



# SOVA-Bench: Benchmarking the Speech Conversation Ability for LLM-based Voice Assistant

Yixuan Hou<sup>1,†</sup>, Heyang Liu<sup>1,2,†</sup>, Yuhao Wang<sup>1,2</sup>, Ziyang Cheng<sup>3</sup>, Ronghua Wu<sup>2</sup>, Qunshan Gu<sup>2,\*</sup>, Yanfeng Wang<sup>1</sup>, Yu Wang<sup>1,\*</sup>

<sup>1</sup>School of Artificial Intelligence, Shanghai Jiao Tong University, China; <sup>2</sup>Ant Group, China;

<sup>3</sup>School of Cyber Science and Engineering, Wuhan University, China

{yixuan303a, liuheyang, colane, wangyanfeng622, yuwangsjtu}@sjtu.edu.cn,  
{r.wu, guqunshan.gqs}@antgroup.com, icelookgoose@gmail.com

## Abstract

Thanks to the steady progress of large language models (LLMs), speech encoding algorithms and vocoder structure, recent advancements have enabled generating speech response directly from a user instruction. However, benchmarking the generated speech quality has been a neglected but critical issue, considering the shift from the pursuit of semantic accuracy to vivid and spontaneous speech flow. Previous evaluation focused on the speech-understanding ability, lacking a quantification of acoustic quality. In this paper, we propose **Speech cOnversational Voice Assistant Benchmark (SOVA-Bench)**, providing a comprehension comparison of the general knowledge, speech recognition and understanding, along with both semantic and acoustic generative ability between available speech LLMs. To the best of our knowledge, SOVA-Bench is one of the most systematic evaluation frameworks for speech LLMs, inspiring the direction of voice interaction systems.

**Index Terms:** speech large language model, voice assistant, evaluation benchmark

## 1. Introduction

Communication through speech outperforms text interaction for its convenience and efficiency. Recent advancements in large language models (LLMs) have led to remarkable breakthroughs in speech LLMs and voice assistants. Starting from an LLM trained on tremendous textual corpora, these models perceive speech using an encoder with adaptors, generating the capability to understand speech flow through supervised fine-tuning, even producing speech response directly with a vocoder. The most promising application of generative speech LLMs lies in the voice assistant. It is proposed to respond to the user's vocal instructions with vivid and spontaneous speech, enabling accurate and efficient interactions. One of most powerful LLM, GPT-4o [1], has already supported speech modality interaction, and then this ability has spread to open-sourced models, such as Mini-Omni [2], LLaMA-Omni [3] and Moshi [4], showing the promising path for speech-central human-machine interaction.

The lack of an evaluation system is one of the main factors limiting the development of generative speech LLMs. Previous work has predominantly adopted benchmarks for the model's understanding capabilities. Dynamic SUPERB provides a platform to compare various models on plenty of speech-processing tasks [5, 6]. AudioBench and AIR-Bench focus on the performance of understanding various audio signals including human speech and natural sound [7, 8]. VoiceBench measures voice assistants for their capability of general knowledge, instruction

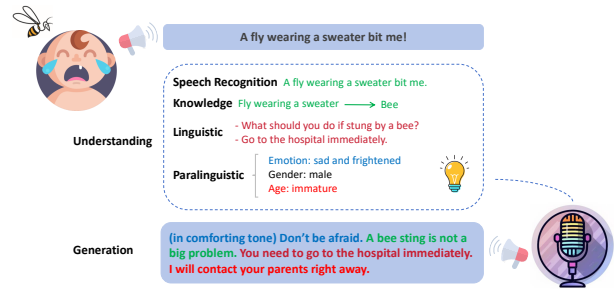


Figure 1: Properties for an ideal voice assistant. The color of the generated response corresponds to distinct aspects of the user's query. For example, if the model receives emotions of sadness and fear, a comforting tone should usually be chosen.

following, robustness and safety alignment [9]. However, the above benchmarks either only measure the model's ability to produce text responses or transform the speech response via an automatic speech recognition (ASR) system. They assess understanding or text-based transformations, but fall short in evaluating the quality of the voice output itself—an essential factor for voice assistants and other interactive systems.

