

---

# FunAudioLLM: Voice Understanding and Generation Foundation Models for Natural Interaction Between Humans and LLMs

---

Tongyi SpeechTeam

Alibaba Group

FunAudioLLM@list.alibaba-inc.com

## Abstract

This report introduces FunAudioLLM, a model family designed to enhance natural voice interactions between humans and large language models (LLMs). At its core are two innovative models: SenseVoice, which handles multilingual speech recognition, emotion recognition, and audio event detection; and CosyVoice, which facilitates natural speech generation with control over multiple languages, timbre, speaking style, and speaker identity. SenseVoice-Small delivers exceptionally low-latency ASR for 5 languages, and SenseVoice-Large supports high-precision ASR for over 50 languages, while CosyVoice excels in multi-lingual voice generation, zero-shot in-context learning, cross-lingual voice cloning, and instruction-following capabilities. The models related to SenseVoice and CosyVoice have been open-sourced on Modelscope and Huggingface, along with the corresponding training, inference, and fine-tuning codes released on GitHub. By integrating these models with LLMs, FunAudioLLM enables applications such as speech-to-speech translation, emotional voice chat, interactive podcasts, and expressive audiobook narration, thereby pushing the boundaries of voice interaction technology. Demos are available at <https://fun-audio-llm.github.io>, and the code can be accessed at <https://github.com/FunAudioLLM>.

## 1 Introduction

In recent years, the advancement in artificial intelligence (AI) has dramatically transformed how humans interact with machines, such as GPT-4o (OpenAI, 2023) and Gemini-1.5 (Reid et al., 2024) and so on (Bai et al., 2023b; Chu et al., 2023). This transformation is particularly evident in the realm of voice processing, where capabilities such as high-precision speech recognition (Radford et al., 2023), emotion recognition (Ma et al., 2024b), and voice generation (Wang et al., 2023a; Du et al., 2024a) are paving the way for more intuitive and human-like interactions. In this report, we introduce FunAudioLLM, an innovative framework designed to facilitate natural voice interactions between humans and large language models (LLMs) (Team, 2023; Bai et al., 2023a; Touvron et al., 2023). At the core of FunAudioLLM are our two groundbreaking models: SenseVoice, for voice understanding, and CosyVoice, for voice generation.

SenseVoice is our state-of-the-art voice understanding model, which excels in multiple domains of voice processing. We offer both SenseVoice-Small and SenseVoice-Large variants. We have open-sourced SenseVoice-Small, which supports multilingual recognition in Chinese, English, Cantonese, Japanese, and Korean, delivering extremely low inference latency by employing a non-autoregressive end-to-end architecture. This design choice results in a performance that is more than 5 times faster than Whisper-small and more than 15 times faster than Whisper-large (Radford et al., 2023). On the other hand, SenseVoice-Large supports speech recognition in over 50 lan-




Figure 1: An overview of our FunAudioLLM models for voice understanding and generation.

guages, with significant advantages in recognizing Chinese and Cantonese. In addition to speech recognition, SenseVoice offers state-of-the-art capabilities in emotion recognition and audio event detection (Mesaros et al., 2021), making it an ideal choice for creating low-latency, human-like voice interaction systems.

Our suite of applications is further enriched by CosyVoice (Du et al., 2024a), a family of fundamental speech generation models designed to produce natural-sounding voices for a variety of contexts. CosyVoice excels in generating multi-lingual voices tailored to specific speakers, zero-shot adaptation to new speakers (Wang et al., 2023a), cross-lingual voice cloning (Zhang et al., 2023), creating emotionally resonant voices (Shin et al., 2022), and offering nuanced control over speech output through instructional text (Ji et al., 2024). CosyVoice supports five languages: Chinese, English, Japanese, Cantonese, and Korean. CosyVoice comes in three open-source models: CosyVoice-base-300M, which specializes in accurately representing speaker identity, zero-shot learning, and cross-lingual voice cloning; CosyVoice-instruct-300M, which focuses on generating emotionally expressive voices and allows for meticulous adjustments via instructional text, extending its capabilities to controllability over various aspects such as speaker identity (Shimizu et al., 2023), speaking style (Ji et al., 2024), and fine-grained paralinguistic features (Kanda et al., 2024); and CosyVoice-sft-300M, which has been fine-tuned on seven multilingual speakers and is ready for immediate deployment.

By integrating SenseVoice, CosyVoice, and LLMs like Qwen (Team, 2023), FunAudioLLM offers a range of rich application demos. These include Speech-to-Speech Translation (Berard et al., 2018), which allows users to speak in foreign languages using their own voice; Emotional Voice Chat (Xue et al., 2024), which enables the model to understand and respond to emotions for more human-like interactions; Interactive Podcast (Laban et al., 2022), wherein users can engage in live discussions with multiple large models; and AudioBook (Chalamandaris et al., 2014), allowing the model to perform expressive, multi-character narration for audiobooks.

Overall, FunAudioLLM leverages the strengths of SenseVoice and CosyVoice to push the boundaries of voice interaction technology, enabling more natural and seamless communication between humans and large language models.

## 2 FunAudioLLM Models

### 2.1 Overview of FunAudioLLM

FunAudioLLM consists of two foundation models for voice understanding and generation, named SenseVoice and CosyVoice, respectively. SenseVoice supports multi-lingual speech recognition, which is trained on over 300k hours. Specifically, SenseVoice-Small is efficient in inference, in which the recognition latency is less than 80ms and is more than 5 and 15 times faster than Whisper-Small and Whisper-large, respectively, and SenseVoice-Large supports high-precision ASR for over 50 languages. Furthermore, SenseVoice supports rich transcription, including state-of-the-art emotion recognition, audio event detection, inverse text normalization (Pusateri et al., 2017) and punctuation (Chen et al., 2020).

Our voice generation model, CosyVoice, can generate multi-lingual speeches, which is trained on over 170k hours and five languages, including Chinese (ZH), English (EN), Japanese (JP), Cantonese (Yue) and Korean (KO). CosyVoice generated samples can achieve a WER of less 2% and speaker

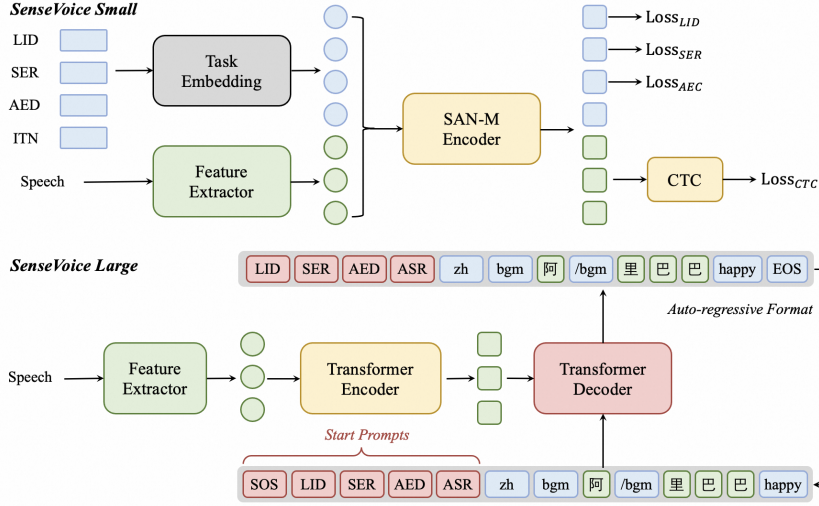


Figure 2: SenseVoice is a comprehensive speech foundation model designed to perform various speech understanding tasks, including Automatic Speech Recognition (ASR), Language Identification (LID), Speech Emotion Recognition (SER), and Audio Event Detection (AED). SenseVoice-Small [Top]: An encoder-only model optimized for rapid speech understanding. It offers high-speed processing while supporting 5 languages. SenseVoice-Large [Bottom]: An encoder-decoder model aimed at achieving more precise speech understanding across a broader range of languages. It excels in accuracy and supports an extensive set of language capabilities.

similarity of over 75%, which achieves the quality level of human parity. CosyVoice supports zero-shot in-context learning, which enables voice cloning with a prompt speech of even 3 seconds. The timbre, emotion, prosody and style can be reproduced within or cross languages. We also released an instruction model, which can control speaker identity, speaking style (e.g., emotion) and other fine-grained paralinguistic features with natural textual instructions. An overview of FunAudioLLM models is shown in Figure 1.

