



DeSTA: Enhancing Speech Language Models through Descriptive Speech-Text Alignment

Ke-Han Lu¹, Zhehuai Chen², Szu-Wei Fu², He Huang², Boris Ginsburg², Yu-Chiang Frank Wang², Hung-yi Lee¹

¹Graduate Institute of Communication Engineering, National Taiwan University, Taiwan
²NVIDIA

{d12942024, hungyilee}@ntu.edu.tw, {zhehuaic, szuweif, heh, bginsburg, frankwang}@nvidia.com

Abstract

Recent speech language models (SLMs) typically incorporate pre-trained speech models to extend the capabilities from large language models (LLMs). In this paper, we propose a **Descriptive Speech-Text Alignment** approach that leverages *speech captioning* to bridge the gap between speech and text modalities, enabling SLMs to interpret and generate comprehensive natural language descriptions, thereby facilitating the capability to understand both linguistic and non-linguistic features in speech. Enhanced with the proposed approach, our model demonstrates superior performance on the Dynamic-SUPERB benchmark, particularly in generalizing to unseen tasks. Moreover, we discover that the aligned model exhibits a zero-shot instruction-following capability without explicit speech instruction tuning. These findings highlight the potential to reshape instruction-following SLMs by incorporating rich, descriptive speech captions.¹

Index Terms: speech language model, instruction tuning, speech caption

1. Introduction

In recent years, the advent and continuous evolution of large language models (LLMs) have revolutionized the landscape of natural language processing (NLP), demonstrating remarkable performance across a diverse array of text generation and understanding tasks [1–5]. The effectiveness of LLMs is significantly enhanced through the process of instruction tuning, which equips these models with the flexibility to adapt to novel tasks by following specific instructions [6, 7].

Motivated by the success of LLMs, recent studies have begun exploring the potential of LLMs in the realm of speech processing [8–17]. These versatile instruction-following speech language models (SLMs) are designed to comprehend textual instructions and perform specific speech processing tasks. As depicted in Figure 1, typically, these models incorporate a pre-trained speech model and instruction-following LLMs as fundamental architecture. Subsequently, these models undergo instruction-tuning across various speech tasks, aiming to harmonize the capabilities of the speech model with the text-based LLM. Nevertheless, this process of aligning speech and text models presents multiple challenges. First, the necessity for curators to explicitly define the task scope is very difficult to standardize when assigning instructions across different tasks. It can also lead to overfitting problems [10] to the training tasks and diminish the emergent capabilities of LLMs. Second, the training targets in current instruction-tuning datasets are usually

¹<https://github.com/kehanlu/Nemo/tree/desta/examples/multimodal/DeSTA>

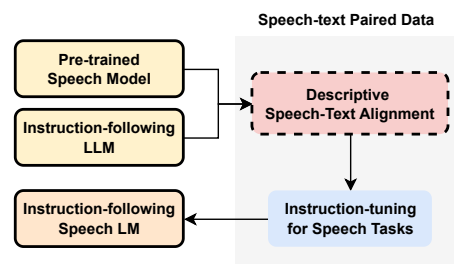


Figure 1: *The Descriptive Speech-Text Alignment framework.*

formulated to perform one task at a time, which does not consider the multifaceted nature of speech. Finally, these datasets are usually formulated as classification tasks requiring only the model to predict one option [18]. This process neglects the inherent long-form generation capability of LLM.

Figure 1 illustrates an overall framework in this paper. To address the highlighted issues, we present a novel descriptive speech-text alignment (DeSTA) stage before instruction tuning phase for SLMs. The speech-text-alignment stage aims at bridging the modality gap between speech and text through *speech captioning*, thereby enables the broad understanding of speech. In this stage, we collect a speech caption dataset that features long-form natural language descriptions that encapsulate the multi-dimensional aspects of speech, by leveraging the metadata from existing datasets. This method mirrors the human ability to interpret and integrate multiple facets of speech into a comprehensive description. For example, “*The woman shouted “I love cats” with joy, her voice loud, quick, and high-pitched...*” which captures not only the spoken words but also the speaking styles conveyed. To access the effectiveness of our speech-text alignment approach, we conduct instruction-tuning on our pre-trained model and evaluate our model on the Dynamic-SUPERB [18] benchmark.

We summarize our contributions as follows.

