

# Soundwave : Less is More for Speech-Text Alignment in LLMs

Yuhao Zhang, Zhiheng Liu, FanBu, Ruiyu Zhang, Benyou Wang\*, Haizhou Li

The Chinese University of Hong Kong, Shenzhen  
yoo hao . zhang @ gmail . com , wang benyou @ cuhk . edu . cn

## Abstract

Existing end-to-end speech large language models (LLMs) usually rely on large-scale annotated data for training, while data-efficient training has not been discussed in depth. We focus on two fundamental problems between speech and text: the representation space gap and sequence length inconsistency. We propose Soundwave, which utilizes an efficient training strategy and a novel architecture to address these issues. Results show that Soundwave outperforms other advanced speech LLMs in speech translation and AIR-Bench speech tasks with only a fraction of the training data. Further analysis shows that Soundwave still retains its intelligence during conversation.

## 1 Introduction

Large language models (LLMs) have profoundly transformed the paradigm of natural language processing (NLP) due to their remarkable abilities in understanding and reasoning (Achiam et al., 2023; Touvron et al., 2023). Recently, multi-modal LLMs have also shown rapid development, with the success of GPT-4o highlighting the potential of speech-focused LLMs (Hurst et al., 2024). A fundamental requirement for achieving seamless communication with LLMs is their ability to accurately interpret speech—essentially enabling LLMs to “hear”.

However, most speech-based LLMs rely on massive labeled datasets and substantial computational resources to enable speech perception (Défossez et al., 2024; Chen et al., 2025). For example, the Qwen2-Audio (Chu et al., 2024) model requires approximately 500,000 hours of data to achieve cross-modal functionality, while 1,000 hours is needed by advanced automatic speech recognition models to achieve comparable results (Gulati et al., 2020). This discrepancy underscores the need for more efficient methods to develop speech-capable LLMs.

\* Corresponding author.

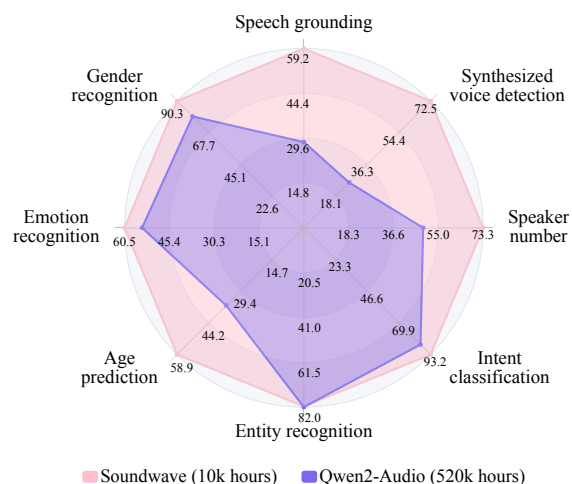


Figure 1: AIR-Bench speech foundation tasks.

We first identify two fundamental challenges to achieve alignment between speech and text (Zhang et al., 2023): (1) the **representation space gap** and (2) **sequence length inconsistency**. The former challenge arises from the two modalities being developed independently, while the second challenge stems from the difference in modeling units—speech is typically represented at the frame level, whereas LLMs operate at the sub-word level. We then decouple the process of building speech LLMs to achieve more efficient training. Specifically, we propose a two-stage training framework designed to efficiently overcome these challenges. The first stage focuses on resolving the representation space gap, while the second stage aims to reduce the sequence length of speech.

Furthermore, to efficiently bridge the gap between speech and text, the quality of alignment data plays a crucial role. To address this, we collect high-quality speech recognition data and manually annotate audio labels to support the first stage. For the second stage, we analyze the proportion of text data to ensure a smooth learning process. During the supervised fine-tuning (SFT) stage, we

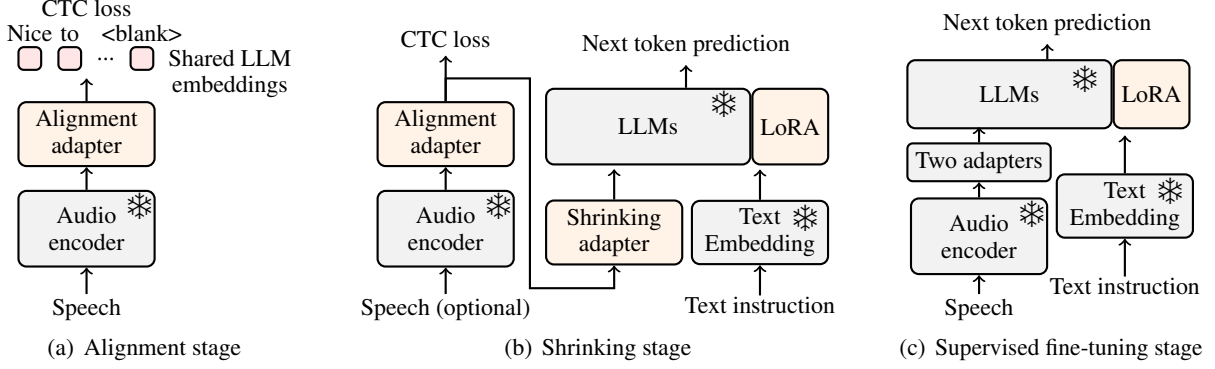


Figure 2: Training progress of Soundwave. The gray modules are frozen while the orange modules are updated.

employ temperature sampling to balance the variety of tasks effectively.

We conduct experiments on various speech-to-text tasks and several sound-related tasks. We also compare our model with the strong system Qwen2-Audio on both closed- and open-ended tasks. Our average results achieve state-of-the-art performance on the AIR-Bench (Yang et al., 2024) speech foundation task as shown in Fig. 1. It also shows comparable results on audio-related tasks. Furthermore, our method exhibits significant performance in zero-shot speech translation, demonstrating that Soundwave unlocks the full potential of LLMs. Soundwave delivers better performance with less training data, lower training costs, and fewer speech sequences.

Our main **contributions** are as follows. 1) We propose an efficient training framework that utilizes only ten thousand hours of training data to achieve state-of-the-art speech understanding performance. 2) We introduce dynamic multi-task learning in the post-training stage to enhance speech modeling and leverage the benefits of text data. 3) We annotate a sound classification dataset to support the alignment between sound and text, and provide high-quality instruction data based on the thinking process for speech LLMs.

## 2 Methodology

### 2.1 Overall Design

The training process consists of three stages, as shown in Fig. 2. Stage I aims to align the representation between speech and text, addressing the *representation space gap* problem. Stage II primarily shrinks the speech sequence and mitigates the *sequence length inconsistency*. The supervised fine-tuning (SFT) stage (Wei et al., 2021) enables

Table 1: The parameters of different modules. The orange represents the number of training parameters.

Modules	#Param.	Training stage	Details
Audio encoder	~635M	-	Whisper Large V3
Alignment adapter	~144M	I&II	One projection layer and Transformer
Shrinking adapter	~67M	II	One cross-attention and layer-norm
LLMs	~8B	-	Llama3.1
LLM adapter	~55M	II&III	LoRA
Total	~9B		

the speech LLMs to generalize across diverse tasks.

The input to the model consists of speech FBank features, which are then processed by the pretrained audio encoder. To efficiently align the representation with that of LLMs, we use the audio encoder that produces semantic features (e.g., Whisper (Radford et al., 2023) or Seamless (Team, 2025)), rather than vector quantization features (Défossez et al., 2022) or self-supervised features (Hsu et al., 2021). We implement an alignment adapter and a shrinking adapter to bridge the gap between speech and text. Additionally, LoRA (Hu et al., 2022) is used to enable efficient fine-tuning. An overview of the modules is provided in Tab. 1.

### 2.2 Stage I: Alignment

We use the auxiliary CTC loss and high-quality data to improve training efficiency.

#### 2.2.1 Auxiliary CTC loss

The audio encoder and LLMs have a gap in their representation spaces due to separate pre-training. One direct approach is to use ASR tasks for alignment. We design an adapter and utilize CTC loss

(Graves et al., 2006) to achieve efficient cross-modal training. Specifically, the adapter consists of a linear layer followed by a Transformer layer (Vaswani, 2017). The linear layer transforms the audio sequence  $A \in \mathbb{R}^{l \times h_a}$  where  $l$  is the length of the speech sequence and  $h_a$  is the hidden size of the audio encoder. We concatenate adjacent features and adjust the dimensionality to match that of the LLMs, resulting in  $A' \in \mathbb{R}^{l/2 \times h_{llm}}$  where  $h_{llm}$  is the hidden size of the LLMs. A Transformer layer then converts the features into the representation space of LLMs. Finally, we use CTC loss to train the adapter, aligning the shared space of the LLMs.

### 2.2.2 High-quality Alignment Data

We believe that *improving data quality is crucial to training efficiency for alignment*. We apply data strategies for two types of data (ASR and sound data), as outlined below. The adapter is trained without the LLMs at this stage, thus the alignment training is fast. Our later experiments in Sec. 5.4 confirm the benefits to training efficiency.

**Verified ASR Data** At this training stage, we use transcriptions from ASR data as the target, which we found to be crucial for improving convergence ratio. The selected high-quality data is all verified by advanced ASR model (Radford et al., 2023) with a Word Error Rate (WER) lower than 10%.

**Standardized Sound Data** Another challenge is processing sound due to the inherent background noise and the diversity of labeling information. To address this, we annotate about 8k pieces of sound category data. We further select clear 20k sound samples, then unify label format and audio length.