## 2.2 Voice Understanding Model: SenseVoice

SenseVoice is a speech foundation model with multiple voice understanding capabilities, including automatic speech recognition (ASR), spoken language identification (LID), speech emotion recognition (SER), and audio event classification (AEC) or audio event detection (AED). Two models with different sizes and architectures are proposed to suit different requirements: SenseVoice-Small, an encoder-only speech foundation model for rapid speech understanding, and SenseVoice-Large, an encoder-decoder (Vaswani et al., 2017) speech foundation model for more accurate speech understanding with more languages supported, as illustrated in Figure 2.

**SenseVoice-Small** is a non-autoregressive encoder-only model for multi-lingual multi-style ASR and multiple speech understanding tasks. Given the input waveform, we first compute the 80-dimensional log-mel filter-bank, and then stack consecutive frames and down-sample them by a factor of 6. The extracted feature is then mapped to the dimension  $D$  of the encoder, denoted as  $\mathbf{X}_{\text{speech}} \in \mathbb{R}^{T \times D}$ , where  $T$  is the length of the down-sampled feature. The encoder is implemented as a memory-equipped self-attention network (SAN-M) (Gao et al., 2020). To specify the task, we prepend four embeddings to the speech feature as the input to the encoder:

$$\mathbf{X} = \text{concat}(\mathbf{e}_{\text{LID}}, \mathbf{e}_{\text{SER}}, \mathbf{e}_{\text{AEC}}, \mathbf{e}_{\text{ITN/NoITN}}, \mathbf{X}_{\text{speech}}) \quad (1)$$

$$\mathbf{P} = \text{Softmax}(\text{Linear}_{D \rightarrow |V'|}(\text{Encoder}(\mathbf{X}))) \quad (2)$$

$\mathbf{X} \in \mathbb{R}^{(T+4) \times D}$  and  $\mathbf{P} \in \mathbb{R}^{(T+4) \times |V'|}$ .  $V'$  is the vocabulary including tokens for ASR and other tasks.  $\mathbf{e}_{\text{LID}}, \mathbf{e}_{\text{SER}}, \mathbf{e}_{\text{AEC}}, \mathbf{e}_{\text{ITN/NoITN}}$  are embeddings of four special tokens:

$\langle \text{LID} \rangle$  indicates the LID task. If  $\langle \text{LID} \rangle$  is prepended, the model is trained to predict the language token, at the corresponding position of the output. In the training stage, we randomly replace  $\langle \text{LID} \rangle$  with the ground truth language token according to probability 0.8 so that the model can either predict the language token, or be configured with a specified language token in the inference stage.

$\langle \text{SER} \rangle$  indicates the SER task. If  $\langle \text{SER} \rangle$  is prepended, the model is trained to predict the speech emotion label, at the corresponding position of the output.

$\langle \text{AEC} \rangle$  indicates the AEC task. If  $\langle \text{AEC} \rangle$  is prepended, the model is trained to predict the audio event label, at the corresponding position of the output.

$\langle \text{ITN} \rangle$  or  $\langle \text{NoITN} \rangle$  specify the transcription style. If  $\langle \text{ITN} \rangle$  is provided, the model is trained to transcribe with inverse text normalization (ITN) and punctuation. If  $\langle \text{NoITN} \rangle$  is provided, the model is trained to transcribe without ITN and punctuation.

In the training stage, the LID, SER, and AEC tasks are optimized using the cross-entropy loss. The ASR task is optimized using the CTC loss (Graves et al., 2006).

**SenseVoice-Large** is an autoregressive encoder-decoder model for multi-lingual ASR and multiple speech understanding tasks. Similar to Whisper (Radford et al., 2023), SenseVoice-Large specifies tasks by a sequence of input tokens to the decoder. Specifically, we specify whether to predict language, speech emotion, and audio events with timestamps by including  $\langle \text{LID} \rangle$ ,  $\langle \text{SER} \rangle$ ,  $\langle \text{AED} \rangle$  tokens respectively. Compared to SenseVoice-Small, the advantage of SenseVoice-Large is the transcription accuracy and supporting for a vast number of languages (50+).

Table 1 gives examples of transcriptions of Whisper, SenseVoice-S, SenseVoice-L, and the ground truth of the ASR task.

<b>Whisper</b>
Absolute shock, but in a great way. Wow. That was awesome. That was awesome. What way to open a song. That was awesome. Awesome. ...
<b>SenseVoice-S</b>
$\langle \text{music} \rangle$ Absolute shocked but in a great way my. $\langle \text{happy} \rangle$ That was awesome, that was awesome what way to open a song that was awesome, awesome, ...
<b>SenseVoice-L</b>
$\langle \text{music} \rangle$ Absolutely shocked but in a great way. That was awesome, $\langle \text{music} \rangle$ that was awesome $\langle \text{happy} \rangle$ what way to open a song, that was awesome, awesome, ...
<b>Ground Truth</b>
Absolutely shocked, but in a great way. Who am I? Wow. That was awesome. That was awesome. What way to open a song. That was awesome. Awesome. ...

Table 1: Examples of transcriptions of Whisper, SenseVoice-S, SenseVoice-L, and the ground truth.

### 2.3 Semantic Speech Tokenizer

A speech tokenizer transforms vocal signals into discrete tokens, enabling their modeling and prediction by autoregressive transformers for speech generation. Our preliminary experiments indicated that the choice of speech tokenizer is pivotal for overall system performance as well as the requirements of both data quality and volume. We evaluated three classes of speech tokenizers: 1) those based on residual quantization like SoundStream (Zeghidour et al., 2022), Encodec (Défossez et al., 2022) and FunCodec (Du et al., 2024b); 2) those utilizing multi-grouped quantization, such as Hifi-Codec (Yang et al., 2023); and 3) “semantic” speech tokens, specifically HuBERT (Hsu et al., 2021). All the above tokenizers are trained in the unsupervised or self-supervised manners. Thus, their association to semantic content is often tenuous, contributing to an unstable synthesis process and a substantial demand for clean training data. Moreover, unsupervised tokenizers are susceptible to data noise, necessitating meticulously curated clean data sets.

Building on the success of SenseVoice models, we introduce a supervised semantic speech tokenizer, denoted as  $\mathcal{S}^3$  (Du et al., 2024a). Using the pre-trained SenseVoice-Large model as a foundation, we incorporate a vector quantizer subsequent to the encoder’s initial six layers, delineated in Figure

Projects	Languages	Zero-shot	Style&Speaker Control	Fine-grained	SFT	Server
Bark	13	✗	✗	✓	✗	✗
ChatTTS	en, zh	✗	✗	✓	✓	WebUI
parler-tts	en	✗	✓	✗	✓	✗
EmotiVoice	en, zh	✗	✓	✗	✓	WebUI
GPT-SoVITS	en, zh, jp	✓	✗	✗	✓	WebUI
OpenVoice	en,sp,fr, zh,jp,kr	✓	✓	✓	✗	✗
CosyVoice	en, zh, jp, yue, kr	✓	✓	✓	✓	WebUI, gRPC

Table 2: Comparison on released features between CosyVoice and other open-sourced projects.

3. Importantly, the integration of an additional positional embedding post-quantization enhances temporal information. The combination of **Encoder**<sub>1</sub> and vector quantizer is considered as the speech tokenizer, employing the index of the closest code vector as speech tokens. The vector quantizer utilizes a solitary codebook with an expansive dictionary containing 4,096 entries. The derived token sequence exhibits a frequency of 50 Hz, thereby reducing the computational load on text-to-token generation within language models.

Since the speech tokenizer is trained to minimize the recognition errors of rich text in an end-to-end manner, the extracted tokens have a strong semantic relationship to textual and paralinguistic information. Furthermore, our  $\mathcal{S}^3$  tokenizer benefits from supervised training, enhancing its robustness to data noise and reducing the reliance on pristine data collection. Consequently, a broader spectrum of data can be utilized for training the model.

