguistic and expressive cues. Such rationales are intended to provide interpretable references for downstream expressive speech evaluation and generation tasks.

3.2 CEAEval-D: Dataset

Evaluating expressive appropriateness under rich contextual settings requires speech data that jointly provide authentic expressive realizations, long-range narrative context, and reliable human judgments. To support this task, we construct CEAEval-D based on narrated Mandarin audiobooks, which naturally exhibit rich discourse structure, diverse speaker roles, and context-dependent expressive variation. Specifically, we collect 84 audiobook works (including 2 high-quality TTS-generated works), resulting in a total of 3,505 hours of performed speech. From this corpus, we strictly select speech segments that are publicly accessible and suitable for release to construct a subset for manual annotation. These annotated segments are drawn from contiguous portions of each work and are accompanied by complete story texts, enabling reliable context construction and expressive assessment. All manually annotated data are curated in accordance with ethical research and data privacy considerations. The data selected for manual annotation comprise 16.1 hours of speech, split into 14.65 hours for training and 1.45 hours for evaluation, with zero overlap in speech samples.

3.2.1 Weak Annotation

To enable context-rich expressive appropriateness evaluation at scale, we generate weak descriptive annotations for the full 3,505-hour corpus using Qwen3-Omni-Captioner (Ma et al., 2025). These captions provide detail captions of speech and are used to distill descriptive and reasoning capabilities into the judge model.

Prior to manual annotation, we apply an auto-

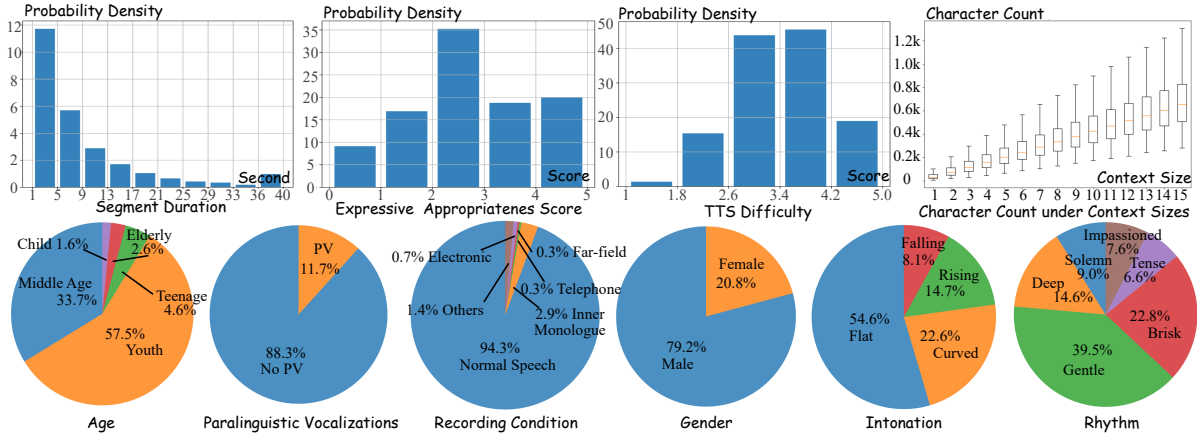


Figure 2: Statistical distribution of annotation categories and attributes in the CEAEval-D dataset.

matic speech recognition (ASR) model (Gao et al., 2023) to pre-segment the 16.1 hours of selected audio and generate preliminary content annotations, which serve as reference material to facilitate and standardize subsequent human annotation.

3.2.2 Manual Annotation

The selected 16.1 hours of speech data are manually annotated to provide reliable supervision for context-rich expressive appropriateness evaluation. Annotation is conducted by 18 native Mandarin-speaking graduate students with backgrounds in speech emotion research, following unified annotation guidelines and a standardized calibration protocol, with inter-annotator agreement verified on a shared calibration subset (see Appendix A).

Within each selected excerpt, speech is further segmented into fine-grained utterances. Each utterance is annotated with a multidimensional set of attributes (detailed in Appendix A), including expressive appropriateness scores, intonation, rhythm, emotion categories, refined textual context, TTS difficulty, recording conditions, background music presence, paralinguistic vocalizations, and sound events. In addition, auxiliary information such as utterance boundaries, refined textual content, and speaker metadata (role name, gender, and age) is also provided. Together, these annotations capture complementary aspects of expressive behavior relevant to appropriateness judgments under rich contextual settings.

As illustrated in Figure 2, we present summary statistics of key annotation dimensions and contextual properties in the dataset, which spans a wide range of context sizes, prosodic patterns, and expressive conditions and thus enables evaluation under diverse discourse settings. In this work, con-

text size (CTS) denotes the number of consecutive dialogue or narrative lines provided as contextual input for a target utterance (see Appendix B), with CTS=0 corresponding to the context-free setting.

3.3 CEAEval-M: Speech-LLM as a Judge

We propose CEAEval-M, a speech-LLM that evaluates expressive appropriateness by jointly reasoning over speech signals and rich textual context. As shown in Figure 3, the model is trained through a three-stage pipeline. First, we distill audio perceptual reasoning abilities from a captioning teacher using 3505-hours data, enabling the model to recognize expressive cues and paralinguistic events in speech. Next, a frozen text-only expressive planner predicts an ideal expressive profile implied by the contextual text, which serves as a reference for appropriate expression. The Speech-LLM judge is then fine-tuned via LoRA to compare the observed speech realization with this planned expressiveness and to produce an appropriateness score in a CoT style, supported by a learnable audio attention bias. Finally, we apply reinforcement learning to further improve scoring robustness and calibration with respect to human judgments.

3.3.1 Expressive Planner

The expressive appropriateness scoring task requires joint reasoning over a speech segment and its narrative context, which may span long sequences of text. As shown in Figure 2, when the context size reaches 15, the accumulated narrative context can exceed 1,200 characters, posing challenges for robust discourse-level modeling. Given the limited amount of textual knowledge available from annotated speech in our 16.1-hour corpus, such supervision is insufficient to bridge the gap be-

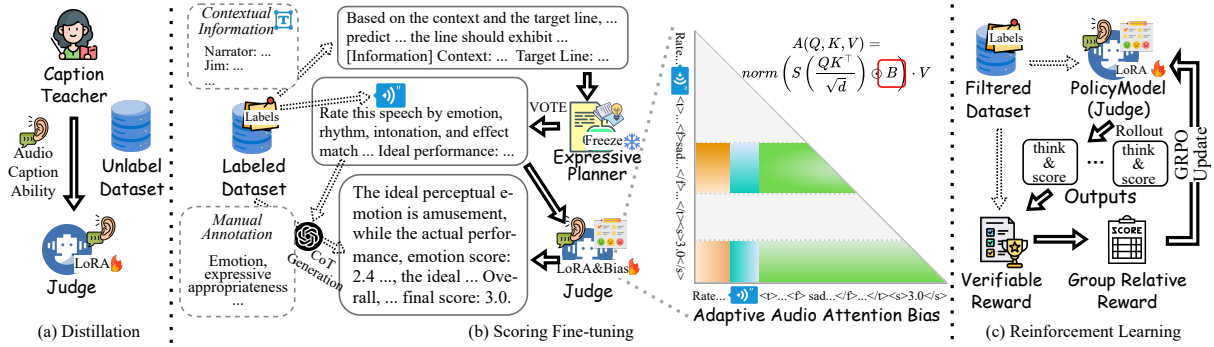


Figure 3: Overview of CEAEval-M, which is trained through a three-stage pipeline for context-rich expressive appropriateness evaluation. Dashed arrows indicate data flow, while solid arrows denote inference or training flow.

tween Omni-style multimodal models and text-only LLMs in long-range textual modeling. We therefore introduce a dedicated text-only large language model, Qwen3-8B (Yang et al., 2025), as an expressive planner to explicitly model narrative context and infer ideal speech expressiveness. The planner takes the narrative context and target utterance as input and predicts an ideal expressive profile covering emotion, rhythm, intonation, and recording condition. To improve robustness under varying context sizes, we construct cumulative context windows ranging from one to fifteen preceding lines and aggregate the resulting predictions through a voting strategy, which mitigates instability caused by context truncation or local ambiguity. Details of the voting procedure are provided in Appendix C.