## 2.3 Stage II: Shrinking

After aligning the data representation, we focus on reducing the length of the speech as detailed in Sec. 2.3.1. Additionally, at this stage, we include various types of foundational audio tasks to better generalize downstream tasks. This introduces a data mixture problem, which is solved by a dynamic data mixture strategy (see Sec. 2.3.2).

### 2.3.1 Dynamic Shrinking

There are two essential aspects to shrinking the audio sequence: *final length determination* and *lossless information retention*.

**Final Length Determination** For the first aspect, we utilize the probability from CTC. The CTC predicts the corresponding word for each position. We

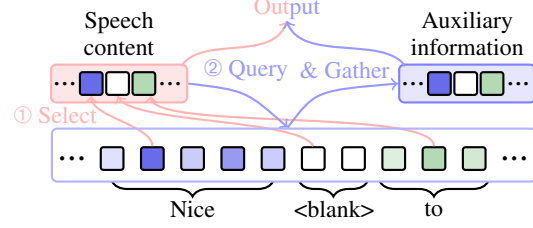


Figure 3: We first select the features based on the peak of CTC prediction. Then, we use these features to query and gather auxiliary information from the original sequence. Finally, we fuse the two features to achieve shrinking. Note that when gathering the information, the selected positions are masked.

then remove duplicate predictions from adjacent positions to obtain the final sequence. Since the sequence has been aligned to text in Stage I, the decoded result can indicate the final length.

**Lossless Information Retention** For the second aspect, we select the content based on the CTC output as the query, and then use attention mechanisms to gather related information, such as tone and pitch, in order to prevent information loss.

Assume the speech features  $x$  have been aligned to the representation space of LLMs, then, we select the features based on the CTC probability to compress the sequence  $x'$ .

$$x_{out} = \text{norm}(x' + \text{cross\_attn}(x', x, x)) \quad (1)$$

where  $\text{norm}$  is the layer norm operation.  $x_{out}$  is the final output of the shrinking adapter.  $x'$  can be viewed as the content feature, while the gathered information, calculated by cross-attention, serves as auxiliary data for the selected features. The whole processing is shown in Fig. 3.

Table 2: The overview of tasks in shrinking stage. The data scales of these data are highly imbalanced.

Task	Input	Output	Size (k)
QA	Text question	Text answer	78
ASR	Speech	Transcription	3,012
ST	Speech	Translation	460
Sound	Mixed speech and sound	Transcription and sound type	25

### 2.3.2 Dynamic Data Mixture

We select both audio data (involving three basic audio tasks) and text data to enable LLMs to generalize to downstream speech understanding. Training with mixed data may be biased by dominant

Table 3: Summary of datasets used in different stages and their total hours.

Dataset	I	II	III	Num.	Hours	Task
GigaSpeech (M) (Chen et al., 2021)	✓	✓		713k	805.11	ASR
TED-LIUM (Hernandez et al., 2018)	✓	✓		144k	244.02	ASR
Multilingual LibriSpeech (En) (Pratap et al., 2020)	✓	✓		985k	4,081.61	ASR
Europarl-ASR (Garcés Díaz-Munío et al., 2021)	✓	✓		719k	418.42	ASR
TextrolSpeech (Ji et al., 2024b)	✓	✓	✓	215k	301.19	ASR, GR <sup>1</sup> , Emotion Recognition
LibriSpeech (Panayotov et al., 2015)	✓	✓	✓	281k	961.05	ASR, Speech Grounding
MUST-C (En-De) (Cattoni et al., 2021)		✓	✓	283k	388.55	Speech Translation
Common Voice (En) (Ardila et al., 2019)		✓	✓	233k	364.64	AP <sup>1</sup> , Speech Translation
Fisher (Cieri et al., 2004)		✓	✓	132k	1,091.42	ASR, Chat
Europarl-ST (Iranzo-Sánchez et al., 2020)			✓	53k	133.16	Language Identification
Common Voice (Ja) (Ardila et al., 2019)			✓	13k	15.00	Language Identification
SLURP (Bastianelli et al., 2020)			✓	141k	101.49	IC <sup>1</sup> , Entity Recognition
CREMA-D (Cao et al., 2014)			✓	7k	5.26	Emotion Recognition
RAVDESS (Livingstone and Russo, 2018)			✓	1k	1.48	Emotion Recognition
IEMOCAP (Busso et al., 2008)			✓	3k	2.16	Emotion Recognition
MELD (Poria et al., 2019)			✓	9k	8.12	Emotion Recognition
VoxCeleb (Nagrani et al., 2017)			✓	156k	435.17	Speaker Num. Verification
FoR (Reimao and Tzerpos, 2019)			✓	54k	47.55	Synthesized Detection
AnyInstruct (Zhan et al., 2024)			✓	107k	206.30	Speech Instruction
VocalSound (Gong et al., 2022)	✓	✓	✓	20k	23.20	Sound Classification
TUT2017 (Duppada and Hiray, 2017)	✓	✓	✓	5k	13.00	Scene Classification
CochlScene (Jeong and Park, 2022)			✓	75k	208.65	Scene Classification
<b>Total</b> <sup>2</sup>				<b>4,349k*</b>	<b>9,856.91*</b>	<b>15</b>

<sup>1</sup> GR is for Gender Recognition, AP is for Age Prediction, and IC is for Intent Classification.

<sup>2</sup> ‘\*’ means that this table is compiled from the perspective of audio, and an audio file may be used multiple times for different tasks. If multiple usages at different tasks are all counted, the number of data samples is 6301k, and the total duration is 14068.77 hours.

tasks due to data imbalance (see Table 2), and existing work has adopted curriculum learning (Das et al., 2024; Tang et al., 2023), though it requires considerable prior knowledge for proper design.

Inspired by temperature-based data sampling, which has previously been used to address multilingual data imbalance (Arivazhagan et al., 2019), we propose a dynamic data mixture guided by sampling temperature. Specifically, the sample rate for each task  $k$  is as follows:

$$p_k = \left( \frac{|\mathcal{D}_k|}{\sum_i |\mathcal{D}_i|} \right)^{\frac{1}{T}} \quad (2)$$

where the  $|\mathcal{D}_k|$  denotes the data size of task  $k$  and  $T$  denotes the temperature.  $T$  is initially set to 1 and gradually increases. This causes the training to start with a *sample-level uniform distribution* and gradually shift to a *task-level uniform distribution*. Training at the former stage might be dominated by rich-source tasks, while at the latter stage, training might be more balanced among tasks, potentially alleviating the over-fitting issue.

Additionally, Chen et al. (2024) shows that text-related tasks aid instruction following for multi-modal LLMs. We also introduce the text task to ensure a smoother cross-modal process. We incorporate the Wizard SFT dataset (Xu et al., 2024) to

help speech LLMs retain their understanding capabilities, thereby enhancing their ability to follow instructions for speech tasks.

## 2.4 Stage III: Supervised Fine-tuning

At this stage, we only fine-tune the parameters of LoRA, as speech and text are already aligned. Our goal is to enable the speech LLMs to handle more complex tasks and respond directly based on the speaker’s speech. Thus, we use both text-based and speech-based instructions during SFT.

## 3 Data Engineering

We introduce the data details for the three stages, respectively, see the summary in Tab. 3. The data shown in the table has been cleaned and filtered, and the details of strategies can be found in App. A.2. We sample some speech from several dataset to control the quality and training cost.

### 3.1 Data During Stage I and II

**ASR Data** We choose high-quality datasets and filter the data with a WER of less than 10%, as tested by Whisper medium. We apply SpecAugment (Park et al., 2019) to enhance the robustness of the model towards speech. To help LLMs understand the conversation and the number of speakers,



Table 4: Examples of three QA formats for building SFT data in the speech emotion task.

	Instruction	Answer
1	What feeling is the speaker trying to convey?	Neutral.
2	What is the emotional theme of the speaker’s message? Pick one answer from: A. neutral B. fear C. surprise D. happiness.	The one answer you should pick is D. happiness.
3	What emotion is the speaker hinting at in their speech?	The speaker is hinting at surprise in their speech.

we splice speech from different speakers. We denote the output format as The first speaker says ... The second speaker says ....

**Sound Data** The sound data is often too short and may be viewed as noise, which causes the model to fail in perceiving it. To address this problem, we embedded environmental sounds into the audio to construct the data. For example: ‘But there was a passenger dropped off for you, a little girl. *<throat\_clearing>* It’s a boy I’ve come for.’ The special token is added to the conversation, and the model needs to transcribe both the sound and speech simultaneously. This method of learning both speech and audio also makes training efficient.

Another problem is that the sound is always a mix of multiple sound categories and exhibits constant repetition. This causes timing and label count issues, which increases the difficulty of learning alignment. Therefore, we standardize the audio to a duration of three seconds based on energy, and manually select and label the data. The detailed process of labeling can be found in our Appendix A.1.

### 3.2 Instruction Data During Stage III

**Text Instructions** The text-based instruction is designed to understand and analyze speech. We have created three types of QA formats, as shown in Tab. 4. The first requires the model to directly answer, which is the most difficult. The second provides detailed choices, and the last requires the model to output the answer in a natural format.

**Speech Instructions** If speech LLMs are to communicate directly with humans, it is essential for them to follow speech instructions. Once the speech is well-aligned, we can achieve this by using text-based dialogue data and synthesizing text into speech. We use AnyInstruct speech subset (Zhan et al., 2024), which is built using this approach.