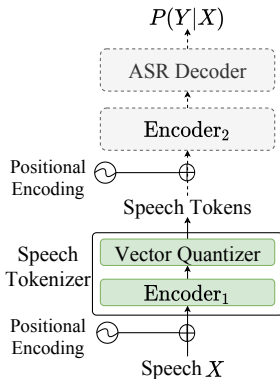


Figure 3: An illustration of our supervised semantic speech tokenizer.

## 2.4 Voice Generation Model: CosyVoice

CosyVoice, a family of fundamental speech generation models (Du et al., 2024a), utilizes  $\mathcal{S}^3$  tokens to synthesize natural-sounding voices suitable for various applications. As a versatile model, CosyVoice excels in tasks such as generating multi-lingual voices tailored to specific speakers, adapting to new speakers without training (zero-shot in-context learning), replicating voices across different languages (cross-lingual voice cloning), creating emotionally resonant voices, and offering nuanced influence over speech output through instructional text. CosyVoice supports five languages, including Chinese (ZH), English (EN), Japanese (JP), Cantonese (Yue) and Korean (KO). We released three open-source models. The first, CosyVoice-base-300M, excels in accurately representing speaker identity, adapting to contexts without any finetuning, and cloning voices across languages. The second, CosyVoice-instruct-300M, is adept in generating emotionally expressive voices and allows for meticulous adjustments via instructional text. Lastly, CosyVoice-sft-300M has been fine-tuned on seven multi-lingual speakers and is ready for immediate deployment. All of them share the common model architecture and learning framework. Compared with other open-sourced projects, CosyVoice released a widest spectrum of supporting features as shown in Table 2.

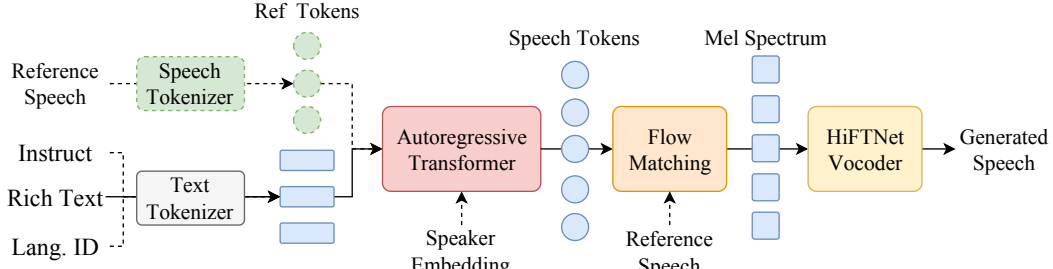


Figure 4: A semantic diagram of CosyVoice models.




Figure 5: Sequence construction for (a) zero-shot in-context learning and (b) cross-lingual voice cloning. LID represents language identifier.

## 2.4.1 System Overview

CosyVoice incorporates an autoregressive Transformer-based language model (LM) to generate speech tokens for the input text. An ordinary differential equation based (ODE-based) diffusion model, flow matching (Lipman et al., 2023), reconstructs Mel spectrum from the generated tokens. Subsequently, a HiFTNet-based vocoder (Li et al., 2023) is followed to synthesize waveforms from the reconstructed Mel spectrum. Dashed models are optional for certain applications, such as cross-lingual cloning and speaker fine-tuned inference.

## 2.4.2 Model Training

At the training stage, the autoregressive language model (LM) is trained using a teacher-forcing paradigm. In this process, tokenized text and a left-shifted version of the speech tokens are provided as input to predict the subsequent speech tokens.

The flow matching model is developed to estimate the conditional probabilities  $P(S|X, v, S_{ref})$ , where  $X$  and  $v$  denote the speech tokens and speaker embeddings (Wang et al., 2023b), respectively.  $S$  and  $S_{ref}$  represent the Mel spectrum of target and reference speech, respectively. A convolutional Transformer U-Net (Mehta et al., 2023) is employed to ascertain the vector field between the prior distribution and the desired one, which is derived from the optimal transport ODE. The straightforward nature of resolving the OT-ODE allows for a significantly reduced number of iterations during the inference stage, typically only five to ten iterations are required to produce a satisfactory Mel spectrogram. We also employ the classifier-free guidance (CFG) (Ho & Salimans, 2022) technique and mask out the 70%~100% preceding feature conditions to boost the in-context learning ability.

For the synthesis of waveforms from the predicted Mel spectrograms, we utilize a vocoder based on HiFTNet (Li et al., 2023). Modifications have been made on HiFTNet to support streaming generation, including the replacement and redesign of certain components. Complete details regarding these adjustments are available in our released code.

### 2.4.3 Zero-shot In-context Learning

CosyVoice models exhibit zero-shot in-context learning capabilities, allowing for the replication of an arbitrary voice with only a brief reference speech sample. This process entails the careful construction of input sequences for the token language model (LM), depicted in Figure 5. For prompt speech and input text in the same language, we merge them to form a unified input, treating the prompt speech tokens as pre-generated. With this input sequence, the autoregressive LM iteratively predicts subsequent tokens until it encounters the “end of sequence” token (E). However, when the prompt speech and input text differ linguistically, we omit the text and tokens associated with the prompt to prevent prosodic characteristics of the original language from influencing the target language. It is important to note that the prompt text, which corresponds to the prompt speech’s content, can be transcribed either through human annotation or ASR models, such as SenseVoice. Similar to the prompt text, the prompt tokens are extracted from the prompt speech with  $S^3$  tokenizer.

After generating the speech tokens, they are appended after the prompt tokens, forming a composite condition for the flow-matching model. Additionally, the speaker embedding and the Mel spectrogram of the prompt speech are incorporated to further enhance timbre and environmental consistency.

### 2.4.4 Instruction Fine-tuning

To enable further controllability on CosyVoice, we experiment with integrating additional instruction fine-tuning (Ji et al., 2023). CosyVoice-instruct extends CosyVoice-base with enhanced instruction-following capabilities. Specifically, it supports controllability over various aspects such as speaker identity (i.e., speaker’s characteristics), speaking style (including emotion, gender, speaking rate, and pitch), and fine-grained paralinguistic features. These features include the ability to insert laughter, breaths, speaking while laughing, and emphasizing certain words. Table 3 shows some examples of speaker identity, speaking style, and fine-grained paralinguistic features.

---

<b>Speaker Identity</b>
1. Selene 'Moonshade', is a <b>mysterious, elegant dancer</b> with a connection to the night. Her movements are both <b>mesmerizing</b> and <b>deadly</b> .<endofprompt>Hope is a good thing.
2. Theo 'Crimson', is a <b>fiery, passionate</b> rebel leader. Fights with fervor for justice, but struggles with <b>impulsiveness</b> .<endofprompt>You don't know about real loss.

---

<b>Speaking Style</b>
1. A <b>happy girl</b> with <b>high tone</b> and <b>quick speech</b> .<endofprompt>The sun is shining brightly today.
2. A <b>sad woman</b> with <b>normal tone</b> and <b>slow speaking speed</b> .<endofprompt>I failed my important exam.

---

<b>Fine-grained Paralinguistics</b>
1. Well that's kind of scary [ <b>laughter</b> ].
2. I don't think I over eat yeah [ <b>breath</b> ] and um I do exercise regularly.
3. Well that pretty much covers < <b>laughter</b> > <b>the subject</b> </laughter> well thanks for calling me.
4. The team's < <b>strong</b> > <b>unity</b> </strong> and < <b>strong</b> > <b>resilience</b> </strong> helped them win the championship.

---

Table 3: Examples of speaker identity, speaking style, and fine-grained paralinguistics.

## 3 Dataset

### 3.1 Training Set for SenseVoice

Figure 6 provides an overview of the dataset utilized for training the SenseVoice models. The SenseVoice-Small model was trained on an extensive audio data corpus of approximately 300,000 hours, covering 5 languages including Chinese, Cantonese, English, Japanese, and Korean. To further enhance the multilingual ability of SenseVoice-Large, an additional 100,000 hours of diverse multilingual data were integrated into the training corpus. To obtain rich transcription labels from speech data, we leveraged open-source models for audio event detection (AED)<sup>1</sup> and speech emo-

<sup>1</sup>[https://github.com/qiuqiangkong/audioset\\_tagging\\_cnn/tree/master](https://github.com/qiuqiangkong/audioset_tagging_cnn/tree/master)




Figure 6: Hours of SenseVoice training data across languages (in log scale).

tion recognition (SER) <sup>2</sup> to generate pseudo labels, yielding an extensive rich transcribe dataset. Specifically, the AED data amounted to 150 million entries, while the SER data comprised 30 million entries.

