

SoundMind: RL-Incentivized Logic Reasoning for Audio-Language Models

Xingjian Diao¹, Chunhui Zhang¹, Keyi Kong², Weiyi Wu¹, Chiyu Ma¹,
Zhongyu Ouyang¹, Peijun Qing¹, Soroush Vosoughi¹, Jiang Gui¹

¹Dartmouth College, ²Shandong University

xingjian.diao.gr@dartmouth.edu

 [SoundMind Dataset](#)

Abstract

While large language models have demonstrated impressive reasoning abilities, their extension to the audio modality, particularly within large audio-language models (LALMs), remains underexplored. Addressing this gap requires a systematic approach that involves a capable base model, high-quality reasoning-oriented audio data, and effective training algorithms. In this work, we present a comprehensive solution for audio logical reasoning (ALR) tasks: we introduce SoundMind, a dataset of 6,446 audio–text annotated samples specifically curated to support complex reasoning. Building on this resource, we propose SoundMind-RL, a rule-based reinforcement learning (RL) algorithm designed to equip audio-language models with robust audio–text reasoning capabilities. By fine-tuning Qwen2.5-Omni-7B on the proposed SoundMind dataset using SoundMind-RL, we achieve strong and consistent improvements over state-of-the-art baselines on the SoundMind benchmark. This work highlights the benefit of combining high-quality, reasoning-focused datasets with specialized RL techniques, and contributes to advancing auditory intelligence in language models. The code and dataset are publicly available at <https://github.com/xid32/SoundMind>.

1 Introduction

Large language models (LLMs) have made remarkable strides in reasoning capabilities through innovations such as Chain-of-Thought (CoT) prompting and specialized reasoning architectures. Recent models such as OpenAI’s o1 model (Jaech et al., 2024) and Deepseek-R1 (Guo et al., 2025) have demonstrated exceptional performance on complex logical tasks (Kimi Team et al., 2025; Zhao et al., 2024; Yuan et al., 2025; Zhang et al., 2024, 2025a,b; Han et al., 2024). Particularly noteworthy is Deepseek-R1’s rule-based reinforcement learning approach, which enables emergent reasoning without relying on traditional frameworks like

Monte Carlo Tree Search (Wan et al., 2024) or process reward models (Lightman et al., 2023). This general reasoning paradigm has been successfully extended to the visual domain, where frameworks such as Visual-CoT (Shao et al., 2024) have significantly improved multimodal models’ cognitive abilities for image and video reasoning.

Despite these advances in text and visual reasoning, there is a significant gap in *audio reasoning and generation capabilities*. While large audio-language models (LALMs) like Audio Flamingo (Kong et al., 2024) and Qwen2.5-Omni-7B (Xu et al., 2025) have made progress in audio understanding, end-to-end audio reasoning remains underdeveloped. This limitation stems primarily from two factors: (1) the simplicity of existing audio reasoning datasets (Suzgun et al., 2023; Kong et al., 2024), which often contain only brief textual labels without proper audio modality annotations, and (2) the technical challenges in maintaining reasoning coherence during long-duration audio generation. Current CoT methods applied to audio often lead to hallucinations and performance degradation when generating extended reasoning sequences. The lack of aligned audio-text annotations further impedes research on reasoning-driven audio generation, creating a bottleneck in developing LALMs with sophisticated reasoning capabilities.

To address these challenges in audio reasoning, we introduce a comprehensive approach with two key components. First, we introduce SoundMind, a benchmark dataset specifically designed for the audio logical reasoning (ALR) task. It contains 6,446 audio–text aligned samples, each containing user content, step-by-step chain-of-thought reasoning, final answers, and the corresponding input–output speech audio (Figure 1). Built upon the LogiQA 2.0–NLI dataset (Liu et al., 2023), SoundMind preserves the full logical structure and augments it with parallel text–audio annotations, enabling models to learn reasoning grounded in audio. Second,

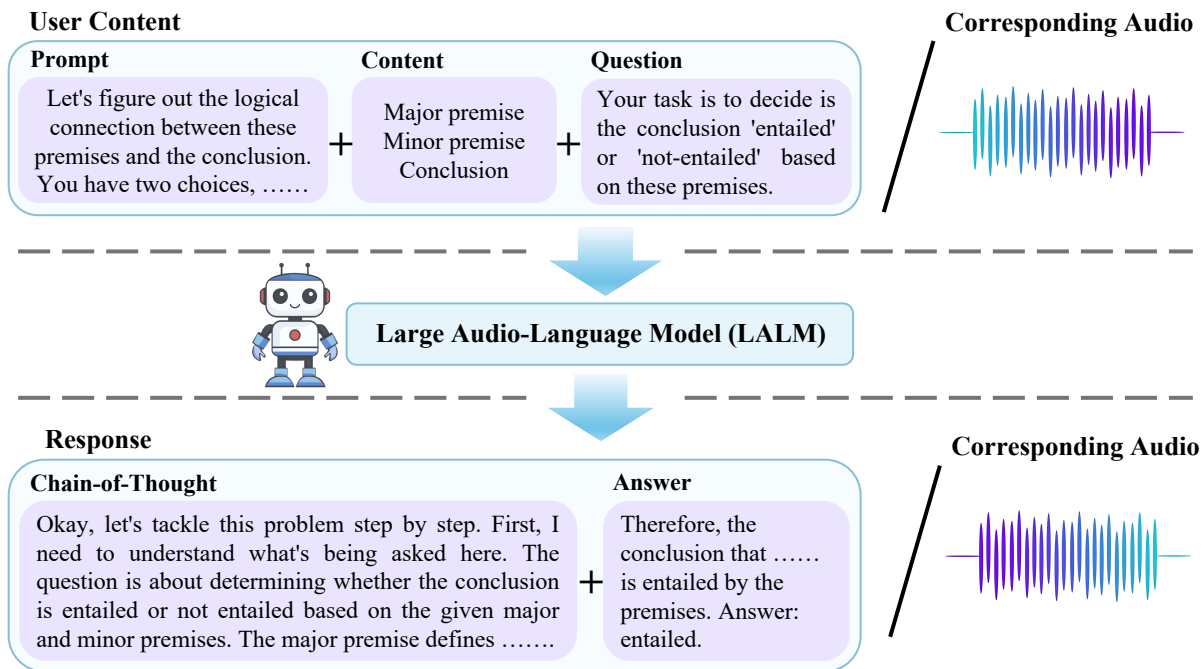


Figure 1: Overview of a SoundMind sample for the audio logical reasoning (ALR) task. Each instance contains **User Content** (top), including a natural-language prompt, structured logical triplets, and a question, as well as **Response** (bottom), which consists of step-by-step chain-of-thought reasoning and the final answer. Both components are provided in text and accompanied by the corresponding input–output speech audio (right). SoundMind therefore offers **complete logical reasoning tasks** with fully aligned content and response across text and audio modalities, supporting multimodal training and evaluation of large audio–language models.

we propose SoundMind-RL, a rule-based reinforcement learning algorithm that addresses the challenges of long-form audio reasoning generation. Drawing inspiration from the Logic-RL framework (Xie et al., 2025a), SoundMind-RL incorporates a strict format-based reward mechanism to prevent shortcut reasoning biased toward the textual modality. It leverages the REINFORCE++ algorithm (Hu et al., 2025) alongside the reward design principles from Deepseek-R1 for effective post-training. Our main contributions are summarized as follows:

- **SoundMind Dataset:** We release a high-quality audio logical reasoning dataset comprising 6,446 samples with user content, CoT reasoning, answers and corresponding audio input/output. It provides deep reasoning annotations in both text and audio modalities, serving as a valuable resource for RL-based training and evaluation of audio–language models.
- **SoundMind-RL Algorithm:** Taking into account the unique challenges of audio reasoning generation, we design a strict format-based reward mechanism that maintains reasoning coherence across modalities. By combining the im-

proved REINFORCE++ algorithm with reward design principles from DeepSeek-R1, we fine-tune Qwen2.5-Omni-7B through reinforcement learning, achieving state-of-the-art results across three input–output modality configurations on the SoundMind benchmark.