**Chain of Thought** To reduce the complexity of some challenging tasks, we built the dataset to enable the model to predict in a manner similar to Chain of Thought (CoT). For example, in the Intent

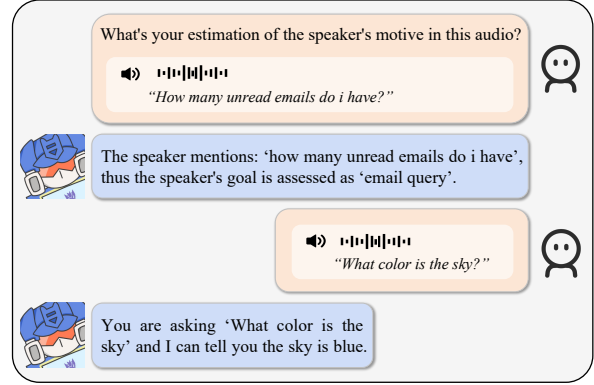


Figure 4: Adding thought processes to address complicated problems and speech instructions.

Classification task, we first had the model output the speech transcription before identifying the intent, as shown in Fig. 4. For the speech instruction, the model initially predicts the transcription and then responds to the question. The reasoning time is slightly increased, but the model can address complicated tasks with limited training data.

## 4 Experiments

### 4.1 Settings

**Training** The audio encoder is Whisper Large V3 (Radford et al., 2023), and the foundation model is Llama-3.1-8B-Instruct (Dubey et al., 2024). The alignment adapter is a projection where the output size is 4096. We apply LoRA to the Attention module, where rank and  $\alpha$  are set to 64 and 16, respectively. Both alignment and shrinking stages consist of 6,000 steps, with the SFT stage set to around 4,000 steps. The sample temperature  $T$  at Stage II starts at 1 and increases by 5 per training epoch. The experiments are conducted on 32 A800 GPUs for training on 10k hours of data. The training time for the two stages is approximately four days, and the SFT requires an additional day. App. C shows more details about training settings.

**Evaluation** We evaluate Soundwave on several basic tasks and the open-ended AIR-Bench. We also remove repeated samples (see App. A.4) before training to avoid data leakage. We primarily

Table 5: Performance on foundation tasks, including ASR, speech translation (ST), Speech Emotion Recognition (SER), Vocal Sound Classification (VSC). ST is evaluated by ScareBLEU (Post, 2018). \* denotes the zero-shot task.

Task	Dateset	Model	Metric	Performance
ASR	Librispeech (Test-clean/Test-other)	SALMONN (Tang et al., 2023)	WER ↓	2.1   4.9
		SpeechVerse (Das et al., 2024)		2.5   4.7
		WavLLM (Hu et al., 2024)		2.0   4.8
		Qwen2-Audio (Chu et al., 2024)		<b>1.6   3.6</b>
		Soundwave		2.1   5.0
ST	CoVoST2 En-De	BLSP (Wang et al., 2023)	BLEU ↑	14.1
		SALMONN (Tang et al., 2023)		18.6
		Qwen2-Audio (Chu et al., 2024)		29.9
		Soundwave		<b>30.6</b>
	MuST-C* (En-Nl En-It En-Ro En-Es)	Qwen2-Audio (Chu et al., 2024)	BLEU ↑	20.7   19.5   11.8   22.1
		Soundwave		<b>27.0   22.2   16.9   26.7</b>
SER	Meld	Qwen2-Audio (Chu et al., 2024)	ACC ↑	0.553
		Soundwave		<b>0.635</b>
VSC	VocalSound	Pengi (Deshmukh et al., 2023)	ACC ↑	0.604
		Qwen2-Audio (Chu et al., 2024)		<b>0.939</b>
		Soundwave		0.905

compare Soundwave with Qwen2-Audio, an advanced model for various audio processing tasks.

## 4.2 Results

**Basic Audio Tasks** We show the results on foundational audio tasks in Tab. 5. We find that our model demonstrates a significant advantage on the ST and SER tasks, which heavily rely on the understanding ability of speech LLMs. We also observe that our model shows strong performance on zero-shot tasks, such as translation tasks in other languages. On the other hand, our model still underperforms the SOTA model on the ASR task, indicating that massive training data is essential for ASR. We only used about 244 hours of sound data, which is dozens of times less than the SOTA, thus there is still a gap on the VSC task.

**AIR-Bench** We compare our model on AIR-Bench across speech foundation, sound foundation, and speech chat tasks. As shown in Tab. 6, our model demonstrates SOTA performance on average speech foundation tasks with only about 10k of training data. Specifically, we outperform the best of previous speech LLMs on six sub-tasks. Since 98.61% of the training data consists of English speech, our model performs worse on the language identification task. This highlights that the proportion of different languages remains important.

Results of the sound foundation task are shown in Tab. 7. Although only around 244 hours of data were used, our model is still superior to other models, except Qwen2-Audio, which is trained with 10k hours. Moreover, our single-encoder architecture performs better than the two-encoder

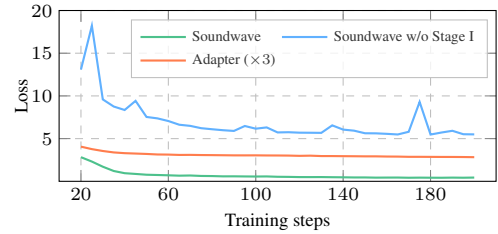


Figure 5: Training curves of different strategies

model (Tang et al., 2023), indicating that fewer encoders can process both speech and sound simultaneously. Our model also performs well in AIR-Bench speech chat task, ranking second only to Qwen2-Audio among open source models.

## 5 Analysis

Considering that analysis based on full data requires massive training cost, we analyze our method based on Librispeech data. The experiments are trained on 8 A800 GPUs with 4,000 steps. We use **Adapter** ( $\times n$ ) to denote that the adapter architecture is the same as Qwen2-Audio, where  $n$  is the down-sampling rate.

### 5.1 Convergence Rate

In Fig. 5, we compare the convergence rate with and without the first alignment stage, and the projection adapter architecture. Soundwave sees a high convergence rate, with the loss rapidly decreasing within the first hundred steps. In contrast, the training process of the other model is much slower without the alignment stage. Furthermore, Soundwave performs worse than other models without stage one, because the shrinking adapter relies on the CTC prediction.

Table 6: Performance on the AIR-Bench speech foundation tasks.

Task	Soundwave	Qwen2-Audio	Qwen-Audio Turbo	SALMONN	BLSP	NExT-GPT	PandaGPT	Whisper +GPT-4
Speech Grounding	<b>59.2</b>	28.3	45.4	25.3	25.0	25.4	23.0	35.0
Language Identification	89.6	93.3	95.9	28.1	30.8	23.7	34.6	<b>96.8</b>
Gender Recognition	<b>90.3</b>	79.3	82.5	35.5	33.2	57.0	66.5	21.9
Emotion Recognition	<b>60.5</b>	54.6	60.0	29.9	27.4	25.7	26.0	59.5
Age Prediction	58.9	36.1	58.8	48.7	51.2	<b>62.4</b>	42.5	41.1
Entity Recognition	81.7	<b>82.0</b>	48.1	51.7	37.2	26.1	34.0	69.8
Intent Classification	<b>93.2</b>	85.8	56.4	36.7	46.6	25.6	28.5	87.7
Speaker num. Verification	<b>73.3</b>	48.8	54.3	34.3	28.1	25.4	43.2	30.0
Synthesized Detection	<b>72.5</b>	25.9	69.3	50.0	50.0	30.8	53.1	40.5
Average	<b>75.5</b>	59.3	63.4	37.8	36.6	33.6	39.0	53.6

Table 7: Performance on the AIR-Bench audio foundation and chat tasks.

Task	Soundwave	Qwen2-Audio	Qwen-Audio Turbo	SALMONN	BLSP	NExT-GPT	PandaGPT	Gemini (1.5-pro)
Sound (average)	62.10	<b>65.10</b>	60.95	32.95	31.43	32.15	43.58	-
Speech Chat	6.51	<b>7.18</b>	7.04	6.16	6.17	3.86	3.58	6.97

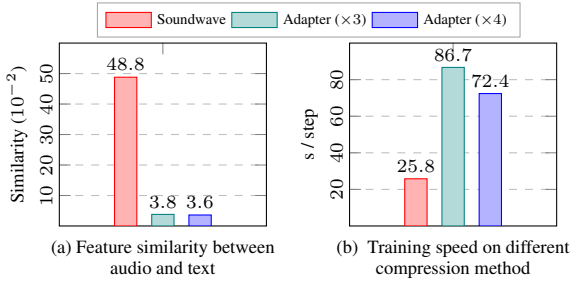


Figure 6: Comparison of alignment effect and speed.

## 5.2 Effect of Alignment

We randomly sampled 200 items from the Librispeech test clean set and then extracted text and speech representations. The similarity of speech and text after average pooling is compared, as shown in Fig. 6. We found that the representation of Soundwave with the alignment adapter is significantly higher than that of other methods. In addition, we further compare the average training speed under the same batch conditions. The training speed in the alignment stage is nearly three times faster than that of other methods. Whether due to the alignment effect or the training method, the alignment adapter shows obvious advantages.