3.3.2 Knowledge Distillation

To focus the model on expressive speech perception, we use Qwen3-Omni-Captioner to generate weak descriptive annotations for 3,505 hours of speech data and perform distillation-based training with Qwen2.5-Omni-7B (Xu et al., 2025) equipped with LoRA (Hu et al., 2022), as shown in Figure 3.

3.3.3 Scoring Model with CoT Supervision

Building on the distilled backbone, we train a judge model for expressive appropriateness evaluation under rich contextual conditions. The judge conditions on planner outputs together with the input speech and performs expressiveness analysis in a CoT manner, covering emotion, recording conditions, rhythm, intonation, paralinguistic vocalizations, and sound events (see Appendix D). To supervise CoT reasoning, we generate CoT annotations using GPT-4o based on ground-truth scores, manually annotated expressive attributes, and the outputs of the expressive planner (details in Appendix E). The judge model is then fine-tuned with LoRA on

the resulting dataset to learn structured reasoning for expressive appropriateness evaluation.

3.3.4 Adaptive Audio Attention Bias

While CoT supervision improves reasoning transparency, it significantly increases textual length. For base models with limited reasoning capacity, this often induces text-dominated shortcut reasoning, causing the model to under-attend to speech signals (Wang et al., 2025a; Sim et al., 2025; Tian et al., 2025). To counter this effect, we introduce an adaptive audio attention bias into the self-attention computation of the Speech-LLM, as illustrated in Figure 3. The modified attention operation is defined as:

$$A(Q, K, V) = \text{norm} \left(S \left(\frac{QK^\top}{\sqrt{d}} \right) \odot B \right) V, \quad (1)$$

where $A(\cdot)$ denotes the attention operation, $S(\cdot)$ is the softmax function, Q , K , and V denote the query, key, and value matrices, d denotes the hidden dimension, \odot denotes element wise multiplication, and B denotes an adaptive attention bias:

$$B = 2 \cdot \sigma(f_p(X)) \cdot M_p + (1 + \sigma(f_a(X))) \cdot M_a + \sigma(f_{\text{CoT}}(X)) \cdot M_{\text{CoT}} + M_{\text{base}}, \quad (2)$$

where X denotes the input hidden state. f_p , f_a , and f_{CoT} denote learnable linear projections that map the input feature to a scalar, and $\sigma(\cdot)$ denotes the sigmoid function. The four binary masks M_p , M_a , M_{CoT} , and M_{base} respectively indicate system prompt regions, audio regions that require focused attention, CoT regions, and remaining regions that remain unchanged. By adapting attention weights through region-specific bias, the proposed model mitigates audio attention dilution caused by increased textual inputs and improves score predic-

tion accuracy under CoT-style supervision. Implementation details of the attention bias construction and mask definitions are provided in Appendix F.

3.3.5 Reinforcement Learning Optimization

Our final objective is accurate and stable expressive appropriateness score prediction. While CoT-supervised training yields reasonable behavior, classification-style supervision does not explicitly model distances between continuous scores, leading to instability. We therefore optimize the judge model using GRPO (Guo et al., 2025) on a filtered and balanced train-set to directly optimize distance-aware score prediction, as shown in Figure 3. Details of the filtering and resampling strategy are provided in Appendix G. We define a reward function that combines regression accuracy (Ji et al., 2025) and bucket-level ordinal consistency:

$$r(\hat{s}, s) = \exp\left(-\frac{|\hat{s} - s|}{\sigma}\right) + \exp(-|b(\hat{s}) - b(s)|), \quad (3)$$

where \hat{s} and s denote the predicted and ground truth scores, wrapped with $\langle s \rangle$ and $\langle /s \rangle$ in the output sequence, $\sigma = 1.0$, and $b(\cdot)$ maps scores to discrete buckets:

$$b(s) = \min(5, \max(0, \lfloor s \rfloor)). \quad (4)$$

The GRPO objective is defined as:

$$\max_{\theta} \mathbb{E}_{y \sim \pi_{\theta}} \left[\text{clip}(r, -\epsilon, \epsilon) - \beta \text{KL}(\pi_{\theta} \parallel \pi_{\text{ref}}) \right], \quad (5)$$

where $\epsilon = 0.1$ and $\beta = 0.01$. The reference policy π_{ref} shares the same backbone as π_{θ} but excludes LoRA parameters.

4 Experiment Setup

4.1 Data

As described in Section 3.3.2, the distillation stage uses 3,505 hours of unlabeled audiobook speech. For context-rich expressive appropriateness scoring, we annotate 16.1 hours of speech, split into 14.65 hours for training and 1.45 hours for testing, with strict story-level separation and no overlap in narratives, characters, or scenes.

4.2 Models and Training Settings

We adopt Qwen3-8B as the expressive planner and Qwen2.5-Omni-7B-Thinker as the backbone of the judge model. All fine-tuning and reinforcement learning stages use LoRA, with the rank set to 32 and the scaling factor alpha set to 64. The learning

rate linearly increases to a peak value of 5×10^{-6} during the first 10% of training steps and then gradually decays to 5×10^{-7} by the end of training. Training runs on eight NVIDIA A40 GPUs, with a per-GPU batch size of 4. To support multilingual instructions and outputs, we design language-specific system prompts for Chinese and English, with details provided in the Appendix D.

4.3 Baselines and Metrics

Since existing public evaluation models don't support rich narrative context as defined in our task, and our approach adopts commonly used strategies such as supervised fine-tuning, CoT reasoning, and reinforcement learning, we focus our comparison on representative speech-capable models. Specifically, we evaluate Gemma-3n (Team et al., 2024), Midashenglm (Dinkel et al., 2025), Phi-4-MM (Abdin et al., 2024), Qwen2.5-Omni-7B (Xu et al., 2025), Qwen3-Omni-30B-Instruct (Yang et al., 2025), Voxtral-Mini (Liu et al., 2025), Step-Audio-R1 (Tian et al., 2025), Kimi-Audio (KimiTeam et al., 2025), GPT-4o-Audio (Hurst et al., 2024), and Gemini-3-Pro (Comanici et al., 2025).

We evaluate score prediction using the Linear Correlation Coefficient (LCC) (Benesty et al., 2009) and a tolerance-based accuracy metric (ACC) (Ito, 1996), where a prediction is considered correct if the absolute difference between the predicted score and the annotated score is within 1. LCC serves as the primary metric for assessing score consistency, while ACC provides a complementary measure of absolute prediction error.

5 Results

5.1 Context-rich Speech Expressiveness Appropriateness Evaluation

We evaluate different models on contextual speech expressiveness appropriateness assessment under varying context sizes (CTS), as shown in Table 2. To account for prompt language sensitivity, we test both Chinese and English prompts and report the best-performing language for each model. We further compare direct scoring with CoT reasoning and analyze performance trends as CTS increases. Details of context construction and prompt designs are provided in Appendix B and H, and a detailed comparison of model parameter counts is provided in Appendix I.