2 Related Work

Open-Source Datasets for Audio Reasoning. Research on audio reasoning remains limited by the scarcity of publicly available benchmarks, making systematic evaluation and comparison difficult. Existing corpora typically treat audio as descriptive input paired with text, yet their supervision remains purely textual, providing final annotations exclusively in text form (Suzgun et al., 2023; Kong et al., 2024; Ghosh et al., 2025). More recently, a number of studies have begun to recognize the critical role of chain-of-thought (CoT) prompting in audio reasoning datasets (Xie et al., 2025b), highlighting the importance of structured reasoning traces. To the best of our knowledge, SoundMind is among the first open-source resources where audio itself serves as the annotated modality for reasoning, providing parallel text–audio representations and en-

Datasets	Type	Transcripts	CoT annotations	Samples	Hours
CoTA	Sound, Speech, Music	✗	✓	1.2M	6K
AudioSkills	Speech	✗	✗	4.2M	9.3K
LongAudio	Audio	✗	✗	263K	8.5K
Big Bench Audio	Speech	✗	✗	1000	2
SoundMind (Ours)	Speech	✓	✓	6446	1K

Table 1: Comparison of publicly available datasets for audio reasoning. Existing resources differ in scope: some provide speech data without transcripts, while others include transcripts but lack step-by-step reasoning annotations aligned with audio. These gaps make systematic study of multimodal reasoning challenging. **SoundMind** combines transcripts, chain-of-thought (CoT) annotations, and speech audio into a unified benchmark, offering a balanced resource for training and evaluating audio–language models.

abling models to learn from and be evaluated on spoken reasoning traces.

Chain-of-Thought Reasoning. LLMs enhance their reasoning ability through in-context learning (ICL), which processes prompts within the surrounding context. CoT techniques further reinforce this capability. Prominent CoT approaches include Tree-of-Thought (ToT) (Yao et al., 2023), manually crafted few-shot CoT (Wei et al., 2022), and various automatic generation strategies (Jin et al., 2024). Recent works have also examined the necessity, theoretical underpinnings, and task-specific effectiveness of CoT reasoning (Sprague et al., 2025). OpenAI’s ol model (Jaech et al., 2024) has reignited interest in CoT prompting and is often paired with reinforcement learning-based training approaches (Hu et al., 2025; Xie et al., 2025a).

Multimodal Chain-of-Thought Reasoning. CoT prompting has seen notable advancements in the multimodal domain. For instance, Visual-CoT (Shao et al., 2024) integrates object detection to assist reasoning, while LLaVA-CoT (Xu et al., 2024) and MAMmoTH-VL (Guo et al., 2024) improve performance via dataset augmentation. However, CoT applications in the audio domain are still in their infancy. Audio-CoT (Ma et al., 2025) demonstrates that zero-shot CoT prompting yields improvements on simple audio tasks, but remains inadequate for complex reasoning.

Although current audio-language models (ALMs) have made progress in comprehension and real-time response, their capability in CoT-style reasoning remains underexplored. Our study addresses this research gap by systematically investigating the application of CoT techniques within ALMs.

3 Dataset Design and Construction

3.1 Key Characteristics

We introduce SoundMind, a unified benchmark for the audio logical reasoning (ALR) task, specifically developed to advance research on reasoning grounded in auditory input. In contrast to most existing datasets that provide only textual supervision, SoundMind offers both text-level and audio-level chain-of-thought (CoT) annotations, allowing models to learn and reason directly from spoken prompts and responses. The dataset emphasizes speech-based scenarios, making it well suited for applications such as dialogue understanding and auditory commonsense reasoning.

As shown in Table 1, existing audio datasets such as CoTA (Xie et al., 2025b), AudioSkills (Ghosh et al., 2025), LongAudio (Ghosh et al., 2025), and Big Bench Audio (Suzgun et al., 2023) provide valuable resources but remain limited in their support for auditory reasoning. Some lack audio-level annotations entirely, while others include audio but omit chain-of-thought (CoT) reasoning traces or detailed step-by-step rationales. For instance, although CoTA offers CoT-style annotations, it does not include audio transcripts, which limits its utility for training and evaluating models that operate directly on speech input. Likewise, AudioSkills and LongAudio supply a large quantity of audio data but no reasoning supervision, making them less suitable for investigating step-by-step inference or developing models that require explicit alignment between input and reasoning output.

The SoundMind dataset is constructed through a structured pipeline that transforms textual logical reasoning tasks into natural spoken audio interactions, thereby enabling the development of audio-based reasoning models. The complete pipeline is