Language	Duration (hr)
ZH	130,000
EN	30,000
Yue	5,000
JP	4,600
KO	2,200

Table 4: Hours of CosyVoice training data across languages.

### 3.2 Training Set for CosyVoice

To train the CosyVoice models, we have amassed a considerable dataset comprising multiple languages. Throughout the collection process, we utilize specialized in-house tools for speech detection, signal-to-noise ratio (SNR) estimation, speaker diarization, and separation. Subsequently, pseudo text labels are generated using SenseVoice-Large and Paraformer. These labels undergo a refinement process with the aid of force-alignment (FA) models, which helps eliminate low-quality data and enhances the accuracy of punctuation. A comprehensive breakdown of the training data’s duration across various languages is presented in Table 4.

For the CosyVoice-instruct model, we fine-tuned CosyVoice-base using instruction training data without incorporating speaker embedding in the autoregressive language model. Table 5 presents the duration of the training data for different types of instructions.

Type	Duration (hr)
Speaker Identity	101
Speaking Style	407
Fine-grained Paralinguistics	48

Table 5: Duration statistics of instruction training data by type.

<sup>2</sup>[https://modelscope.cn/models/iic/emotion2vec\\_plus\\_large](https://modelscope.cn/models/iic/emotion2vec_plus_large)






Figure 7: Comparison of SenseVoice and Whisper on Common Voice, with or without LID

## 4 Experimental Results

### 4.1 Multilingual Speech Recognition

**Metrics.** We use Character Error Rate (CER) to evaluate the models in five languages: Chinese, Cantonese, Japanese, Korean, and Thai, and use the Word Error Rate (WER) for all other languages. Both the ground truth transcriptions and the recognition outputs are standardized using text normalization before the error rate calculation, in alignment with the methodology used by Whisper. All Chinese characters were converted into the simplified Chinese version, together with an additional text normalization pipeline<sup>3</sup>.

Results in Table 6 show the comparison of Whisper, SenseVoice and Paraformer (Gao et al., 2022, 2023; Shi et al., 2024) on popular open speech recognition benchmark datasets, including AISHELL-1 (Bu et al., 2017), AISHELL-2 (Du et al., 2018), WenetSpeech (Zhang et al., 2022), Librispeech (Panayotov et al., 2015), and Common Voice (Ardila et al., 2019). It can be seen that SenseVoice-S and SenseVoice-L outperform their Whisper counterparts by a significant margin in most test sets except Librispeech.

Figure 7 illustrates the comparative performance of SenseVoice-Large and Whisper-Large-V3 on a broader range of languages, with or without ground truth LID as input. While SenseVoice-Large performs comparably with Whisper-Large-V3 in general, SenseVoice-Large obtains significantly better performance in languages like Cantonese (Yue), Catalan (CA), and Marathi (MR).

The evaluation of inference efficiency is shown in Table 7. The Real-time factor (RTF, the ratio of the transcribing time to the audio length) and 10s Audio Latency (the average time cost when transcribing a 10s audio.) are benchmarked on an A800 machine, with a decoding batch size of 1. For the encoder-decoder-based model (Whisper-S, Whisper-L-V3, and SenseVoice-L), we perform beam search in decoding with a beam size of 5. Owing to its non-autoregressive architecture, SenseVoice-S obtains extremely low inference latency—more than 5 times faster compared to Whisper-small and more than 15 times faster compared to Whisper-L-V3. SenseVoice-L shows close performance with Whisper-L-V3.

### 4.2 Speech Emotion Recognition

We evaluate the SER ability of the SenseVoice on 7 popular emotion recognition datasets, including CREMA-D(Cao et al., 2014), MELD(Poria et al., 2019), IEMOCAP(Busso et al., 2008), MSP-

<sup>3</sup>[https://github.com/speechio/chinese\\_text\\_normalization/blob/master/python/cn\\_tn.py](https://github.com/speechio/chinese_text_normalization/blob/master/python/cn_tn.py)

	Whisper-S	Whisper-L-V3	SenseVoice-S	SenseVoice-L	Paraformer-zh
AISHELL-1 test	10.04	5.14	2.96	2.09	<b>1.95</b>
AISHELL-2 test_ios	8.78	4.96	3.80	3.04	<b>2.85</b>
WenetSpeech test_meeting	25.62	18.87	7.44	<b>6.73</b>	6.97
WenetSpeech test_net	16.66	10.48	7.84	<b>6.01</b>	6.74
LibriSpeech test_clean	3.13	<b>1.82</b>	3.15	2.57	-
LibriSpeech test_other	7.37	<b>3.50</b>	7.18	4.28	-
CommonVoice zh-CN	19.60	12.55	10.78	<b>7.68</b>	10.30
CommonVoice en	14.85	9.39	14.71	<b>9.00</b>	-
CommonVoice yue	38.97	10.41	7.09	<b>6.78</b>	-
CommonVoice ja	19.51	10.34	11.96	<b>9.19</b>	-
CommonVoice ko	10.48	5.59	8.28	<b>5.21</b>	-
CommonVoice 5 lang. Average	20.68	9.66	10.56	<b>7.57</b>	-

Table 6: Performance comparisons among different models on Chinese and English Open Corpus.

Model	Framework	Parameters	Support Language	RTF	10s Audio Latency(ms)
Whisper-S	Autoregressive	224M	50+	0.042	518
Whisper-L-V3	Autoregressive	1550M	50+	0.111	1281
Paraformer-zh	Non-autoregressive	220M	zh	0.009	100
SenseVoice-S	Non-autoregressive	234M	zh,yue,en,ja,ko	0.007	70
SenseVoice-L	Autoregressive	1587M	50+	0.110	1623

Table 7: Comparison of model architecture, parameter scale, supported languages, and inference efficiency of SenseVoice, Paraformer, and Whisper.

Podcast(Martinez-Lucas et al., 2020), CASIA(Zhang & Jia, 2008), MER2023(Lian et al., 2023) and ESD(Zhou et al., 2021). These corpora cover both Chinese and English, and scenarios like acts, TV dramas, and daily conversation. We report unweighted average accuracy (UA), weighted average accuracy (WA), macro F1 Score (F1), and weighted average F1 (WF1), and compare them with some recently published SER benchmarks (EmoBox (Ma et al., 2024a), Emo-Superb(Wu et al., 2024) and MerBench (Lian et al., 2024)) from literature in Table 8. We show that SenseVoice achieves a good performance on all test sets and all metrics even without fine-tuning on the target domain.

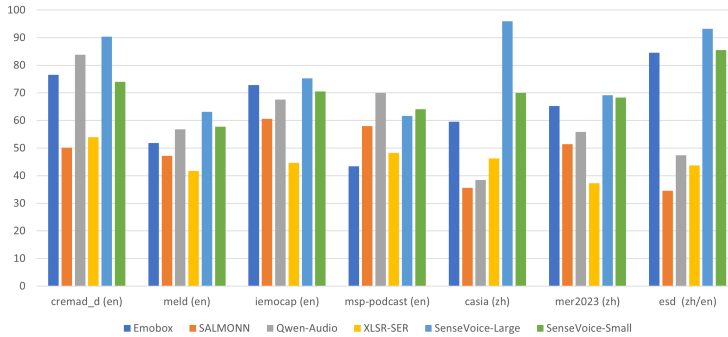


Figure 8: Weighted Average Accuracy (WA(%)) comparison with other open source SER models.

We further compare SenseVoice with some open-sourced SER models. Results are shown in Figure 8. XLSR-SER is the most popular SER model on HuggingFace<sup>4</sup>, and Qwen-Audio(Chu et al., 2023) and SALMONN(Tang et al., 2024) are two Audio-LLM models which can recognize speech emotion with natural language prompt. Results from EmoBox are also involved in the figure as references. SenseVoice-Large achieves the best results on almost all datasets while the SenseVoice-Small also outperforms other baseline models on the majority datasets.