Across all evaluated models, the context-free baseline (CTS=0) consistently underperforms set-

Model	w/o CoT								w CoT							
	LCC \uparrow				ACC % \uparrow				LCC \uparrow				ACC % \uparrow			
	CTS	0	5	10	15	0	5	10	15	0	5	10	15	0	5	10
Qwen2.5-Omni	0.09	0.16	0.15	0.07	28.01	27.85	27.69	28.99	0.03	0.14	0.11	0.09	28.11	31.43	30.94	28.83
Kimi-Audio	-0.01	0.17	0.19	0.13	35.67	36.32	33.71	31.60	0.04	0.06	0.04	0.03	30.13	29.48	26.22	26.71
Phi-4-MM	-0.01	0.27	0.18	0.10	33.39	33.22	25.20	32.08	0.00	0.07	0.06	0.06	32.74	34.69	30.78	32.90
Gemma-3n	0.15	0.21	0.10	0.06	30.78	40.07	29.15	28.34	0.14	0.10	0.04	0.07	31.27	34.04	34.85	32.84
Step-Audio-R1	0.16	0.12	0.16	0.11	28.83	27.69	26.22	27.20	0.08	0.17	0.03	0.07	22.80	25.73	24.27	24.43
Midashenglm	0.09	0.24	0.16	0.17	22.80	28.99	32.90	29.80	-0.04	0.19	0.19	0.09	21.82	28.83	34.04	35.71
GPT-4o-Audio	0.08	0.09	0.13	0.06	29.99	31.82	31.66	32.73	0.11	0.18	0.19	0.22	23.16	22.76	26.71	27.41
Gemini-3-Pro	0.11	0.14	0.16	0.10	29.41	27.69	26.23	26.41	0.09	0.17	0.14	0.12	22.41	20.43	18.73	21.17
Voxtral-Mini	0.20	0.33	0.23	0.22	32.74	31.11	31.27	30.94	0.04	0.02	0.15	0.06	33.39	32.57	31.76	29.80
Qwen3-Omni	0.21	0.27	0.25	0.28	35.34	34.85	35.18	33.55	0.21	0.30	0.29	0.22	24.42	29.49	32.74	30.13
CEAEval-M	0.54	0.58	0.61	0.61	59.96	62.00	64.11	65.47	0.61	0.69	0.71	0.72	64.33	68.47	70.12	70.80

Table 2: Performance comparison on contextual speech expressiveness appropriateness assessment, evaluated with and without CoT across different context sizes (CTS).

tings with contextual input. As CTS increases from 0 to moderate values, most models show clear improvements in both LCC and ACC, indicating that narrative context plays a critical role in aligning speech expressiveness with communicative intent. This trend underscores the importance of contextual information for expressive appropriateness evaluation, even when existing models cannot fully exploit long-range context.

As CTS increases, most baseline models exhibit a rise-then-decline performance pattern. Moderate context initially improves alignment between speech expressiveness and narrative intent, whereas longer contexts degrade performance. As illustrated in Fig. 2, when CTS exceeds 5, the contextual text rapidly grows beyond 300 characters, causing multimodal models to become increasingly text-dominated and less sensitive to acoustic cues, which degrades evaluation performance (Wang et al., 2025a; Liu et al., 2024; Tian et al., 2025). In contrast, our method achieves consistently higher and more stable performance across all context sizes, reaching an ACC of 70.80% and an LCC of 0.72. By decoupling context modeling from speech scoring via an expressive planner and explicitly rebalancing audio attention in the judge model, our approach enables CoT-style reasoning without overwhelming the speech modality.

5.2 Evaluation Baselines with Planner

To isolate the effect of the expressive planner, we evaluate all baseline models under the planner-assisted setting, as shown in Table 3. Compared to direct contextual conditioning in Table 2, planner-assisted evaluation consistently improves performance and reduces variance across models, indi-

Model	w/o CoT		w CoT	
	LCC \uparrow	ACC % \uparrow	LCC \uparrow	ACC % \uparrow
Gemma-3n	0.21	48.70	0.15	32.74
Kimi-Audio	0.20	48.21	0.09	42.35
Phi-4-MM	0.26	46.58	0.15	46.25
Midashenglm	0.27	44.30	0.19	44.46
Step-Audio-R1	0.21	40.72	0.16	42.18
Qwen2.5-Omni	0.24	49.89	0.27	42.35
Voxtral-Mini	0.32	54.37	0.24	45.86
GPT-4o-Audio	0.25	47.04	0.29	44.95
Gemini-3-Pro	0.22	51.21	0.25	48.81
Qwen3-Omni	0.36	58.17	0.30	49.47
CEAEval-M	0.61	65.47	0.72	70.80

Table 3: Planner-assisted evaluation with and without CoT.

cating that abstracting narrative context into structured expressive plans leads to more reliable expressive appropriateness evaluation. This effect becomes particularly evident under long-context settings: when the number of context segments reaches 15, raw contextual text often exceeds 600 characters, a regime in which baseline models tend to exhibit unstable or text-dominated reasoning. By converting long narrative inputs into semantically grounded plans through multi-context voting, the expressive planner alleviates the burden of long-text processing in scoring models. Notably, the planner operates as a text-only module, allowing speech-centric scoring models to focus on acoustic perception while still benefiting from rich contextual information.

5.3 Ablation Study

We conduct ablation experiments to analyze the contribution of each proposed component, with results summarized in Table 4 and Figure 4. Each

ID	Distill.	CoT	Planner	AttenBias	RL	LCC \uparrow	ACC % \uparrow
(0)	Baseline-Qwen2.5-Omni(w/o-SFT)					0.09	28.83
(1)	No	No	No	No	No	0.45	48.49
(2)	Yes	No	No	No	No	0.53	56.55
(3)	Yes	No	Only15	No	No	0.58	63.47
(4)	Yes	No	VOTE	No	No	0.61	64.11
(5)	Yes	No	GPT4o	No	No	0.63	65.03
(6)	Yes	No	VOTE	No	Yes	0.65	66.86
(7)	Yes	Yes	VOTE	No	No	0.40	49.09
(8)	Yes	Yes+No	VOTE	No	No	0.41	50.17
(9)	Yes	Yes+No	VOTE	No	Yes	0.47	54.44
(10)	Yes	Yes	VOTE	Yes	No	0.61	64.07
(11)	Yes	Yes+No	VOTE	Yes	No	0.64	67.33
(12)	Yes	Yes+No	VOTE	Yes	Yes	0.72	70.80

Table 4: Ablation study on the effects of distillation, chain-of-thought (CoT), planner, attention bias, and reinforcement learning (RL).

model configuration is indexed by its ID in the table. For configurations involving CoT supervision, inference is also performed in a CoT manner to ensure consistency between training and evaluation.

Knowledge distillation. We first examine the effect of knowledge distillation. Comparing ID (1) and ID (2), distilling expressive perception knowledge from Qwen3-Omni into Qwen2.5-Omni leads to a clear improvement in both LCC and ACC. This result indicates that initializing the scoring model with stronger audio caption capability is beneficial for expressive appropriateness evaluation. Notably, Qwen3-Omni already demonstrates competitive performance in our task (Table 2), motivating its use as the teacher model for distillation.

Expressive planner and context size. We analyze the effect of the expressive planner under increasing context length in Figure 4. Without the planner, both Qwen2.5-Omni and our model exhibit unstable performance as context size grows, following a rise-then-decline pattern (green and brown lines). Introducing the expressive planner fundamentally changes this behavior: with planner-based abstraction, performance increases more steadily and gradually converges. This trend is further supported by the quantitative comparison between ID (2) and ID (3) in Table 4, indicating that summarizing long narrative context into structured expressiveness plans reduces the burden of long-context reasoning in the speech-LLM model. To further mitigate variability across context configurations, we introduce a voting mechanism. Comparing ID (3) and ID (4), voting consistently improves stability and overall performance, which is also reflected by smoother trends in Figure 4.




Figure 4: Performance trends under increasing context size.

We additionally evaluate GPT-4o as an alternative expressive planner without voting (ID (5)). Although it yields higher performance, the Qwen3-8B-based voting planner achieves comparable results. We therefore adopt the voting-based planner as the default setting, balancing performance, reproducibility, and deployment cost.

Reinforcement learning. We further introduce reinforcement learning to directly optimize score prediction accuracy. Comparisons between ID (4) and ID (6), ID (8) and ID (9), as well as ID (11) and ID (12), show consistent performance gains from reinforcement learning. These results indicate that distance-aware optimization further improves numerical stability on top of supervised fine-tuning.

CoT and attention bias. Finally, we examine the interaction between CoT-style supervision and audio attention. As shown by ID (7), ID (8), and ID (9), introducing CoT-style reasoning alone leads to noticeable performance degradation from ID (4). This is caused by the increased amount of textual content produced during CoT reasoning, which, given the limited reasoning capacity of the base model (Qwen2.5-Omni), shifts its attention away from the speech modality. To address this issue, we introduce an adaptive attention bias mechanism. Comparisons between ID (7) and ID (10), as well as between ID (8) and ID (11), show that attention bias effectively counteracts the negative impact of CoT supervision and restores performance. Moreover, when CoT and non-CoT supervision are jointly applied during training (ID (11)), the resulting model outperforms its non-CoT counterpart (ID (4)), indicating that CoT provides complementary benefits when its modality imbalance is properly controlled. The introduction of the expressive planner and CoT reasoning not only improves scoring accuracy but also provides high interpretability. A concrete example of this reasoning process is presented in Appendix J, which compares the ideal expressive plan with the actual speech realization.

6 Conclusion

This paper introduces expressive appropriateness evaluation for Mandarin speech under rich narrative context and presents the first systematic study of this problem from the perspectives of task formulation, data construction, and model design. We define expressive appropriateness as the alignment between speech realization and the latent communicative intent implied by contextual narratives, rather than isolated acoustic attributes. To support this task, we construct a real-speech dataset comprising 16.1 hours of contextualized audiobook speech, annotated with 15 carefully designed dimensions and exhibiting high inter-annotator consistency. Building on this dataset, we propose a context-rich evaluation framework that integrates knowledge distillation, planner–scorer decoupling, adaptive audio attention bias, and reinforcement learning. This design enables robust reasoning over long-range narrative context while preserving sensitivity to speech signals, resulting in the predicted expressive appropriateness scores that closely align with human judgments. Beyond audiobook, the proposed framework provides a general diagnostic tool for expressive speech generation in dialogue systems, offering a principled way to assess whether generated speech appropriately reflects contextual intent.

Limitations

While this work presents a comprehensive framework for context-rich expressive appropriateness evaluation, several limitations remain. First, as expressive appropriateness is shaped by language-specific and cultural factors, our current study focuses on Mandarin speech. In future work, we plan to extend the proposed framework to additional languages and cultural contexts and scale up the annotated dataset size, with appropriate adaptation to account for cross-linguistic and cross-cultural variation in expressive intent. Second, we primarily model contextual information from textual narratives, and incorporating speech-level context across neighboring utterances may further enhance expressive evaluation. Finally, although human annotation improves reliability, expressive appropriateness remains inherently subjective, and automatic scores should be interpreted with caution rather than used as the sole criterion for real-world decision-making.

Ethical Statement

Human annotation and fair compensation. All human annotators involved in data creation were native Mandarin speakers with academic backgrounds in speech emotion or affective speech research. Annotators were either legally employed graduate students supported by formal scholarships, and all participated as co-authors of this work. Annotation work was compensated in accordance with local minimum wage regulations and institutional guidelines. This compensation scheme aligns with ACL requirements regarding fair treatment and remuneration of human participants.

Data privacy and consent. We will not release the full 3,505 hours of speech data. Only the manually annotated subset will be made available. All annotated speech segments are carefully reviewed and are derived from publicly accessible, user-uploaded audio content on platforms such as Bilibili. Each released audiobook segment will be shorter than 10 minutes, which is substantially shorter than preview excerpts commonly provided by commercial audiobook platforms, thereby minimizing potential copyright and privacy risks. The released data will not contain any personally identifiable information or sensitive user data.

Licensing and responsible use. The manually annotated dataset subset will be released under a CC-BY-NC license. This license explicitly restricts usage to non-commercial academic research and is consistent with ACL guidelines on ethical dataset release and respect for copyright. Users of the dataset are required to adhere to the license terms and applicable regulations.

Model release for reproducibility. While the full 3,505 hours of speech data will not be publicly released, the distilled model checkpoints and final model parameters trained on this data will be made publicly available. The released models do not contain or expose raw audio, transcripts, or identifiable user information, and are provided solely to support reproducibility and further academic research. Releasing model parameters without distributing the underlying audio data does not constitute data redistribution and is consistent with common practice in speech and language research.

Diversity and representativeness. The annotated dataset includes a diverse range of expressive speech. More than one quarter of the annotated

samples are produced by female speakers, and the data cover a wide variety of expressive and contextual settings. While the dataset focuses on Mandarin speech, this composition reflects a deliberate effort to mitigate representational bias within the targeted linguistic and narrative domain.

Environmental and safety considerations. This work focuses on the evaluation of expressive appropriateness in speech and does not involve the deployment of generative systems in real-world settings. As such, it does not introduce direct safety, security, or environmental risks. Nevertheless, as with other automatic speech evaluation frameworks, there is a potential risk of misuse if the proposed metrics or models are applied in high-stakes decision-making scenarios without appropriate human oversight. We emphasize that CEAEval is intended as a research benchmark and analysis tool, rather than a standalone decision-making system.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. U23B2053). This work was also supported by Tencent and the Tencent-NTU Joint Research Laboratory (CENTURY), Nanyang Technological University, Singapore.

References

1996. Methods for subjective determination of transmission quality. Technical Report ITU-T Recommendation P.800, International Telecommunication Union.
- Marah Abdin, Jyoti Aneja, Harkirat Behl, Sébastien Bubeck, Ronen Eldan, Suriya Gunasekar, Michael Harrison, Russell J Hewett, Mojan Javaheripi, Piero Kauffmann, and 1 others. 2024. Phi-4 technical report. *arXiv preprint arXiv:2412.08905*.
- Philip Anastassiou, Jiawei Chen, Jitong Chen, Yuanzhe Chen, Zhuo Chen, Ziyi Chen, Jian Cong, Lelai Deng, Chuang Ding, Lu Gao, and 1 others. 2024. Seed-tts: A family of high-quality versatile speech generation models. *arXiv preprint arXiv:2406.02430*.
- Jacob Benesty, Jingdong Chen, Yiteng Huang, and Israel Cohen. 2009. *Pearson Correlation Coefficient*, pages 1–4. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, Luke Marris, Sam Petulla, Colin Gaffney, Asaf Aharoni, Nathan Lintz, Tiago Cardal Pais, Henrik Jacobsen, Idan Szpektor, Nan-Jiang Jiang, and 3416 others. 2025. *Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities*. *Preprint, arXiv:2507.06261*.
- Jian Cong, Shan Yang, Na Hu, Guangzhi Li, Lei Xie, and Dan Su. 2021. Controllable context-aware conversational speech synthesis. *arXiv preprint arXiv:2106.10828*.
- Heinrich Dinkel, Gang Li, Jizhong Liu, Jian Luan, Yadong Niu, Xingwei Sun, Tianzi Wang, Qiyang Xiao, Junbo Zhang, and Jiahao Zhou. 2025. Midashenglm: Efficient audio understanding with general audio captions. *arXiv preprint arXiv:2508.03983*.
- Szu-Wei Fu, Yu Tsao, Hsin-Te Hwang, and Hsin-Min Wang. 2018. Quality-net: An end-to-end non-intrusive speech quality assessment model based on blstm. *arXiv preprint arXiv:1808.05344*.
- Zhifu Gao, Zerui Li, Jiaming Wang, Haoneng Luo, Xian Shi, Mengzhe Chen, Yabin Li, Lingyun Zuo, Zhihao Du, and Shiliang Zhang. 2023. *Funasr: A fundamental end-to-end speech recognition toolkit*. In *Interspeech 2023*, pages 1593–1597.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shiron Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, and 1 others. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Chae-Bin Im, Sang-Hoon Lee, Seung-Bin Kim, and Seong-Wan Lee. 2022. Emoq-tts: Emotion intensity quantization for fine-grained controllable emotional text-to-speech. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6317–6321. IEEE.
- Shengpeng Ji, Yifu Chen, Minghui Fang, Jialong Zuo, Jingyu Lu, Hanting Wang, Ziyue Jiang, Long Zhou, Shujie Liu, Xize Cheng, and 1 others. 2024. Wavchat: A survey of spoken dialogue models. *arXiv preprint arXiv:2411.13577*.
- Shengpeng Ji, Tianle Liang, Yangzhuo Li, Jialong Zuo, Minghui Fang, Jinzheng He, Yifu Chen, Zhengqing Liu, Ziyue Jiang, Xize Cheng, and 1 others. 2025. Wavreward: Spoken dialogue models with generalist reward evaluators. *arXiv preprint arXiv:2505.09558*.

- Changhao Jiang, Jiajun Sun, Yifei Cao, Jiabao Zhuang, Hui Li, Baoyu Fan, Tao Ji, Tao Gui, and Qi Zhang. 2025. Speechrole: A large-scale dataset and benchmark for evaluating speech role-playing agents. *arXiv preprint arXiv:2508.02013*.
- KimiTeam, Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen, Wei Song, Xu Tan, Heyi Tang, Zhengtao Wang, Chu Wei, Yifei Xin, Xinran Xu, Jianwei Yu, Yutao Zhang, Xinyu Zhou, Y. Charles, and 21 others. 2025. [Kimi-audio technical report](#). *Preprint*, arXiv:2504.18425.
- Alexander H Liu, Andy Ehrenberg, Andy Lo, Clément Denoix, Corentin Barreau, Guillaume Lamplé, Jean-Malo Delignon, Khyathi Raghavi Chandu, Patrick von Platen, Pavankumar Reddy Muddireddy, and 1 others. 2025. Voxtral. *arXiv preprint arXiv:2507.13264*.
- Shi Liu, Kecheng Zheng, and Wei Chen. 2024. Paying more attention to image: A training-free method for alleviating hallucination in vlms. In *European Conference on Computer Vision*, pages 125–140. Springer.
- Chen-Chou Lo, Szu-Wei Fu, Wen-Chin Huang, Xin Wang, Junichi Yamagishi, Yu Tsao, and Hsin-Min Wang. 2019. Mosnet: Deep learning-based objective assessment for voice conversion. *Interspeech 2019*.
- Ziyang Ma, Ruiyang Xu, Zhenghao Xing, Yunfei Chu, Yuxuan Wang, Jinzheng He, Jin Xu, Pheng-Ann Heng, Kai Yu, Junyang Lin, Eng Siong Chng, and Xie Chen. 2025. [Omni-captioner: Data pipeline, models, and benchmark for omni detailed perception](#). *Preprint*, arXiv:2510.12720.
- 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.
- Georgia Maniati, Alexandra Vioni, Nikolaos Ellinas, Karolos Nikitaras, Konstantinos Klapsas, June Sig Sung, Gunu Jho, Aimilios Chalamandaris, and Pirros Tsiakoulis. 2022. Somos: The samsung open mos dataset for the evaluation of neural text-to-speech synthesis. *arXiv preprint arXiv:2204.03040*.
- Ruskin Raj Manku, Yuzhi Tang, Xingjian Shi, Mu Li, and Alex Smola. 2025. Emergenttts-eval: Evaluating tts models on complex prosodic, expressiveness, and linguistic challenges using model-as-a-judge. *arXiv preprint arXiv:2505.23009*.
- Kenneth O McGraw and Seok P Wong. 1996. Forming inferences about some intraclass correlation coefficients. *Psychological methods*, 1(1):30.
- Mary L McHugh. 2012. Interrater reliability: the kappa statistic. *Biochemia medica*, 22(3):276–282.
- Gabriel Mittag, Babak Naderi, Assmaa Chehadi, and Sebastian Möller. 2021. Nisqa: A deep cnn-self-attention model for multidimensional speech quality prediction with crowdsourced datasets. *arXiv preprint arXiv:2104.09494*.
- Saif M Mohammad. 2025. Nrc vad lexicon v2: Norms for valence, arousal, and dominance for over 55k english terms. *arXiv preprint arXiv:2503.23547*.
- Kyeongman Park, Seongho Joo, and Kyomin Jung. 2025. Multiactor-audiobook: Zero-shot audiobook generation with faces and voices of multiple speakers. *arXiv preprint arXiv:2505.13082*.
- Chandan KA Reddy, Vishak Gopal, and Ross Cutler. 2021. Dnsmos: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6493–6497. IEEE.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Flávio Ribeiro, Dinei Florêncio, Cha Zhang, and Michael Seltzer. 2011. Crowdmos: An approach for crowdsourcing mean opinion score studies. In *2011 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 2416–2419. IEEE.
- Jiatong Shi, Jionghao Han, Yichen Lu, Santiago Pascual, Pengfei Wu, Chenye Cui, Shinji Watanabe, Chao Weng, and Cong Zhou. 2025. Speech-drame: A framework for human-aligned benchmarks in speech role-play. *arXiv preprint arXiv:2511.01261*.
- Mong Yuan Sim, Wei Emma Zhang, Xiang Dai, and Biaoan Fang. 2025. Can vlms actually see and read? a survey on modality collapse in vision-language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 24452–24470.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun MA, and Chao Zhang. 2024. [SALMONN: Towards generic hearing abilities for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Ashish Tawari and Mohan Manubhai Trivedi. 2010. Speech emotion analysis: Exploring the role of context. *IEEE Transactions on multimedia*, 12(6):502–509.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, and 1 others. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295*.

Fei Tian, Xiangyu Tony Zhang, Yuxin Zhang, Haoyang Zhang, Yuxin Li, Daijiao Liu, Yayue Deng, Donghang Wu, Jun Chen, Liang Zhao, Chengyuan Yao, Hexin Liu, Eng Siong Chng, Xuerui Yang, Xiangyu Zhang, Daxin Jiang, and Gang Yu. 2025. *Step-audio-r1 technical report*. *Preprint*, arXiv:2511.15848.

Junyu Wang, Ziyang Ma, Zhengding Luo, Tianrui Wang, Meng Ge, Xiaobao Wang, and Longbiao Wang. 2025a. Pay more attention to audio: Mitigating imbalance of cross-modal attention in large audio language models. *arXiv preprint arXiv:2509.18816*.

Siyin Wang, Wenyi Yu, Xianzhao Chen, Xiaohai Tian, Jun Zhang, Lu Lu, Yu Tsao, Junichi Yamagishi, Yuxuan Wang, and Chao Zhang. 2025b. Qualispeech: A speech quality assessment dataset with natural language reasoning and descriptions. *arXiv preprint arXiv:2503.20290*.

Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, 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: A comprehensive benchmark for end-to-end spoken dialogue models. *arXiv preprint arXiv:2502.17810*.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Jixun Yao, Guobin Ma, Huixin Xue, Huakang Chen, Chunbo Hao, Yuepeng Jiang, Haohe Liu, Ruibin Yuan, Jin Xu, Wei Xue, and 1 others. 2025. Songeval: A benchmark dataset for song aesthetics evaluation. *arXiv preprint arXiv:2505.10793*.

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*.

Song Zhang. 2003. *Chinese broadcasting announcing*. Communication University of China Press.

Xueyao Zhang, Chaoren Wang, Huan Liao, Ziniu Li, Yuancheng Wang, Li Wang, Dongya Jia, Yuanzhe Chen, Xiulin Li, Zhuo Chen, and 1 others. 2025. Speechjudge: Towards human-level judgment for speech naturalness. *arXiv preprint arXiv:2511.07931*.

Kun Zhou, Berrak Sisman, Rajib Rana, Björn W Schuller, and Haizhou Li. 2022. Emotion intensity and its control for emotional voice conversion. *IEEE Transactions on Affective Computing*, 14(1):31–48.

A Data Annotation and Inter-Annotator Reliability

We recruit 18 native Mandarin-speaking graduate students with backgrounds in speech emotion and speech perception research to annotate 16.1 hours of speech data following unified annotation guidelines. The annotator group consists of 11 male and 7 female participants. Before annotation, all annotators receive detailed instructions on the annotation task and interface, illustrated in Figure 5, Table 5, and Table 6. Annotators are informed that the data are used solely for scientific research purposes. Each annotator labels 4 to 5 stories, and the complete annotation process takes approximately two weeks per participant. The annotation interface supports synchronized playback of speech and its surrounding textual context, enabling annotators to consider discourse-level narrative context when making judgments.

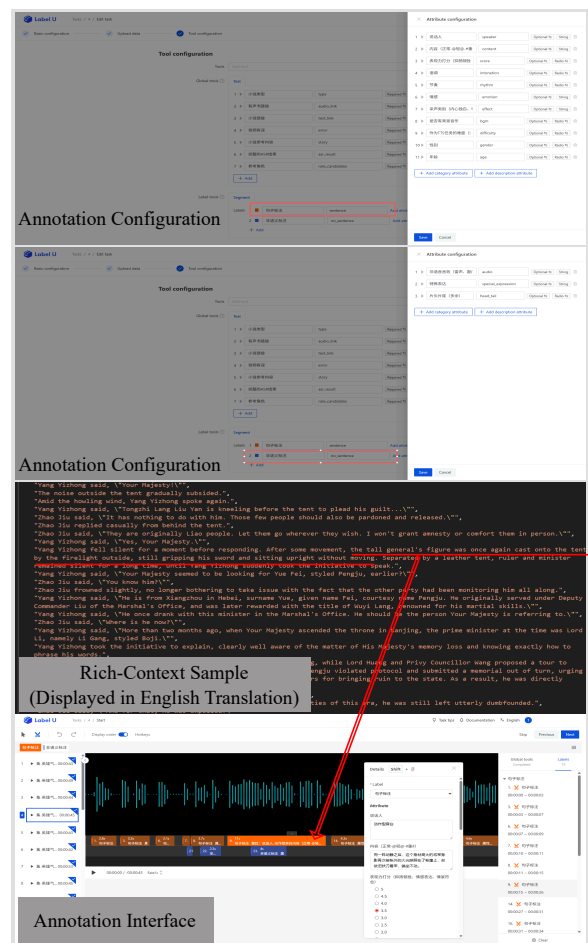


Figure 5: Annotation interface and configuration used in this work. The figure shows the annotation configuration panels, an example of the rich narrative context presented to annotators, and the annotation interface.

Score	Description
0-1	Clearly inappropriate. The expressive realization is severely misaligned with the narrative context. Emotion, prosody, or paralinguistic behavior is clearly incorrect or contradictory, seriously disrupting comprehension or immersion.
1-2	Weak or unnatural. Partial expressive intent is present, but major mismatches remain. Problems such as inappropriate emotional tone, unnatural intonation or rhythm, or conflicting expressive cues limit contextual appropriateness.
2-3	Somewhat appropriate with issues. The overall expressive direction is roughly correct, but inconsistencies in emotion, prosodic realization, or expressive emphasis reduce coherence with the narrative context.
3-4	Generally appropriate with room for improvement. Expressive realization largely matches the narrative context and communicative intent. Minor issues in emotional nuance, timing, or prosodic smoothness remain.
4-5	Highly appropriate, natural, and expressive. Emotion, prosody, and paralinguistic cues are well coordinated and fully aligned with the narrative context, resulting in a fluent and convincing expressive realization.

Table 5: Scoring criteria for expressive appropriateness.

Score	Description
0-1	Very easy. The utterance requires minimal expressive variation and neutral delivery. Correct rendering can be easily achieved.
1-2	Easy. Limited expressive control is required, such as slight emphasis or mild emotional coloring. Overall delivery remains straightforward.
2-3	Moderate difficulty. The utterance involves noticeable expressive elements, such as clear emotional cues or prosodic variation. Careful but manageable control is required.
3-4	Difficult. The utterance demands precise expressive modulation, including nuanced emotion, timing, or intonation. Achieving a natural rendering is challenging.
4-5	Very difficult. The utterance requires complex and fine-grained expressive control, such as subtle emotional shifts, layered prosody, or strong context dependence. It is hard to render naturally.

Table 6: Scoring criteria for TTS difficulty.

Annotation dimensions and their relation to expressive appropriateness. Expressive appropriateness is annotated as an integrated perceptual judgment rather than as a deterministic combination of isolated attributes. Annotators are instructed to assess whether the expressive realization of a speech utterance appropriately reflects the commu-

nicative intent implied by its contextual narrative. Following established principles in Chinese broadcast speech and reading aesthetics (Zhang, 2003), several expressive attributes are annotated to provide structured support for this holistic judgment. As summarized in Table 7, the dataset provides fine-grained annotations across 15 distinct dimensions, covering appropriateness, prosody, emotion, text, speaker metadata, and environmental factors.

Category	Annotation Dimensions
Perceptual Judgment	1. Overall Expressive Score 2. TTS Difficulty
Acoustic & Prosody	3. Intonation 4. Rhythm
Emotion & Intent	5. Emotion 6. Paralinguistic Vocalizations
Context & Text	7. Refined Textual Context 8. Refined Textual Content 9. Utterance Boundaries
Speaker Metadata	10. Speaker Role Name 11. Speaker Age 12. Speaker Gender
Environment	13. Recording Conditions 14. Background Music Presence 15. Sound Events

Table 7: Overview of the 15 annotation dimensions in CEAEval-D.

Emotional expression is annotated using open-ended textual descriptions. Annotators are allowed to freely describe perceived emotions (e.g., happy, angry, sad) as well as compound or dynamic emotional states (e.g., calm turning into excitement), reflecting the continuous and evolving nature of expressive affect in narrative speech. Prosodic realization is characterized along two dimensions, intonation and rhythm, following the taxonomy in (Zhang, 2003). Intonation is categorized into four types: flat, rising, curved, and falling, capturing overall pitch movement patterns. Rhythm is categorized into six types: brisk, heavy, low-paced, high-energy, relaxed, and tense, reflecting differences in speech tempo, energy, and stress distribution. Recording conditions are annotated using open-ended textual descriptions to capture perceptual factors such as far-field or telephone that may influence expressive perception. Paralinguistic vocalizations and sound events are also annotated in free-form text, covering non-verbal cues such as laughter, gasps, sighs, breath noises, or other expressive sounds. In addition, TTS difficulty is annotated to indicate the degree of expressive control required to render an utterance appropriately un-

der its context. Unlike expressive appropriateness, which reflects perceptual outcome, TTS difficulty captures expressive complexity from a production perspective.

Calibration and reliability analysis. Before large-scale annotation, we conduct a calibration phase in which all annotators independently label the same 14.8-minute subset of data under identical guidelines. Centralized feedback is provided to align annotators’ interpretations of the scoring criteria and reduce subjective variability. Inter-annotator agreement statistics computed on this calibrated subset are summarized in Table 8.

For continuous annotations, including expressive appropriateness scores, TTS difficulty, and emotion ratings, we measure inter-annotator reliability using ICC(2,1) (McGraw and Wong, 1996), which captures absolute agreement under a two-way random-effects model. Emotion agreement is computed by mapping categorical emotion descriptions into the Valence–Arousal–Dominance (VAD) space (Mohammad, 2025), computing ICC separately for each dimension, and averaging the results. For categorical attributes such as intonation, rhythm, speaker age, background music presence, and speaker gender, we report percent agreement (McHugh, 2012) based on the majority label. For paralinguistic vocalizations annotated in free-form text, agreement is quantified using an embedding-based semantic similarity (Reimers and Gurevych, 2019) measure, defined as the average pairwise cosine similarity among annotators’ textual descriptions.

Overall, the agreement scores indicate a high level of consistency across annotation dimensions. Expressive appropriateness scoring achieves an **ICC of 0.87**, and emotion annotations exhibit an average **ICC of 0.93** in VAD space. Most categorical attributes exceed **0.9 in percent agreement**, demonstrating that the annotation interface, training procedure, and calibration protocol together support reliable multi-dimensional annotation for context-rich expressive appropriateness evaluation.

B Context Construction and Context Size

We construct a local context window composed of multiple neighboring dialogue or narrative lines. The context is represented as an ordered list of text lines, where each item corresponds to a single utterance or narration segment in the surrounding story, as shown in Figure 5. The target line itself is

Type	Annotation	Metric	Value \uparrow
Numeric	Expr. App. Score	ICC(2,1)	0.867
	TTS Difficulty	ICC(2,1)	0.810
	Emotion (VAD)	ICC(2,1)	0.934
Categorical	Intonation	Pct. Agr.	0.831
	Rhythm	Pct. Agr.	0.915
	Age	Pct. Agr.	0.981
	BGM	Pct. Agr.	0.990
	Gender	Pct. Agr.	0.994
Textual	Recording Cond.	Agreement	0.990
Textual	Paraling. Vocal.	Agreement	0.907

Table 8: Inter-annotator agreement on a 14.8-minute calibration set annotated by 18 annotators. ICC(2,1) is reported for numeric annotations, percent agreement (Pct. Agr.) for closed-set categorical annotations, and embedding-based agreement for textual annotations.

treated separately from the contextual input.

The parameter context size (CTS) specifies the number of surrounding context lines provided in addition to the target line. When CTS is set to 0, the input consists solely of the target line, corresponding to the context-free setting. For CTS > 0 , the context contains exactly C neighboring text lines, with a preference for lines immediately preceding the target line. If sufficient preceding lines are available, the context includes the C lines immediately before the target line. When the target line appears near the beginning of a dialogue or narrative and fewer than C preceding lines exist, the context window is expanded forward to include subsequent lines, ensuring that the total number of context lines remains fixed at C . By varying the context size, we control the amount of discourse-level information available to the evaluation model, ranging from isolated utterance evaluation (CTS=0) to richer narrative-level conditioning.

C Contextual Prompting and Voting Strategy for the Expressive Planner

This section describes the contextual prompting and voting strategy used by the expressive planner.

Contextual prompting. For each target utterance, the planner predicts an ideal expressive profile conditioned on the textual narrative context. For a given target line, multiple planner inputs are constructed by varying the context size (CTS) from 1 to 15. Each CTS configuration, together with the same target line, is provided as an independent input to the planner using a fixed prompt template as below:

System Instruction: Expressive Planning from Narrative Context

You are an expert in speech expressiveness analysis for narrative-driven speech. Your task is to infer the ideal expressive realization of a target utterance based solely on the provided textual narrative context.

Given the surrounding narrative context and the target utterance, analyze the implied communicative intent, character state, and discourse role, and predict how the utterance should be expressed in speech.

Please consider:

1. Narrative progression, character relationships, and situational context.
2. Implied emotional state and possible emotional shifts.
3. Expressive delivery style and recording condition, including speaking distance, inner monologue, and sound-related delivery effects (e.g., phone speech, distant speech, intermittent effects).

[Input]

Narrative Context: %s

Target Utterance: %s

[Output Requirements]

Output exactly one expressive plan in the following JSON format.

The fields emotion and recording condition should be described in concise natural language (free-form), while rhythm and intonation should be selected from the predefined categories.

```
{  
  "emotion": "<free-form description, e.g., calm,  
  angry, frightened, tender, conflicted>",  
  "rhythm": "brisk / heavy / low-paced / high-  
  energy / relaxed / tense",  
  "intonation": "flat / rising / curved / falling",  
  "recording condition": "<freeform description,  
  e.g., normal speaking, phone speech, distant  
  voice, inner monologue>"  
}
```

Do not include any explanation or additional text outside the JSON object.

Voting strategy. For each target utterance, the planner produces 15 expressive plans. Voting is performed over the joint expressive plan rather than over individual attributes. Specifically, each output is treated as a four-element combination of emo-

tion, rhythm, intonation, and recording condition. Identical combinations predicted under different context spans are grouped and counted. The final expressive plan is selected as the combination with the highest frequency across all context variants. In the event of a tie, the plan predicted under the longest context span is selected. This strategy favors expressive plans that are stable across varying amounts of narrative context.

D Simplified Prompt with Expressive Planner

Under our framework, the system prompt for expressive appropriateness evaluation is substantially simplified by conditioning the judge model on the output of the expressive planner. Rather than requiring the model to infer the expected expressive intent directly from a long narrative context, the planner provides a structured description of the ideal expressive realization for the target utterance. The judge then focuses on comparing the actual speech signal against this planned expressiveness. During both training and inference, the judge can be prompted either to directly output a score (w/o CoT) or to generate an explicit reasoning process before producing the final score (w/ CoT), enabling controlled evaluation of the effect of chain-of-thought reasoning.

Simplified system prompt conditioned on expressive planner output

You are given a speech segment and an expressive plan that describes how this utterance should ideally be expressed in context.

Your task is to evaluate whether the actual speech realization matches the planned expressiveness.

Please analyze the speech with respect to the following aspects: emotion, rhythm, intonation, paralinguistic vocalizations, recording condition, and overall naturalness.

Assign a single expressive appropriateness score between 0.0 and 5.0.

Scoring Guidelines:

0–1: Clearly mismatched with the expressive plan.

1–2: Weak, unnatural, or largely monotone expression.

2–3: Roughly matches the plan but with noticeable issues.

3–4: Generally appropriate with minor flaws.
 4–5: Highly appropriate, vivid, and natural.
 # Ideal Expressive Plan:
 emotion: {ideal_emotion}
 rhythm: {ideal_rhythm}
 intonation: {ideal_intonation}
 recording_condition:
 {ideal_recording_condition}
 % w/o CoT setting:
 % Give the final expressive appropriateness
 score only.
 % w/ CoT setting:
 % Think step by step and explain your analysis
 before giving the final score.

E Chain-of-Thought Generation Prompt

Before training, we generate CoT supervision using GPT-4o, conditioned on ground-truth expressive scores, manually annotated expressive attributes, and the voted outputs of the expressive planner. The final score is provided as a condition, and the model is instructed to generate a reasoning process that explains expressive alignment and mismatch leading to this score. The prompt used for CoT generation is shown below.

Prompt for CoT generation conditioned on ground-truth score

You are given a spoken utterance together with its expressive evaluation result. Your task is to generate a clear and coherent chain-of-thought that explains why the speech receives the given expressive appropriateness score.
 The final score is already provided and should be used only as a condition. Do NOT predict, restate, or emphasize the score itself.
 In your reasoning, compare the ideal expressive intent with the actual speech realization, and analyze expressive alignment or mismatch from the following perspectives: emotion, rhythm, intonation, recording condition, and any relevant paralinguistic or sound-related effects.
 After comparing the two performances along each dimension, assign a score for that dimension.
 # Scoring Explanation:
 0–1: Clearly mismatched with the ideal expression.

1–2: Weak, unnatural, or largely monotone expression.
 2–3: Roughly matches the ideal expression but with noticeable issues.
 3–4: Generally appropriate with minor flaws.
 4–5: Highly appropriate, vivid, and natural.
 [Input]
 Target utterance: {target_line}
 Ideal expressive plan:
 emotion={ideal_emotion}
 rhythm={ideal_rhythm}
 intonation={ideal_intonation}
 recording_condition={ideal_record_condition}
 Actual expressive attributes (human annotation):
 emotion={actual_emotion}
 rhythm={actual_rhythm}
 intonation={actual_intonation}
 recording_condition={actual_record_condition}
 Ground-truth expressive score: {actual_score}
 (Optional) Paralinguistic or sound-related events: {linguistic_sounds}
 [Output]
 Generate a natural-language chain-of-thought that explains the expressive appropriateness of the speech under the given context, explicitly providing a score for each expressive dimension. Do not output the final score.

Here, the ideal expressive attributes are provided by the expressive planner and aggregated via voting. By conditioning CoT generation on both ideal and actual expressive attributes, the model is encouraged to explicitly reason about expressive consistency and deviation across dimensions. The generated CoT texts are translated into Chinese and used together with the original texts in bilingual training.

F Audio Attention Bias

As discussed in the main text, the proposed audio attention bias mechanism modulates attention strength according to the different token regions, with the goal of mitigating text-dominant reasoning under CoT-style supervision. This section details the construction of the region masks M_p , M_a , M_{CoT} , and M_{base} in Eq. (1), as well as their dynamic activation during autoregressive inference.

Sequence structure and region annotation. Following the sequence formulation illustrated in Fig-

ure 3, the input to the Speech-LLM is augmented with special boundary tokens that explicitly mark semantically distinct regions:

$$[P_0, \dots, P_n, \langle a \rangle, S_0, \dots, S_m, \langle /a \rangle, \langle bos \rangle, \langle t \rangle, T_0, \dots, \langle f \rangle, T_i, \dots, \langle /f \rangle, \dots, \langle /t \rangle, \langle s \rangle, 3.0, \langle /s \rangle]$$

where P_0, \dots, P_n denote system prompt tokens, S_0, \dots, S_m denote audio tokens enclosed by $\langle a \rangle$ and $\langle /a \rangle$, and T_0, \dots denote chain-of-thought tokens enclosed by $\langle t \rangle$ and $\langle /t \rangle$. Within the chain-of-thought region, the token pair $\langle f \rangle$ and $\langle /f \rangle$ marks expressive focus spans explicitly refer to speech-dependent attributes (e.g., actual emotion, intonation, or paralinguistic cues) and therefore require increased attention to the audio modality. The score value is enclosed by $\langle s \rangle$ and $\langle /s \rangle$ to indicate the score prediction stage.

Region masks. Based on the above sequence structure, we define four mutually exclusive binary region masks according to the explicit boundary tokens. The system prompt mask M_p covers the system prompt region $[P_0, \dots, P_n]$. The audio mask M_a corresponds exclusively to the audio token region $[\langle a \rangle, S_0, \dots, S_m, \langle /a \rangle]$. The CoT mask M_{CoT} marks the reasoning region $[\langle t \rangle, T_0, \dots, \langle /t \rangle]$, including all internal expressive analysis tokens. All remaining tokens are assigned to the base mask M_{base} .

Dynamic bias activation. During autoregressive inference, region-specific attention bias terms in Eq. (1) are activated according to the boundary tokens encountered in the input sequence. When the model enters a region marked by a start token (e.g., $\langle f \rangle$ or $\langle t \rangle$), the corresponding bias component becomes effective for subsequent positions. When the matching end token is reached, the bias is deactivated and attention reverts to the base setting governed by M_{base} . The magnitude of each bias component is dynamically predicted from the current hidden representation via the learnable projections f_p , f_a , and f_{CoT} , allowing the model to adapt attention strength based on contextual needs.

Figure 6 visualizes the resulting attention bias matrices at different transformer layers under CoT-style inference, illustrating how the proposed mechanism dynamically rebalances attention toward audio-related regions during expressive appropriateness scoring.

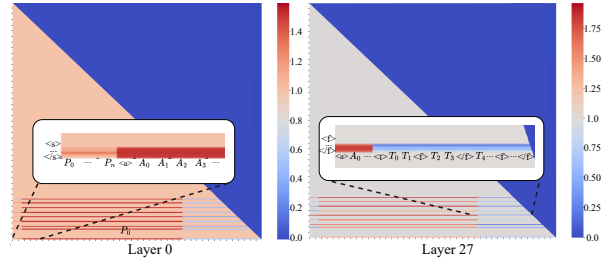


Figure 6: Visualization of 0th and 27th transformer layer attention bias matrices during inference.

G Filtered Training Set for Reinforcement Learning

For reinforcement learning optimization, we apply a filtering and resampling strategy to improve training stability. Speech samples shorter than 1 second or longer than 45 seconds are removed to exclude unreliable or extreme-duration cases. To mitigate score imbalance during policy optimization, we further perform score-balanced resampling, where samples are grouped into integer score bins and resampled to ensure approximately uniform bin frequencies. This strategy reduces variance in reward estimation and stabilizes GRPO training.

H Multilingual System Prompts for Baselines

To ensure reproducibility, we provide the English system prompt used for baseline expressive appropriateness evaluation below. The Chinese prompt follows the same structure, scoring criteria, and output format, and is obtained via a direct sentence-level translation of the English version.

English system prompt for expressive appropriateness evaluation

You are given a speech segment, its surrounding textual context, and a target utterance. Your task is to evaluate the expressive appropriateness of the speech, that is, whether the expressive realization of the utterance appropriately matches the intended communicative setting implied by the context.

Please consider the following aspects in an integrated manner: emotion, rhythm, intonation, recording condition, and any relevant paralinguistic or sound-related cues.

Assign a single expressive appropriateness score in the range of 0.0 to 5.0.

Scoring Guidelines:

0–1: Clearly inappropriate; expressive realization is severely mismatched with the context.
 1–2: Weak or unnatural; major expressive issues remain.
 2–3: Somewhat appropriate with noticeable problems.
 3–4: Generally appropriate with minor flaws.
 4–5: Highly appropriate, natural, and expressive.

Output Format:
 Return only the final score in the following format: <score>1.5</score>
 # Reasoning Mode (optional):
 The model may either directly output the score (w/o CoT), or generate an explicit reasoning process before the final score (w/ CoT), depending on the experimental setting.

[Context]
 {context}
 [Target Utterance]
 {target line}

I Model Parameter Counts

To facilitate a clearer comparison of model capacity as discussed in Section 5.1, we summarize the parameter counts of all evaluated models in Table 9. Note that for the planner-assisted baselines, the 8B parameters of the expressive planner (Qwen3-8B) are effectively utilized during inference.

Model	Model Size
Qwen2.5-Omni	7B
Kimi-Audio	7B
Phi-4-MM	5.6B
Gemma-3n	6B
Step-Audio-R1	32B
Midashenglm	7B
GPT-4o-Audio	-
Gemini-1.5-Pro	-
Voxtral-Mini	3B
Qwen3-Omni	30B
Ours	7B (Judge) + 8B (Planner)

Table 9: Parameter counts of the models evaluated in this work.

J Case Study: Planner and Judge Output

To illustrate how CEAEval-M performs context-rich expressive appropriateness evaluation in prac-

tice, we present a representative case study below. The example demonstrates how the text-only Expressive Planner infers the ideal expressive profile from a multi-turn narrative context, and how the Speech-LLM Judge leverages this profile to perform step-by-step chain-of-thought (CoT) reasoning before predicting the final score.

Case Study: Context-Rich Evaluation

[Input Narrative Context (CTS = 4)]

- *Line 1:* Lin Nianlei said: “Ah, it’s really warm.”
- *Line 2:* Qin Yu said nothing more and continued walking forward at a steady pace. Wearing gloves, Lin Nianlei brushed her hands over the icicles hanging from the windowsills of the houses, then suddenly shouted.
- *Line 3:* Lin Nianlei said: “You big dummy.”
- *Line 4:* Qin Yu said: “What?”

[Target Utterance (Line 5)]

Lin Nianlei said: “Cheer up. What’s already happened can’t be changed by anyone. But since we live in such a time, we should always believe that tomorrow will be beautiful.”

[Input Speech]

target_utterance_speech.wav

[Expressive Planner Output]

```
The ideal performance in this
context is: { "emotion":
"gentle", "rhythm": "soothing",
"intonation": "rising
intonation", "recording
condition": "normal speech" }
```

[Speech-LLM Judge Output (CoT + Score)]

The ideal emotion is gentleness, while the actual performance conveys comfort, encouragement, and calmness. The emotional direction is consistent and naturally expressive, resulting in an emotion score of 4.2. In terms of rhythm, the ideal style is soothing, while the actual delivery is light and lively, with smooth pacing and well-placed pauses, yielding a rhythm score of 4.0. The ideal intonation is rising, while the actual intonation is curved. Although the intonation pattern does not exactly match, it is rich and appropriate to the context, resulting in an intonation score of 4.0. The ideal recording condition is normal speech, and the actual

performance is also normal speech, with good overall clarity, leading to a recording condition score of 4.0. In addition, inhalation sounds appear in the speech and are naturally integrated into the context, further enhancing emotional expression. Overall, the performance is natural and fluent, with minor shortcomings in detail. I assign a final score of <score>4.0</score>.

[Human Expressive Appropriateness Score]

Ground Truth: 4.0