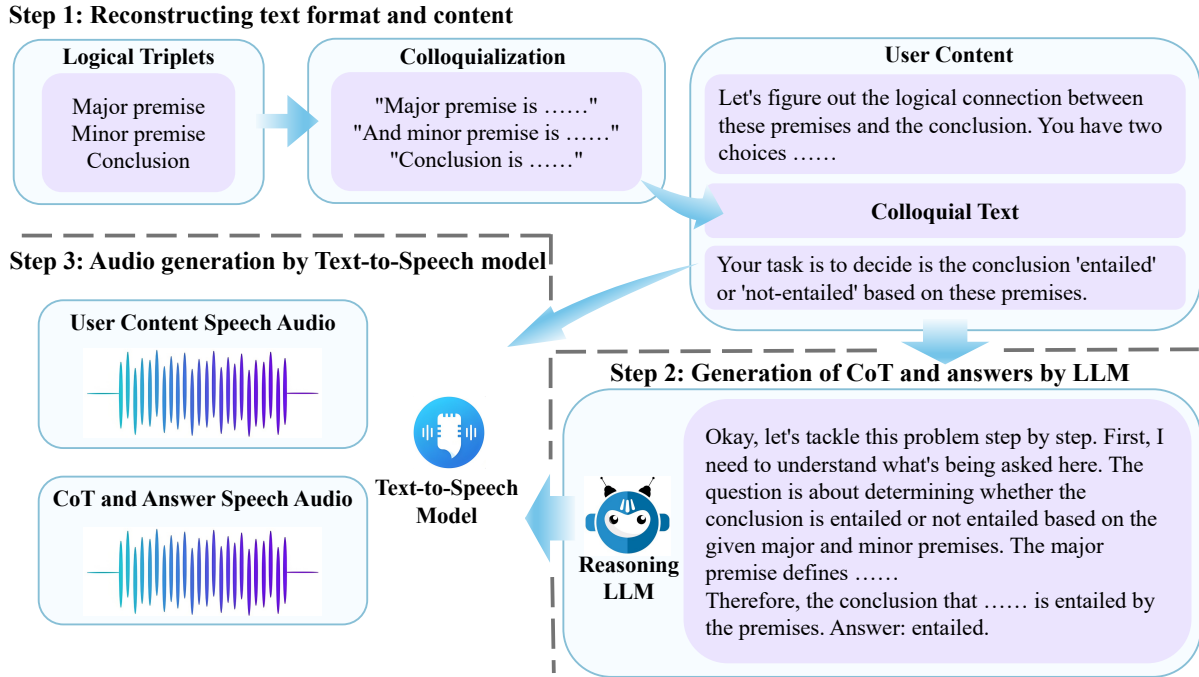


Figure 2: Three-step pipeline for constructing SoundMind samples. **Step 1:** Structured logical triplets are converted into natural, conversational prompts through a Colloquialization module. **Step 2:** A large language model generates detailed chain-of-thought reasoning and final answers, providing rich supervision for reasoning tasks. **Step 3:** Both the user content and the generated reasoning–answer pairs are converted into speech using a text-to-speech (TTS) model, yielding two carefully aligned audio segments that together capture the full reasoning interaction.

illustrated in Figure 2.

Although SoundMind contains fewer samples (6,446) than large-scale datasets such as AudioSkills, it is specifically designed for logic-oriented reasoning and provides over 1,074 hours of carefully annotated speech audio with both text-level and audio-level CoT supervision. This comprehensive annotation makes SoundMind a valuable benchmark for training and evaluating models that require structured, multimodal reasoning grounded in audio perception.

3.2 Data Generation Pipeline

Our pipeline takes as input structured logical triplets consisting of a major premise, a minor premise, and a conclusion. These are first processed by a **Colloquialization** module that rewrites the formal statements into natural, conversational prompts (e.g., “The major premise is ...”, “Let’s figure out the logical connection ...”). This step enhances the fluency of the resulting audio and more closely reflects realistic user queries.

The colloquialized content and instructions are merged into a single **User Content** block and synthesized into speech audio using the high-fidelity text-to-speech model MegaTTS 3 (Jiang et al.,

2025), forming the input side of the dataset. The prompts used in this process are listed in Table 2.

To generate the output, we employ the large language model DeepSeek-R1 to produce step-by-step CoT reasoning and final answers, providing richer supervision for model training. The generated reasoning and answers are then synthesized into speech audio with MegaTTS 3 (Jiang et al., 2025), yielding the corresponding spoken response.

As a result, each SoundMind instance consists of two carefully aligned audio segments: one representing the user query and the other presenting the detailed step-by-step reasoning and the corresponding final answer, as illustrated in Figure 3.

4 SoundMind-RL Algorithm

We improve our system through iterative optimization of rule-based rewards. All λ in the following equations are hyperparameters.

Answer Format Correctness Evaluation. To ensure the correctness of answer formatting, we first require that the token “Answer:” must appear within the last five characters of the model response. Considering that the SoundMind dataset covers both text and audio modalities, we design

Position	Prompt
System	Your task is to decide if the conclusion is "entailed" or "not-entailed" based on these premises. You are a wise person who answers two-choice questions, "entailed" or "not entailed". Use plain text for thought processes and answers, not markdown or LaTeX. The thought process and response style should be colloquial, which I can then translate directly into audio using the TTS model. The final output is the Answer, nothing else, and the format is Answer: YOUR ANSWER. For example: "Answer: entailed." or "Answer: not entailed." The final answer must contain nothing else! The thought process should be very complete, careful, and cautious. When you think and generate a chain of thought, you need to test your answer from various angles.
Before the major premise	Let's figure out the logical connection between these premises and the conclusion. You have two choices: "entailed" means the conclusion must be true based on the given premises, or "not-entailed" means the conclusion can't be true based on the premises. Here's the setup:
Behind the conclusion	Your task is to decide is the conclusion is "entailed" or "not-entailed" based on these premises.

Table 2: Instructional prompts used at different positions during chain-of-thought generation. The *system* prompt specifies the overall task, output format, and encourages a careful, colloquial reasoning style suitable for speech synthesis. Additional *contextual prompts* placed before the major premise and after the conclusion provide further guidance, helping the model interpret logical relations and produce well-structured reasoning traces.

two specific format scoring methods (denoted as S_{format}), calculated as follows:

$$S_{\text{format}}^{(1)} = \begin{cases} \lambda_1, & \text{if correct for text tokens} \\ 0, & \text{otherwise,} \end{cases} \quad (1)$$

$$S_{\text{format}}^{(2)} = \begin{cases} \lambda_2, & \text{if correct for audio tokens} \\ 0, & \text{otherwise.} \end{cases} \quad (2)$$

Answer Correctness Evaluation. Once the format compliance is verified, this module evaluates the factual accuracy of the model response. Specifically, the answer score (S_{answer}) is computed based on the consistency between the model's predicted answer and the ground truth answer, using the following formulation:

$$S_{\text{answer}} = \begin{cases} \lambda_3, & \text{if answer = ground truth} \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

Reasoning Length Evaluation. We additionally introduce a reward evaluation based on the reasoning length of the model response. This score (S_{len}) is computed by comparing the ratio of the model output length to the reference reasoning length. Given that the SoundMind dataset provides supervision in both text and audio modalities, we design two distinct length evaluation methods, defined as follows:

$$S_{\text{len}}^{(1)} = \lambda_4 \times \min \left(1, \frac{L_{\text{model}}}{L_{\text{annotation}}} \right), \quad (4)$$

$$S_{\text{len}}^{(2)} = \lambda_5 \times \min \left(1, \frac{T_{\text{model}}}{T_{\text{annotation}}} \right). \quad (5)$$

where L denotes the length of text tokens, and T denotes the length of audio tokens.

REINFORCE++ Policy Optimization. In our setting, the policy π_θ corresponds to a large-scale ALM, which receives an audio question as input and produces a reasoning response in either text or audio form. To optimize the model with the above composite reward, we adopt the REINFORCE++ (Hu et al., 2025)—a clipped policy-gradient method that eliminates the need for a value (critic) network while leveraging PPO-style stability and sample efficiency. Specifically, REINFORCE++ updates the policy by maximizing the following objective:

$$\mathcal{J}_{\text{REINFORCE++}}(\theta) = \mathbb{E}_{(x,y) \sim \mathcal{D}, o_{\leq t} \sim \pi_{\theta_{\text{old}}}} \left[\min \left(r_t(\theta) \hat{A}_t, \text{clip} \left(r_t(\theta), 1 - \epsilon, 1 + \epsilon \right) \hat{A}_t \right) \right] \quad (6)$$

where x is the audio input, y is the generated response, and $r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}$ is the importance sampling ratio. \hat{A}_t is the normalized advantage, given by: $A_t = R(x, y) - \beta \sum_{i=t}^T \text{KL}(i)$, $\hat{A}_t = \frac{A_t - \mu_A}{\sigma_A}$. Here, $R(x, y)$ is the cumulative reward, composed of the weighted sum of the above reward terms (S_{format} , S_{answer} , S_{len}), and $\text{KL}(i)$ denotes the token-level KL divergence between the current policy and the reference SFT model. The KL term is defined as:

$$\text{KL}(i) = \log \frac{\pi_\theta(a_i|s_i)}{\pi_{\text{SFT}}(a_i|s_i)}, \quad (7)$$

where π_θ is the current policy being optimized, π_{SFT} is the fixed reference (supervised fine-tuned) model, β is the KL penalty coefficient, and μ_A ,

User Content

Let's figure out the logical connection between these premises and the conclusion. You have two choices, 'entailed' means the conclusion must be true based on the given premises, or 'not-entailed' means the conclusion can't be true based on the premises.