## 5.3 Effect of Shrinking

We compare the performance and compression ratios of different strategies on ASR tasks. We found that our approach compresses significantly based on text length. Our method maintains stable performance with 2.5% compression ratios. However, the compression method leads to performance degrada-

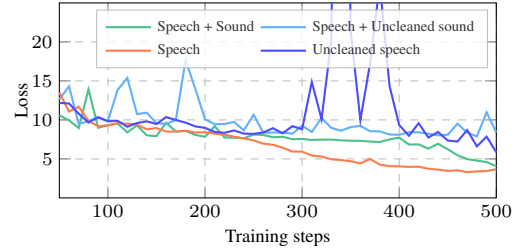


Figure 7: Training curves with cleaned and uncleaned data at the alignment stage.

tion on other test tasks without the aid of auxiliary information. This demonstrates that auxiliary information can compensate for missing features, allowing the LLMs to receive complete information.

We exhibit the inference speed in Tab. 8, using Time To First Token (TTFT) as the metric. Our method shows a speed-up of about 15% and 25% compared to Adapter ( $\times 3$ ) and Adapter ( $\times 4$ ) methods, respectively. This demonstrates that our method uses fewer tokens while achieving greater inference speed-up. We found that the shrinking adapter does not incur significant computational cost, proving it is both lightweight and effective.

## 5.4 Data Quality

The training loss for Stage I, with and without cleaning the speech and sound data, is compared in Fig. 7. When uncleaned speech is used, the training process becomes unstable. Additionally, if the sound data is not properly processed, it significantly worsens the overall training. Given that the alignment stage only trains a few parameters to align the two pretrained large models, the quality

Table 8: Comparison of different shrinking methods on the Librispeech ASR dataset.

Method	Test clean	Test other	TTFT (ms)	Compression ratio
Shrinking adapter	3.1	6.6	72	2.5%
w/o auxiliary info.	3.1	7.1	72	2.5%
Adapter ( $\times 3$ )	3.8	6.5	95	33.3%
Adapter ( $\times 4$ )	4.3	7.8	85	25.0%

of the training data is crucial.

## 5.5 Data Scaling

We compare the performance from 1k to 10k hours of data, and the results are shown in Fig. 8. Our model, using only 1k hours of data, achieves performance comparable to previous speech LLMs. Note that we use only the ASR task as the SFT data, yet our model demonstrates decent capability in instruction following. This demonstrates that the speech representation is well aligned with the text representation. When we further scale up the training data, all tasks show consistent improvements.

## 5.6 Knowledge-Based QA

We present a case of using the speech instruction to ask complex questions in Fig. 9. We find that Soundwave inherits the rich knowledge of LLMs during the conversation. For more examples of performance in physics, chemistry, finance, mathematics, and other fields, refer to App. E.

## 6 Related Work

Speech contains rich non-semantic information compared to text (Wang et al., 2024; Bu et al., 2024; Huang et al., 2024). For LLMs to achieve an accurate understanding of audio, they must have a comprehensive perception of speech rather than relying solely on text (Ji et al., 2024a; Ao et al., 2024). As a result, many researchers have studied how to build end-to-end speech LLMs (Hu et al., 2024; Tang et al., 2023; Chu et al., 2024; Ghosh et al., 2024; Fang et al., 2024; Geng et al., 2025).

Some studies have found the *less is more* phenomenon in LLMs with respect to data usage (Zhou et al., 2024; Song et al., 2025), meaning that efficient use of data can also achieve good performance. However, for speech LLMs, data efficiency has not been fully explored. Therefore, this work addresses this issue by focusing on the key challenge of speech-text alignment.

The acoustic features and text features differ significantly in both their representation space

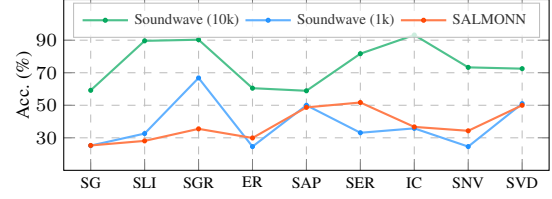


Figure 8: Comparison of scaling effect in AIR-Bench speech foundation tasks.

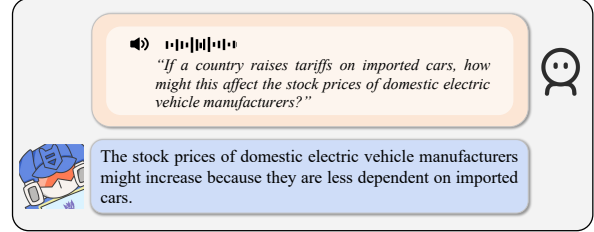


Figure 9: A case of answering the spoken question.

and length. To address this issue, Chu et al. (2024, 2023) employ convolution network to down-sample the speech, while others opt for solutions with more learnable parameters, such as Q-Former (Tang et al., 2023) and linear layers (Hu et al., 2024). Unlike previous work, the proposed Soundwave implements two adapters to address differences in representation and length, which also make training more efficient.

Speech LLMs are primarily designed for two capabilities: Speech and Sound. Tang et al. (2023); Hu et al. (2024) combine Whisper with other feature extractors, such as BEATs (Chen et al., 2023) and WavLM (Chen et al., 2022), to process sound features. Chu et al. (2024) show that a fully fine-tuned encoder can also capture sound information. Our work demonstrates that a frozen encoder can efficiently process both types of features when provided with the proper data and training strategy.

## 7 Conclusion

Speech understanding is a core capability for multi-modal LLMs, yet current speech LLMs often rely on enormous amounts of training data, putting them out of reach for most academic researchers due to the high costs involved. To address this, we developed a more data-efficient solution: a three-stage training strategy paired with a model architecture that incorporates two adapters. This approach effectively tackles the mismatches in representation and length between speech and text. The trained Soundwave delivers top-tier performance on the AIR-Bench speech tasks, while requiring significantly less training data.



## Acknowledgement

This work was supported by the Shenzhen Science and Technology Program (JCYJ20220818103001002), Shenzhen Doctoral Startup Funding (RCBS20221008093330065), Tianyuan Fund for Mathematics of National Natural Science Foundation of China (NSFC) (12326608), Shenzhen Science and Technology Program (Shenzhen Key Laboratory Grant No. ZDSYS20230626091302006), and Shenzhen Stability Science Program 2023, Shenzhen Key Lab of Multi-Modal Cognitive Computing. The authors would like to thank anonymous reviewers for their insightful comments.

## Limitations

Our work still has some limitations, specifically in the following three aspects:

- We have not verified the feasibility of our approach on larger models with more parameters.
- Due to time and manpower limitations, the amount of sound data we have labeled from the scene dataset is still relatively small. As a result, we are unable to conduct in-depth experiments to determine the optimal amount of sound data to include.
- Due to the lack of relevant data, our model does not perform well in music understanding tasks and has limited support for multiple languages.

Next, we will expand the parameter size of our model to verify the feasibility of our approach on larger models. We will also incorporate music understanding and multilingual data to enhance these capabilities. In addition, we will continue annotating the sound data to further validate the optimal data ratio. We also hope that other researchers in the community will conduct related studies.

## Ethical Considerations

**Use of Artifacts** Our study employs Whisper Large V3 as the audio encoder to extract and process speech input data and utilizes Llama-3.1-8B-Instruct as the foundation model for downstream tasks. In using these models, we adhere to academic standards and have cited their original papers

and relevant documentation to ensure proper scholarly attribution. Additionally, Whisper is released under [MIT License](#), while Llama-3.1-8B-Instruct is subject to [Llama 3.1 Community License](#). We have ensured that our application of the model does not violate any of the specified restrictions, thereby maintaining compliance with the license terms.

**Data Collection** All the datasets used in our study are publicly released open-source datasets, and we strictly adhere to the corresponding open-source license agreements to ensure the legality and compliance of the data sources. In addition, the supplementary data annotation work we conducted did not involve any data privacy or sensitive information. Detailed procedures and workflow of the data annotation work can be found in Section [A.1](#). The content related to Statistics For Data can be found in Section [A.3](#).

**Computational Experiment Design and Execution** In Section [4.1](#), we detail the number of parameters of the base model used, the total computational budget, and the computing infrastructure employed. Tab. [13](#) in Section [C](#) lists the hyperparameters used during training and other related configuration details. In Sections [4.2](#) and [5](#), we present the final training results and a comparative analysis of experiments.

**Data Annotation and Ethical Compliance** In the Section [A.1](#), we provide detailed explanations of the manual data annotation work. Section [A.1.1](#) details our data processing methods. We display the complete instruction text given to participants in Fig. [10](#), and explain our volunteer recruitment methods, salaries, and annotator characteristics in Section [A.1.3](#) and [A.1.4](#). Throughout the entire data processing procedure, no ethical risks to personal privacy or data security were posed, and therefore no ethics committee review was required.