<sup>4</sup><https://huggingface.co/ehcalabres/wav2vec2-lg-xlsr-en-speech-emotion-recognition>

Test set	EmoBox			Emo-Superb	MerBench	SenseVoice-L				SenseVoice-S			
	UA	WA	F1	F1	WF1	UA	WA	F1	WF1	UA	WA	F1	WF1
CASIA	59.6	59.6	56.3	–	–	96.0	96.0	95.5	95.5	70.0	70.0	70.3	70.3
CREMA-D	76.8	76.5	76.6	67.7	–	90.1	90.4	89.8	89.9	73.1	74.0	65.7	65.7
ESD	84.6	84.6	84.3	–	–	93.2	93.2	92.2	92.2	85.5	85.5	81.0	81.0
IEMOCAP	73.5	72.9	73.1	–	69.7	73.9	75.3	73.2	72.8	70.5	65.7	67.9	67.8
MELD	31.5	51.9	32.9	–	46.5	58.7	63.1	50.9	65.7	50.8	57.8	44.6	57.7
MER2023	61.2	65.2	62.3	–	67.5	70.9	69.2	55.6	57.4	69.0	68.3	52.8	56.6
MSPPodcast	21.4	43.4	21.5	38.4	–	46.0	61.7	45.0	58.9	49.4	64.1	46.4	63.1

Table 8: SER performance comparisons on different evaluation benchmarks.

### 4.3 Audio Event Detection

Both SenseVoice-Small and SenseVoice-Large models can classify the audio event in the speech, including music, applause, and laughter. The SenseVoice-L can further predict the start and end position of the audio event, while the SenseVoice-Small can only predict what happened in the audio, with at most one event per utterance. SenseVoice-Small can detect more kinds of events, such as coughing, sneezing, breathing, and crying which could occur during human-machine interaction.




Figure 9: F1(%) Score comparison of the SenseVoice with the audio event detection models BEATS and PANNs on different audio event detection tasks.

We compare SenseVoice with the SOTA audio event detection models BEATS(Chen et al., 2023a) and PANNs(Kong et al., 2020) on different tasks, including environment sound classification (ESC-50)(Piczak, 2015), baby cry/laugh detection<sup>5</sup>, coughing detection (Coswara)(Sharma et al., 2020)<sup>6</sup> and in-home talkshow event detection. As SenseVoice only predicts the event of our interest, which may not include event categories in other models, we use the F1 score on each event for evaluation. Qwen-audio is also evaluated for comparison.

We find that SenseVoice serves as a good audio event classification or detection model, though BEATS and PANNs may have better F1 scores, which may be attributed to two reasons. Firstly, BEATS and PANNs can modify the detection threshold to trade-off the accuracy and recall rate to obtain a higher F1 score, but threshold modification is much more difficult for SenseVoice and Qwen-Audio (An interesting discovery is that SenseVoice and Qwen-Audio always have a much higher accuracy than the recall rate, which could be more friendly for the human-machine interaction). Secondly, SenseVoice is trained with ASR data with AED pseudo labeling rather than AED-specific data.

<sup>5</sup>[https://github.com/giulbia/baby\\_cry\\_detection/tree/master](https://github.com/giulbia/baby_cry_detection/tree/master)

<sup>6</sup><https://github.com/iiscleap/Coswara-Data/tree/master>

#### 4.4 Preserving Semantic Information by $S^3$ Tokenizer

To assess the  $S^3$  tokenizer’s ability to preserve semantic information, we compared the recognition performance of the quantizer-augmented SenseVoice-L against its original version and the Whisper-Large V3 model. The models underwent evaluation using the Common Voice zh-CN and en benchmarks, with the findings detailed in Table 9.

From the table, we can see that our  $S^3$  tokens demonstrate robust recognition performance in both the Chinese and English test sets. Notably, on the common\_voice\_zh-CN set,  $S^3$  tokens surpass the performance of the Whisper-Large V3 model, achieving a 4.14% relative reduction in error rate. This suggests a substantial correlation between  $S^3$  tokens and semantic content. It is worth noting that there is only a single codebook in the  $S^3$  tokenizer with a dictionary size of 4,096 entries.

Test set	Whisper-L-V3		SenseVoice-L		$S^3$ tokens	
	w/o lid	w/ lid	w/o lid	w/ lid	w/o lid	w/ lid
common_voice_zh-CN	12.82	12.55	8.76	8.68	12.24	12.06
common_voice_en	13.55	9.39	9.79	9.77	15.43	15.38

Table 9: The evaluation on  $S^3$  tokens’ capability to preserve semantic information. We employ character and word error rates for zh-CN and en languages on the Common Voice benchmarks. Please note that the SenseVoice-L model in this table is an intermediate version, and is not identical to the one presented in Table 6.

#### 4.5 Evaluation on Generation Quality of CosyVoice

We evaluate the quality of CosyVoice’s speech synthesis by examining content consistency and speaker similarity. The “test-clean” subset of LibriTTS (Zen et al., 2019) and the test set of AISHELL-3 (Shi et al., 2021) are employed to construct evaluation set for English and Chinese, respectively. For each text in these sets, we randomly select a prompt speech. Content consistency was evaluated using Whisper-Large V3 (Radford et al., 2023) for English and Paraformer (Gao et al., 2022) for Chinese recognition. Speaker similarity was quantified by calculating the cosine similarity between speaker embeddings of the generated and prompt speeches, extracted using ERes2Net (Chen et al., 2023b).

Similar to other autoregressive language models, we employ a random sampling decoding strategy for our token LM and assessed the synthesis process using five different random seed values: 0, 7, 42, 123, and 1,337. The resultant evaluation metrics were averaged to determine the mean and standard deviation. Additionally, we conducted an ASR re-ranking to demonstrate potential performance improvements in offline mode.

Tables 10 and 11 present the results for English and Chinese, respectively. On the English dataset, CosyVoice attained human-level performance with similar content recognition and higher speaker similarity. ASR re-ranking notably enhanced content consistency, yielding a reduced word error rate (WER) of 1.51%. CosyVoice outperformed ChatTTS in WER and the number of insertion and deletion errors, indicating superior content consistency. We did not assess speaker similarity for ChatTTS as it doesn’t release voice cloning capabilities.

Model	WER (%)	#Ins.&Del.	SS
Original	2.66	92	69.67
ChatTTS	8.32	441	-
CosyVoice	2.89±0.18	88.60±3.88	74.30±0.15
+ 5× re-ranking	1.51	47	74.30

Table 10: The comparison of original and CosyVoice generated speeches on the LibriTTS test-clean set in terms of word error rate (WER) and speaker similarity (SS). “±” joins the mean and standard deviation for each evaluation metric.

As for the results on Chinese, the generated utterances of CosyVoice achieves a comparable CER as well as the errors of insertion and deletion compared with the original utterances. It seems

Model	CER (%)	#Ins.&Del.	SS
Original	2.52	25	74.15
ChatTTS	3.87	111	-
CosyVoice	3.82±0.24	24.4±2.24	81.58±0.16
+ 5× re-ranking	1.84	11	81.58

Table 11: The comparison of original and CosyVoice generated speeches on the AISHELL-3 test set in terms of character error rate (CER) and speaker similarity (SS). “±” joins the mean and standard deviation for each evaluation metric.

that ChatTTS has a better generation ability on Chinese than English in terms of CER. Although ChatTTS and CosyVoice achieves a similar CER, ChatTTS produces more insertion and deletion errors, This is due to the problem of speaker leaking, where modal particles of another speaker is generated unexpectedly. On the contrary, CosyVoice doesn’t suffer this problem with much less insertion and deletion errors. With ASR re-ranking, CosyVoice reached a remarkably low CER of 1.84%. As seen with English, CosyVoice also exhibited greater speaker similarity than the original utterances, showcasing its effective voice-cloning proficiency.

#### 4.6 Evaluation on Emotion Controllability of CosyVoice

To verify the emotion controllability, we use the public speech emotion recognition model emo2vec<sup>7</sup> (Ma et al., 2024b). We generate and evaluate 100 English utterances for each of the six emotions: happy, angry, sad, surprised, fearful, and disgusted. The content of the synthesized text is designed to match the target emotion. We then measure the accuracy of the predicted emotions from the synthesized speech for each emotion.