Here's the setup:

Major premise is "Scientists have found that lack of exposure to bacteria and germs in childhood is one of the main reasons for the sharp increase in diabetes cases among people under the age of five in recent years. Children living in rural areas are less likely to suffer from diseases such as allergies, asthma and eczema because they are exposed to bacteria earlier and have more opportunities to spend time with pets".

Minor premise is "Most of the bacteria on human skin play a protective role and create a unique body odor for everyone".

Conclusion is "an environment that destroys bacteria too thoroughly may do harm to children's health".

Your task is to decide is the conclusion 'entailed' or 'not-entailed' based on these premises?

Part of Chain-of-thought

Okay, let's dive into this problem. So, the task is to determine whether the conclusion is entailed by the premises provided. The conclusion is that "an environment that destroys bacteria too thoroughly may do harm to children's health." The major and minor premises are about the role of bacteria exposure in childhood and the protective role of bacteria on human skin.

First, I need to parse each premise carefully. The major premise states two main points: one, that lack of exposure to bacteria and germs in childhood is a main reason for the increase in diabetes cases in kids under five. Two, children in rural areas have less allergies, asthma, and eczema because they're exposed to bacteria earlier and have more contact with pets. The minor premise adds that most bacteria on human skin are protective and contribute to unique body odor.

The conclusion suggests that overly thorough destruction of bacteria (like excessive cleaning or use of antibacterial products) might be harmful to children's health. The question is whether this conclusion logically follows from the premises.

Let's start by breaking down the major premise. It links reduced early bacterial exposure to higher diabetes rates in young children. This seems to be referencing the "hygiene hypothesis", which posits that early exposure to microbes helps train the immune system, reducing the risk of autoimmune diseases. The premise also connects rural living (with more bacteria and pet exposure) to lower rates of allergies, asthma, and eczema, which are immune-related conditions. So, the major premise is suggesting that early bacterial exposure is protective against certain diseases.....

Answer

Thus, based on the premises, the conclusion that overly sterile environments may harm children's health is entailed.

Answer: entailed.

Figure 3: An illustrative example from the SoundMind dataset. Each sample contains three components: (1) **User Content**, which includes a natural-language prompt and a structured logical triplet (major premise, minor premise, conclusion); (2) **Chain-of-Thought**, a step-by-step reasoning trace generated by a large language model, explaining the entailment decision; and (3) the final **Answer**. All components are provided in both text form and synthesized speech audio, enabling fully multimodal training and evaluation of audio–language models.

σ_A are the batch mean and standard deviation for normalization. ϵ is the PPO clip parameter. REINFORCE++ combines the stability of PPO's clipped surrogate loss with the efficiency of critic-free Monte Carlo policy gradient updates, using a token-level KL penalty and batch-normalized advantages. It trains audio-language models for end-to-end reasoning over audio questions.

5 Experiments

5.1 Setup

We fine-tune the Qwen2.5-Omni-7B model on the SoundMind dataset using the proposed SoundMind-RL algorithm. All experiments were conducted under a consistent hardware environment, consisting of an Intel(R) Xeon(R) Platinum 8468 CPU and 8 NVIDIA H800 GPUs, each equipped with 80 GB of memory. The weighting coefficients for the reward components were set as follows: $\lambda_1 = 1.0$, $\lambda_2 = 0.5$, $\lambda_3 = 2.0$, $\lambda_4 = 1.0$, and $\lambda_5 = 0.75$.

The training procedure was executed for 50,000 steps. All other hyperparameter settings followed those used in Logic-RL (Xie et al., 2025a).

To fully exploit the availability of both text and audio modalities in the SoundMind dataset, we evaluate our approach under three input–output configurations: ① Table 4: audio-only input with text-based reasoning output, benchmarked against five strong multimodal LLMs: MiniCPM-o (Yao et al., 2024), Gemini-Pro-V1.5 (Team et al., 2024), Baichuan-Omni-1.5 (Li et al., 2025), Qwen2-Audio (Chu et al., 2024), and Qwen2.5-Omni-7B (Xu et al., 2025). ② Table 5: text-only input with audio-based reasoning output. Due to the lack of comparable open-source multimodal models, we use Qwen2.5-Omni-7B as the reference baseline. ③ Table 6: audio-only input with audio-based reasoning output. Together, these three configurations allow us to comprehensively assess SoundMind-RL across text and audio modalities, covering fully cross-modal reasoning scenarios.

Split	# Entailed	# Not-ent.	Avg. Inp. Tok.	Avg. Out. Tok.	Avg. Inp. Dur. (s)	Avg. Out. Dur. (s)
Training	2326	2858	182	1683	62.33	608.42
Test	296	360	158	1424	57.90	586.51
Validation	264	342	155	1426	57.85	586.39

Table 3: Detailed statistics of the SoundMind dataset across training, validation, and test splits. “# Entailed” and “# Not-ent.” denote the number of samples labeled as entailed and not-entailed. “Avg. Inp./Out. Tok.” reports the mean token counts for the input and outputs, respectively. “Avg. Inp./Out. Dur. (s)” provides the average duration (in seconds) of the input audio and the corresponding output audio. Together, these statistics highlight the dataset’s scale, near-balanced class distribution, and focus on long-form reasoning supervision.