- We introduce a descriptive speech-text alignment training approach designed to train SLMs on both linguistic and non-linguistic information through a comprehensive speech caption. We will release our implementation and speech caption dataset in the future.
- With the proposed speech-text alignment training, our fine-tuned models outperform baselines systems on the Dynamic-SUPERB benchmark, particularly in tasks not covered during the training phase. This highlights the generalizability of our proposed framework.
- We discover that, although the pre-trained model is trained for speech captioning, it demonstrates zero-shot instruction-

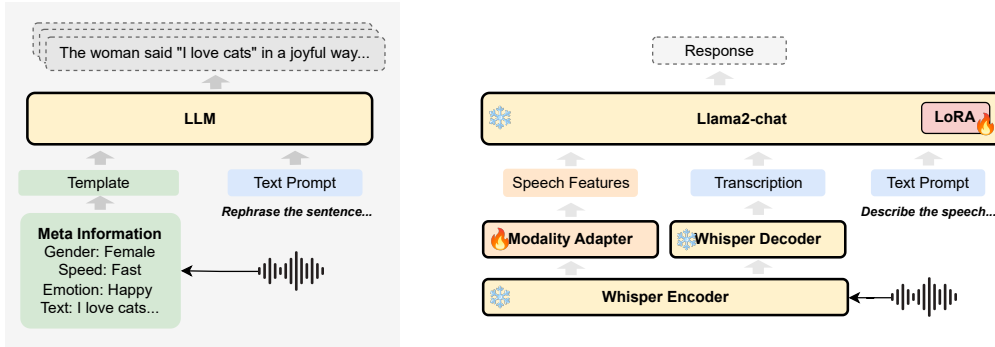


Figure 2: (Left) Data generation pipeline. (Right) Model architecture for speech-text alignment training.

following capabilities derived from LLM by utilizing LoRA-scaling [10] at testing time.

2. Related work

Recently, there has been growing interest in integrating speech models with LLM to enable instruction-following for multitask speech processing. To this end, many recent innovations focus on creating speech instruction dataset that features a wide range of speech processing tasks for the instruction-tuning stage on Figure 1. A common approach is to leverage the power of LLMs to automatically curate diverse instructions for speech tasks [8–10, 12]. Although LLMs can not understand speech directly, it can understand and interpret the text description with strong knowledge of speech [8]. For example, LTU [8, 9] generate large amounts of open-ended question-answering pairs from GPT. On the other hand, several studies decompose the learning process into multiple stage to enhance the learning efficiency and effectiveness [10, 19, 20]. For instance, SALMONN [10] presents a pre-training stage using speech recognition and audio captioning data and Qwen-audio [19] conduct multitask pre-training before instruction finetuning.

3. Descriptive Speech-Text Alignment

In this work, we introduce a descriptive speech-text alignment training stage that aims at training the SLM to generate comprehensive speech captions. As demonstrate on Figure 2(Left), we utilize a large language model to generate speech captions with metadata from the audio. Then the SLM learns the multifaceted speech concept from the caption data accordingly.

3.1. Speech caption

The primary objective of creating a speech caption dataset is to accurately capture the complex nature of speech and translate it into comprehensive, natural language descriptions. As demonstrated in Figure 2(Left), following the methodologies of [8, 9], we collect meta-information from existing datasets. This meta-information includes attributes such as speaking style (e.g., pitch, volume, and speaking speed), speaker information (e.g., gender), and the actual spoken content. We then empirically curate initial templates based on these attributes and formulate them into multiple natural language sentences. For example, a template might be structured as: "A [gender] speaker says [text] with [emotion] emotion." Next, we employ a LLM to generate diverse captions that reflect various writing

styles and tones based on the provided sentences. Specifically, through prompt engineering, we instruct the LLM to accurately reflect the original spoken content while creatively incorporating the speech attributes to avoid hallucination. For each audio, we generate several captions based on different templates and prompts to ensure the dataset represents a wide range of expressiveness while avoiding repetition.

3.2. Model architecture

As depicted in Figure 2(Right), our architecture incorporates a pre-trained Whisper model [21] alongside an instruction-following Llama2-chat model [4]. Throughout the training phase, these pre-trained models remain unchanged. A randomly initialized modality adapter is employed to map the speech feature from the Whisper encoder into the input representation space of Llama. Additionally, the transcribed text is fed into the language model as supplementary input. Low-rank adapters [22] are attached to the LLM to enhance training efficiency.