To address this gap, we introduce an exhaustive benchmark, **Speech cOnversational Voice Assistant Benchmark (SOVA-Bench)**, to quantify the performance of generative speech LLMs, especially voice assistants. As shown in Figure 1, an ideal voice assistant should respond to user queries using preserved general knowledge based on instruction content, linguistic, and paralinguistic information. SOVA-Bench evaluates the models with distinct subsets. It ingests raw audio from multiple datasets and presents in the format of question-answering (QA) and instruction querying. By encompassing both understanding and generation, including measures for tone, emotion, and other paralinguistic cues, this method provides a more comprehensive evaluation framework that mirrors real-world communication requirements. Ultimately, it addresses a significant gap in the field, paving the way for the development of advanced generative speech models that deliver not only accurate information but also a reliable, engaging, and human-like auditory experience. Our main contributions are as follows:

- We propose SOVA-Bench, an evaluation system for generative speech LLMs and voice assistants, quantifying the performance of general knowledge, as well as the ability to recognize, understand and generate speech flow.
- SOVA-Bench compares multiple generative speech models under the same evaluation framework. Based on the comparison experiments, we identify the limitations of existing voice interaction models.
- SOVA-Bench provides a standard for the evaluation of voice

<sup>†</sup>These authors contributed equally.

\*Corresponding Authors

Table 1: Comparison with other available speech LLM evaluation benchmarks.

Benchmark	Num of Tasks	Knowledge	Recognition	Understanding		Generation	Evaluation Subject
				Linguistic	Paralinguistic		
Dynamic SUPERB Phase-2	180	✗	✓	✓	✓	✗	Universal Speech Models
AudioBench	8	✗	✓	✓	✓	✗	AudioLLMs
AIR-Bench	20	✗	✓	✓	✓	✗	AudioLLMs
VoiceBench	6	✓	✗	✓	✗	✗	LLM-based Voice Assistants
OpenAudioBench	3	✓	✗	✓	✗	✗	Audio-focused LLMs
<b>SOVA-Bench (Ours)</b>	8	✓	✓	✓	✓	✓	Generative Speech LLMs

interactive systems and facilitates subsequent researchers to compare model performance.

## 2. Related Works

### 2.1. Speech LLM

Speech LLM extends the understanding capability to speech flow, performing modality alignment between speech and text via an encoder with adaptors. Earlier work focused on generating textual responses from speech instructions, optionally adopting a TTS module for a cascaded speech-in-speech-out paradigm, such as SALMONN [10], SpeechGPT [11] and Qwen2-audio [12]. These models preserve the ability to understand speech signals while show a significant difference from human-like real-time conversations and voice assistants. The real-time API introduced by GPT-4o supports end-to-end speech generation without an explicit TTS module. It can produce text and speech responses at the same time, enabling fluent and efficient communication with users. By utilizing speech vocoder methods, several open-sourced models have also achieved speech response from discrete acoustic tokens. Mini-Omni [2] adopted SNAC [13] to quantify the model’s response speech during training with a transformer-based streaming audio decoder to covert the predicted embedding to speech flow. LLaMA-Omni [3] alternatively discrete continuous speech to units using pre-trained HuBERT [14], and a HiFi-GAN vocoder [15, 16] for unit-to-speech. Moshi [4] proposed a novel speech encoder, Mimi, incorporating both semantic and acoustic speech embeddings for information extraction and generative speech wave construction. Other speech LLMs that enable direct speech response include Westlake-Omni [17], Freeze-Omni [18], SyncLLM [19], OmniFlatten [20], etc. These models mainly follow a similar schema - generate discrete audio tokens using LLM before converting to speech signals. The generative speech LLMs support real-time communications, showing great potential for voice assistants.

### 2.2. Speech LLM evaluation

The evaluation of speech models is a contentious issue. For specialized models, the comparison between reference and output is widely adopted, such as automatic speech recognition (ASR) and speech emotion recognition (SER) models. For speech LLMs with multiple capabilities, comprehensive evaluation benchmarks have been proposed to systematically assess a model’s performance, as shown in Table 1. Dynamic SUPERB is a universal speech processing benchmark consisting of 33 tasks to evaluate speech models in a zero-shot fashion, and then extend to 180 tasks in the 2nd phase [5, 6]. AudioBench presents a set of measurements targeting at linguistic and paralinguistic speech understanding, and audio scene recognition [7]. AIR-Bench encompasses a series of tasks to measure the performance of understanding human speech, natural sound, and music [8]. It preserves two dimensions - the foundation

benchmark for specialized ability and the chat benchmark for general performance. These two benchmarks are originally designed for Audio LLMs. VoiceBench aims at a multifaceted evaluation for LLM-based voice assistants [9]. It distills proper instances to access the model’s capability of general knowledge, instruction following and safety alignment, along with the robustness of speaker, environment and content variation. OpenAudioBench is designed to assess the capabilities of multimodal and audio-focused language models, consisting of subsets of logical reasoning, general knowledge and open-ended questions [21]. The above-mentioned benchmarks show the same limitation, for they all concentrate on the understanding capability while ignoring the generative performance.