**Use of AI Tools** In the course of this research project, AI tools were only utilized in specific aspects, such as assisting with coding and providing grammar checks and language refinement in the writing of the paper, to enhance efficiency and textual quality. Beyond these applications, the core research content, data processing, experimental design, analysis, and conclusions were all independently conducted by the research team, without any other form of artificial intelligence involvement, ensuring the rigor and originality of the study.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Junyi Ao, Yuancheng Wang, Xiaohai Tian, Dekun Chen, Jun Zhang, Lu Lu, Yuxuan Wang, Haizhou Li, and Zhizheng Wu. 2024. Sd-eval: A benchmark dataset for spoken dialogue understanding beyond words. *arXiv preprint arXiv:2406.13340*.
- Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. 2019. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*.
- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, et al. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges. *arXiv preprint arXiv:1907.05019*.
- Emanuele Bastianelli, Andrea Vanzo, Pawel Swietojanski, and Verena Rieser. 2020. [SLURP: A spoken language understanding resource package](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 7252–7262. Association for Computational Linguistics.
- Fan Bu, Yuhao Zhang, Xidong Wang, Benyou Wang, Qun Liu, and Haizhou Li. 2024. Roadmap towards superhuman speech understanding using large language models. *arXiv preprint arXiv:2410.13268*.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Abe Kazemzadeh, Emily Mower, Samuel Kim, Jeanette N Chang, Sungbok Lee, and Shrikanth S Narayanan. 2008. Iemocap: Interactive emotional dyadic motion capture database. *Language resources and evaluation*, 42:335–359.
- Houwei Cao, David G Cooper, Michael K Keutmann, Ruben C Gur, Ani Nenkova, and Ragini Verma. 2014. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE transactions on affective computing*, 5(4):377–390.
- Roldano Cattoni, Mattia Antonino Di Gangi, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2021. [Must-c: A multilingual corpus for end-to-end speech translation](#). *Comput. Speech Lang.*, 66:101155.
- Guiming Hardy Chen, Shunian Chen, Ruifei Zhang, Junying Chen, Xiangbo Wu, Zhiyi Zhang, Zhihong Chen, Jianquan Li, Xiang Wan, and Benyou Wang. 2024. Allava: Harnessing gpt4v-synthesized data for a lite vision-language model. *arXiv preprint arXiv:2402.11684*.
- Guoguo Chen, Shuzhou Chai, Guan-Bo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, Mingjie Jin, Sanjeev Khudanpur, Shinji Watanabe, Shuaijiang Zhao, Wei Zou, Xiangang Li, Xuchen Yao, Yongqing Wang, Zhao You, and Zhiyong Yan. 2021. Gigaspeech: An evolving, multi-domain ASR corpus with 10, 000 hours of transcribed audio. In *Proceeding of Inter-speech, 2021*.
- Qian Chen, Yafeng Chen, Yanni Chen, Mengzhe Chen, Yingda Chen, Chong Deng, Zhihao Du, Ruize Gao, Changfeng Gao, Zhifu Gao, et al. 2025. Minmo: A multimodal large language model for seamless voice interaction. *arXiv preprint arXiv:2501.06282*.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Jian Wu, Michael Zeng, Xiangzhan Yu, and Furu Wei. 2022. [Wavlm: Large-scale self-supervised pre-training for full stack speech processing](#). *IEEE J. Sel. Top. Signal Process.*, 16(6):1505–1518.
- Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, Wanxiang Che, Xiangzhan Yu, and Furu Wei. 2023. [Beats: Audio pre-training with acoustic tokenizers](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 5178–5193. PMLR.
- 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*.
- 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. *arXiv preprint arXiv:2311.07919*.
- Christopher Cieri, David Miller, and Kevin Walker. 2004. The fisher corpus: A resource for the next generations of speech-to-text. In *LREC*, volume 4, pages 69–71.
- Nilaksh Das, Saket Dingliwal, Srikanth Ronanki, Rohit Paturi, Zhaocheng Huang, Prashant Mathur, Jie Yuan, Dhanush Bekal, Xing Niu, Sai Muralidhar Jayanthi, et al. 2024. Speechverse: A large-scale generalizable audio language model. *arXiv preprint arXiv:2405.08295*.
- Alexandre Défossez, Jade Copet, Gabriel Synnaeve, and Yossi Adi. 2022. High fidelity neural audio compression. *arXiv preprint arXiv:2210.13438*.
- Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou, Edouard Grave, and Neil Zeghidour. 2024. Moshi: a speech-text foundation model for real-time dialogue. *arXiv preprint arXiv:2410.00037*.

- Soham Deshmukh, Benjamin Elizalde, Rita Singh, and Huaming Wang. 2023. Pengi: An audio language model for audio tasks. In *Proceeding of NeurIPS 2023*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Venkatesh Duppada and Sushant Hiray. 2017. Ensemble of deep neural networks for acoustic scene classification. *arXiv preprint arXiv:1708.05826*.
- Qingkai Fang, Shoutao Guo, Yan Zhou, Zhengrui Ma, Shaolei Zhang, and Yang Feng. 2024. Llama-omni: Seamless speech interaction with large language models. *arXiv preprint arXiv:2409.06666*.
- Joel Frank and Lea Schönherr. 2021. [Wavefake: A data set to facilitate audio deepfake detection](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.
- Gonçal Garcés Díaz-Munío, Joan Albert Silvestre Cerdà, Javier Jorge-Cano, Adrián Giménez Pastor, Javier Iranzo-Sánchez, Pau Baquero-Arnal, Nahuel Roselló, Alejandro Manuel Pérez-González de Martos, Jorge Civera Saiz, José Alberto Sanchis Navarro, et al. 2021. Europarl-asr: A large corpus of parliamentary debates for streaming asr benchmarking and speech data filtering/verbatimization. *Proc. Interspeech 2021*, pages 3695–3699.
- Xuelong Geng, Kun Wei, Qijie Shao, Shuiyun Liu, Zhennan Lin, Zhixian Zhao, Guojian Li, Wenjie Tian, Peikun Chen, Yangze Li, et al. 2025. Osum: Advancing open speech understanding models with limited resources in academia. *arXiv preprint arXiv:2501.13306*.
- Sreyan Ghosh, Sonal Kumar, Ashish Seth, Chandra Kiran Reddy Evuru, Utkarsh Tyagi, Sakshi Singh, Oriol Nieto, Ramani Duraiswami, and Dinesh Manocha. 2024. GAMA: A large audio-language model with advanced audio understanding and complex reasoning abilities. *arXiv preprint arXiv:2406.11768*.
- Yuan Gong, Jin Yu, and James R. Glass. 2022. [Vocal-sound: A dataset for improving human vocal sounds recognition](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2022, Virtual and Singapore, 23-27 May 2022*, pages 151–155. IEEE.
- Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, pages 369–376.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Esteve. 2018. Ted-lium 3: Twice as much data and corpus repartition for experiments on speaker adaptation. In *Speech and Computer: 20th International Conference, SPECOM 2018, Leipzig, Germany, September 18–22, 2018, Proceedings 20*, pages 198–208. Springer.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM transactions on audio, speech, and language processing*, 29:3451–3460.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [Lora: Low-rank adaptation of large language models](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Shujie Hu, Long Zhou, Shujie Liu, Sanyuan Chen, Hongkun Hao, Jing Pan, Xunying Liu, Jinyu Li, Sunit Sivasankaran, Linquan Liu, et al. 2024. Wavllm: Towards robust and adaptive speech large language model. *arXiv preprint arXiv:2404.00656*.
- Chien-Yu Huang, Ke-Han Lu, Shih-Heng Wang, Chi-Yuan Hsiao, Chun-Yi Kuan, Haibin Wu, Siddhant Arora, Kai-Wei Chang, Jiatong Shi, Yifan Peng, Roshan S. Sharma, Shinji Watanabe, Bhiksha Ramakrishnan, Shady Shehata, and Hung-Yi Lee. 2024. [Dynamic-superb: Towards a dynamic, collaborative, and comprehensive instruction-tuning benchmark for speech](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2024, Seoul, Republic of Korea, April 14-19, 2024*, pages 12136–12140. IEEE.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Javier Iranzo-Sánchez, Joan Albert Silvestre-Cerdà, Javier Jorge, Nahuel Roselló, Adrià Giménez, Albert Sanchis, Jorge Civera, and Alfons Juan. 2020. [Europarl-st: A multilingual corpus for speech translation of parliamentary debates](#). In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, pages 8229–8233. IEEE.
- Il-Young Jeong and Jeongsoo Park. 2022. [Cochlscene: Acquisition of acoustic scene data using crowdsourcing](#). *CoRR*, abs/2211.02289.