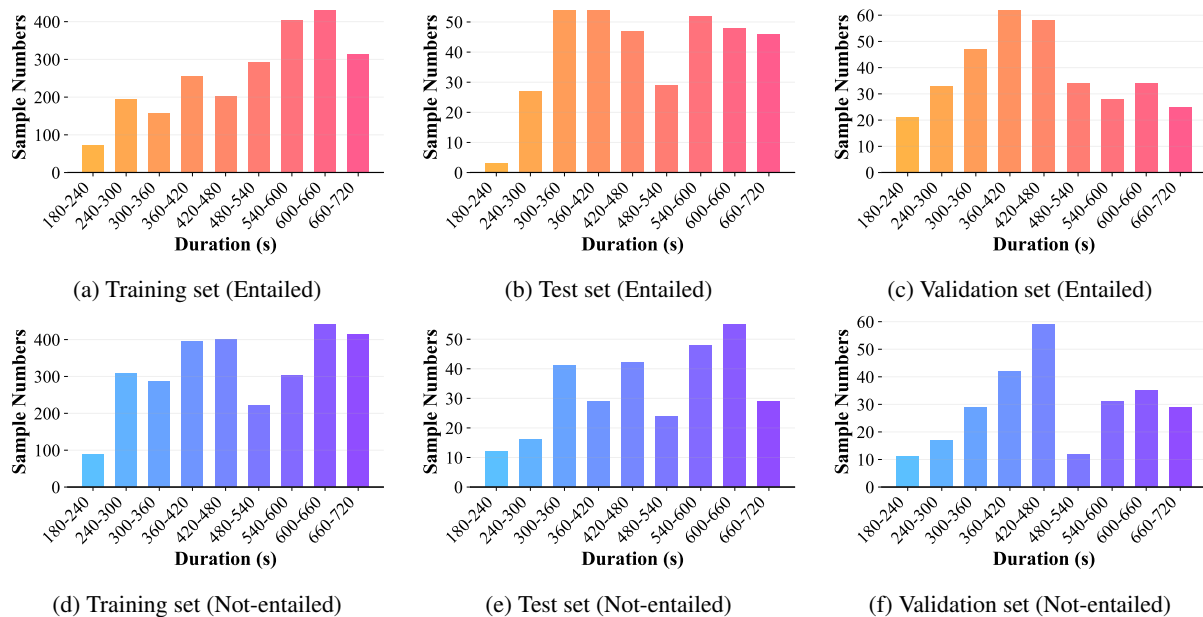


Figure 4: Distribution of audio durations across dataset splits and entailment labels. Subfigures (a–c) show histograms for “Entailed” samples in the training, test, and validation sets, while (d–f) show the corresponding distributions for “Not-entailed” samples. These plots illustrate the broad coverage of audio durations and the consistent duration profiles across splits, highlighting that the SoundMind dataset offers both diversity and balance for evaluating model performance across different temporal contexts.

5.2 Dataset Analysis

Table 3 summarizes key statistics of the SoundMind dataset, clearly illustrating its overall scale, class balance, and linguistic diversity across training, validation, and test splits. The corpus contains 6,446 samples with a near-balanced distribution of entailed and not-entailed instances, thereby ensuring representative coverage of both classes for robust entailment prediction.

The user content averages 160–180 tokens across all splits, whereas the generated CoT responses are much longer, exceeding 1,400 tokens on average. This substantial length gap reflects the dataset’s focus on detailed, step-by-step reasoning, providing rich supervision that encourages faithful and transparent reasoning in audio-based tasks.

In terms of audio representation, the user content speech averages roughly one minute per sample, whereas the CoT reasoning audio extends to nearly ten minutes on average. This substantial duration provides a challenging and realistic benchmark for assessing models’ ability to perform long-form audio comprehension and sustained logical reasoning.

As shown in Figure 4, the SoundMind dataset covers a broad range of audio durations, capturing important variability for developing and evaluating models for audio-based logical reasoning. Most segments fall within 3–12 minutes, with 28.6% of training samples concentrated in the 540–720s range, which provides ample contextual information while keeping computational cost tractable. Notably, the dataset deliberately excludes very short clips to ensure that each sample contains suf-

ficiently meaningful reasoning content.

The dataset exhibits a class distribution of 44.9% entailed and 55.1% not-entailed instances overall, closely resembling patterns in many real-world reasoning scenarios. This skew becomes more pronounced in the mid-range durations (240–420s), suggesting that longer and acoustically richer segments may introduce additional difficulty for reliable entailment classification. Notably, the 480–540s range in the training split shows a reversed distribution, with 56.9% entailed samples, potentially reflecting distinctive acoustic patterns or systematic properties of the original source data.

The SoundMind dataset is split to maintain consistent entailment ratios across training (42.8%), test (54.9%), and validation (56.3%) sets, thereby enabling reliable evaluation while preserving the inherent task complexity. Each split carefully preserves proportional coverage across all duration bins, effectively minimizing potential duration-related bias and ensuring representative diversity. The test set is nearly balanced (360 entailed vs. 296 not-entailed), providing a solid basis for rigorous and fair model assessment.

Overall, this distribution profile aligns with the dataset’s goal of advancing audio–language reasoning research by offering diverse but systematically controlled interaction scenarios. Such a design compels models to acquire robust cross-modal reasoning skills, discouraging reliance on superficial duration-based correlations and encouraging deeper semantic understanding.

5.3 Results and Analysis

5.3.1 Audio-to-Text Reasoning

Model	Accuracy (%) [↑]
MiniCPM-o	73.17
Gemini-Pro-V1.5	74.54
Baichuan-Omni-1.5	70.58
Qwen2-Audio	58.23
Qwen2.5-Omni-7B	77.59
Qwen2.5-Omni-7B (SoundMind-RL)	81.40

Table 4: Accuracy (%) of evaluated models on the SoundMind benchmark under the audio-to-text configuration. In this setting, models receive audio-only inputs and produce textual outputs containing the final answer. The table compares several strong multimodal LLM baselines with the reinforcement-tuned Qwen2.5-Omni-7B (SoundMind-RL).

Qwen2.5-Omni-7B fine-tuned with SoundMind-RL achieves state-of-the-art performance on the challenging audio-to-text reasoning task. Table 4 reports results for generating textual outputs from audio-only inputs. Our model attains an accuracy of 81.40%, clearly outperforming all evaluated baselines. Compared to its counterpart Qwen2.5-Omni-7B (77.59%), the reinforcement-tuned variant yields an absolute improvement of 3.81%, clearly demonstrating that the SoundMind reward framework substantially enhances the model’s ability to produce accurate text-based logical conclusions from speech audio input.

Among the remaining baseline models, Gemini-Pro-V1.5 reaches 74.54%, followed closely by MiniCPM-o (73.17%) and Baichuan-Omni-1.5 (70.58%). Qwen2-Audio performs the weakest, achieving only 58.23%, clearly underscoring that generic multimodal alignment is insufficient for reasoning-intensive tasks without specialized optimization. Overall, Qwen2.5-Omni-7B (SoundMind-RL) further improves upon Gemini-Pro by 6.86% and surpasses MiniCPM-o by 8.23%, establishing a consistently clear lead across models of comparable or larger capacity. These findings strongly confirm that our reinforcement learning approach strengthens reasoning accuracy and generalizes effectively across audio reasoning tasks.

5.3.2 Text-to-Audio Reasoning

Model	WER (%) [↓]	Acc. (%) [↑]
Qwen2.5-Omni-7B	2.18	80.79
Qwen2.5-Omni-7B (SoundMind-RL)	6.99	83.84

Table 5: Performance of Qwen2.5-Omni-7B and its SoundMind-RL fine-tuned variant on the SoundMind benchmark under the text-to-audio setting. In this configuration, models receive text-only inputs and generate audio outputs. Results are reported using Word Error Rate (WER) to assess speech fidelity and accuracy (%) to measure reasoning correctness.

SoundMind-RL significantly improves reasoning performance in the text-to-audio setting. As shown in Table 5, Qwen2.5-Omni-7B fine-tuned with SoundMind-RL achieves an accuracy of 83.84%, yielding a 3.05% improvement over the Qwen2.5-Omni-7B baseline. These results clearly indicate that the reward mechanism effectively encourages the model to produce coherent and consistent spoken reasoning, even without audio input. The observed accuracy gain is accompanied by a

higher WER (6.99% vs. 2.18%), which can be attributed to the generation of longer and more detailed reasoning sequences under reinforcement learning, thereby increasing the chance of recognition errors while preserving semantic correctness.