Table 12 shows the comparison of emotion control accuracy between CosyVoice-base and CosyVoice-instruct. For CosyVoice-instruct, the input consists of content text accompanied by a speaking style instruction (e.g., “Happy.<endofprompt>Content Text”). In contrast, CosyVoice-base only receives the content text as input. The results indicate that CosyVoice-instruct with emotional instructions demonstrates a significant improvement over both CosyVoice-base and CosyVoice-instruct without emotional instructions.

Model	Happy	Sad	Angry	Surprised	Fearful	Disgusted
CosyVoice-base	1.00±0.00	0.45±0.05	0.59±0.03	0.26±0.02	0.88±0.01	0.46±0.06
CosyVoice-instruct	1.00±0.00	0.98±0.02	0.83±0.04	0.64±0.03	0.87±0.03	0.93±0.02
w/o instruction	0.98±0.01	0.77±0.04	0.49±0.12	0.28±0.06	0.83±0.04	0.45±0.16

Table 12: Comparison of emotion control accuracy between CosyVoice-base and CosyVoice-instruct. “±” joins the mean and standard deviation for each evaluation metric.

#### 4.7 CosyVoice as a Data Generator

A straightforward application of CosyVoice is as a data generator to augment the training data of other tasks, such as ASR, speech-to-speech translation (S2ST). Taking the ASR task an example, we conduct an experiment on the Librispeech corpus to evaluate CosyVoice’s capability in generating high-quality data. The experimental results are shown in Table 13, where “Librispeech” denotes the original 960-hour data. “Syn on LS text” and “Syn on LS text” denote the generated data with the text from Librispeech and MLS training sets, respectively. From the table, we can see that only training on the synthesized data, the ASR model can achieve a comparable result than the original Librispeech training set. Upon integration of them, a notable enhancement in recognition accuracy is observed. An interesting finding is that involving the synthesized data on the MLS text significantly improve the recognition performance. This may indicates that the text diversity is more critical for ASR task than the duration of speech itself. This improvement can be attributed to the varied

<sup>7</sup>[https://modelscope.cn/models/iic/emotion2vec\\_base\\_finetuned](https://modelscope.cn/models/iic/emotion2vec_base_finetuned)

linguistic content introduced by CosyVoice synthesized samples. The findings from our evaluation underscore the high quality of the samples generated by CosyVoice.

Training Data	dev_clean	dev_other	test_clean	test_other
Librispeech	2.77	5.84	2.79	5.97
Syn on LS text	2.79	6.37	3.00	6.59
Librispeech + Syn on LS text	2.44	5.52	2.56	5.68
Librispeech + Syn on LS text $\times 2$	2.51	5.23	2.68	5.26
Librispeech + Syn on LS, MLS text	<b>1.93</b>	<b>4.43</b>	<b>2.04</b>	<b>4.53</b>

Table 13: Evaluation on CosyVoice generation quality by treating it as a data generator. Word error rates (%) on the human-uttered test sets are employed as the evaluation metrics.

## 5 Applications

The FunAudioLLM is an innovative framework designed to facilitate natural voice interactions between humans and large language models (LLMs). By integrating SenseVoice, CosyVoice, and LLMs, FunAudioLLM offers a variety of rich application demos, including speech-to-speech translation, emotional voice chat, interactive podcasts, and expressive audiobook narration. The demos are available at <https://fun-audio-llm.github.io>.

By combining SenseVoice, LLMs, and CosyVoice, we can effortlessly perform speech-to-speech translation (S2ST), as illustrated in Figure 10. SenseVoice is used to recognize the input speech in its original language, the LLM translates the source language to the target language, and CosyVoice synthesizes the target speech with cross-lingual voice cloning. This allows users to speak in foreign languages using their own voice.




Figure 10: A diagram of Speech-to-Speech Translation.

By integrating SenseVoice, LLMs, and CosyVoice, we can develop an Emotional Voice Chat application, as depicted in Figure 11. SenseVoice recognizes the input speech and its emotion and audio event, the LLM generates the response content with a speaking style description, and CosyVoice produces emotional speech following the given speaking style description.




Figure 11: A diagram of Emotional Voice Chat.

By leveraging SenseVoice, an LLM-based multi-agent system with real-time world knowledge, and CosyVoice, we can create an interactive podcast, as shown in Figure 12. We can use an LLM plugin to fetch real-time daily knowledge, which a content-generation agent then transforms into a podcast script. The Multi-Agent system matches podcast roles, and CosyVoice synthesizes the voices. Users can also insert themselves into the podcast for interactive dialogues with the Multi-Agent system.

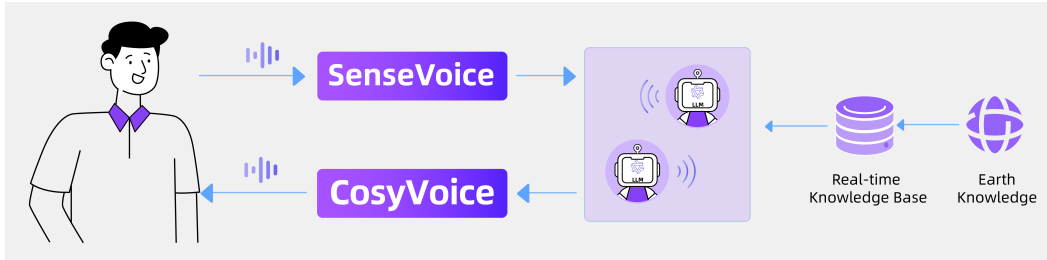


Figure 12: A diagram of Interactive Podcast.

Through the analytical capabilities of LLMs to structure and identify emotions within books, and synthesizing this with CosyVoice, we achieve audiobooks with enhanced expressiveness, as illustrated in Figure 13. The LLM is used for narrative and dialogue analysis, character analysis, and fine-grained sentiment analysis, while CosyVoice synthesizes the speech with enhanced expressiveness.




Figure 13: A diagram of Expressive Audiobook.

## 6 Limitations

SenseVoice has certain limitations that need to be addressed. Firstly, the ASR performance generally remains much lower for under-resourced languages. Secondly, SenseVoice is not designed for streaming transcription. Therefore, future work may focus on developing streamable voice understanding models based on SenseVoice.

CosyVoice also has several limitations. Firstly, it supports a limited number of languages. While it can express emotions and speaking styles based on explicit instructions, it cannot infer the appropriate emotion or style based on the semantic content of the text. Additionally, CosyVoice does not perform well when tasked with singing. There’s still room for improvement in achieving expressive emotional changes while maintaining the original timbre of the voice.

Another limitation is that the two innovative models within FunAudioLLM are not trained end-to-end with LLMs. This pipeline approach may introduce error propagation, which could affect overall performance.

## 7 Authors (alphabetical order of family name)

- Keyu An
- Qian Chen
- Chong Deng
- Zhihao Du
- Changfeng Gao
- Zhifu Gao
- Yue Gu
- Ting He
- Hangrui Hu
- Kai Hu
- Shengpeng Ji
- Yabin Li
- Zerui Li
- Heng Lu
- Xiang Lv
- Bin Ma
- Ziyang Ma
- Chongjia Ni
- Changhe Song
- Jiaqi Shi
- Xian Shi
- Hao Wang
- Wen Wang
- Yuxuan Wang
- Zhangyu Xiao
- Zhijie Yan
- Yexin Yang
- Bin Zhang
- Qinglin Zhang
- Shiliang Zhang
- Nan Zhao
- Siqi Zheng

## 8 Acknowledgment

We extend our heartfelt appreciation to the developers and contributors of the following open-source projects: FunASR, FunCodec, Whisper, ESPNet, WeNet, SLAM-LLM, Matcha-TTS, and Tortoise. Their innovative efforts and valuable code contributions have significantly inspired our work and facilitated our research. We are also grateful to numerous other projects not explicitly mentioned here, which have equally provided considerable assistance and played an instrumental role in the success of our endeavors.