- Shengpeng Ji, Yifu Chen, Minghui Fang, Jialong Zuo, Jingyu Lu, Hanting Wang, Ziyue Jiang, Long Zhou, Shujie Liu, Xize Cheng, et al. 2024a. Wavchat: A survey of spoken dialogue models. *arXiv preprint arXiv:2411.13577*.
- Shengpeng Ji, Jialong Zuo, Minghui Fang, Ziyue Jiang, Feiyang Chen, Xinyu Duan, Baoxing Huai, and Zhou Zhao. 2024b. [Textrolspeech: A text style control speech corpus with codec language text-to-speech models](#). In *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2024, Seoul, Republic of Korea, April 14-19, 2024*, pages 10301–10305. IEEE.
- Steven R Livingstone and Frank A Russo. 2018. The ryerson audio-visual database of emotional speech and song (ravdess): A dynamic, multimodal set of facial and vocal expressions in north american english. *PloS one*, 13(5):e0196391.
- Arsha Nagrani, Joon Son Chung, and Andrew Zisserman. 2017. [Voxceleb: A large-scale speaker identification dataset](#). In *18th Annual Conference of the International Speech Communication Association, Interspeech 2017, Stockholm, Sweden, August 20-24, 2017*, pages 2616–2620. ISCA.
- Olympusmons. 2021. [Librispeech asr test clean word timestamp](#). Accessed: 2024-12-23.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An ASR corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2015, South Brisbane, Queensland, Australia, April 19-24, 2015*, pages 5206–5210. IEEE.
- Daniel S Park, William Chan, Yu Zhang, Chung-Cheng Chiu, Barret Zoph, Ekin D Cubuk, and Quoc V Le. 2019. SpecAugment: A simple data augmentation method for automatic speech recognition. *arXiv preprint arXiv:1904.08779*.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. 2019. [MELD: A multimodal multi-party dataset for emotion recognition in conversations](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 527–536. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. [MLS: A large-scale multilingual dataset for speech research](#). In *21st Annual Conference of the International Speech Communication Association, Interspeech 2020, Virtual Event, Shanghai, China, October 25-29, 2020*, pages 2757–2761. ISCA.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *Proceedings of ICML, 2023*.
- Ricardo Reimao and Vassilios Tzerpos. 2019. [For: A dataset for synthetic speech detection](#). In *2019 International Conference on Speech Technology and Human-Computer Dialogue (SpeD)*, pages 1–10.
- S Sakshi, Utkarsh Tyagi, Sonal Kumar, Ashish Seth, Ramaneswaran Selvakumar, Oriol Nieto, Ramani Duraiswami, Sreyan Ghosh, and Dinesh Manocha. 2024. Mmau: A massive multi-task audio understanding and reasoning benchmark. *arXiv preprint arXiv:2410.19168*.
- Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li. 2020. Aishell-3: A multi-speaker mandarin tts corpus and the baselines. *arXiv preprint arXiv:2010.11567*.
- Dingjie Song, Wenjun Wang, Shunian Chen, Xidong Wang, Michael X. Guan, and Benyou Wang. 2025. [Less is more: A simple yet effective token reduction method for efficient multi-modal LLMs](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 7614–7623, Abu Dhabi, UAE. Association for Computational Linguistics.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2023. [Salmonn: Towards generic hearing abilities for large language models](#). *arXiv preprint arXiv:2310.13289*.
- SEAMLESS Communication Team. 2025. Joint speech and text machine translation for up to 100 languages. *Nature*, 637(8046):587–593.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*.
- A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Bin Wang, Xunlong Zou, Geyu Lin, Shuo Sun, Zhuohan Liu, Wenyu Zhang, Zhengyuan Liu, AiTi Aw, and Nancy F Chen. 2024. [Audiobench: A universal benchmark for audio large language models](#). *arXiv preprint arXiv:2406.16020*.
- Changhan Wang, Anne Wu, Jiatao Gu, and Juan Pino. 2021. Covost 2 and massively multilingual speech translation. In *Interspeech*, pages 2247–2251.
- Chen Wang, Minpeng Liao, Zhongqiang Huang, Jinliang Lu, Junhong Wu, Yuchen Liu, Chengqing



Zong, and Jiajun Zhang. 2023. Blsp: Bootstrapping language-speech pre-training via behavior alignment of continuation writing. *arXiv preprint arXiv:2309.00916*.

Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*.

Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, Qingwei Lin, and Daxin Jiang. 2024. Wizardlm: Empowering large pre-trained language models to follow complex instructions. In *Proceeding of ICLR 2024*.

Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald. 2019. [Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit \(version 0.92\)](#).

Qian Yang, Jin Xu, Wenrui Liu, Yunfei Chu, Ziyue Jiang, Xiaohuan Zhou, Yichong Leng, Yuanjun Lv, Zhou Zhao, Chang Zhou, et al. 2024. Airbench: Benchmarking large audio-language models via generative comprehension. *arXiv preprint arXiv:2402.07729*.

Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J. Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu. 2019. [Libritts: A corpus derived from librispeech for text-to-speech](#). In *20th Annual Conference of the International Speech Communication Association, Interspeech 2019, Graz, Austria, September 15-19, 2019*, pages 1526–1530. ISCA.

Jun Zhan, Junqi Dai, Jiasheng Ye, Yunhua Zhou, Dong Zhang, Zhigeng Liu, Xin Zhang, Ruibin Yuan, Ge Zhang, Linyang Li, et al. 2024. Anygpt: Unified multimodal llm with discrete sequence modeling. *arXiv preprint arXiv:2402.12226*.

Yuhao Zhang, Chen Xu, Bei Li, Hao Chen, Tong Xiao, Chunliang Zhang, and Jingbo Zhu. 2023. Rethinking and improving multi-task learning for end-to-end speech translation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 10753–10765.

Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia Efrat, Ping Yu, Lili Yu, et al. 2024. Lima: Less is more for alignment. *Advances in Neural Information Processing Systems*, 36.

## A Data Construction and Preparation

### A.1 Sound Re-annotation

Since a scene segment may contain multiple sounds, we allow users to select multiple sound labels. In the end, we labeled 7,863 audio files, totaling 7.30 hours, with 5,197 files being single-label, totaling 4.81 hours. Specific details can be found in the following sections. We will release all of these labeled data.

#### A.1.1 Data Splitting

We divide the original 10s data into combinations of 3s-3s-4s for annotation.

#### A.1.2 Data Annotation

In accordance with the pre-existing scene labels, we established detailed sound annotations, such as waves and birdsong for a beach scenario. Subsequently, we recruited a number of volunteers to perform data annotation tasks. The interface utilized for this process is illustrated in Fig. 10.

### Annotation

Figure 10: Page for sound annotation.

#### A.1.3 Volunteer Sources and Salaries

We extensively recruited volunteers for this project, comprising 70% undergraduate students, 25% graduate students, and 5% individuals who have already graduated. Each participant was compensated with a one-time payment of 200 RMB, which aligns with the prevailing wage levels in mainland China.

#### A.1.4 Volunteer Authorization

All volunteers have agreed to the public release of their labeled data to promote academic research within the community.

### A.2 Data Process

The process from raw data to final application in this paper includes two steps: Data Selection and Data Filtration.

**Data Selection** We performed data selection on the following dataset:

- **TED-LIUM**, we selected the speaker adaptation part as our dataset, as it is more balanced and representative in characteristics (number of speakers, gender, duration) (Hernandez et al., 2018).
- **GiGaSpeech** It contains five data sizes: XL, L, M, S, and XS. We noticed that the WER limit for XL is relatively loose, so we did not choose this size of data. At the same time, due to the large number of data in L, it would cause GigaSpeech’s data to occupy too high a proportion, while the data in S and XS are too small. Therefore, we ultimately chose M.
- **Common Voice (En)**, Common Voice contains multiple versions. Since we mainly use it to construct the data required for Covost2 (Wang et al., 2021), we selected the matching version, which is Common Voice En 15.0. At the same time, we also use it to construct the data required for Age Prediction, which does not have specific version requirements. For convenience, we consistently use Common Voice En 15.0.
- **Common Voice (Ja)**, Common Voice contains multiple versions. Since we are using this data to build a Language Identification task, and in order to balance with data from other languages, we need to select around 15 hours of data. Smaller versions do not provide enough data, and larger versions would result in unnecessary overhead, so we ultimately chose Common Voice ja 7.0.

**Data Filtration** To ensure data quality, we performed filtering on the dataset:

- **Duration**, it cannot exceed 30 seconds for the Whisper encoder, and it cannot be less than 3 seconds for more balanced training.

- **Lenth** We excluded transcriptions that were longer than 200 to ensure the stability of the training process. We limit the frame length to no more than 100 times the text length, as exceeding this indicate the speech contains excessive noise.

- **WER** To ensure stable training, we only retained the data with a WER of less than 10% in the Whisper recognition results.

**Instruction Data Generation** During the SFT data preparation phase, we used the GPT-4o-mini model to automatically generate question-answer pairs for all the training tasks. For example about the Fisher conversation ASR data, the process involved extracting dialogue segments from the ASR dataset and inputting them into the GPT-4o-mini model, which was instructed to generate contextually relevant questions framed from an objective third-person perspective. For each question, the model also generated corresponding answers directly derived from the conversation. These question-answer pairs were then incorporated into the training dataset for the chat task.

### A.3 Data Statistics

We present the usage of audio data and the total amount for each dataset in Table 3. It is important to note that in this statistical process, the same data is counted only once across different stages, only once across different tasks, and only once even if constructed using different methods within the same task. If you are interested in the specific data usage for each stage and task, please refer to the subsequent section.

#### A.3.1 Alignment Stage Data

In Tab. 9, we present the datasets used during the alignment stage, along with their respective quantities and durations.

Table 9: Summary of datasets, their size, and duration used in the alignment stage.

Dataset	Size	Duration (h)
Multilingual LibriSpeech	984,559	4081.61
GigaSpeech	713,394	805.11
TED-LIUM	143,641	244.02
TUT w./ libritts	4,698	18.48
Vocalsound w./ libritts	19,737	82.42
Europarl-ASR	718,663	418.42
LibriSpeech	281,241	961.05
TextrolSpeech	170,502	246.85
<b>SUM</b>	<b>3,036,435</b>	<b>6,857.97</b>

As described in Section 3, the sound and scene data used in this stage are both constructed with LibriTTS (Zen et al., 2019) inserted in between. "TUT w./ LibriTTS" refers to the dataset where TUT is sandwiched between LibriTTS, while "TUT w./ LibriTTS" refers to the dataset where VocalSound is sandwiched between LibriTTS.

### A.3.2 Shrinking Stage Data

In Tab. 9, we present the datasets used during the shrinking stage, along with their respective quantities and durations.

Table 10: Summary of datasets, their size, and duration in the shrinking stage.