5.3.3 Audio-to-Audio Reasoning

Model	WER (%)↓	Acc. (%)↑
Qwen2.5-Omni-7B	2.23	77.59
Qwen2.5-Omni-7B (SoundMind-RL)	8.95	81.40

Table 6: Performance of Qwen2.5-Omni-7B and its SoundMind-RL fine-tuned variant on the SoundMind benchmark under the audio-to-audio reasoning setting. In this configuration, models receive audio-only inputs and generate audio reasoning outputs.

In the fully audio-based setting, our SoundMind-RL framework delivers notable performance gains under the most challenging condition. Table 6 presents results where models must derive and verbalize reasoning solely from acoustic input. Qwen2.5-Omni-7B fine-tuned with SoundMind-RL achieves 81.40% accuracy, improving substantially over the baseline (77.59%) and showing that the model can extract task-relevant logical dependencies directly from audio signals. The result is accompanied by a higher WER (8.95% vs. 2.23%), which we attribute to the generation of longer, more elaborate reasoning sequences and the added difficulty of learning from raw acoustic input. Our analysis attributes this to the absence of explicit fluency constraints during reinforcement learning, which instead prioritizes structured reasoning and semantic correctness.

Takeaway. Across all evaluated settings, SoundMind-RL consistently enhances logical reasoning performance for audio–language models. Qwen2.5-Omni-7B fine-tuned with SoundMind-RL achieves substantial gains across audio-to-text, text-to-audio, and audio-to-audio configurations, outperforming all baseline models. These results demonstrate that reward-guided reinforcement learning can reliably improve multimodal reasoning accuracy by encouraging faithful answer formatting, factual correctness, and sufficiently detailed reasoning traces. For speech-based outputs, we observe a moderate increase in WER relative to the baseline, which correlates with the production of longer, information-rich reasoning segments and suggests that the model prioritizes semantic completeness over surface fluency. Future work can

explore fluency-aware or prosody-sensitive reward components to better balance reasoning fidelity with the naturalness of generated speech.

6 Conclusion

We introduce SoundMind-RL, a novel rule-based reinforcement learning framework that empowers large-scale audio-language models with advanced logical reasoning capabilities across both audio and textual modalities. To support such training, we construct SoundMind, a dataset for the task of audio logical reasoning, containing 6,446 carefully curated audio–text pairs, each annotated with modality-specific chain-of-thought reasoning. Experimental results demonstrate that our method significantly improves performance and establishes state-of-the-art results on the SoundMind benchmark, consistently outperforming strong baselines across three reasoning settings: text-to-audio, audio-to-text, and audio-to-audio. We hope this work provides a useful resource and framework for advancing research in reasoning-oriented audio–language modeling and reinforcement-based multimodal learning.

Limitations

While this study shows promising results, several limitations remain. (1) The rule-based reward design enforces reasoning format consistency but may introduce rigidity, and its generalization to more diverse or open-ended tasks remains to be examined. (2) The dataset relies on synthetic speech and automatically generated chain-of-thought annotations; despite careful curation, subtle artifacts or biases may still be present and could influence model behavior. (3) The increase in Word Error Rate (WER) observed in audio generation suggests a possible trade-off between reasoning depth and fluency, motivating further work on prosody modeling and decoding strategies. We leave these challenges for future work, aiming to advance the development of audio reasoning systems that are more robust, generalizable, and aligned with human needs.

Ethical Considerations

We have not identified any ethical concerns directly related to this study.

Acknowledgment

This study is supported by the Department of Defense grant HT9425-23-1-0267.

References

- Martijn Bartelds, Nay San, Bradley McDonnell, Dan Jurafsky, and Martijn Wieling. 2023. Making more of little data: Improving low-resource automatic speech recognition using data augmentation. *arXiv preprint arXiv:2305.10951*.
- Jean Carletta, Simone Ashby, Sebastien Bourban, Mike Flynn, Mael Guillemot, Thomas Hain, Jaroslav Kadlec, Vasilis Karaiskos, Wessel Kraaij, Melissa Kronenthal, et al. 2005. The ami meeting corpus: A pre-announcement. In *International Workshop on Machine Learning for Multimodal Interaction*.
- Ming Cheng, Xingjian Diao, Shitong Cheng, and Wenjun Liu. 2024. Saic: Integration of speech anonymization and identity classification. In *AI for Health Equity and Fairness: Leveraging AI to Address Social Determinants of Health*, pages 295–306. Springer.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, et al. 2024. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*.
- Xingjian Diao, Chunhui Zhang, Tingxuan Wu, Ming Cheng, Zhongyu Ouyang, Weiyi Wu, and Jiang Gui. 2025. Learning musical representations for music performance question answering. *arXiv preprint arXiv:2502.06710*.
- Sreyan Ghosh, Zhifeng Kong, Sonal Kumar, S Sakshi, Jaehyeon Kim, Wei Ping, Rafael Valle, Dinesh Manocha, and Bryan Catanzaro. 2025. Audio flamingo 2: An audio-language model with long-audio understanding and expert reasoning abilities. *arXiv preprint arXiv:2503.03983*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Jarvis Guo, Tuney Zheng, Yuelin Bai, Bo Li, Yubo Wang, King Zhu, Yizhi Li, Graham Neubig, Wenhu Chen, and Xiang Yue. 2024. Mammoth-vl: Eliciting multimodal reasoning with instruction tuning at scale. *arXiv preprint arXiv:2412.05237*.
- Xiaotian Han, Yiren Jian, Xuefeng Hu, Haogeng Liu, Yiqi Wang, Qihang Fan, Yuang Ai, Huaibo Huang, Ran He, Zhenheng Yang, et al. 2024. Infiwebmath-40b: Advancing multimodal pre-training for enhanced mathematical reasoning. *arXiv preprint arXiv:2409.12568*.
- Jian Hu, Jason Klein Liu, Haotian Xu, and Wei Shen. 2025. Reinforce++: An efficient rlhf algorithm with robustness to both prompt and reward models. *arXiv preprint arXiv:2501.03262*.
- Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Ye Jia, Yifan Ding, Ankur Bapna, Colin Cherry, Yu Zhang, Alexis Conneau, and Nobuyuki Morioka. 2022. Leveraging unsupervised and weakly-supervised data to improve direct speech-to-speech translation. *arXiv preprint arXiv:2203.13339*.
- Ziyue Jiang, Yi Ren, Ruiqi Li, Shengpeng Ji, Boyang Zhang, Zhenhui Ye, Chen Zhang, Bai Jionghao, Xiaoda Yang, Jialong Zuo, et al. 2025. Megatts 3: Sparse alignment enhanced latent diffusion transformer for zero-shot speech synthesis. *arXiv preprint arXiv:2502.18924*.
- Feihu Jin, Yifan Liu, and Ying Tan. 2024. Zero-shot chain-of-thought reasoning guided by evolutionary algorithms in large language models. *arXiv preprint arXiv:2402.05376*.
- Kimi Team, Angang Du, Bofei Gao, Bowei Xing, Changjiu Jiang, Cheng Chen, Cheng Li, Chenjun Xiao, Chenzhuang Du, Chonghua Liao, et al. 2025. Kimi k1. 5: Scaling reinforcement learning with llms. *arXiv preprint arXiv:2501.12599*.
- Zhifeng Kong, Arushi Goel, Rohan Badlani, Wei Ping, Rafael Valle, and Bryan Catanzaro. 2024. Audio flamingo: A novel audio language model with few-shot learning and dialogue abilities. In *International Conference on Machine Learning*.
- Bin Li, Bin Sun, Shutao Li, Encheng Chen, Hongru Liu, Yixuan Weng, Yongping Bai, and Meiling Hu. 2024. Distinct but correct: generating diversified and entity-revised medical response. *Science China Information Sciences*.
- Yadong Li, Jun Liu, Tao Zhang, Song Chen, Tianpeng Li, Zehuan Li, Lijun Liu, Lingfeng Ming, Guosheng Dong, Da Pan, et al. 2025. Baichuan-omni-1.5 technical report. *arXiv preprint arXiv:2501.15368*.
- Hunter Lightman, Vineet Kosaraju, Yuri Burda, Harrison Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. In *International Conference on Learning Representations*.
- Hanmeng Liu, Jian Liu, Leyang Cui, Zhiyang Teng, Nan Duan, Ming Zhou, and Yue Zhang. 2023. Logiqa 2.0—an improved dataset for logical reasoning in natural language understanding. *Transactions on Audio, Speech, and Language Processing*.
- Ziyang Ma, Zhuo Chen, Yuping Wang, Eng Siong Chng, and Xie Chen. 2025. Audio-cot: Exploring chain-of-thought reasoning in large audio language model. *arXiv preprint arXiv:2501.07246*.
- OpenAI. 2024. Whisperspeech. <https://github.com/WhisperSpeech/WhisperSpeech>.