## References

- Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M Tyers, and Gregor Weber. Common voice: A massively-multilingual speech corpus. *arXiv preprint arXiv:1912.06670*, 2019.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *CoRR*, abs/2309.16609, 2023a.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *CoRR*, abs/2308.12966, 2023b.
- Alexandre Berard, Laurent Besacier, Ali Can Kocabiyikoglu, and Olivier Pietquin. End-to-end automatic speech translation of audiobooks. In *ICASSP*, pp. 6224–6228. IEEE, 2018.
- Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao Zheng. Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline. In *2017 20th conference of the oriental chapter of the international coordinating committee on speech databases and speech I/O systems and assessment (O-COCOSDA)*, pp. 1–5. IEEE, 2017.
- Carlos Busso, Murtaza Bulut, Chi-Chun Lee, Ebrahim (Abe) Kazemzadeh, Emily Mower Provost, Samuel Kim, Jeannette N. Chang, Sungbok Lee, and Shrikanth S. Narayanan. Iemocap: interactive emotional dyadic motion capture database. *Language Resources and Evaluation*, 42: 335–359, 2008.
- Houwei Cao, David G. Cooper, Michael K. Keutmann, Ruben C. Gur, Ani Nenkova, and Ragini Verma. Crema-d: Crowd-sourced emotional multimodal actors dataset. *IEEE Transactions on Affective Computing*, 5(4):377–390, 2014. doi: 10.1109/TAFFC.2014.2336244.
- Aimilios Chalamandaris, Pirros Tsiakoulis, Sotiris Karabetsos, and Spyros Raptis. Using audio books for training a text-to-speech system. In *LREC*, pp. 3076–3080. European Language Resources Association (ELRA), 2014.
- Qian Chen, Mengzhe Chen, Bo Li, and Wen Wang. Controllable time-delay transformer for real-time punctuation prediction and disfluency detection. In *ICASSP*, pp. 8069–8073. IEEE, 2020.
- Sanyuan Chen, Yu Wu, Chengyi Wang, Shujie Liu, Daniel Tompkins, Zhuo Chen, Wanxiang Che, Xiangzhan Yu, and Furu Wei. Beats: Audio pre-training with acoustic tokenizers. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pp. 5178–5193. PMLR, 2023a.
- Yafeng Chen, Siqi Zheng, Hui Wang, Luyao Cheng, Qian Chen, and Jiajun Qi. An enhanced res2net with local and global feature fusion for speaker verification. In *Interspeech*. ISCA, 2023b.
- Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models, 2023.



- Jiayu Du, Xingyu Na, Xuechen Liu, and Hui Bu. Aishell-2: Transforming mandarin asr research into industrial scale. *arXiv preprint arXiv:1808.10583*, 2018.
- Zhihao Du, Qian Chen, Shiliang Zhang, Kai Hu, Heng Lu, Yexin Yang, Yue Gu, Ziyang Ma, and Zhijie Yan. Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer based on supervised semantic tokens. *arxiv*, 2024a.
- Zhihao Du, Shiliang Zhang, Kai Hu, and Siqi Zheng. Funcodec: A fundamental, reproducible and integrable open-source toolkit for neural speech codec. In *ICASSP*, 2024b.
- Alexandre Défossez, Jade Copet, Gabriel Synnaeve, and Yossi Adi. High fidelity neural audio compression. *arXiv:2210.13438*, 2022.
- Zhifu Gao, Shiliang Zhang, Ming Lei, and Ian McLoughlin. SAN-M: memory equipped self-attention for end-to-end speech recognition. In *21st Annual Conference of the International Speech Communication Association, Interspeech 2020, Virtual Event, Shanghai, China, October 25-29, 2020*, pp. 6–10. ISCA, 2020.
- Zhifu Gao, Shiliang Zhang, Ian McLoughlin, and Zhijie Yan. Paraformer: Fast and accurate parallel transformer for non-autoregressive end-to-end speech recognition. In *Interspeech*, pp. 2063–2067. ISCA, 2022.
- Zhifu Gao, Zerui Li, Jiaming Wang, Haoneng Luo, Xian Shi, Mengzhe Chen, Yabin Li, Lingyun Zuo, Zhihao Du, Zhangyu Xiao, and Shiliang Zhang. Funasr: A fundamental end-to-end speech recognition toolkit. In *INTERSPEECH*, 2023.
- Alex Graves, Santiago Fernández, Faustino J. Gomez, and Jürgen Schmidhuber. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Machine Learning, Proceedings of the Twenty-Third International Conference (ICML 2006), Pittsburgh, Pennsylvania, USA, June 25-29, 2006*, volume 148, pp. 369–376. ACM, 2006.
- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *CoRR*, abs/2207.12598, 2022.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE ACM Trans. Audio Speech Lang. Process.*, 29:3451–3460, 2021.
- Shengpeng Ji, Jialong Zuo, Minghui Fang, Ziyue Jiang, Feiyang Chen, Xinyu Duan, Baoxing Huai, and Zhou Zhao. Textrolspeech: A text style control speech corpus with codec language text-to-speech models. *CoRR*, abs/2308.14430, 2023.
- Shengpeng Ji, Jialong Zuo, Minghui Fang, Siqi Zheng, Qian Chen, Wen Wang, Ziyue Jiang, Hai Huang, Xize Cheng, Rongjie Huang, and Zhou Zhao. Controlspeech: Towards simultaneous zero-shot speaker cloning and zero-shot language style control with decoupled codec. 2024. URL <https://arxiv.org/abs/2406.01205>.
- Naoyuki Kanda, Xiaofei Wang, Sefik Emre Eskimez, Manthan Thakker, Hemin Yang, Zirun Zhu, Min Tang, Canrun Li, Chung-Hsien Tsai, Zhen Xiao, Yufei Xia, Jinzhu Li, Yanqing Liu, Sheng Zhao, and Michael Zeng. Making flow-matching-based zero-shot text-to-speech laugh as you like. *CoRR*, abs/2402.07383, 2024.
- Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D. Plumbley. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2880–2894, 2020. doi: 10.1109/TASLP.2020.3030497.
- Philippe Laban, Elicia Ye, Srujay Korlakunta, John F. Canny, and Marti A. Hearst. Newspod: Automatic and interactive news podcasts. In *IUI*, pp. 691–706. ACM, 2022.
- Yinghao Aaron Li, Cong Han, Xilin Jiang, and Nima Mesgarani. Hiftnet: A fast high-quality neural vocoder with harmonic-plus-noise filter and inverse short time fourier transform. *CoRR*, abs/2309.09493, 2023.