Modality adapter The modality adapter is designed to extract meaningful representations from speech inputs. Specifically, the adapter processes the hidden outputs from the intermediate layers of the Whisper encoder to obtain high-level speech features. In this work, we explore two architectures: CNN and Qformer [23]. While the CNN preserves temporal information from the speech, the Qformer provides more flexibility with its attention mechanism. Next, these layer-wise representations are combined through a weighted summation using learnable weights to obtain the final representations. Finally, a projection layer is employed to map the continuous representations into the embedding space of Llama.

Large language model The speech features from the modality adapter and the transcribed text are concatenated with text prompts. In the speech-text alignment phase, we use a set of prompts for speech captioning, such as "Describe the speech" and "What can be inferred from this audio?", to maintain the input structure of instruction tuning. As a result, the model is trained to generate speech captions with the next-token-prediction loss in this stage.

4. Experiment

4.1. Speech caption dataset

In our study, we created a speech caption dataset by combining the LibriTTS [24], IEMOCAP [25], and PromptTTS [26] datasets, all known for their expressive speech attributes. These

Table 1: *Dynamic-SUPERB* results. STA denotes the model is pre-trained with speech-text alignment. †We specify the seen and unseen tasks respect to the *Dynamic-SUPERB* training and evaluation set, while ASR+ChatGPT does not require training.

Model	STA	Seen						Unseen					All	
		CON 9	SEM 2	PAR 4	DEG 6	SPK 3	Avg 24	CON 2	SEM 4	PAR 3	DEG 13	SPK 2	Avg 24	Avg 48
ASR+ChatGPT†		67.11	43.75	38.75	47.50	35.67	52.78	50.50	71.63	5.17	40.31	47.75	42.60	47.10
ImageBind-LLM [18]		64.39	48.25	59.88	79.17	51.50	65.07	16.75	22.50	21.00	48.65	43.50	37.75	51.06
Whisper-LLM [18]		76.94	56.50	68.00	91.67	92.17	79.28	8.00	21.50	6.83	59.69	59.25	42.38	60.85
CNN		90.33	63.25	68.88	49.33	46.83	70.70	10.75	50.25	23.83	48.00	49.25	42.35	55.58
CNN	✓	95.44	65.50	63.13	50.83	57.83	73.52	71.25	73.25	11.00	50.08	45.00	50.40	61.05
Qformer		96.39	65.00	77.25	44.33	55.33	74.87	87.75	80.50	4.17	43.19	46.25	48.50	60.47
Qformer	✓	95.00	67.50	74.38	71.25	59.00	80.15	74.50	75.75	20.33	57.54	46.50	56.42	67.63

Table 2: *Statistics of speech caption dataset.*

Pretrain statics	LibriTTS	IEMOCAP	PromptTTS
# Audios	20,807	4,262	23,544
# Captions	62,309	12,781	70,439
Duration(hours)	88.3	16.2	132.6
Avg. length(tokens)	60.8	62.4	66.3

datasets include meta information such as gender, pitch, volume, speaking speed, and text transcriptions. We utilized publicly available curated metadata [9] for LibriTTS and IEMOCAP, adding emotion labels for IEMOCAP. Additionally, we enriched our dataset with the PromptTTS dataset, which features emotional speech and human-annotated style description. We employed the Zephyr-7b-beta [27]² to generate captions by randomly combining three prompts and five templates for each audio. Ultimately, as detailed in Table 2, our dataset consists of 48,613 audio clips and provides 145,529 audio-caption pairs, totaling 237.1 hours, with an average length of 60 tokens each.

4.2. Descriptive speech-text alignment

We use the NeMo toolkits [28] to implement the proposed method. We utilize publicly available checkpoints of Llama2-7b-chat³ and Whisper-large-v3⁴, which have 1.5 billion parameters, as our fundamental architecture. The LoRA adapters (rank=32) are injected into the query, key, and value projection layers of the attention mechanisms within Llama. A scaling factor α is set to control the impact of these adapters at testing time [10]. For processing speech inputs, the Whisper encoder generates 1,500 hidden representations for each audio. The Qformer architecture consists of a stack of two Transformer decoder blocks [29], coupled with 64 learnable query vectors. Similarly, for the CNN settings, two CNN layers are designed to downsample the encoder outputs into 60 representations for each audio. The speech is transcribed by the Whisper beforehand to enhance training efficiency. The total number of trainable parameters in our model is 56.3M for the Qformer architecture and 46.4M for the CNN architecture, respectively. We train the model on the speech caption dataset using the Adam optimizer with a cosine annealing scheduler for 5 epochs. We use 4 V100 GPUs and the global batch size is 12 with learning rate of 1e-4.

