



Leveraging LLM and Self-Supervised Training Models for Speech Recognition in Chinese Dialects: A Comparative Analysis

Tianyi Xu^{1,2}, Hongjie Chen¹, Qing Wang¹, Lv Hang¹, Jian Kang¹, Jie Li¹, Zhennan Lin², Yongxiang Li¹, Lei Xie²

¹Institute of Artificial Intelligence (TeleAI), China Telecom, China

²Audio, Speech and Language Processing Group (ASLP@NPU), School of Computer Science, Northwestern Polytechnical University, China

xutianyi@mail.nwpu.edu.cn; lxie@nwpu.edu.cn

Abstract

Large-scale training corpora have significantly improved the performance of ASR models. Unfortunately, due to the relative scarcity of data, Chinese accents and dialects remain a challenge for most ASR models. Recent advancements in self-supervised learning have shown that self-supervised pre-training, combined with large language models (LLM), can effectively enhance ASR performance in low-resource scenarios. We aim to investigate the effectiveness of this paradigm for Chinese dialects. Specifically, we pre-train a Data2vec2 model on 300,000 hours of unlabeled dialect and accented speech data and do alignment training on a supervised dataset of 40,000 hours. Then, we systematically examine the impact of various projectors and LLMs on Mandarin, dialect, and accented speech recognition performance under this paradigm. Our method achieved SOTA results on multiple dialect datasets, including Kespeech. We will open-source our work to promote reproducible research.¹

Index Terms: speech recognition, LLM, speech foundation model

1. Introduction

Large-scale training corpora have significantly improved the performance of automatic speech recognition (ASR) models [1, 2, 3, 4]. However, labeled dialect data remains prohibitively expensive, making it challenging to obtain in large quantities. As a result, most pre-trained models in Chinese speech processing communities, such as WeNet [5], K2 [6], and FunASR [7], exhibit limited capabilities in dialect recognition. Fortunately, advances in self-supervised learning have demonstrated the potential of large-scale training to further improve model performance. Self-supervised large language models (LLM) [8, 9, 10, 11] have shown exceptional capabilities to understand and generate human language, and self-supervised speech models such as Wav2Vec2 [12], HuBERT [13] and Data2Vec2 [14] have demonstrated strong proficiency in modeling speech. Given that ASR traditionally relies on both acoustic and language modeling, the advantages of self-supervised pre-trained acoustic models and LLMs—particularly their scalability in terms of training data and model size—present promising opportunities for low-resource ASR systems.

Research efforts to integrate LLMs with ASR systems generally follow two primary approaches. The first is the cascaded approach [15, 16, 17], where LLMs are employed to refine the n -best hypotheses generated by ASR systems. Among studies adopting this approach, MMGER [18] stands out as a closely related work, achieving the previous SOTA performance

in Chinese accents on the Kespeech [19] dataset. However, this approach has a notable limitation: for complex cases, the n -best hypotheses may fail to provide sufficiently informative input to the LLM, making it difficult for the model to infer correct results from erroneous predictions. This limitation can sometimes result in semantic distortions during error correction. The second approach involves training audio-text cross-modal LLMs [20, 21, 22], where a speech encoder processes audio signals to generate embeddings that are subsequently fed into an LLM for decoding. This method aims to integrate acoustic features more effectively with linguistic context to enhance ASR accuracy. Specifically, SALMONN [20] employs Whisper [1] to extract semantic content while utilizing BEATs [23] to enable comprehensive perception of speech, music, and sound events. Qwen-Audio and Qwen2-Audio [22, 24] use Whisper as an encoder and optimize task performance across various domains with structured task instructions. Similarly, SLAM-ASR [21] achieved SOTA results on the 960-hour English LibriSpeech dataset by training only the projector. In the context of Chinese ASR, Geng et al. [25] achieved SOTA performance on AISHELL, Test-net, and Test-meeting by leveraging HuBERT encoder and Transformer projection layers.