- Zheng Lian, Haiyang Sun, Licai Sun, Kang Chen, Mingyu Xu, Kexin Wang, Ke Xu, Yu He, Ying Li, Jinming Zhao, Ye Liu, Bin Liu, Jiangyan Yi, Meng Wang, Erik Cambria, Guoying Zhao, Björn W. Schuller, and Jianhua Tao. Mer 2023: Multi-label learning, modality robustness, and semi-supervised learning, 2023.
- Zheng Lian, Licai Sun, Yong Ren, Hao Gu, Haiyang Sun, Lan Chen, Bin Liu, and Jianhua Tao. Merbench: A unified evaluation benchmark for multimodal emotion recognition, 2024.
- Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *ICLR*. OpenReview.net, 2023.
- Ziyang Ma, Mingjie Chen, Hezhao Zhang, Zhisheng Zheng, Wenxi Chen, Xiquan Li, Jiabin Ye, Xie Chen, and Thomas Hain. Emobox: Multilingual multi-corpus speech emotion recognition toolkit and benchmark. In *Proc. INTERSPEECH*, 2024a.
- Ziyang Ma, Zhisheng Zheng, Jiabin Ye, Jinchao Li, Zhifu Gao, Shiliang Zhang, and Xie Chen. emotion2vec: Self-supervised pre-training for speech emotion representation. *Proc. ACL Findings*, 2024b.
- Luz Martinez-Lucas, Mohammed Abdelwahab, and Carlos Busso. The MSP-Conversation Corpus. In *Proc. Interspeech 2020*, pp. 1823–1827, 2020.
- Shivam Mehta, Ruibo Tu, Jonas Beskow, Éva Székely, and Gustav Eje Henter. Matcha-tts: A fast TTS architecture with conditional flow matching. *CoRR*, abs/2309.03199, 2023.
- Annamaria Mesaros, Toni Heittola, Tuomas Virtanen, and Mark D. Plumbley. Sound event detection: A tutorial. *IEEE Signal Process. Mag.*, 38(5):67–83, 2021.
- OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 5206–5210. IEEE, 2015.
- Karol J. Piczak. ESC: dataset for environmental sound classification. In *ACM Multimedia*, pp. 1015–1018. ACM, 2015.
- Soujanya Poria, Devamanyu Hazarika, Navonil Majumder, Gautam Naik, Erik Cambria, and Rada Mihalcea. MELD: A multimodal multi-party dataset for emotion recognition in conversations. In Anna Korhonen, David Traum, and Lluís Màrquez (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 527–536, 2019.
- Ernest Pusateri, Bharat Ram Ambati, Elizabeth Brooks, Ondrej Plátek, Donald McAllaster, and Venki Nagesha. A mostly data-driven approach to inverse text normalization. In *INTER\_SPEECH*, pp. 2784–2788. ISCA, 2017.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pp. 28492–28518. PMLR, 2023.
- Machel Reid, Nikolay Savinov, Denis Teplyashin, Dmitry Lepikhin, Timothy P. Lillicrap, Jean-Baptiste Alayrac, Radu Soricut, Angeliki Lazaridou, Orhan Firat, Julian Schrittwieser, Ioannis Antonoglou, Rohan Anil, Sebastian Borgeaud, Andrew M. Dai, Katie Millican, Ethan Dyer, Mia Glaese, Thibault Sottiaux, Benjamin Lee, Fabio Viola, Malcolm Reynolds, Yuanzhong Xu, James Molloy, Jilin Chen, Michael Isard, Paul Barham, Tom Hennigan, Ross McIlroy, Melvin Johnson, Johan Schalkwyk, Eli Collins, Eliza Rutherford, Erica Moreira, Kareem Ayoub, Megha Goel, Clemens Meyer, Gregory Thornton, Zhen Yang, Henryk Michalewski, Zaheer Abbas, Nathan Schucher, Ankesh Anand, Richard Ives, James Keeling, Karel Lenc, Salem Haykal, Siamak Shakeri, Pranav Shyam, Aakanksha Chowdhery, Roman Ring, Stephen Spencer, Eren Sezener, and et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *CoRR*, abs/2403.05530, 2024.

- Neeraj Sharma, Prashant Krishnan, Rohit Kumar, Shreyas Ramoji, Srikanth Raj Chetupalli, Nirmala R., Prasanta Kumar Ghosh, and Sriram Ganapathy. Coswara — a database of breathing, cough, and voice sounds for covid-19 diagnosis. In *Interspeech 2020*. ISCA, 2020.
- Xian Shi, Yexin Yang, Zerui Li, Yanni Chen, Zhifu Gao, and Shiliang Zhang. Seaco-paraformer: A non-autoregressive asr system with flexible and effective hotword customization ability. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 10346–10350. IEEE, 2024.
- Yao Shi, Hui Bu, Xin Xu, Shaoji Zhang, and Ming Li. AISHELL-3: A multi-speaker mandarin TTS corpus. In *Interspeech*, pp. 2756–2760. ISCA, 2021.
- Reo Shimizu, Ryuichi Yamamoto, Masaya Kawamura, Yuma Shirahata, Hironori Doi, Tatsuya Komatsu, and Kentaro Tachibana. Prompttts++: Controlling speaker identity in prompt-based text-to-speech using natural language descriptions. *CoRR*, abs/2309.08140, 2023.
- Yookyung Shin, Younggun Lee, Suhee Jo, Yeongtae Hwang, and Taesu Kim. Text-driven emotional style control and cross-speaker style transfer in neural TTS. In *INTERSPEECH*, pp. 2313–2317. ISCA, 2022.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun MA, and Chao Zhang. SALMONN: Towards generic hearing abilities for large language models. In *The Twelfth International Conference on Learning Representations*, 2024.
- Qwen Team. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971, 2023.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 5998–6008, 2017.
- Chengyi Wang, Sanyuan Chen, Yu Wu, Ziqiang Zhang, Long Zhou, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, Lei He, Sheng Zhao, and Furu Wei. Neural codec language models are zero-shot text to speech synthesizers. *CoRR*, abs/2301.02111, 2023a.
- Hui Wang, Siqi Zheng, Yafeng Chen, Luyao Cheng, and Qian Chen. CAM++: A fast and efficient network for speaker verification using context-aware masking. In *INTERSPEECH*, pp. 5301–5305. ISCA, 2023b.
- Haibin Wu, Huang-Cheng Chou, Kai-Wei Chang, Lucas Goncalves, Jiawei Du, Jyh-Shing Roger Jang, Chi-Chun Lee, and Hung-Yi Lee. Emo-superb: An in-depth look at speech emotion recognition, 2024.
- Hongfei Xue, Yuhao Liang, Bingshen Mu, Shiliang Zhang, Mengzhe Chen, Qian Chen, and Lei Xie. E-chat: Emotion-sensitive spoken dialogue system with large language models. *CoRR*, abs/2401.00475, 2024.
- Dongchao Yang, Songxiang Liu, Rongjie Huang, Jinchuan Tian, Chao Weng, and Yuexian Zou. Hifi-codec: Group-residual vector quantization for high fidelity audio codec. *CoRR*, abs/2305.02765, 2023.
- Neil Zeghidour, Alejandro Luebs, Ahmed Omran, Jan Skoglund, and Marco Tagliasacchi. Soundstream: An end-to-end neural audio codec. *IEEE ACM Trans. Audio Speech Lang. Process.*, 30: 495–507, 2022.
- Heiga Zen, Viet Dang, Rob Clark, and et al. Libritts: A corpus derived from librispeech for text-to-speech. *arXiv:1904.02882*, 2019.

Binbin Zhang, Hang Lv, Pengcheng Guo, Qijie Shao, Chao Yang, Lei Xie, Xin Xu, Hui Bu, Xiaoyu Chen, Chenchen Zeng, et al. Wenetspeech: A 10000+ hours multi-domain mandarin corpus for speech recognition. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6182–6186. IEEE, 2022.

JTFLM Zhang and Huibin Jia. Design of speech corpus for mandarin text to speech. In *The blizzard challenge 2008 workshop*, 2008.

Ziqiang Zhang, Long Zhou, Chengyi Wang, Sanyuan Chen, Yu Wu, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, Lei He, Sheng Zhao, and Furu Wei. Speak foreign languages with your own voice: Cross-lingual neural codec language modeling. *CoRR*, abs/2303.03926, 2023.

Kun Zhou, Berrak Sisman, Rui Liu, and Haizhou Li. Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset, 2021.

## **A Auxiliary Results of SenseVoice on Common Voice.**

Language	Whisper-L-V3		SenseVoice-L	
	w/o lid	w lid	w/o lid	w lid
zh-CN	12.82	12.55	7.92	7.68
en	13.55	9.39	14.30	9.00
yue	40.42	10.51	7.08	6.78
ja	11.18	10.34	9.58	9.19
ko	5.59	5.59	5.23	5.21
fr	11.13	10.77	8.67	8.45
es	5.00	4.74	5.37	4.63
it	5.93	5.46	5.74	5.16
ru	6.16	5.67	6.60	5.23
id	8.98	7.22	12.80	6.97
th	9.73	5.80	4.36	4.12
de	6.06	5.70	6.94	6.57
ca	16.76	13.20	5.90	5.62
nl	5.51	4.28	6.65	5.23
pt	6.90	5.92	9.05	6.88
pl	7.26	5.95	10.01	7.47
cs	10.99	9.04	11.45	9.70
hi	46.17	16.88	48.85	10.06
tr	14.65	12.04	14.10	11.09
ro	13.43	10.84	18.21	12.01
hu	13.89	13.40	12.53	12.27
da	15.59	12.49	17.41	13.23
bg	17.05	14.24	18.56	13.25
mr	38.14	31.13	20.80	13.51
el	15.58	13.73	25.39	16.98
uk	15.89	11.60	12.43	20.75
az	36.32	25.21	72.65	28.63
sw	54.10	50.43	26.21	25.85
fa	37.44	34.86	32.40	40.33
bn	42.25	40.15	44.10	43.80

Table 14: Performance comparisons among different models with and without language id.