## 3. SOVA-Bench

As shown in Table 2, SOVA-Bench aims to comprehensively evaluate and compare speech LLMs’ abilities to perceive instructions and generate responses both in speech formats, therefore contributing to the evolution of voice assistants.

### 3.1. Voice Assistant Properties

Voice assistants are intelligent systems designed to enhance human-computer interaction through natural speech processing, understanding and generation. An ideal voice assistant shall preserve such properties:

- **General Knowledge:** Voice assistants accumulate massive and diverse general knowledge far beyond the level of normal human mastery. When the user proposes an inquiry, the system responds accurately with high credit and confidence.
- **Speech Recognition:** Voice assistants enable speech recognition, which is not only a common task benefiting user experience but also contributes to speech understanding.
- **Speech Understanding:** Voice assistants should understand speech in both linguistic and paralinguistic dimensions. For the former part, the system distills the semantic information and responds accordingly. For the paralinguistic part, it should generate messages such as the user’s identification and emotion, therefore producing personalized responses.
- **Speech Generation:** Voice assistants generate vivid and spontaneous speech responses when being evoked. The generated speech should preserve high accuracy, fluency, and clarity, and be consistent with its text transcription. The speech response quality should be evaluated from the perspective of both semantic and acoustic.

### 3.2. Source Dataset

**General knowledge:** To evaluate the general knowledge performance of speech LLMs, SOVA-Bench incorporates TriviaQA, an open-domain question-answering (QA) dataset with multiple reference resources [22]. We utilize the unfiltered dev set and remove the reference to form a knowledge QA dataset.

Table 2: Statistics for SOVA-Bench

Dimension	Task	Source Data	Synthesized	Format	Input Modality		Num of Instances	Evaluation	
					Text	Speech		Method	Metrics
Knowledge	Knowledge QA	TriviaQA	✓	Open-domain	✗	✓	6840	GPTEval	Accuracy
Recognition	Speech Recognition	LibriSpeech	✗		✓	✓		Normal	WER
Linguistic	Spoken QA	LibriSQA	✗	Multi-choice (4)	✓	✓	2619	Normal	Accuracy
		Spoken SQuAD	✓	Open-domain	✓	✓	3243	GPTEval	Accuracy
Paralinguistic	Emotion Recognition	IEMOCAP	✗	Multi-choice (5)	✓	✓	5819	GPTEval	Accuracy
	Gender Recognition	Common Voice	✗	Multi choice (2)	✓	✓	2469	GPTEval	Accuracy
	Age Recognition	Common Voice	✗	Multi-choice (5)	✓	✓	2511	GPTEval	Accuracy
Generation	Consistency	Alpaca	✓	Open-domain	✗	✓	3745	Normal	WER
	Semantic	Alpaca	✓	Open-domain	✗	✓	3745	GPTEval	GPTScore
	Acoustic	Alpaca	✓	Open-domain	✗	✓	3745	UTMOSv2	Score

**Speech recognition:** For ASR performance, we evaluate the model’s performance with LibriSpeech [23]. All of the four evaluation sets are tested for recognition accuracy.

**Speech understanding:** To measure the semantic understanding of speech information, SOVA-Bench consists of two formats of spoken QA - LibriSQA for multi-choice form and Spoken SQuAD for open-ended test [24, 25]. LibriSQA is developed based on LibriSpeech with the help of ChatGPT, sharing the same speech content. The comparison between the ASR and SQA performance can reveal the inner process of models to perceive and understand speech semantic information. In our benchmark, we include the second part of LibriSQA to form a simpler spoken language understanding task. Spoken SQuAD is an open-domain QA dataset, while the right answer is included in the speech transcription. We preserve a single question for each valid sentence. For those with multiple answers, we preserve the most complete one if they are in an inclusion relationship, otherwise all answers are considered correct. For paralinguistic understanding, SOVA-Bench includes emotion recognition, gender recognition and age prediction. These tasks cover the necessary user information required by voice assistants. SOVA-Bench incorporates IEMOCAP for emotion recognition [26]. The original source includes 10 emotions, and we delete neutral kinds and merge indistinguishable categories to form a 5-choice selection task: angry, happy, fearful, frustrated and sad, excited and surprised. For gender and age recognition, we use the latest version of Common Voice [27]. Only those with speaker information annotations are preserved.