Our study extends these investigations by exploring the potential of the speech foundation encoder with the LLM-Decoder paradigm on large-scale Chinese dialect datasets. Specifically, we utilize two datasets: an unlabeled dataset consisting of approximately 300,000 hours of dialect and accented speech and a labeled dataset comprising 40,000 hours. Using these resources, we systematically investigate the impact of various speech encoders, projectors, and LLMs on the ASR performance of Chinese dialects. Our experimental results reveal several important findings:

- First, among speech encoders, Data2Vec2 is a good speech encoder for large language models and can have competitive performance when downsampled to 25Hz. Combined with the 4-time downsampling from the projector, it can achieve competitive performance at 6.25 Hz framerate for the LLM, greatly reducing the computational overhead.
- Second, fully connected layers outperform other projectors, particularly during the early stages of training. Notably, differences in the effectiveness of projectors are most prominent during the first training stage but diminish as the parameters of the pre-trained models are progressively unfrozen in subsequent stages.
- Third, a multi-stage training strategy proves highly effective in enhancing model performance, with the addition of a fourth training stage, building upon conventional three-stage approaches, yields further improvements.

Compared to previous work [21, 25], our approach employs

¹Work done during internship at China Telecom.

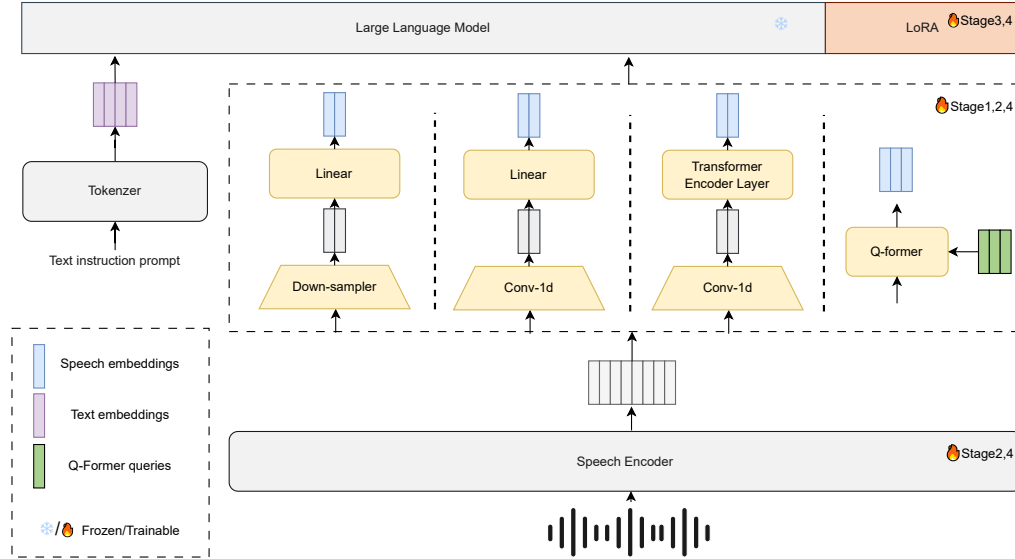


Figure 1: Overall model structure.

Data2Vec2 as the speech encoder and incorporates convolutional layers as the projector. Combined with a four-stage training strategy, the proposed method achieves SOTA performance on multiple Chinese dialect and accent test sets, including Kespeech (with CER of 6.48% vs. 7.52%), even with a small LLM (0.5B) and a low frame rate (6.25Hz). To promote reproducibility, we will release all associated training, inference, and evaluation scripts, as well as pre-trained models and training logs.

2. Method

As illustrated in Figure 1, the model architecture consists of an audio encoder and a LLM. For each sample during training, we define the text prompt (e.g., "Transcribe the following speech"), the speech utterance, and the corresponding transcription as P , S and T , respectively. The text prompt and transcription are tokenized using the LLM's tokenizer and embedded into feature vectors E_p and E_t as:

$$E_p = \text{Embedding}(\text{Tokenizer}(P)), \quad (1)$$

$$E_t = \text{Embedding}(\text{Tokenizer}(T)). \quad (2)$$

For the input speech S , features are first extracted using the speech encoder, producing the encoder output H_s , represented as:

$$H_s = \text{Encoder}(S). \quad (3)$$

The encoder output H_s is then passed through a projector, followed by a linear layer, to generate a feature sequence E_s with the same dimensionality as the input to the LLM, denoted as:

$$E_s = \text{Linear}(\text{Projector}(H_s)), \quad (4)$$

Here, the projector preserves the dimensionality of the features output by the speech encoder, while the linear layer maps these features to the embedding space of the LLM. Finally, E_s , E_p , and E_t are concatenated into a unified feature sequence and fed into the LLM. The model generates the transcription Y , as the output, defined as:

$$Y = \text{LLM}(\text{Regulation}(E_p, E_s, E_t)). \quad (5)$$

3. Experimental Setup

3.1. Datasets

Training Set: The training data comprises a diverse set of internal dialect and accent datasets, including samples from Anhui, Gansu, Hebei, Shandong, Shanxi, Tianjin, Cantonese, Henan, Sichuan, Chongqing, Northeast China, Shaanxi, Hubei, Fujian, Guizhou, Hangzhou, Hunan, Jiangxi, Shanghai, Suzhou, Yunnan, Chongqing and Hakka. Additionally, public Mandarin datasets, such as Wenetspeech [26], Aishell [27], and AliMeeting, are incorporated. For training the unsupervised models, the total dataset size is approximately 300,000 hours, while the supervised dataset is 40,000 hours.

Testing Set: Due to the lack of open-source test datasets for many dialects and the limitations of existing datasets—such as overly short utterances that fail to capture distinct dialectal characteristics—we employ internally annotated test sets for evaluation. Specifically, we conduct evaluations on four dialects: Shanghai, Hunan, Henan, and Cantonese, each with approximately 2,000 audio samples. Additionally, we utilize the open-source Kespeech dataset [19] to assess performance on accented speech. Kespeech includes Mandarin as well as dialects from regions such as Beijing, Jiang-Huai, Jiao-Liao, Ji-Lu, Lan-Yin, Northeastern, Southwestern, and Zhongyuan. However, we observe that the accents in this dataset are not particularly pronounced, likely because regional labels for speakers were directly assigned as accent labels, even when many speakers exhibit minimal accent features. Finally, to evaluate performance on Mandarin, we use the Test-net and Test-meeting subsets from the Wenetspeech dataset.

3.2. Components

3.2.1. LLM

To ensure robust recognition performance across various Chinese accents and dialects, we require an LLM with good performance in Chinese-related tasks. Given the complexity of LLMs, the academic community currently lacks standardized evaluation metrics. Therefore, we refer to the CLiB² leaderboard,

²<https://github.com/jeinlee1991/chinese-LLM-benchmark?tab=readme-ov-file>

which assesses the capabilities of Chinese LLMs across multiple tasks. Among open-source models of under 7B parameters, the Qwen 2 series [28] achieved the highest scores. Consequently, our study utilizes Qwen 2 of 0.5B, 1.5B, 7B, and Qwen 2.5 3B (as Qwen 2.5 was only released during the later stages of our research). Notably, our testing reveals that the Qwen series also demonstrates the capability to handle conversational responses in dialectal and colloquial Chinese expressions.

3.2.2. Speech foundation encoder

In prior work [29], the performance of several self-supervised speech encoders was investigated, including Wav2vec2 [12], Hubert [13], BEST-RQ [30], DinoSR [31], and Data2vec2 [14]. Our findings showed that, in Wenetspeech pre-training, Data2vec2 outperformed the other models. Subsequently, we trained Data2vec2 on a larger dataset consisting of 300,000 hours of unlabeled multilingual and accented speech. For the current study, we use the TeleSpeechPT variant of Data2vec2 with a model size of 24 layers and 301 million parameters. The model was trained at a sampling rate of 25Hz, which is lower than the commonly used 50Hz rate of the Whisper encoder.