Dataset	Size	Duration (h)
ASR (First Stage)	3,012,000	6757.06
Covost2 (en-de)	232,953	364.89
MUST-C (en-de)	226,810	395.10
Fisher	153,102	1091.42
TUT w./ LibriTTS	4,937	19.33
VocalSound w./ LibriTTS	19,737	82.42
<b>SUM</b>	<b>3,649,539</b>	<b>8,710.23</b>

Here, we have two points that need clarification: 1) Why is the ASR data reduced compared to the first stage? 2) Why has the "TUT w./ LibriTTS" data increased compared to the first stage?

Here, we will explain why the ASR data in this stage is reduced. In this stage, the LoRA parameters of the large model need to be trained. To ensure stable training, it is important to avoid overlap between tasks executed by different instructions. As shown in Figure 1, we have a dedicated task for Sound, which primarily involves predicting special Sound tokens. To prevent overlap, we removed the data containing Sound special tokens from the ASR dataset.

Here, we explain the reason for the increased amount of "TUT w./ LibriTTS" data in this stage. In the previous training phase, we removed data that was deemed to have a high level of noise during manual annotation in order to ensure the stability of the training process. However, in this phase, we included these data to enhance the model’s robustness.

### A.3.3 SFT Data

Our SFT data statistic details are shown in Tab. 11.

## A.4 Mitigating Data Leakage Risks

In this section, we will discuss the risks of data leakage in several parts. For one set of tasks, we

used non-homogeneous training data, while for another set, although we employed homogeneous data, we rigorously considered the issue of data leakage.

### A.4.1 Non-homogeneous Training Data

**Speech Gender Recognition** The task involved in the test set is AIR-Bench (Yang et al., 2024), which uses Common Voice (Ardila et al., 2019) and MELD (Poria et al., 2019) to construct the data. We use TextrolSpeech (Ji et al., 2024b) to construct the data, which is considered non-homogeneous data in comparison.

**Spoken Language Identification** This task involves a total of 7 languages: Chinese, English, Italian, German, French, Spanish, and Japanese. AIR-Bench (Yang et al., 2024) uses Covost2 (Wang et al., 2021) in its construction, which is sourced from Common Voice (Ardila et al., 2019). The construction of data in English, Italian, German, French, and Spanish, we used Europarl-ASR (Garcés Díaz-Munío et al., 2021), while for Chinese data, we used AISHELL3 (Shi et al., 2020). These sources are different from Common Voice, so there is no data leakage. For Japanese data construction, we used Common Voice, which is the same source as AIR-Bench, so we paid special attention to potential leakage issues. We noticed that there were only two Japanese samples in AIR-Bench, so we manually removed these two entries.

### A.4.2 Homogeneous Training Data

**Speech Grounding** Since our training set is constructed using the same dataset as in the AIR-Bench test, we made sure that the test set was not included in the training. We removed data where the same word in the same position was queried in the audio. Specifically, due to the difficulty of ensuring that randomly selected data doesn’t overlap during the selection process, we adopted a post-processing approach where we deleted training data with the same filename and identical queries.

**Emotion Recognition** During the construction of our training dataset, we utilized the TextrolSpeech (Ji et al., 2024b), RAVDESS (Livingstone and Russo, 2018), CREMA-D (Cao et al., 2014), IEMOCAP (Busso et al., 2008), and MELD (Poria et al., 2019) datasets. Notably, TextrolSpeech is composed of multiple datasets, including ESD, MEAD, MESS, SAVEE, and TESS. Given that AIR-Bench incorporates data from IEMOCAP and MELD, we have entirely excluded these datasets

Table 11: Summary of datasets, their size, and duration used in the SFT stage.

Task	Dataset	Size	Duration (hours)
ASR	LibriSpeech	281241	961.05
Translation EN-DE	Covost2, MuST-C	455648	753.18
Speech Grounding	LibriSpeech	23828	51.36
Spoken Language Identification	Common Voice, Europarl-ST	269485	293.15
Speaker Gender Recognition	TextrolSpeech	319986	466.09
Emotion Recognition	TextrolSpeech, RAVDESS, CREMA-D, IEMOCAP	258541	347.18
Speaker Age Prediction	Common Voice	77239	120.04
Speech Entity Recognition	SLURP	114792	91.62
Intent Classification	SLURP	212473	151.52
Speaker Num. Verification	VoxCeleb	156033	435.17
Synthesized Voice Detection	FoR	53696	47.55
Vocal Sound Classification	VocalSound	29962	34.77
Acoustic Scene Classification	TUT2017, CofCScene	159590	443.31
Chat	Fisher	132209	956.25
Speech Instruction	AnyInstruct	106770	206.30
<b>SUM</b>	<b>18</b>	<b>2651493</b>	<b>5358.54</b>

from our training data. Additionally, the test set of MELD has also been removed to ensure data integrity and prevent potential data leakage.

**Speech Entity Recognition and Intent Classification** We used the same source data as AIR-Bench for construction, both utilizing SLURP (Bastianelli et al., 2020), so we paid special attention to data leakage issues. Since AIR-Bench retained the original file names for all its files, we directly removed this portion of the data from our dataset.

**Speaker Number Verification** We used the same source data as AIR-Bench for construction, both utilizing VoxCeleb (Nagrani et al., 2017). We used a fully random selection method, choosing a series of speech pairs to form our training set. Since the random selection process is hard to control, we removed any data that had already appeared in AIR-Bench after the selection.

**Synthesized Voice Detection** We used FoR (Reimao and Tzerpos, 2019) to construct our training set. We noticed that FoR overlaps with AIR-Bench. However, AIR-Bench does not provide detailed records of the specific sources of these data, so we removed the repeated test and development sets from FoR.

### A.5 Dataset Licence

The paper and license for the dataset we used are listed in Table 15. There are several points regarding the usage of data that need to be clarified.

**Fisher** Our access to the Fisher dataset is authorized by the National University of Singapore. We will strictly adhere to the usage agreement and will not, in any way, release, broadcast, disclose, display, copy, reproduce, or redistribute this data to

anyone outside of our team.

## B Performance on MMAU benchmark

The following results on MMAU benchmark (Sakshi et al., 2024) in Tab 12. We can find that Soundwave still outperforms Qwen2-Audio in Speech tasks and shows decent performance on Sound tasks.

Table 12: Results on MMAU speech and Sound tasks.

Model	Speech Tasks	Sound Tasks	Avg.
Soundwave	<b>46.25</b>	46.25	46.25
Qwen2-Audio	42.04	54.95	<b>48.50</b>
Qwen-Audio	30.03	<b>55.25</b>	42.64
SALAMONN	25.50	41.00	33.25
GAMA	18.91	41.44	30.18

## C Training Configurations

The training settings of different stages are shown in Tab. 13. For all training and decoding processes, we set ‘*You are a helpful language and speech assistant. You are able to understand the speech content that the user provides and assist the user with a variety of tasks using natural language.*’ as the system prompt.

Table 13: Overview of training parameters at different stages.

Settings	Stage 1	Stage 2	Stage 3
Batch	32	16	8
Learning rate	1e-4	3e-5	3e-5
Accumulation steps	8	8	4
Training param.	144M	266M	122M



Table 14: Comparison of training parameters and data for different models.

Model	#Params	#Data (Hours)	AIR-Bench Speech Avg.(%)
Qwen2-Audio	~8B	~520K	59.3
Qwen-Audio Turbo	~8B	~147K	63.4
SALMONN	~14B	~9K	37.8
BLSP	~7B	~13K	36.6
NExT-GPT	~12B	-	33.6
PandaGPT	~14B	-	39.0
Soundwave	~9B	~10K	75.5

All in all, we trained a total of 266M parameters in our three-stage process. To better highlight the advantages of our model’s parameters compared to others, we plotted the relationship between AIR-Bench performance and training parameters, which is shown in Fig. 11. We also compare the size of model parameters and training data in Tab. 14.

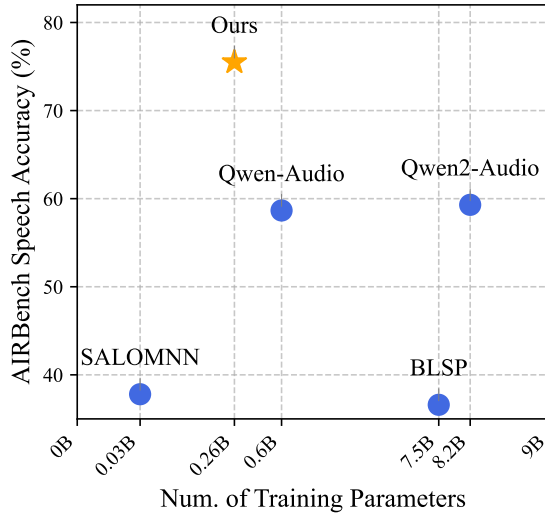


Figure 11: AIR-Bench speech accuracy with number of training parameters.

## D AIR-Bench Sound and Music Foundation Tasks

The AIR-Bench sound and music tasks evaluate models on various auditory capabilities. Sound tasks focus on identifying, classifying, and reasoning with environmental sounds, while music tasks involve classifying musical elements, analyzing pitch and velocity, and understanding emotional content. As Tab. 16 shown, Soundwave demonstrates exceptional performance in vocal sound and acoustic scene classification, achieving impressive accuracy. Though few sound data is used, sound average score of Soundwave still ranks second. This

highlights its strong performance across sound-related tasks, even with limited data.

## E Speech Instruction

We demonstrate our model’s ability to follow voice commands from two aspects: generation tasks and knowledge question-answering tasks.