²<https://huggingface.co/HuggingFaceH4/zephyr-7b-beta>

³<https://huggingface.co/meta-llama/Llama-2-7b-chat>

⁴<https://huggingface.co/openai/whisper-large-v3>

4.3. Instruction tuning

The speech-text aligned model is further refined by utilizing the *Dynamic-SUPERB* [18] training set. This dataset includes a wide range of instruction-guided speech processing tasks, which are categorized into five dimensions: content (CON), semantic (SEM), paralinguistic (PAR), degradation (DEG), and speaker (SPK). The training set comprises 22 instances, resulting in a total of 107K instruction pairs. Empirical evidence indicates that tasks within the content dimension are generally straightforward for LLMs to address. Therefore, we employ random selection to adjust our training data by incorporating fewer training samples from the content dimension, leading to a dataset with 76.5K samples. Additionally, we encountered an issue with task overfitting when employing a multi-task training approach. To mitigate this, we reduced the dataset size by 80%, yielding a total of 15K data samples. Unless otherwise stated, we use this dataset configuration in our experiments. In the inference stage, we employ a reduced version of *Dynamic-SUPERB* evaluation set, containing 48 speech-related instances, each with 200 samples. We adhere to the evaluation settings to calculate the exact match accuracy for all instances. The complete task list is available at <https://github.com/dynamic-superb/dynamic-superb>.

5. Results

5.1. Results on *Dynamic-SUPERB* benchmark

Table 1 presents the results from *Dynamic-SUPERB*, categorizing them into *seen* and *unseen* categories. The ASR+ChatGPT model is a cascaded system in which ChatGPT processes text transcribed by Whisper-large-v3. Both ImageBind-LLM and Whisper-LLM are baseline systems proposed alongside the *Dynamic-SUPERB* benchmark [18]. ImageBind-LLM [30] is a multimodal language model that leverages an ImageBind [31] encoder to integrate features from audio, vision, and text into a unified representation space. On the other hand, Whisper-LLM, based on the architecture of ImageBind-LLM, replaces the encoder with a Whisper encoder. Both models are further instruction-finetuned on the *Dynamic-SUPERB* training set, enabling them to follow instructions for executing speech processing tasks.

At first glance, the speech-text aligned Qformer model surpasses existing baseline systems, achieving an overall accuracy of 67.63%. Notably, it excels in both seen categories and makes significant strides in the unseen category, where it achieves an average accuracy of 56.42%. While CNN architectures are less effective than the Qformer, they exhibit superior performance in unseen categories compared to baseline models. It is im-

Table 3: *Abalation studies on instruction-finetuning data size.*

Model	STA	Data	Seen	Unseen	All
ImageBind-LLM		107K	65.07	37.75	51.06
Whisper-LLM		107K	79.28	42.38	60.85
CNN		15K	70.70	42.35	55.58
	✓	15K	73.52	50.40	61.05
		75.6K	71.22	45.06	57.11
	✓	75.6K	74.37	48.63	61.30
Qformer		15K	74.87	48.50	60.47
	✓	15K	80.15	56.42	67.63
		75.6K	68.02	41.33	53.66
	✓	75.6K	77.00	52.42	64.83

portant to highlight that there are no tasks explicitly related to the metadata employed during the speech-text alignment phase, with the exception of two emotion-related tasks in the seen category. Additionally, compared with the results to those without speech-text alignment, enhancements are evident in both seen and unseen categories. This indicates that the proposed method is not only good at adapting to new training tasks but also maintains the generalization capabilities of Llama.

Upon comparing performance across different dimensions, we observed that the Whisper-LLM significantly outperforms performance in the DEG and SPK dimensions. In contrast, our model demonstrates enhanced performance in the CON, SEM, and PAR dimensions. This disparity may be attributed to our method of incorporating transcribed text as input to the LLM, prompting the model to process both linguistic and non-linguistic features concurrently. However, We found that this may negatively impact the performance in tasks involving overlapping or noisy speech. As a consequence, our architecture faces challenges in such environments.

5.2. Abalation studies on instruction-finetuning data size

Table 3 demonstrates the performance of Dynamic-SUPERB across different training sizes. Our model exhibits superior generalization capabilities on both seen and unseen categories when trained with 15K data points. Surprisingly, models trained with 75.6K data samples do not surpass the performance of those trained with a smaller dataset, despite the significant increase in data quantity. This phenomenon may be attributed to the structure of Dynamic-SUPERB, which is designed as a multiple-choice question task (i.e., a classification task with predefined options). Therefore, the model might readily identify patterns, which could lead to an overfitting issue, especially with specific choices such as 'yes' or 'no'.