**Speech generation:** SOVA-Bench is developed for voice assistants, following an instruction-inquiry scenery in the generation performance evaluation. We use a cleaned version of Alpaca [28]. We select a small portion and ensure the preserved instances without additional input. This part includes user questions and instructions. The speech generation performance is evaluated across 3 dimensions: the consistency between text and speech response, and the semantic and acoustic quality. The consistency is measured by using text response as transcription and comparing with speech output. The semantic performance is evaluated with GPT, and we implement one advanced MOS predictor, UTMOSv2 [29], to quantify the acoustic score.

It should be noted that neither TriviaQA nor Alpaca has speech audio. We implement an advanced TTS model, CosyVoice [30], to synthesize speech audio in a zero-shot fashion. The speech prompt is derived from Common Voice, and we select instances with speaker labels and ensure the gender-ratio to 1:1. The synthesized speech is processed by Whisper-large-v3 [31] to compare with the original transcription, and only those with high accuracy are preserved. As illustrated in Figure 2, this synthesized method ensures the diversity of speak-




Figure 2: Creation pipeline for dataset without speech modality. A multi-speaker prompt set is adopted for zero-shot TTS.

ers and acoustic environments, thereby more robustly evaluating the performance.

## 4. Evaluation Experiment

### 4.1. Evaluated Models

We have tested various widely used voice assistants on SOVA-bench, as shown in Table 3, including: Mini-Omni [2], Mini-Omni2 [32], LLaMA-Omni [3], Freeze-Omni [18], Moshi [4], and GLM-4-Voice [33]. Additionally, we developed a cascade model for comparative evaluation, which combines Whisper-large-v3 as the ASR model and GPT-4o-mini as the LLM. Among these models, only LLaMA-Omni supports synchronous inputs of both text and speech modalities, while Mini-Omni and Mini-Omni2 support either text or speech at a time.

Table 3: The summarization of the tested speech LLMs. The ‘Modality’ column represents the input modalities, where ‘T’ and ‘S’ represent text and speech respectively.

Speech LLM	Architecture		Modalities
	Speech Encoder	Base LLM	
Mini-Omni	Whisper-small	Qwen2-0.5B	T ∨ S
Mini-Omni2	Whisper-small	Qwen2-0.5B	T ∨ S
Moshi	Mimi	Helium-7B	S
LLaMA-Omni	Whisper-large-v3	LLaMA-3.1-8B-Instruct	T ∧ S
Freeze-Omni	CNN + Transformer	Qwen2-7B-Instruct	S
GLM-4-Voice	VQ + Whisper	GLM-4-9B	S
Cascade	Whisper-large-v3	GPT-4o-mini	S

It is important to note that for Moshi and GLM-4-Voice, we randomly sampled 500 instances from each subset. Since Moshi typically begins a conversation with an introductory statement (e.g. “Hey, how is everything going?”), we include 2.5 seconds of blank audio before each input to avoid confusion.

Table 4: Evaluation performance for speech LLMs on SOVA-Bench

Model	LLM size	Knowledge $\uparrow$	Recognition $\downarrow$		Understanding $\uparrow$					Generation $\uparrow$		
			dev_clean/other	test_clean/other	LibriSQA	Spoken SQuAD	Emotion	Age	Gender	Consistency	Semantic	Acoustic
Mini-Omni	0.5B	2.42	5.18 / 12.03	8.68 / 12.75	-	-	28.53	11.79	27.22	10.04	1.44	3.62
Mini-Omni2	0.5B	0.37	7.54 / 11.58	3.99 / 13.44	-	-	24.52	16.05	16.73	25.92	1.68	<b>3.66</b>
Moshi	7B	4.00	-	-	-	-	-	-	-	5.94	1.57	2.43
LLaMA-Omni	8B	24.91	>100	>100	<b>50.67</b>	<b>43.07</b>	37.74	19.43	<b>78.41</b>	<b>4.71</b>	<b>3.28</b>	3.61
Freeze-Omni	7B	<b>27.27</b>	<b>3.29 / 7.40</b>	3.24 / 7.68	-	-	<b>38.31</b>	17.80	39.13	10.96	3.20	3.46
GLM-4-Voice	7B	25.58	-	<b>2.82 / 7.66</b>	-	-	30.16	<b>37.84</b>	69.40	10.02	3.25	3.09
Cascade	-	<b>58.98</b>	<b>1.92 / 3.77</b>	<b>1.81 / 3.62</b>	<b>81.11</b>	<b>93.05</b>	-	-	-	-	<b>3.88</b>	-




Figure 3: Performance for paralinguistic understanding.