3.2.3. Projector

The speech modality is more sparsely represented compared to the text modality, leading to increased computational complexity. To enhance inference efficiency, previous studies have frequently employed projection layers to align text and speech modalities. Since the community lacks a consensus conclusion on the most suitable projection layer for speech-based LLMs, we perform a comparative evaluation in the context of Chinese dialects. Specifically, we compare four different projection layers: fully connected (Linear), 1-dimension Convolutional (Conv1d), Transformer, and Q-Former [32]. For the Linear and Conv1d layers, we follow the configuration from SLAM and conduct experiments with downsampling rates of 1, 2, 4, and 8. The Q-Former configuration also follows SLAM, with the number of queries set to 64. The Transformer configuration follows the setup in [25].

3.3. Training strategy

LLM-based models are typically very large, and training the entire model simultaneously can sometimes lead to suboptimal performance. Previous works [20, 22, 25] used a three-stage fine-tuning strategy to ensure convergence between different modules. However, we hypothesize that this fine-tuning approach might cause the model to converge to local optima within individual modules rather than achieving a global optimum. Therefore, building on the three-stage training framework, we introduce an additional fourth stage, which involves fine-tuning the entire model. Our experiments show that this modification improves ASR performance. Specifically, the training process is shown in Figure 1: (1) First, only the projector is trained while all other components are frozen. (2) Second, only the speech encoder is trained. (3) Third, the encoder and projector are frozen, and LoRA is used to fine-tune the LLM. (4) Finally, both the encoder and the projector are unfrozen for further optimization.

3.4. Implementation details

We train our model using 8 40G A100 GPUs, and fine-tuning for the 7B version takes about 30 days for four stages. We use the framework from SLAM [21]. We use the AdamW opti-

Table 1: Comparison of CER (%) on Test-net/Test-meeting for Different Projection Layers after Pre-trained only Data2vec2 encoder vs CTC fine-tuned Data2vec2 encoder and Qwen2 0.5B LLM during the First Training Stage.

	Frame Rate	Pretrained	Finetuned
Conv1d	6.25Hz	18.19/16.64	31.46/27.87
Linear	6.25Hz	16.47/8.80	14.09/11.09
Transformer	6.25Hz	17.32/12.37	14.34/11.32
Q-former	-	27.00/17.76	23.72/17.53

Table 2: Comparison of CER (%) on Test-meeting/Test-net after fine-tuning on Wenetspeech with Multiple Projection Layers at Different Sampling Rates.

	25Hz	12.5Hz	6.25Hz	3.125Hz
Conv1D		20.44/14.10	31.46/27.87	38.97/41.66
Linear	16.48/10.79	18.15/14.89	14.09/11.09	24.62/25.00
Transformer	17.32/12.37	20.03/14.57	19.06/11.87	27.65/26.94

mizer [33] with the following hyperparameters: lr = 1.0e-05, beta = (0.9, 0.99), eps = 1.0e-06, and weight_decay = 0.01. To mitigate potential issues related to gradient explosion during training, we apply gradient clipping [34] with a threshold value set to 5. This ensures that gradients over 5 or below -5 are clipped to 5 or -5, respectively. Additionally, we use a gradient accumulation of 20 to increase the effective batch size to 100. Regarding the training approach, when training the LLM, we freeze the LLM body and only update the LLM using LoRA fine-tuning [35]. We configure LoRA with alpha = 32, rank = 12. The alpha parameter controls the weight of the LoRA matrix, and the rank parameter determines the dimension of the low-rank matrix. All experiments in this work follow these configurations unless otherwise specified.

4. Experimental Results

4.1. Comparison of Projector Architectures

We evaluate the effectiveness of different projection layers through small-scale experiments. Specifically, four types of projection layers—Linear, Conv1d, Transformer, and Q-Former are compared using the Data2Vec2 model pre-trained on 300,000 hours of data and tested on the Test-meeting and Test-net datasets. In these experiments, the LLM is fixed as Qwen2 0.5B, and training is limited only to the first stage for 1,000,000 steps. The Data2vec2 encoder, originally at a sampling rate of 25Hz, is further downsampled by a factor of 4, resulting in an embedding sampling rate of 6.25Hz at the LLM input. For the Q-Former projection layer, a fixed embedding length of 64 is used instead of a predefined sampling rate. The results, presented in Table 2, indicate that among the four projection layers, the fully connected layer achieves the best overall performance.