- Paribesh Regmi, Arjun Dahal, and Basanta Joshi. 2019. Nepali speech recognition using rnn-ctc model. *International Journal of Computer Applications*.
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2024. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning. In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.
- Zayne Rea Sprague, Fangcong Yin, Juan Diego Rodriguez, Dongwei Jiang, Manya Wadhwa, Prasann Singhal, Xinyu Zhao, Xi Ye, Kyle Mahowald, and Greg Durrett. 2025. To cot or not to cot? chain-of-thought helps mainly on math and symbolic reasoning. In *The Thirteenth International Conference on Learning Representations*.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc Le, Ed Chi, Denny Zhou, et al. 2023. Challenging big-bench tasks and whether chain-of-thought can solve them. In *Findings of the Association for Computational Linguistics: ACL*.
- Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.
- Ziyu Wan, Xidong Feng, Muning Wen, Stephen Marcus McAleer, Ying Wen, Weinan Zhang, and Jun Wang. 2024. Alphazero-like tree-search can guide large language model decoding and training. In *International Conference on Machine Learning*.
- Wengran Wang, Yudong Rao, Rui Zhi, Samiha Marwan, Ge Gao, and Thomas W Price. 2020. Step tutor: Supporting students through step-by-step example-based feedback. In *ACM conference on Innovation and Technology in Computer Science Education*.
- Xinsheng Wang, Mingqi Jiang, Ziyang Ma, Ziyu Zhang, Songxiang Liu, Linqin Li, Zheng Liang, Qixi Zheng, Rui Wang, Xiaoqin Feng, et al. 2025. Spark-tts: An efficient llm-based text-to-speech model with single-stream decoupled speech tokens. *arXiv preprint arXiv:2503.01710*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in neural information processing systems*.
- Tian Xie, Zitian Gao, Qingnan Ren, Haoming Luo, Yuqian Hong, Bryan Dai, Joey Zhou, Kai Qiu, Zhirong Wu, and Chong Luo. 2025a. Logic-rl: Unleashing llm reasoning with rule-based reinforcement learning. *arXiv preprint arXiv:2502.14768*.
- Zhifei Xie, Mingbao Lin, Zihang Liu, Pengcheng Wu, Shuicheng Yan, and Chunyan Miao. 2025b. Audio-reasoner: Improving reasoning capability in large audio language models. *arXiv preprint arXiv:2503.02318*.
- Guowei Xu, Peng Jin, Li Hao, Yibing Song, Lichao Sun, and Li Yuan. 2024. Llava-o1: Let vision language models reason step-by-step. *arXiv preprint arXiv:2411.10440*.
- Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, et al. 2025. Qwen2. 5-omni technical report. *arXiv preprint arXiv:2503.20215*.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. In *Advances in neural information processing systems*.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. 2024. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*.
- Wenhao You, Xingjian Diao, Chunhui Zhang, Keyi Kong, Weiyi Wu, Zhongyu Ouyang, Chiyu Ma, Tingxuan Wu, Noah Wei, Zong Ke, et al. 2025. Music’s multimodal complexity in avqa: Why we need more than general multimodal llms. *arXiv preprint arXiv:2505.20638*.
- Xiangchi Yuan, Chunhui Zhang, Zheyuan Liu, Dachuan Shi, Leyan Pan, Soroush Vosoughi, and Wenke Lee. 2025. Superficial self-improved reasoners benefit from model merging. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*.
- Rodolfo Zevallos. 2022. Text-to-speech data augmentation for low resource speech recognition. *arXiv preprint arXiv:2204.00291*.
- Xiao Zhan, Noura Abdi, William Seymour, and Jose Such. 2024. Healthcare voice ai assistants: factors influencing trust and intention to use. *ACM on Human-Computer Interaction*.
- Chunhui Zhang, Yiren Jian, Zhongyu Ouyang, and Soroush Vosoughi. 2024. Working memory identifies reasoning limits in language models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.
- Chunhui Zhang, Zhongyu Ouyang, Kwonjoon Lee, Nakul Agarwal, Sean Dae Houlihan, Soroush Vosoughi, and Shao-Yuan Lo. 2025a. Overcoming multi-step complexity in multimodal theory-of-mind reasoning: A scalable bayesian planner. In *International Conference on Machine Learning*.

Chunhui Zhang, Sirui Wang, Zhongyu Ouyang, Xiangchi Yuan, and Soroush Vosoughi. 2025b. Growing through experience: Scaling episodic grounding in language models. In *Annual Meeting of the Association for Computational Linguistics*.

Yu Zhao, Huifeng Yin, Bo Zeng, Hao Wang, Tianqi Shi, Chenyang Lyu, Longyue Wang, Weihua Luo, and Kaifu Zhang. 2024. Marco-o1: Towards open reasoning models for open-ended solutions. *arXiv preprint arXiv:2411.14405*.

Guolong Zhong, Hongyu Song, Ruoyu Wang, Lei Sun, Diyuan Liu, Jia Pan, Xin Fang, Jun Du, J. Zhang, and Lirong Dai. 2022. External text based data augmentation for low-resource speech recognition in the constrained condition of openasr21 challenge. In *Interspeech*.

Zyphra. 2024. Zonos. <https://github.com/Zyphra/Zonos>.