5.3. Results on zero-shot instruction following

We discovered that our models, pre-trained with speech-text alignment prior to instruction-based fine-tuning, exhibit the ability to follow instructions in a zero-shot manner. To demonstrate this zero-shot capability, we assessed the performance of the model, which had undergone only speech-text alignment training, by evaluating it on a set of 100 randomly chosen audio samples from the PromptTTS test dataset. We tasked our Qformer model with a variety of questions related to gender, pitch, volume, and speaking speed, using both multiple-choice and yes/no formats. Motivated by the findings in [10], we scaled LoRA factors to trade-off between textual knowledge and speech-text alignment learning. Table 4 demonstrates the

Table 4: *Zero-shot instruction following results based on different LoRA-scaling factor. Success rate denotes the model follow the instruction and respond correctly, while accuracy indicates the performance when model follows the instruction.*

α	Success Rate	Accuracy	Following Rate
1.00	N/A	N/A	N/A
0.75	38.25	71.00	53.88
0.50	69.75	70.72	98.63
0.25	68.83	68.63	100.00
0.00	62.63	62.63	100.00

Table 5: *Generation examples from a sentence spoke by a female speaker in a slow pace.*

α	Response
<i>Transcription: "Of course not," said he, jumping up from the throne, where he had seated himself.</i>	
Q: What is the gender of the speaker? Answer 'male' or 'female'.	
1.00	The speaker's voice is soft and gentle,... (<i>Description</i>)
0.75	The speaker's gender is identified as female.
0.50	The speaker is a female.
0.25	Male
0.00	Male
Q: Does the speaker talk in a fast pace? Answer 'yes' or 'no'.	
1.00	The speaker's volume is soft, which...(<i>Description</i>)
0.75	The speaker's voice is soft and low, with a slow pace that allows for clear enunciation of each word.
0.50	No.
0.25	No.
0.00	No

results, and Table 5 shows the generated examples based on different LoRA-scaling factors. Initially, the original model ($\alpha = 1.0$), as expected, could not follow the instructions because it was trained for speech captioning. However, when we lowered the LoRA-scaling factor, the model began to follow instructions and responded in a more precise way. Surprisingly, even when we completely removed the LoRA adapter ($\alpha = 0.0$), the model still performed well. This shows that the proposed speech-text alignment training can effectively bring speech models to the LLM modeling space through the learned modality adapter, and reducing the LoRA-scaling factor makes the LLM behave more like its original form (i.e., instruction-following LLM). These findings imply that with the proposed speech-text alignment approach, we can utilize the instruction-following ability derived from LLMs without explicitly curating the instruction set for specific tasks by applying pruning methods during the testing time.

6. Conclusion

In conclusion, our study introduces a descriptive speech-text alignment approach by leveraging speech captioning. This method bridges the modality gap between speech and text, enabling SLMs to comprehend a wide range of speech features. The experimental results on the Dynamic-SUPERB benchmark and the zero-shot evaluation demonstrate the effectiveness of our approach, particularly in its ability to generalize to new tasks. This indicates the potential to harness the general capabilities from instruction-following LLMs with our approach.