### E.1 Generation Tasks

We selected some commonly used generation tasks in daily life, which demonstrate our model’s ability to assist in handling everyday affairs. Our presentation results are shown in Figures 12, 13, and 14.

### E.2 Knowledge-based QA Tasks

We will categorize the knowledge to showcase our model’s QA ability across various domains of knowledge. Specifically, for math, please refer to Fig. 15. For physics, please refer to Fig. 16. For biology, please refer to Fig. 17. For history, please refer to Fig. 18.

Table 15: The summary for the dataset


DataSet Name	Citation	License
Librispeech w./ timestamp	Olympusmons (2021)	Apache 2.0
Librispeech	Panayotov et al. (2015)	CC BY 4.0
CREMA-D	Cao et al. (2014)	DbCL-1.0
TED-LIUM	Hernandez et al. (2018)	CC BY-NC-ND 3.0
MLS	Pratap et al. (2020)	CC BY 4.0
Europarl-ASR	Garcés Díaz-Munío et al. (2021)	CC BY 4.0
TextrolSpeech	Ji et al. (2024b)	MIT License
LibriTTS	Zen et al. (2019)	CC BY 4.0
VCTK	Yamagishi et al. (2019)	ODC-By 1.0
TUT2017	Duppada and Hiray (2017)	Non-Commercial
VocalSound	Gong et al. (2022)	CC BY-SA 4.0
MUST-C	Cattoni et al. (2021)	CC BY-NC-ND 4.0
Europarl-ST	Iranzo-Sánchez et al. (2020)	CC BY-NC 4.0
Common Voice	Ardila et al. (2019)	CC0 1.0 Universal
CochlScene	Jeong and Park (2022)	CC BY-SA
SLURP	Bastianelli et al. (2020)	CC BY-NC 4.0
RAVDESS	Livingstone and Russo (2018)	CC BY-NC-SA 4.0
IEMOCAP	Busso et al. (2008)	IEMOCAP License
MELD	Poria et al. (2019)	GPL-3.0
Gigaspeech	Chen et al. (2021)	Apache 2.0
Covost2	Wang et al. (2021)	CC0 1.0 Universal
VoxCeleb	Nagrani et al. (2017)	CC BY 4.0
WaveFake	Frank and Schönherr (2021)	MIT License
For	Reimao and Tzerpos (2019)	LGPL-3.0
AnyInstruct	Zhan et al. (2024)	
Fisher	Cieri et al. (2004)	LDC License

Table 16: Performance on the AIR-Bench sound and music foundation tasks.



Task	Soundwave	Qwen2-Audio	Qwen-Audio Turbo	SALMONN	BLSP	NExT-GPT	PandaGPT
Audio Grounding	23.1	34.9	41.6	24.0	34.6	<b>62.2</b>	38.3
Vocal Sound Classification	<b>91.7</b>	89.3	78.1	45.3	29.8	23.5	31.6
Acoustic Scene Classification	<b>83.8</b>	67.4	61.3	34.1	25.2	24.1	55.7
Sound Question Answering	49.7	<b>68.8</b>	62.8	28.4	36.1	18.8	48.7
Sound avg.	62.1	<b>65.1</b>	61.0	33.0	31.4	32.2	43.6
Music Instruments Classification	37.1	<b>65.8</b>	59.6	41.3	22.8	24.3	47.7
Music Genre Classification	49.5	<b>78.8</b>	77.1	45.3	26.1	28.1	39.8
Music Note Analysis-Pitch	27.7	28.7	<b>30.1</b>	26.4	23.5	25.1	26.4
Music Note Analysis-Velocity	23.2	26.2	25.1	22.8	24.9	23.1	<b>27.2</b>
Music Question Answering	65.0	<b>65.7</b>	62.5	54.6	31.0	47.1	50.7
Music Emotion Detection	38.3	<b>46.9</b>	39.0	32.2	28.3	25.4	36.7
Music avg.	40.1	<b>52.0</b>	48.9	37.1	26.1	28.9	38.1

Table 17: Performance on the AIR-Bench chat tasks.

Task	Soundwave	Qwen2-Audio	Qwen-Audio Turbo	SALMONN	BLSP	NExT-GPT	PandaGPT	Whisper +GPT-4
<b>Speech</b>	6.41	7.18	7.04	6.16	6.17	3.86	3.58	7.54
<b>Sound</b>	5.33	6.99	6.59	6.28	5.55	4.76	5.46	/
<b>Music</b>	5.10	6.79	5.98	5.95	5.08	4.18	5.06	/
<b>Mixed Audio</b>	4.98	6.77	5.77	6.08	4.52	2.92	2.93	/
<b>Average</b>	5.46	6.93	6.34	6.11	5.33	4.13	4.25	/




*"For a new wireless earbud product, create three brief advertisement slogans. Each slogan must mention exactly two features from 'noise cancellation', 'battery life', and 'comfortable fit'. And each slogan must begin with a gerund and end with a word. Each slogan should be fewer than 20 words."*






"Enjoying music without interruptions" - Wireless earbuds with active noise cancellation and long battery life.  
 "Relaxing with a comfortable fit" - Wireless earbuds with long battery life and comfortable fit.  
 "Listening to music without distractions-free" - Wireless earbuds with noise cancellation and long battery life.

Figure 12: Generate task show case 1.




*"Hey, please come up with three daily workout suggestions. Each suggestion should be under 25 words and include a specific time (for example, '7 AM'), an exercise (for example, 'jogging'), and a workout method (for example, 'outdoor')."*






Sure, here are three daily workout suggestions:  
 1. Morning jog: Start your day with a 30-minute jog at 7 am to boost your energy and set a positive tone for the day.  
 2. Evening yoga: End your day with a 45-minute yoga session at 8 pm to relax and unwind.  
 3. Lunchtime walk: Take a 30-minute walk during your lunch break at 12 pm to 1 pm to stay active and refreshed.

Figure 13: Generate task show case 2.





*"I got up at 6 AM to exercise, then had breakfast at 7 AM. At 9 AM, I went to a meeting at work, had lunch with my team at 12 PM, returned to the office at 2 PM to check emails, and attended a training session at 4 PM. I finished work at 6 PM, had dinner with my family at 7 PM, and relaxed by reading at 9 PM. Please generate a schedule for me."*






Sure, I can help you with that. Here's a suggested schedule for you:  
 6:00 AM - Wake up  
 7:00 AM - Breakfast  
 9:00 AM - Meeting at work  
 12:00 PM - Lunch with team  
 2:00 PM - Check emails  
 4:00 PM - Training session  
 6:00 PM - Finish work  
 7:00 PM - Dinner with family  
 9:00 PM - Relaxation time (reading, etc.)  
 Please let me know if this schedule meets your needs or if you need any adjustments.

Figure 14: Generate task show case 3.






*"What is the square root of 144?"*







The square root of 144 is 12.



*"Two girls each got 1/6 of the 24 litres of water. Then a boy got 6 litres of water. How many litres of water were left?"*







There were 24 litres of water initially. Two girls each got  $\frac{1}{6}$  of the water, so they got  $2 \times 4 = 8$  litres. Then a boy got 6 litres. So, there were  $24 - 8 - 6 = 10$  litres of water left.

Figure 15: Knowledge-based QA about math.






*"Please explain the quantum tunneling effect and describe how this effect is applied in semiconductor devices."*







Quantum tunneling is a phenomenon where a particle can pass through a barrier even if it doesn't have enough energy to overcome the barrier. This effect is applied in semiconductor devices to create a tunneling current.

*"Quantum tunneling is a phenomenon where a particle can pass through a barrier even if it doesn't have enough energy to overcome the barrier. This effect is applied in semiconductor devices to create a tunneling current."*








The buoyant forces experienced by the glass balls and plastic balls of the same volume are equal.


Figure 16: Knowledge-based QA about physics.








“What is required for physical mapping analysis of DNA fragments? Please listen to the options. A. Exonuclease. B. DNaseI. C. Restriction Endonuclease. D. DNA polymerase.”







The correct answer is C. Restriction Endonuclease. Restriction endonucleases are enzymes that cut DNA at specific sequences, which is necessary for physical mapping analysis of DNA fragments.



“In enzyme separation and purification, which experimental result is the most ideal? Please listen to the options. A. High purification fold and high recovery rate. B. High protein recovery rate. C. Low recovery rate, but high purification fold. D. Maximum specific activity.”







The most ideal experimental result in enzyme separation and purification is option A. High purification fold and high recovery rate. This means that the enzyme is highly purified and also recovered in high amounts.

Figure 17: Knowledge-based QA about biology.






“In the process of establishing the constitutional monarchy in Britain, the Act of Settlement, enacted in 1701, played an important role. The law stipulated that any decision made by the king must be signed by a minister who agrees with the decision in order to take effect. The purpose of this provision was to?”







The purpose of this provision is to ensure that the king's decisions are not arbitrary and that they are subject to scrutiny and approval by a responsible official.

“In the early stages of the French Revolution, to overthrow the long-standing feudal system and aristocratic privileges, what was the purpose of issuing the Declaration of the Rights of Man and of the Citizen, which stipulated that all citizens enjoy equal fundamental rights?”





The purpose of issuing the Declaration of the Rights of Man and of the Citizen was to overthrow the long-standing feudal system and aristocratic privileges, and to stipulate that all citizens enjoy equal fundamental rights.

Figure 18: Knowledge-based QA about history.