A Baselines

- **MiniCPM-o** (Yao et al., 2024): An open-source multimodal language model supporting vision, speech, and real-time streaming. It emphasizes efficiency and deployability, making it suitable for interactive multimodal understanding and generation across a variety of general-purpose and resource-constrained scenarios.
- **Gemini-Pro-V1.5** (Team et al., 2024): A multimodal model developed by Google that processes text, image, audio, and video inputs in a unified way. It is designed to handle extended context reasoning and cross-modal alignment, enabling coherent understanding and response in complex, multi-step tasks across modalities.
- **Baichuan-Omni-1.5** (Li et al., 2025): A multimodal model that combines text, vision, and audio understanding with audio generation. It aims to enable coherent cross-modal interaction and supports a wide range of multimodal and multitask applications in interactive settings.
- **Qwen2-Audio** (Chu et al., 2024): A speech language model that accepts both audio and text as input for natural voice-based interaction. It can perform detailed audio analysis, transcription, and multilingual understanding, serving as a general-purpose audio–text reasoning system.
- **Qwen2.5-Omni-7B** (Xu et al., 2025): A multimodal model that handles text, image, audio, and video inputs and generates both text and speech outputs. It supports streaming interactions and

aims to provide robust multimodal reasoning and natural communication in real-time applications.

B Text-to-Speech Model

Constructing a fully human-recorded version of the audio logical reasoning dataset—covering thousands of user prompts, chain-of-thought reasoning traces, and final answers in both input and output modalities—would require extensive scripting, careful speaker coordination, and thousands of hours of professional narration. Such a large-scale recording effort is prohibitively expensive and logistically infeasible given the diversity and duration of reasoning sequences we target. Consistent with established practice in speech and multimodal learning research (Jia et al., 2022; Zevallos, 2022; Regmi et al., 2019; Cheng et al., 2024; Bartelds et al., 2023; Zhong et al., 2022), we synthesize the audio modality using a high-fidelity text-to-speech (TTS) system. This strategy enables scalable and reproducible generation of paired input–output audio segments, while preserving fine control over speaking rate, prosody, and acoustic quality—key factors for training models to perform multi-minute reasoning without performance degradation.

To ensure that the synthesized audio faithfully reflects the complexity of logical reasoning, we conducted a systematic evaluation of several state-of-the-art open-source TTS systems. We compared MegaTTS-3 (Jiang et al., 2025), WhisperSpeech (OpenAI, 2024), Spark-TTS (Wang et al., 2025), and Zonos (Zyphra, 2024) across multiple dimensions, including perceived naturalness, prosody stability over long contexts, noise robustness, and intelligibility of extended reasoning chains. Each candidate system was benchmarked by synthesizing representative samples from our dataset and evaluated via multi-rater human listening studies. MegaTTS-3 consistently achieved the highest mean opinion scores, particularly in naturalness and prosodic consistency, qualities that we consider essential for faithfully preserving the chain-of-thought signal.

On the basis of these findings, we select MegaTTS-3 as the synthesis engine for both user prompts and model reasoning traces. Its ability to produce long, coherent, and acoustically stable speech makes it well suited for the extended reasoning sequences characteristic of the ALR task. The resulting corpus provides a controlled yet realistic training signal that aligns with our objective of enabling robust, audio-native logical reasoning

Configuration	λ_1 (TextFmt)	λ_2 (AudioFmt)	λ_3 (Answer)	λ_4 (TextLen)	λ_5 (AudioLen)	Accuracy (%)	Δ vs. full (pp)
(1) w/o audio rewards	1.0	0	2.0	1.0	0	70.82	-10.58
(2) w/o text rewards	0	1.0	2.0	0	1.0	48.84	-32.56
(3) w/o answer reward	1.0	0.5	0	1.0	0.75	60.24	-21.16
(4) full (Ours)	1.0	0.5	2.0	1.0	0.75	81.40	—

Table 7: Ablation of the composite reward in SoundMind-RL, evaluated on the test set. λ_1/λ_2 enforce text/audio format compliance, λ_3 supervises *answer correctness*, and λ_4/λ_5 encourage *reasoning length* in text/audio space. Accuracy is reported in %, and Δ indicates *percentage-point (pp)* difference relative to the full model.

in large audio-language models. We further acknowledge the open-source speech community for providing high-quality TTS systems, which have been instrumental in enabling reproducible, large-scale construction of the audio datasets.

C Potential Applications

Reasoning-capable ALMs can enable applications that demand explainability, transparency, and natural interaction. We highlight key domains where audio-native reasoning offers clear value and outline open challenges for future research.

Education and Tutoring. Spoken step-by-step explanations can help learners understand complex problems by exposing intermediate reasoning rather than presenting only a final answer (Wang et al., 2020). This is particularly relevant for subjects such as logic, mathematics, and music, where reasoning over multimodal signals is essential for mastery. Recent studies on music-related question answering and multimodal reasoning further highlight the need for models that can integrate audio and symbolic information to support learning and feedback (You et al., 2025; Diao et al., 2025).

Accessibility and Screen-Free Interaction. For users with visual impairments, hands-busy scenarios, or clinical settings, verbal reasoning offers an inclusive interface to access complex information. Replayable and summarizable reasoning can support comprehension and patient-centered decisions, consistent with efforts on trustworthy medical response generation (Li et al., 2024) and healthcare voice assistants (Zhan et al., 2024). Ensuring natural prosody and coherence over multi-minute speech remains crucial for high-stakes applications.

Conversational Agents and Voice Assistants. In applications such as troubleshooting, procedural guidance, and eligibility checking, users often ask “why” in addition to “what.” Models capable of structured, spoken justifications could reduce repeated follow-ups and improve user trust.

Meetings, Lectures, and Knowledge Capture.

Beyond transcription, audio-native reasoning models could verbalize intermediate inferences, summarize discussions, and surface entailment relations in real time, turning speech streams into structured knowledge. Robustness to overlapping speech, noise, and speaker variability remains a key challenge, as exemplified by the AMI Meeting Corpus (Carletta et al., 2005).

D Ablation Analysis

Quantitative impact. Table 7 reports the results of ablating individual reward components. Removing *text-format/length* supervision (Row 2) causes the most severe degradation, with accuracy dropping to 48.84% (-32.56% relative to the full model). This sharp decline underscores that structural guidance on the text side is essential for producing well-formed and logically verifiable reasoning traces. Excluding the *answer-correctness* reward (Row 3) reduces accuracy to 60.24% (-21.16%), indicating that explicit factual supervision is necessary to ensure generated reasoning converges to the correct entailment decision rather than drifting into irrelevant but superficially coherent explanations. Removing *audio-format/length* rewards (Row 1) still lowers accuracy to 70.82% (-10.58%), suggesting that modality-consistent formatting and duration constraints regularize the shared policy and enhance robustness. Overall, these results demonstrate that each reward term contributes meaningfully to model performance, and that their combination is crucial for maximizing reasoning accuracy.

Key takeaways. Table 7 highlights three main takeaways: (1) Text-side format and length supervision is the dominant factor for producing reliable, verifiable reasoning traces. (2) Answer-correctness reward is essential to align reasoning with ground truth labels and convert fluent explanations into correct decisions. (3) Audio-side rewards serve as powerful regularizers, improving policy robustness and benefiting reasoning quality across modalities.