4.2. Projector Sampling Rate

We explore the impact of varying down-sampling rates in the projector. Experiments were conducted on projection layers implemented using fully connected layers, convolutional layers, and Transformer layers, as shown in table 2. For Q-Former, which compresses speech representations into a fixed-length embedding, comparisons with specific sampling rates are not feasible. Specifically, we test four down-sampling rates—1, 2, 4, and 8—corresponding to frame rates of 25 Hz, 12.5 Hz, 6.25 Hz, and 3.125 Hz for the speech encoder output. In contrast, Q-Former is fixed to produce embeddings with a length of

Table 3: Performance CER (%) on different LLMs through 4-stage fine-tuning on a 40,000-hour multi-accent, multi-dialect dataset.

LLM/ Baseline Model	CTC Finetuned encoder	Projector	LLM finetune method	Frame rate	He Nan	Shang Hai	Cantonese	Hu Nan	Kespeech	Test-net	Test-meeting
Qwen2-Audio-Instruct	-	-	-	50Hz	45.06	49.89	17.72	66.19	10.88	13.01	26.12
Whisper-Large-V3	-	-	-	50Hz	87.78	87.78	55.51	83.61	31.06	8.20	9.08
MMGER	-	-	-	-	-	-	-	-	7.52	-	-
Qwen2 0.5b	True	Conv1D	LoRA	12.5Hz	20.73	21.06	11.61	26.93	7.85	7.39	8.05
Qwen2 0.5b	True	Conv1D	LoRA	6.25Hz	20.48	19.81	11.03	25.16	7.64	7.29	7.69
Qwen2 0.5b	True	Conv1D	Full	6.25Hz	19.97	14.58	11.73	23.89	7.20	6.80	6.98
Qwen2 1.5b	True	Conv1D	LoRA	6.25Hz	20.64	18.56	11.40	23.93	7.62	6.93	7.44
Qwen2.5 3b	False	Linear	LoRA	6.25Hz	28.06	21.09	11.06	27.54	8.27	6.91	7.23
Qwen2.5 3b	True	Linear	LoRA	6.25Hz	27.72	21.78	11.79	25.44	7.78	7.22	7.19
Qwen2.5 3b	True	Conv1D	LoRA	6.25Hz	19.95	17.33	11.31	24.31	6.84	6.59	6.99
Qwen2 7b	True	Conv1D	LoRA	6.25Hz	18.95	18.30	10.98	23.31	6.48	7.54	6.77

Table 4: The improvement in CER (%) of the model across multiple dialectal datasets in 4-stage fine-tuning.

Stage	Projector	He Nan	Shang Hai	Hu Nan	Cantonese	kespeech	Test net	Test meeting
1	Linear	58.59	84.44	115.68	25.32	25.74	17.68	20.82
	Conv1d	61.13	58.16	169.03	21.35	24.86	21.89	22.41
2	Linear	60.23	61.25	72.37	21.89	21.01	14.73	15.60
	Conv1d	56.49	50.22	176.89	17.00	18.38	16.71	15.14
3	Linear	32.75	25.44	23.90	12.22	8.90	8.04	8.05
	Conv1d	27.37	25.16	20.45	11.73	8.42	7.48	8.01
4	Linear	27.72	21.78	21.78	11.30	7.90	6.96	7.36
	Conv1d	25.98	20.19	18.56	11.39	7.40	6.84	7.12

64. The results reveal that lower down-sampling rates preserve more information, resulting in lower CER. However, lower down-sampling rates also increase the computational load on the LLM, requiring more resources during both training and inference. Subsequent experiments use a down-sampling rate of 4 to balance recognition accuracy and computational efficiency. Combined with the encoder’s intrinsic down-sampling rate of 4, this configuration has a frame rate of 6.25 Hz for the LLM input, offering a practical trade-off between performance and resource consumption.