7. References

- [1] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [2] R. Anil, A. M. Dai, O. Firat, M. Johnson, D. Lepikhin, A. Passos, S. Shakeri, E. Taropa, P. Bailey, Z. Chen *et al.*, “Palm 2 technical report,” *arXiv preprint arXiv:2305.10403*, 2023.
- [3] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [4] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale *et al.*, “Llama 2: Open foundation and fine-tuned chat models,” *arXiv preprint arXiv:2307.09288*, 2023.
- [5] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, “Qwen technical report,” *arXiv preprint arXiv:2309.16609*, 2023.
- [6] J. Wei, M. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le, “Finetuned language models are zero-shot learners,” in *International Conference on Learning Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=gEzrGCozdqR>
- [7] S. Zhang, L. Dong, X. Li, S. Zhang, X. Sun, S. Wang, J. Li, R. Hu, T. Zhang, F. Wu *et al.*, “Instruction tuning for large language models: A survey,” *arXiv preprint arXiv:2308.10792*, 2023.
- [8] Y. Gong, H. Luo, A. H. Liu, L. Karlinsky, and J. Glass, “Listen, think, and understand,” *arXiv preprint arXiv:2305.10790*, 2023.
- [9] Y. Gong, A. H. Liu, H. Luo, L. Karlinsky, and J. Glass, “Joint audio and speech understanding,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2023.
- [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,” *arXiv preprint arXiv:2310.13289*, 2023.
- [11] Z. Kong, A. Goel, R. Badlani, W. Ping, R. Valle, and B. Catanzaro, “Audio flamingo: A novel audio language model with few-shot learning and dialogue abilities,” *arXiv preprint arXiv:2402.01831*, 2024.
- [12] Y. Shu, S. Dong, G. Chen, W. Huang, R. Zhang, D. Shi, Q. Xiang, and Y. Shi, “Llasm: Large language and speech model,” *arXiv preprint arXiv:2308.15930*, 2023.
- [13] J. Wu, Y. Gaur, Z. Chen, L. Zhou, Y. Zhu, T. Wang, J. Li, S. Liu, B. Ren, L. Liu *et al.*, “On decoder-only architecture for speech-to-text and large language model integration,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [14] S. Deshmukh, B. Elizalde, R. Singh, and H. Wang, “Pengi: An audio language model for audio tasks,” *Advances in Neural Information Processing Systems*, vol. 36, 2024.
- [15] M. Wang, W. Han, I. Shafran, Z. Wu, C.-C. Chiu, Y. Cao, N. Chen, Y. Zhang, H. Soltan, P. K. Rubenstein *et al.*, “Slm: Bridge the thin gap between speech and text foundation models,” in *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2023, pp. 1–8.
- [16] J. Pan, J. Wu, Y. Gaur, S. Sivasankaran, Z. Chen, S. Liu, and J. Li, “Cosmic: Data efficient instruction-tuning for speech in-context learning,” *arXiv preprint arXiv:2311.02248*, 2023.
- [17] R. Huang, M. Li, D. Yang, J. Shi, X. Chang, Z. Ye, Y. Wu, Z. Hong, J. Huang, J. Liu *et al.*, “Audiogpt: Understanding and generating speech, music, sound, and talking head,” *arXiv preprint arXiv:2304.12995*, 2023.
- [18] 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,” *arXiv preprint arXiv:2309.09510*, 2023.
- [19] Y. Chu, J. Xu, X. Zhou, Q. Yang, S. Zhang, Z. Yan, C. Zhou, and J. Zhou, “Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models,” *arXiv preprint arXiv:2311.07919*, 2023.
- [20] S. Liu, A. S. Hussain, C. Sun, and Y. Shan, “Music understanding llama: Advancing text-to-music generation with question answering and captioning,” *arXiv preprint arXiv:2308.11276*, 2023.
- [21] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, “Robust speech recognition via large-scale weak supervision,” 2022.
- [22] E. J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen *et al.*, “Lora: Low-rank adaptation of large language models,” in *International Conference on Learning Representations*, 2021.
- [23] J. Li, D. Li, S. Savarese, and S. Hoi, “BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models,” in *Proceedings of the 40th International Conference on Machine Learning*, 2023.
- [24] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, “LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech,” in *Proc. Interspeech 2019*, 2019, pp. 1526–1530.
- [25] 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.
- [26] Z. Guo, Y. Leng, Y. Wu, S. Zhao, and X. Tan, “Prompttts: Controllable text-to-speech with text descriptions,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [27] L. Tunstall, E. Beeching, N. Lambert, N. Rajani, K. Rasul, Y. Belkada, S. Huang, L. von Werra, C. Fourrier, N. Habib *et al.*, “Zephyr: Direct distillation of lm alignment,” *arXiv preprint arXiv:2310.16944*, 2023.
- [28] O. Kuchaiev, J. Li, H. Nguyen, O. Hrinchuk, R. Leary, B. Ginsburg, S. Kriman, S. Beliaev, V. Lavrukhin, J. Cook *et al.*, “Nemo: a toolkit for building ai applications using neural modules,” *arXiv preprint arXiv:1909.09577*, 2019.
- [29] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [30] J. Han, R. Zhang, W. Shao, P. Gao, P. Xu, H. Xiao, K. Zhang, C. Liu, S. Wen, Z. Guo *et al.*, “Imagebind-llm: Multi-modality instruction tuning,” *arXiv preprint arXiv:2309.03905*, 2023.
- [31] R. Girdhar *et al.*, “Imagebind: One embedding space to bind them all,” in *CVPR*, 2023.