#### 4.2. Main Results

The overall performance of tested speech LLMs is presented in Table 4. We observe considerable variation in performance, with no single model consistently outperforming across all tasks.

Speech models with larger LLM parameters (7B, 8B) usually outperform smaller models (0.5B) on knowledge, speech understanding and the semantic quality of generated responses. These tasks largely depend on the capabilities preserved from its LLM backbone. The cascade system preserves 58.98% accuracy on knowledge QA, while the best speech LLM, Freeze-Omni, achieves 27.27%. LLaMA-Omni is the only model supporting the simultaneous input of both text and speech, but exhibits a performance deficit of 30% and 50% in both spoken QA formats compared to the cascade method. A strong performance deterioration is observed when extending to speech modality. Moshi achieves an inferior score, as its streaming modeling results in frequent interruptions and irrelevant responses.

Most speech LLMs demonstrate comparable performance in speech recognition. The results of Freeze-Omni and GLM-4-Voice are borrowed from their papers, as the intermediate model and specific prompts are not available. LLaMA-Omni could not perform speech recognition, for it does not involve pre-training on the ASR task, resulting in irrelevant responses given ASR instructions. Influenced by the capability of instruction following, the paralinguistic part of speech understanding tasks is also affected. Besides the overall accuracy, as shown in Figure 3, we simultaneously evaluated the model’s answer rate (AR) and conditioned accuracy (CA), with the latter representing the proportion of correct responses among the answered samples. LLaMA-Omni consistently exhibits high performance in instruction following; however, the best model in terms of conditioned accuracy is not the same across the three tasks. Most speech LLMs lack sufficient capability to capture specific paralinguistic information. LLaMA-Omni, Freeze-Omni, and GLM-4-Voice demonstrated clearly above-chance accuracy in gender prediction, while Mini-Omni and GLM-4-Voice showed

notable performance in age prediction. Emotion recognition presents a special case for evaluation, as it is also significantly influenced by semantic content. All models demonstrate positive performance on the emotion recognition task, with Mini-Omni correctly identifying nearly half of the emotional labels.

Regarding the quality of generated speech, Moshi and LLaMA-Omni exhibit strong speech-text consistency. In contrast, Mini-omni2 demonstrates the poorest performance, primarily due to premature truncation occurring during the speech generation process. LLaMA-Omni and GLM-4-Voice exhibit the highest performance in terms of the semantic quality of generated responses. Theoretically, a larger parameter LLM backbone is conducive to generating higher-quality responses, while Moshi once again received a lower score due to an excessive amount of irrelevant output. As for the acoustic quality, Mini-Omni2 achieves the highest score, a little higher than Mini-Omni and LLaMA-Omni. The acoustic performance is affected by the speech token modeling method.

#### 4.3. Limitations for current speech LLMs

Based on the experimental results from SOVA-Bench, we summarize the current limitations of speech LLMs as follows:

- Limited input modalities: Only LLAMA-Omni supports simultaneous input of both speech and text. However, in voice assistant scenarios, additional text information may be necessary for optimal performance.
- Instruction following variability: Most models typically struggle to effectively follow instructions, particularly when dealing with tasks not involved in their training process.
- Severe performance degradation: Introducing speech modalities into LLMs often leads to a marked decline in general knowledge, semantic understanding, and generation capabilities. A critical challenge is to preserve the performance of textual LLMs while supporting speech interactions.
- Unsatisfactory output speech quality: Some speech LLMs exhibit limited speech-text consistency and speech quality. These issues are not only related to the LLM used but also to the speech tokenization and training schemes.