4.3. Comparison of ASR-Finetuned and Non-Finetuned Encoders

We investigate whether ASR-specific finetuning of the encoder, using the CTC objective, benefits the semantic alignment of the encoder with the LLM or introduces conflicts with the knowledge required during LLM finetuning. To this end, we pre-train a Data2Vec2 encoder and subsequently finetune it with CTC to acquire speech recognition capabilities. Specifically, the Data2Vec2 model is finetuned for 10 epochs using the same supervised 40,000-hour dataset. The resulting ASR-finetuned encoder is then used to replace the pre-trained encoder. Each setup follows the same 1,000,000-step finetuning process. As shown in Table 1, the results indicate that ASR-finetuned encoders perform worse across all four projection layer types compared to their non-finetuned counterparts. Notably, the performance gap was most significant in the case of Conv1d projection layers. This suggests that the knowledge obtained through CTC finetuning may conflict with the semantic information required during LLM finetuning, ultimately hindering the overall recognition performance.

4.4. Analysis of Multi-Stage Finetuning

We extend the investigation to large-scale experiments with multi-stage training. Specifically, we select the configurations with the best and worst (Linear and Conv1d) performance after the first training stage, as identified in Table 1, and use them for the 4-stage finetuning process using the larger supervised dataset containing 40,000 hours of multi-dialect speech data. The results are presented in Table 4. The findings indicate that

although the performance gap between the best and worst configurations is substantial after the first stage (31.46/27.87 vs. 16.47/8.80 on the Test-net and Test-meeting datasets), this disparity progressively decreases during the subsequent training stages and is even reversed by the final stage. Based on these results, we hypothesize that as additional model components are unfrozen during multistage fine-tuning, the differences in the projector’s parameters and structures become increasingly negligible.

4.5. Comparative Analysis of LLMs with Different Sizes

We evaluate the performance of Qwen series LLMs of 0.5B, 1.5B, 3B, and 7B during the 4-stage fine-tuning on the 40,000-hour dataset. The results are summarized in Table 3. For Qwen2 0.5B, we compare LoRA fine-tuning versus full fine-tuning in LLM. We observe a significant performance increase, which suggests that the LoRA parameter is not sufficient for acoustic representation training. In particular, the 3B model belongs to the Qwen2.5 series, which became available only later in the study. Its performance is compared with that of other Qwen2 models. Although Qwen2.5 3B is trained using more text tokens (7T vs. 18T) and has an MMLU close to Qwen2 7B (65.6 vs. 70.3), we observed that this has no significant benefit. In all experiments, larger LLMs demonstrated superior overall performance, highlighting the benefits of scaling up the model size. However, even the smallest model, the 0.5B LLM, achieved competitive results, surpassing the current SOTA MMGER [18] in Kespeech, and outperformed Qwen2-Audio and Whisper-Large-V3 in both internal test sets and public test sets.

5. Conclusion

In this study, we investigate the performance of various structural configurations within the paradigm of self-supervised training encoders paired with LLM decoders, leveraging a large-scale dataset comprising 40,000 hours of Chinese dialect and accented speech. For the encoder, we employ Data2Vec2, pre-trained on 300,000 hours of unlabeled dialect and accented speech data. Regarding the projection layer, our finding reveals that fully connected layers demonstrated a performance advantage over other types of projectors. In terms of training strategies, we observe that multi-stage training significantly improved model performance. Extending the conventional 3-stage training approach by incorporating a fourth stage further enhanced the results. The distinction in the effectiveness of the projector is most pronounced during the first stage of training, while these differences diminish progressively as the parameters of the pre-trained model unfreeze throughout the multi-stage training process. We achieve SOTA performance on multiple Chinese dialect test sets including Kespeech, and we will open-source our recipes and pre-trained models.

6. References

- [1] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust Speech Recognition via Large-Scale Weak Supervision," in *ICML*, 2023.
- [2] Y. Zhang, W. Han, J. Qin, Y. Wang, A. Bapna, Z. Chen, N. Chen, B. Li, V. Axelrod, G. Wang *et al.*, "Google usm: Scaling automatic speech recognition beyond 100 languages," *arXiv preprint arXiv:2303.01037*, 2023.
- [3] V. Pratap, A. Tjandra, B. Shi, P. Tomasello, A. Babu, S. Kundu, A. Elkahky, Z. