

# CORD: Bridging the Audio–Text Reasoning Gap via Weighted On-policy Cross-modal Distillation

Anonymous ACL submission

## Abstract

Large Audio Language Models (LALMs) have garnered significant research interest. Despite being built upon text-based large language models (LLMs), LALMs frequently exhibit a degradation in knowledge and reasoning capabilities. We hypothesize that this limitation stems from the failure of current training paradigms to effectively bridge the acoustic-semantic gap within the feature representation space. To address this challenge, we propose CORD, a unified alignment framework that performs on-line cross-modal self-distillation. Specifically, it aligns audio-conditioned reasoning with its text-conditioned counterpart within a unified model. Leveraging the text modality as an internal teacher, CORD performs multi-granularity alignment throughout the audio rollout process. At the token level, it employs on-policy reverse KL divergence with importance-aware weighting to prioritize early and semantically critical tokens. At the sequence level, CORD introduces a judge-based global reward to optimize complete reasoning trajectories via Group Relative Policy Optimization (GRPO). Empirical results across multiple benchmarks demonstrate that CORD consistently enhances audio-conditioned reasoning and substantially bridges the audio–text performance gap with only 80k synthetic training samples, validating the efficacy and data efficiency of our on-policy, multi-level cross-modal alignment approach.

## 1 Introduction

Large Language Models (LLMs) (Qwen et al., 2025; DeepSeek-AI et al., 2025; OpenAI, 2025) have demonstrated exceptional semantic understanding capabilities, sparking research to extend this intelligence to multimodal domains through end-to-end processing. In the audio domain, most state-of-the-art Large Audio-Language Models (LALMs) (Chu et al., 2023, 2024; Wu et al., 2025; Ding et al., 2025; Zeng et al., 2024; Fang et al., 2024) are built upon pretrained text-based LLMs

by incorporating an audio encoder (Radford et al., 2022) and a modality alignment module. Raw audio signals are first encoded into acoustic representations, which are then projected into the LLM’s embedding space through audio-text paired supervision. This paradigm implicitly assumes that training on audio-text interleaved data is sufficient to align audio and text into a unified semantic space.

However, recent studies (Wang et al., 2024; Cuervo et al., 2025) have revealed a persistent performance disparity between the two modalities. Despite receiving semantically equivalent inputs, LALMs often exhibit markedly inferior performance on audio-conditioned tasks compared to their text-conditioned counterparts. This gap is particularly pronounced in data-constrained regimes (Chu et al., 2024; Wu et al., 2025; Xu et al., 2025), where limited training samples expose the limitations of existing alignment mechanisms and highlight the need for more data-efficient cross-modal alignment.

Several recent approaches have attempted to improve audio-conditioned reasoning performance, primarily through supervised fine-tuning with labeled speech data or knowledge distillation from external text-based teachers. Despite their progress, these methods suffer from three fundamental limitations. **(1) Limited scalability.** Supervised fine-tuning (He et al., 2024; Minixhofer et al., 2025) relies on large-scale, high-quality annotated speech data, which is expensive to collect and difficult to scale across diverse tasks and domains. **(2) Off-policy distillation and distribution mismatch.** Teacher-based distillation methods (Cuervo et al., 2025; Wang et al., 2024; Tseng et al., 2025) typically provide supervision along the teacher’s text-generation trajectories, rather than the student’s actual audio-conditioned inference states. This off-policy supervision leads to distribution mismatch and limits the ability to correct accumulated audio-specific reasoning errors. **(3) Uniform token-level**

**supervision.** Conventional KL-based distillation treats all tokens equally (Wang et al., 2025; Li et al., 2025; Tseng et al., 2025), failing to emphasize semantically critical tokens that drive cross-modal misalignment, and lacking sequence-level constraints to explicitly regulate global reasoning trajectories.

To address these challenges, we propose **CORD** (Cross-modal Weighted On-policy Reward-guided Distillation), a unified alignment framework that performs online cross-modal self-distillation without relying on any external teacher. CORD leverages the model’s internal text modality as an in-model teacher and conducts multi-granularity distillation along the student model’s real audio-conditioned reasoning trajectories. By constructing on-policy alignment objectives based on audio modality rollouts, CORD directly aligns audio-conditioned reasoning behavior with its text-conditioned counterpart within a unified model architecture.

At the token level, CORD introduces a fine-grained weighting mechanism that prioritizes tokens exhibiting high cross-modal divergence as well as those appearing at early stages of reasoning, where semantic deviations are more likely to propagate and dominate the final outcome. This weighted reverse KL objective enables targeted correction of modality-specific reasoning errors. At the sequence level, CORD formulates a cross-modal reward function and employs Group Relative Policy Optimization (GRPO) (Shao et al., 2024) to align complete reasoning trajectories, encouraging the audio modality to follow reasoning policies consistent with the text modality.

Extensive experiments on multiple reasoning benchmarks demonstrate that CORD substantially improves audio-conditioned reasoning performance. In particular, CORD reduces the audio-text performance gap by an average of **41.6%** on Qwen2-Audio-7B-Instruct and **44.8%** on Step-Audio2-mini, significantly outperforming conventional distillation baselines. On several tasks, CORD nearly eliminates the modality gap altogether, indicating that on-policy cross-modal self-distillation with an internal teacher provides an effective and scalable solution for mitigating reasoning degradation in LALMs.

Our contributions can be summarized as follows:

- **Internal on-policy cross-modal self-distillation.** We propose **CORD**, a fully

in-model, weighted on-policy cross-modal self-distillation framework that aligns audio-conditioned reasoning with text-conditioned behavior without relying on any external teacher, avoiding off-policy mismatch and architecture-induced noise.

- **Multi-granularity alignment of reasoning trajectories.** CORD jointly enforces token-level and sequence-level alignment by emphasizing semantically critical and early reasoning tokens with weighted reverse KL, while regulating global reasoning trajectories via a reward-guided GRPO objective.
- **Effective reduction of the audio-text reasoning gap.** We demonstrate that CORD consistently reduces the audio-text gap by over 40% across various backbones, nearly reaching parity with text-conditioned performance in several reasoning tasks.

## 2 Related Work

### 2.1 Audio-Text Alignment in LALMs

While cascaded ASR-LLM pipelines largely preserve text-domain performance, they discard speaker and paralinguistic cues essential for speech interaction (Maimon et al., 2025), prompting recent work on end-to-end LALMs (Tang et al., 2024; Chu et al., 2024; Xie and Wu, 2024). However, despite architectural advances, such models consistently underperform text-based LLMs on language understanding and reasoning benchmarks, revealing a persistent *text-speech understanding gap* (Cui et al., 2025). Existing approaches mainly attempt to reduce this gap via representation-level cross-modal alignment or large-scale synthetic speech data augmentation (Held et al., 2025), but often show limited gains on broad reasoning tasks or rely on massive proprietary datasets that hinder reproducibility (Zeng et al., 2025).

### 2.2 On Policy Distillation

On-policy distillation (OPD) trains a student using supervision computed on the *student’s own rollouts*, rather than on teacher-generated trajectories. Agarwal et al. (2024) introduces this paradigm for language models, showing that supervising at the student’s prefixes can correct error cascades that would not be exposed under teacher trajectories. MiniLLM (Gu et al., 2025) further develops

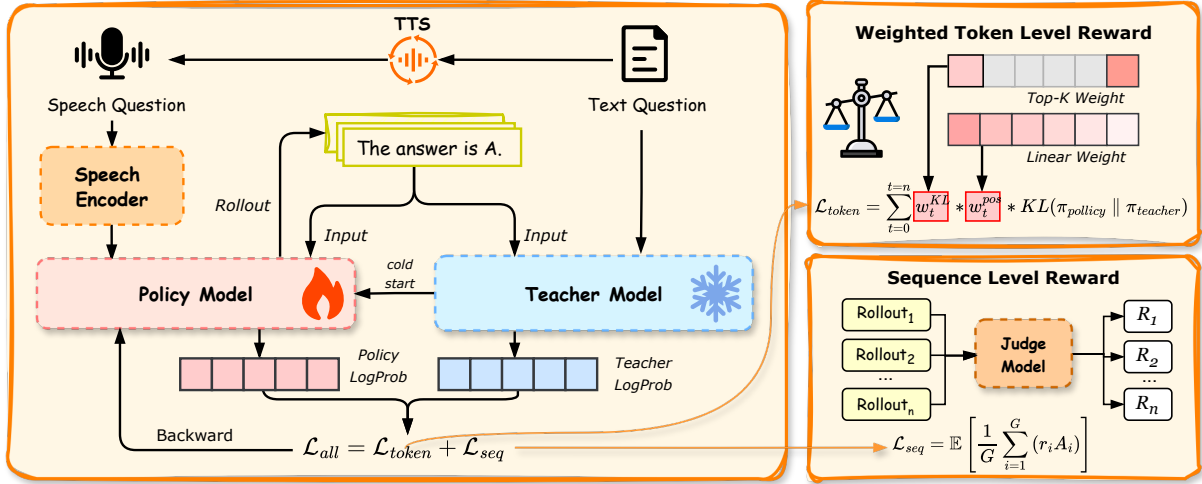


Figure 1: The overall framework of **CORD**. Given semantically equivalent audio and text inputs, CORD performs on-policy cross-modal self-distillation within a single model. Audio-conditioned trajectories are aligned to text-conditioned behaviors at two levels: (i) a token-level objective that applies importance-aware and position-aware reverse KL weighting along audio rollouts, and (ii) a sequence-level objective that uses a judge-based reward optimized via GRPO to enforce global reasoning consistency.

on-policy knowledge distillation for compressing LLMs, emphasizing that matching behaviors on student trajectories can yield strong performance with a smaller student. Recently, some works (Yang et al., 2025; Lu and Lab, 2025; Chen et al., 2025) reports that on-policy distillation can reach comparable reasoning performance to RL-style alignment at a fraction of the training cost under their training recipe. Beyond white-box settings, Ye et al. (2025) study *black-box* on-policy distillation where only teacher API access is available, proposing mechanisms to stabilize on-policy supervision without full teacher internals.

### 3 Method

In this section, we present CORD (Cross-modal Weighted On-policy Reward-guided Distillation), a framework designed to bridge the performance gap in LALMs by aligning audio-conditioned reasoning trajectories with their text-conditioned counterparts. Unlike static or off-policy supervision, CORD enforces alignment along the model’s actual inference paths using an in-model teacher strategy. As shown in Figure 1, CORD employs a multi-granularity alignment approach consisting of two complementary objectives:

- **Token-level alignment** corrects fine-grained semantic deviations at each decoding step using importance-weighted KL.
- **Sequence-level alignment** enforces global

consistency across entire trajectories via GRPO.

By combining local step-wise correction with global trajectory regulation, CORD provides a unified framework for cross-modal alignment. In the following, we detail the problem formulation, the specific alignment objectives, and provide a comprehensive analysis of the underlying mechanisms that ensure training stability and semantic fidelity.

#### 3.1 Problem Formulation

Given semantically equivalent audio and text inputs  $(x^a, x^t)$ , a LALM  $p_\theta$  induces autoregressive distributions over output sequences  $y$ :

$$p_\theta(y | x) = \prod_{t=1}^T p_\theta(y_t | y_{<t}, x), \quad x \in \{x^a, x^t\}. \quad (1)$$

Empirically, LALMs often exhibit divergent behaviors across modalities despite identical input semantics. CORD aims to minimize this discrepancy by explicitly constraining the *audio-conditioned inference trajectories*  $y \sim p_\theta(\cdot | x^a)$ , rather than merely matching marginal output distributions.

#### 3.2 On-policy Cross-modal Distillation

As shown in Figure 1, CORD aligns the model’s cross-modal behavior along trajectories sampled from its current policy. For each  $x^a$ , we sample  $y \sim p_\theta(\cdot | x^a)$  to obtain an on-policy decoding trajectory. At each step  $t$ , the model induces two

distributions over the vocabulary  $\mathcal{V}$  conditioned on the same prefix  $y_{<t}$ :  $p_\theta(\cdot | y_{<t}, x^a)$  and  $p_\theta(\cdot | y_{<t}, x^t)$ .

To quantify the discrepancy at each state, we employ **Reverse Kullback–Leibler (KL) Divergence**:

$$\begin{aligned} D_t &= \text{KL}(p_\theta(\cdot | y_{<t}, x^a) \parallel p_\theta(\cdot | y_{<t}, x^t)) \\ &= \sum_{v \in \mathcal{V}} p_\theta(v | y_{<t}, x^a) \log \frac{p_\theta(v | y_{<t}, x^a)}{p_\theta(v | y_{<t}, x^t)}. \end{aligned} \quad (2)$$

Compared to forward KL, reverse KL places stronger emphasis on high-probability tokens under the text-conditioned distribution, encouraging the audio-conditioned policy to recover critical reasoning decisions made by the text modality. When applied *on-policy* along audio-conditioned trajectories, this formulation enables targeted correction of semantic deviations that arise during audio inference, rather than enforcing global distributional matching.

### 3.3 Analysis of Cross-modal Discrepancy

To motivate the design of CORD, we first investigate the reasoning behavior of LALMs under audio and text modalities. Let  $y = (y_1, \dots, y_T)$  denote a sequence sampled from the audio-conditioned policy  $p_\theta(\cdot | x^a)$ . At each decoding step  $t$ , given the same prefix  $y_{<t}$ , the model induces two conditional token distributions over the vocabulary  $\mathcal{V}$ : one conditioned on audio input  $x^a$ , and the other on text input  $x^t$ . We quantify their discrepancy using the token-level Kullback–Leibler divergence  $D_t$ .

A common practice is to uniformly average token-level KL divergences,

$$\frac{1}{T} \sum_{t=1}^T D_t. \quad (3)$$

This approach implicitly assumes that all tokens contribute equally to cross-modal alignment. However, our empirical analysis on the MMSU benchmark reveals that the distribution of  $D_t$  is highly skewed and non-uniform.

As illustrated in Figure 2, the distribution of  $D_t$  is highly skewed: most tokens exhibit very small divergence, while only a small number of critical tokens have substantially larger KL values. Specifically, we visualize the reverse KL distribution on the MMSU benchmark to investigate this phenomenon. As illustrated in Figure 2 (top-left), the distribution follows a heavy-tailed pattern, where the 80th percentile corresponds to a

remarkably low divergence of only 0.23. This indicates that the vast majority of tokens are already well-aligned across modalities. Furthermore, the scatter plot (bottom-left) shows that these high-discrepancy states are concentrated in early decoding stages ( $r = -0.139$ ). Misalignments at these pivotal early states often trigger a cascade effect, leading to cumulative reasoning failures.

In contrast, the word clouds in Figure 2(a) reveal that high-KL tokens are predominantly concentrated on semantically critical reasoning words and multiple-choice options (e.g., A, B). The misalignment at these pivotal states is a primary driver of incorrect responses; moreover, errors occurring at these early high-discrepancy tokens tend to trigger a cascade effect, leading to cumulative failures in subsequent decoding steps.

As a result, uniform averaging leads to a low overall loss magnitude, causing gradients from high-KL tokens to be diluted by numerous low-KL tokens. This weakens corrective updates precisely at positions that dominate semantic misalignment. To address this issue, we introduce an importance-aware token-level weighting scheme.

### 3.4 Importance-aware Token-level Alignment

To address the gradient dilution issue, CORD introduces a multi-dimensional weighting scheme that prioritizes semantically critical and early-stage tokens.

**Top- $K$  KL-based Importance Weighting.** We first select the  $K$  tokens with the largest divergence values, where  $K = 20$  in all experiments. Let  $\mathcal{I}_K^*$  denote the index set of the top- $K$  tokens ranked by  $D_t$ . We define a KL-based importance weight

$$w_t^{\text{KL}} = \begin{cases} \alpha, & t \in \mathcal{I}_K^*, \\ 1, & \text{otherwise,} \end{cases} \quad (4)$$

where  $\alpha > 1$  is a hyperparameter that controls the strength of emphasis on high-divergence tokens. This hard selection prevents low-divergence tokens from dominating the optimization and ensures that gradient updates focus on tokens that exhibit significant cross-modal semantic mismatch.

**Sequential Decay Weighting.** We further observe that early decoding errors are more detrimental than later ones: once an incorrect semantic decision is made at an early step, subsequent tokens are unlikely to fully correct the reasoning trajectory. To emphasize early alignment, we introduce a



the text-conditioned answer, and  $r_{\text{seq}} = 0$  otherwise. This reward captures global reasoning alignment beyond local token-wise similarity.

**GRPO Optimization.** To optimize the model under this sequence-level reward, we adopt Group Relative Policy Optimization (GRPO). For each audio input  $x^a$ , we sample a group of  $N$  on-policy trajectories  $\{y^{(i)}\}_{i=1}^N$  from the current policy  $p_\theta(\cdot | x^a)$ . Each trajectory is evaluated by the judge model, yielding rewards  $\{r_{\text{seq}}^{(i)}\}_{i=1}^N$ .

GRPO computes a relative advantage by comparing each trajectory’s reward to the group average:

$$A^{(i)} = r_{\text{seq}}^{(i)} - \frac{1}{N} \sum_{j=1}^N r_{\text{seq}}^{(j)}. \quad (9)$$

Following the approach in DAPO (Yu et al., 2025), we omit the explicit KL divergence penalty. Consequently, the sequence-level optimization objective is defined as:

$$\mathcal{L}_{\text{seq}} = -\mathbb{E}_{\{y^{(i)}\}} \left[ \frac{1}{N} \sum_{i=1}^N A^{(i)} \log p_\theta \left( y^{(i)} | x^a \right) \right]. \quad (10)$$

By optimizing this objective, the model increases the likelihood of audio-conditioned trajectories that achieve higher global alignment rewards relative to other on-policy rollouts. Crucially, GRPO operates entirely on-policy, ensuring that global supervision is applied to the exact inference states encountered during audio-conditioned reasoning.

Together with token-level alignment, this judge-guided GRPO objective explicitly constrains global reasoning outcomes and mitigates failure cases where locally aligned tokens still result in semantically inconsistent answers.

### 3.6 Overall Objective

CORD jointly optimizes local and global alignment objectives:

$$\mathcal{L}_{\text{CORD}} = \mathcal{L}_{\text{tok}} + \mathcal{L}_{\text{seq}}, \quad (11)$$

where  $\mathcal{L}_{\text{seq}}$  denotes the GRPO-based sequence-level loss.

## 4 Experiments

### 4.1 Experimental Setting

**Baselines.** In our experiments, we adopt Qwen2-Audio-7B-Instruct (Chu et al., 2024) and Step-Audio-2-Mini (Wu et al., 2025) as

base models. Following the previous distillation method (Cuervo et al., 2025), We implement two distillation objectives using teacher rollouts:

- **Supervised Fine-Tuning (SFT):** Given a teacher rollout  $y \sim p_\theta(\cdot | x^t)$ , the student is optimized to maximize the likelihood of the audio-conditioned trajectory:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{y \sim p_\theta^t} \left[ \sum_{t=1}^{|y|} \log p_{\theta,t}^a \right] \quad (12)$$

- **Forward KL Divergence:** We minimize the distribution discrepancy between modalities over teacher rollouts:

$$\mathcal{L}_{\text{FKL}} = \text{KL}(p_\theta(\cdot | y_{<t}, x^t) \parallel p_\theta(\cdot | y_{<t}, x^a)) \quad (13)$$

**Dataset.** We curate 80,000 instances from **NuminaMath** (LI et al., 2024) to construct our training data. By employing the Kokoro (hexgrad, 2024) model to synthesize audio for each text instruction, we generate semantically equivalent pairs across the *text* and *audio* modalities.

**Judge Model.** Our judge model for GRPO was developed by distilling the evaluation outputs of proprietary frontier models on millions of text-based instruction-following samples. This judge model exhibits exceptional performance, with self-evaluation accuracy consistently exceeding 99%.

**Evaluation Benchmarks.** To evaluate the model’s knowledge-based question answering capabilities, we utilize MMSU and OpenBookQA (OBQA) from VoiceBench (Chen et al., 2024). GSM8K (Cobbe et al., 2021) is employed to assess the model’s mathematical reasoning performance. Furthermore, we adopt MMAU (Sakshi et al., 2024) to evaluate acoustic and paralinguistic reasoning; this benchmark is further categorized into three distinct subsets: *Music*, *Speech*, and *Sound*.

**Experimental Details.** Please refer to Appendix A for more details.

### 4.2 Main Results

Tables 1 and 2 report the performance of CORD across multiple benchmarks using two representative LALMs, Qwen2-Audio-7B-Instruct and Step-Audio2-mini. The results show that CORD substantially **reduces the performance gap** between audio

Method	MMSU		OBQA		GSM8K		AVG.
	Acc.	$\Delta_{\text{base}}$	Acc.	$\Delta_{\text{base}}$	Acc.	$\Delta_{\text{base}}$	
<b>Qwen2-Audio-7B-Instruct</b>	36.04	8.42	51.20	17.59	20.73	19.74	15.25
+ SFT	36.47	7.99 $\downarrow 0.43$	49.49	19.30 $\uparrow 1.71$	33.68	6.79 $\downarrow 12.95$	11.36
+ Forward KL	36.77	7.69 $\downarrow 0.73$	48.21	20.58 $\uparrow 2.99$	36.05	4.42 $\downarrow 15.32$	10.90
+ CORD ( <i>ours</i> )	<b>38.06</b>	<b>6.40</b> $\downarrow 2.02$	<b>52.77</b>	<b>16.02</b> $\downarrow 1.57$	<b>36.20</b>	<b>4.27</b> $\downarrow 15.47$	<b>8.90</b>
<b>Step-Audio2 mini</b>	52.31	8.16	72.30	7.70	43.75	16.72	10.86
+ SFT	54.68	5.79 $\downarrow 2.37$	74.72	5.28 $\downarrow 2.42$	27.89	32.58 $\uparrow 15.86$	14.55
+ Forward KL	53.12	7.35 $\downarrow 0.81$	72.08	7.92 $\uparrow 0.22$	46.57	13.90 $\downarrow 2.82$	9.72
+ CORD ( <i>ours</i> )	<b>57.63</b>	<b>2.84</b> $\downarrow 5.32$	<b>77.74</b>	<b>2.26</b> $\downarrow 5.44$	<b>47.56</b>	<b>12.91</b> $\downarrow 3.81$	<b>6.00</b>

Table 1: Comparison of audio-conditioned reasoning performance across different backbones. **Base Models** (first row of each block) refer to the original instruction-tuned Qwen2-Audio-7B and Step-Audio2-mini.  $\Delta_{\text{base}}$  represents the modality gap for each method, defined as  $\text{Acc}_{\text{text}}^{\text{Base}} - \text{Acc}_{\text{audio}}^{\text{Method}}$ , where  $\text{Acc}_{\text{text}}^{\text{Base}}$  is the fixed text-conditioned accuracy of the Base Model. For Qwen2-Audio, these text baselines are 44.46, 68.79, and 40.47; for Step-Audio2, they are 60.47, 80.00, and 60.47. **AVG.** is the mean of the modality gaps. Arrows indicate the reduction ( $\downarrow$ ) or increase ( $\uparrow$ ) in the gap compared to the Base Model’s initial  $\Delta_{\text{base}}$ .

and text modalities, while effectively **preserving auxiliary audio capabilities** beyond speech.

**Cross-modal gap reduction.** As shown in Table 1, CORD achieves consistent and substantial reductions in the audio–text performance gap across both backbone models. Across all benchmarks, CORD consistently outperforms SFT and Forward-KL, demonstrating the effectiveness of on-policy, trajectory-level alignment. Specifically, on Qwen2-Audio-7B-Instruct, CORD reduces the average audio–text gap by **41.6%** relative to the base model, whereas Forward-KL yields only a 28.5% reduction. On Step-Audio2-mini, CORD achieves a **44.8%** gap reduction, while Forward-KL leads to only a 10.5% reduction.

In addition, by comparing results between different models, we observe a clear trend: Step-Audio2-mini benefits more from CORD (44.8%  $\rightarrow$  41.6%), exhibiting larger gap reductions and stronger alignment. This suggests that a stronger text-conditioned teacher leads to more effective cross-modal alignment. This observation indicates that CORD naturally scales with the quality of the base LALM and can more effectively exploit improvements in text reasoning to enhance audio-conditioned performance.

**Emergent Cross-domain Generalization.** Although CORD is trained exclusively on a math-focused dataset, both models exhibit more pronounced improvements on general-domain benchmarks such as MMSU and OBQA than on the math-intensive GSM8K benchmark. This observa-

Method	Music	Sound	Speech	Avg.
<b>Base Model</b>	58.98	64.74	58.73	60.81
+ SFT	56.29	64.44	51.51	57.39
+ Forward KL	55.99	61.70	53.01	56.90
+ CORD ( <i>ours</i> )	60.18	64.44	55.42	60.01

Table 2: Fine-grained performance on the MMAU benchmark across music, sound, and speech categories based on Qwen2-Audio-7B-Instruct.

tion indicates that CORD does not merely acquire domain-specific knowledge, but instead learns a transferable meta-capability of cross-modal alignment.

**Preserving Auxiliary Audio Capabilities.** Table 2 shows that CORD effectively preserves general audio understanding capabilities while aligning for complex reasoning, whereas baseline methods suffer from significant performance degradation. Specifically, the Forward-KL baseline exhibits a noticeable performance tax, with scores dropping by 2.99 points in *music* and 3.04 points in *sound*, suggesting that conventional distillation may inadvertently lead to the catastrophic forgetting of non-speech acoustic patterns. In contrast, CORD demonstrates remarkable robustness: it not only maintains near-parity with the base model in the *sound* category but even yields a slight improvement in *music*. These results indicate that trajectory-level on-policy alignment effectively mitigates collateral damage to auxiliary audio modalities, ensuring a more stable and balanced cross-

Method	Step	MMSU	OBQA	GSM8K
Base Model (i)	–	36.04	51.20	20.73
<i>Stability Analysis (GRPO only)</i>				
+ GRPO (ii)	500	36.92	50.54	35.59
+ GRPO (ii)	1000	27.87	36.48	19.89
<i>Incremental Ablation (Cumulative)</i>				
+ GRPO + OPD (iii)	3000	37.41	51.20	36.12
+ Full (iv)	3000	38.06	52.77	36.20

Table 3: Ablation study of individual components and training stability. We evaluate the incremental contributions of: (i) the base model, (ii) sequence-level **GRPO**, (iii) On-Policy Distillation (**OPD**), and (iv) the **Full** CORD framework, which integrates all previous modules with **token-level importance weighting** for fine-grained alignment. **Red values** denote model collapse where performance falls below the base baseline.

509 modal alignment that retains pre-trained general  
510 audio intelligence.

### 511 4.3 Ablation Studies

512 We conduct ablation studies to analyze the contribu-  
513 tion of each component in CORD and to investigate  
514 the sensitivity of key hyperparameters. All abla-  
515 tions are performed on Qwen2-Audio-7B-Instruct  
516 under the same training setup as the main experi-  
517 ments.

#### 518 4.3.1 Component-wise Ablation

519 In Table 3, we conduct an incremental ablation  
520 study to evaluate the efficacy of each module within  
521 the CORD framework.

522 We observe that while initial reinforcement learn-  
523 ing via GRPO yields performance gains at 500  
524 steps, it suffers from severe model collapse as train-  
525 ing extends to 1000 steps. This degradation is most  
526 pronounced on GSM8K (35.59  $\rightarrow$  19.89), where  
527 performance falls even below the base model base-  
528 line. Notably, our OPD module acts as a powerful  
529 regularizer that eliminates this instability, ensuring  
530 sustained optimization for up to 3000 steps with-  
531 out quality loss (35.59  $\rightarrow$  36.12 in GSM8K). Fi-  
532 nally, by augmenting OPD with cross-modal token-  
533 level weighting, the complete CORD framework  
534 achieves the best overall performance, with an av-  
535 erage improvement of 6.35 over the base model (i).  
536 This validates the synergy between stable sequence-  
537 level reinforcement learning, on-policy anchoring,  
538 and fine-grained alignment.




Figure 3: Sensitivity analysis of the weighting intensity  $\alpha$  and  $\beta$ . To reduce hyperparameter complexity, we set  $\alpha = \beta$ . The scores represent relative improvements over the baseline ( $\alpha = \beta = 1.0$ ). A value of 2.0 yields the most consistent gains across all tasks.

#### 539 4.3.2 Sensitivity to Weighting Intensity

540 To efficiently navigate the hyperparameter space  
541 while balancing the contributions of token-level im-  
542 portance ( $\alpha$  in Equation. 4) and positional weight-  
543 ing ( $\beta$  in Equation. 5), we couple these parameters  
544 by setting  $\alpha = \beta$ . As illustrated in Figure 3, the  
545 performance across all benchmarks exhibits a char-  
546 acteristic bell-shaped trend, achieving an optimal  
547 trade-off at  $\alpha = \beta = 2.0$ . This configuration suffi-  
548 ciently accentuates pivotal reasoning states without  
549 over-concentrating gradients on a sparse subset of  
550 tokens, which could otherwise compromise opti-  
551 mization stability. In contrast, smaller values (e.g.,  
552 1.0) revert toward uniform KL optimization, while  
553 excessively large values (e.g., 2.5) lead to marginal  
554 performance drops due to the over-suppression of  
555 long-tail semantic information. Consequently, we  
556 adopt  $\alpha = \beta = 2.0$  as the default setting for all  
557 primary experiments.

## 558 5 Conclusion

559 We introduce **CORD**, a weighted on-policy cross-  
560 modal self-distillation framework designed to  
561 bridge the audio–text gap in LALMs. By align-  
562 ing audio-conditioned reasoning with text base-  
563 lines at both token and sequence levels, CORD  
564 employs an importance-aware KL objective and a  
565 judge-guided GRPO objective to ensure local se-  
566 mantic accuracy and global trajectory consistency.  
567 Extensive benchmarks demonstrate that CORD sig-  
568 nificantly narrows the modality gap, highlighting  
569 on-policy trajectory alignment as a robust paradigm  
570 for cross-modal alignment.

## 6 Limitations

Despite the significant performance gains achieved by **CORD**, our study has several limitations. First, due to computational resource constraints, our framework was primarily evaluated on a dataset of 80,000 instances focused on mathematical reasoning. While this scale is sufficient to demonstrate the efficacy of our method, exploring the behavior of **CORD** under larger and more diverse multimodal data remains an important direction for future work. Second, the current experiments are largely centered on audio understanding benchmarks. The generalizability of our importance-aware weighting and on-policy distillation to broader, more diverse domains, such as general audio-visual scene understanding.

## References

Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos, Matthieu Geist, and Olivier Bachem. 2024. **On-policy distillation of language models: Learning from self-generated mistakes**. *Preprint*, arXiv:2306.13649.

Howard Chen, Noam Razin, Karthik Narasimhan, and Danqi Chen. 2025. **Retaining by doing: The role of on-policy data in mitigating forgetting**. *Preprint*, arXiv:2510.18874.

Yiming Chen, Xianghu Yue, Chen Zhang, Xiaoxue Gao, Robby T Tan, and Haizhou Li. 2024. **Voicebench: Benchmarking llm-based voice assistants**. *arXiv preprint arXiv:2410.17196*.

Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, and 1 others. 2024. **Qwen2-audio technical report**. *arXiv preprint arXiv:2407.10759*.

Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. 2023. **Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models**. *Preprint*, arXiv:2311.07919.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. **Training verifiers to solve math word problems**. *arXiv preprint arXiv:2110.14168*.

Santiago Cuervo, Skyler Seto, Maureen de Seyssel, Richard He Bai, Zijin Gu, Tatiana Likhomanenko, Navdeep Jaitly, and Zakaria Aldeneh. 2025. **Closing the gap between text and speech understanding in llms**. *arXiv preprint arXiv:2510.13632*.

Wenqian Cui, Xiaoqi Jiao, Ziqiao Meng, and Irwin King. 2025. **Voxeval: Benchmarking the knowledge understanding capabilities of end-to-end spoken language models**. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 16735–16753, Vienna, Austria. Association for Computational Linguistics.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, and 181 others. 2025. **Deepseek-v3 technical report**. *Preprint*, arXiv:2412.19437.

Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen, Wei Song, Xu Tan, Heyi Tang, and 1 others. 2025. **Kimi-audio technical report**. *arXiv preprint arXiv:2504.18425*.

Qingkai Fang, Shoutao Guo, Yan Zhou, Zhengrui Ma, Shaolei Zhang, and Yang Feng. 2024. **Llama-omni: Seamless speech interaction with large language models**. *Preprint*, arXiv:2409.06666.

Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. 2025. **Minillm: Knowledge distillation of large language models**. *Preprint*, arXiv:2306.08543.

Haorui He, Zengqiang Shang, Chaoren Wang, Xuyuan Li, Yicheng Gu, Hua Hua, Liwei Liu, Chen Yang, Jiaqi Li, Peiyang Shi, Yuancheng Wang, Kai Chen, Pengyuan Zhang, and Zhizheng Wu. 2024. **Emilia: An extensive, multilingual, and diverse speech dataset for large-scale speech generation**. *Preprint*, arXiv:2407.05361.

William Held, Yanzhe Zhang, Minzhi Li, Weiyang Shi, Michael J. Ryan, and Diyi Yang. 2025. **Distilling an end-to-end voice assistant without instruction training data**. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 7876–7891, Vienna, Austria. Association for Computational Linguistics.

hexgrad. 2024. **Kokoro: Open-weight tts model with 82 million parameters**. <https://github.com/hexgrad/kokoro>.

Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang, Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann Fleureau, Guillaume Lample, and Stanislas Polu. 2024. **Numinamath**. [<https://huggingface.co/AI-MO/NuminaMath-CoT>]([https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina\\_dataset.pdf](https://github.com/project-numina/aimo-progress-prize/blob/main/report/numina_dataset.pdf)).

Yang Li, Zhichen Dong, Yuhan Sun, Weixun Wang, Shaopan Xiong, Yijia Luo, Jiashun Liu, Han Lu, Jiamang Wang, Wenbo Su, Bo Zheng, and Junchi Yan. 2025. **Attention illuminates llm reasoning: The preplan-and-anchor rhythm enables fine-grained policy optimization**. *Preprint*, arXiv:2510.13554.