### 5. Conclusion

In this paper, we introduce a comprehensive benchmark, SOVA-Bench, for evaluating generative speech LLMs from both semantic and acoustic perspectives. Evaluation experiments on novel speech-interactive models reveal variations in the quality of generated speech responses and a substantial degradation in performance compared to the cascade model, particularly in general knowledge and speech understanding. Considering the deficiencies in unified evaluation standards within the field of speech LLMs, especially voice assistants, SOVA-Bench provides a systematic framework for comparing model performance.

## 6. Acknowledgements

This work was supported by the Shanghai Jiao Tong University - Ant Group Intelligent Multimedia Joint Laboratory.

## 7. References

- [1] A. Hurst, A. Lerer, A. P. Goucher, A. Perelman, A. Ramesh, A. Clark, A. Ostrow, A. Welihinda, A. Hayes, A. Radford *et al.*, “GPT-4o system card,” *arXiv preprint arXiv:2410.21276*, 2024.
- [2] Z. Xie and C. Wu, “Mini-Omni: Language models can hear, talk while thinking in streaming,” *arXiv preprint arXiv:2408.16725*, 2024.
- [3] Q. Fang, S. Guo, Y. Zhou, Z. Ma, S. Zhang, and Y. Feng, “LLaMA-Omni: Seamless speech interaction with large language models,” *arXiv preprint arXiv:2409.06666*, 2024.
- [4] A. Défossez, L. Mazaré, M. Orsini, A. Royer, P. Pérez, H. Jégou, E. Grave, and N. Zeghidour, “Moshi: a speech-text foundation model for real-time dialogue,” *arXiv preprint arXiv:2410.00037*, 2024.
- [5] C.-y. Huang, K.-H. Lu, S.-H. Wang, C.-Y. Hsiao, C.-Y. Kuan, H. Wu, S. Arora, K.-W. Chang, J. Shi, Y. Peng *et al.*, “Dynamic-SUPERB: Towards a dynamic, collaborative, and comprehensive instruction-tuning benchmark for speech,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 12 136–12 140.
- [6] C.-y. Huang, W.-C. Chen, S.-w. Yang, A. T. Liu, C.-A. Li, Y.-X. Lin, W.-C. Tseng, A. Diwan, Y.-J. Shih, J. Shi *et al.*, “Dynamic-SUPERB Phase-2: A collaboratively expanding benchmark for measuring the capabilities of spoken language models with 180 tasks,” *arXiv preprint arXiv:2411.05361*, 2024.
- [7] B. Wang, X. Zou, G. Lin, S. Sun, Z. Liu, W. Zhang, Z. Liu, A. Aw, and N. F. Chen, “AudioBench: A universal benchmark for audio large language models,” *arXiv preprint arXiv:2406.16020*, 2024.
- [8] Q. Yang, J. Xu, W. Liu, Y. Chu, Z. Jiang, X. Zhou, Y. Leng, Y. Lv, Z. Zhao, C. Zhou *et al.*, “AIR-Bench: Benchmarking large audio-language models via generative comprehension,” *arXiv preprint arXiv:2402.07729*, 2024.
- [9] Y. Chen, X. Yue, C. Zhang, X. Gao, R. T. Tan, and H. Li, “VoiceBench: Benchmarking llm-based voice assistants,” *arXiv preprint arXiv:2410.17196*, 2024.
- [10] C. Tang, W. Yu, G. Sun, X. Chen, T. Tan, W. Li, L. Lu, Z. Ma, and C. Zhang, “SALMONN: Towards generic hearing abilities for large language models,” in *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*.
- [11] D. Zhang, S. Li, X. Zhang, J. Zhan, P. Wang, Y. Zhou, and X. Qiu, “SpeechGPT: Empowering large language models with intrinsic cross-modal conversational abilities,” in *Findings of the Association for Computational Linguistics: EMNLP 2023*, 2023, pp. 15 757–15 773.
- [12] Y. Chu, J. Xu, Q. Yang, H. Wei, X. Wei, Z. Guo, Y. Leng, Y. Lv, J. He, J. Lin *et al.*, “Qwen2-audio technical report,” *arXiv preprint arXiv:2407.10759*, 2024.
- [13] H. Siuzdak, F. Grötschla, and L. A. Lanzendörfer, “SNAC: Multi-scale neural audio codec,” in *Audio Imagination: NeurIPS 2024 Workshop AI-Driven Speech, Music, and Sound Generation*.
- [14] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, “Hubert: Self-supervised speech representation learning by masked prediction of hidden units,” *IEEE/ACM transactions on audio, speech, and language processing*, vol. 29, pp. 3451–3460, 2021.
- [15] J. Kong, J. Kim, and J. Bae, “HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis,” *Advances in neural information processing systems*, vol. 33, pp. 17 022–17 033, 2020.
- [16] A. Polyak, Y. Adi, J. Copet, E. Kharitonov, K. Lakhotia, W.-N. Hsu, A. Mohamed, and E. Dupoux, “Speech resynthesis from discrete disentangled self-supervised representations,” in *INTER-SPEECH 2021-Annual Conference of the International Speech Communication Association*, 2021.
- [17] Xinchen-AI, “Westlake-Omni,” <https://github.com/xinchen-ai/Westlake-Omni>, 2024.
- [18] X. Wang, Y. Li, C. Fu, L. Xie, K. Li, X. Sun, and L. Ma, “Freeze-Omni: A smart and low latency speech-to-speech dialogue model with frozen llm,” *arXiv preprint arXiv:2411.00774*, 2024.
- [19] B. Veluri, B. Peloquin, B. Yu, H. Gong, and S. Gollakota, “Beyond turn-based interfaces: Synchronous llms as full-duplex dialogue agents,” in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, 2024, pp. 21 390–21 402.
- [20] Q. Zhang, L. Cheng, C. Deng, Q. Chen, W. Wang, S. Zheng, J. Liu, H. Yu, and C. Tan, “OmniFlatten: An end-to-end GPT model for seamless voice conversation,” *arXiv preprint arXiv:2410.17799*, 2024.
- [21] Y. Li, J. Liu, T. Zhang, S. Chen, T. Li, Z. Li, L. Liu, L. Ming, G. Dong, D. Pan *et al.*, “Baichuan-Omni-1.5 technical report,” *arXiv preprint arXiv:2501.15368*, 2025.
- [22] M. Joshi, E. Choi, D. S. Weld, and L. Zettlemoyer, “TriviaQA: A large scale distantly supervised challenge dataset for reading comprehension,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2017, pp. 1601–1611.
- [23] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, “LibriSpeech: an asr corpus based on public domain audio books,” in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [24] Z. Zhao, Y. Jiang, H. Liu, Y. Wang, and Y. Wang, “LibriSQA: A novel dataset and framework for spoken question answering with large language models,” *IEEE Transactions on Artificial Intelligence*, 2024.
- [25] C.-H. Lee, S.-L. Wu, C.-L. Liu, and H. yi Lee, “Spoken SQUAD: A study of mitigating the impact of speech recognition errors on listening comprehension,” in *Interspeech 2018*, 2018, pp. 3459–3463.
- [26] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, “IEMOCAP: Interactive emotional dyadic motion capture database,” *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.
- [27] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. Tyers, and G. Weber, “Common Voice: A massively-multilingual speech corpus,” in *Proceedings of the Twelfth Language Resources and Evaluation Conference*, 2020, pp. 4218–4222.
- [28] R. Taori, I. Gulrajani, T. Zhang, Y. Dubois, X. Li, C. Guestrin, P. Liang, and T. B. Hashimoto, “Stanford alpaca: an instruction-following llama model (2023),” *URL [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca)*, vol. 1, no. 9, 2023.
- [29] K. Baba, W. Nakata, Y. Saito, and H. Saruwatari, “The t05 system for the voicemos challenge 2024: Transfer learning from deep image classifier to naturalness mos prediction of high-quality synthetic speech,” *arXiv preprint arXiv:2409.09305*, 2024.
- [30] Z. Du, Q. Chen, S. Zhang, K. Hu, H. Lu, Y. Yang, H. Hu, S. Zheng, Y. Gu, Z. Ma *et al.*, “Cosyvoice: A scalable multilingual zero-shot text-to-speech synthesizer based on supervised semantic tokens,” *arXiv preprint arXiv:2407.05407*, 2024.
- [31] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [32] Z. Xie and C. Wu, “Mini-Omni2: Towards open-source GPT-4o with vision, speech and duplex capabilities,” *arXiv preprint arXiv:2410.11190*, 2024.
- [33] A. Zeng, Z. Du, M. Liu, K. Wang, S. Jiang, L. Zhao, Y. Dong, and J. Tang, “GLM-4-Voice: Towards intelligent and human-like end-to-end spoken chatbot,” *arXiv preprint arXiv:2412.02612*, 2024.