678	Kevin Lu and Thinking Machines Lab. 2025. <a href="#">On-policy distillation</a> . <i>Thinking Machines Lab: Connectionism</i> . <a href="https://thinkingmachines.ai/blog/on-policy-distillation">https://thinkingmachines.ai/blog/on-policy-distillation</a> .	733
679		734
680		735
681		
682	Gallil Maimon, Amit Roth, and Yossi Adi. 2025. <a href="#">Salmon: A suite for acoustic language model evaluation</a> . In <i>ICASSP 2025 – 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)</i> , pages 1–5. IEEE.	736
683		737
684		738
685		739
686		
687	Christoph Minixhofer, Ondrej Klejch, and Peter Bell. 2025. <a href="#">Scaling laws for synthetic speech for model training</a> . In <i>Proceedings of Interspeech 2025</i> . ISCA.	740
688		741
689		742
690	OpenAI. 2025. <a href="#">gpt-oss-120b gpt-oss-20b model card</a> . <i>Preprint</i> , arXiv:2508.10925.	743
691		744
692	Qwen, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jingren Zhou, Junyang Lin, and 24 others. 2025. <a href="#">Qwen2.5 technical report</a> . <i>Preprint</i> , arXiv:2412.15115.	745
693		746
694		747
695		748
696		749
697		750
698		751
699	Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. <a href="#">Robust speech recognition via large-scale weak supervision</a> . <i>Preprint</i> , arXiv:2212.04356.	752
700		753
701		754
702		755
703	S Sakshi, Utkarsh Tyagi, Sonal Kumar, Ashish Seth, Rameswaran Selvakumar, Oriol Nieto, Ramani Duraiswami, Sreyan Ghosh, and Dinesh Manocha. 2024. <a href="#">Mmau: A massive multi-task audio understanding and reasoning benchmark</a> . <i>arXiv preprint arXiv:2410.19168</i> .	756
704		757
705		758
706		759
707		760
708		761
709	Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, and 1 others. 2024. <a href="#">Deepseekmath: Pushing the limits of mathematical reasoning in open language models</a> . <i>arXiv preprint arXiv:2402.03300</i> .	762
710		763
711		764
712		765
713		766
714		
715	Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2024. <a href="#">SALMONN: Towards generic hearing abilities for large language models</a> . In <i>International Conference on Learning Representations (ICLR)</i> .	767
716		768
717		769
718		770
719		771
720	Liang-Hsuan Tseng, Yi-Chang Chen, Kuan-Yi Lee, Da-Shan Shiu, and Hung-Yi Lee. 2025. <a href="#">Taste: Text-aligned speech tokenization and embedding for spoken language modeling</a> . <i>Preprint</i> , arXiv:2504.07053.	772
721		773
722		774
723		775
724	Chen Wang, Minpeng Liao, Zhongqiang Huang, and Jiajun Zhang. 2024. <a href="#">Blsp-kd: Bootstrapping language-speech pre-training via knowledge distillation</a> . <i>Preprint</i> , arXiv:2405.19041.	776
725		777
726		778
727		779
728	Shenzhi Wang, Le Yu, Chang Gao, Chujie Zheng, Shixuan Liu, Rui Lu, Kai Dang, Xionghui Chen, Jianxin Yang, Zhenru Zhang, Yuqiong Liu, An Yang, Andrew Zhao, Yang Yue, Shiji Song, Bowen Yu, Gao Huang, and Junyang Lin. 2025. <a href="#">Beyond the 80/20</a>	780
729		781
730		782
731		783
732		784
		785
		786
		787
		788
		789
		790
		791
		792
		793
		794
		795
		796
		797
		798
		799
		800
		801
		802
		803
		804
		805
		806
		807
		808
		809
		810
		811
		812
		813
		814
		815
		816
		817
		818
		819
		820
		821
		822
		823
		824
		825
		826
		827
		828
		829
		830
		831
		832
		833
		834
		835
		836
		837
		838
		839
		840
		841
		842
		843
		844
		845
		846
		847
		848
		849
		850
		851
		852
		853
		854
		855
		856
		857
		858
		859
		860
		861
		862
		863
		864
		865
		866
		867
		868
		869
		870
		871
		872
		873
		874
		875
		876
		877
		878
		879
		880
		881
		882
		883
		884
		885
		886
		887
		888
		889
		890
		891
		892
		893
		894
		895
		896
		897
		898
		899
		900
		901
		902
		903
		904
		905
		906
		907
		908
		909
		910
		911
		912
		913
		914
		915
		916
		917
		918
		919
		920
		921
		922
		923
		924
		925
		926
		927
		928
		929
		930
		931
		932
		933
		934
		935
		936
		937
		938
		939
		940
		941
		942
		943
		944
		945
		946
		947
		948
		949
		950
		951
		952
		953
		954
		955
		956
		957
		958
		959
		960
		961
		962
		963
		964
		965
		966
		967
		968
		969
		970
		971
		972
		973
		974
		975
		976
		977
		978
		979
		980
		981
		982
		983
		984
		985
		986
		987
		988
		989
		990
		991
		992
		993
		994
		995
		996
		997
		998
		999
		1000

787 For rollout sampling, we use a temperature of 1.0  
788 for CORD and Forward-KL, while a higher tem-  
789 perature of 1.5 is adopted for GRPO to encourage  
790 trajectory diversity. For the token-level alignment  
791 objective, we fix the KL-based importance scaling  
792 factor and the positional decay factor to  $\alpha = 2$  and  
793  $\beta = 2$ , respectively, in all experiments.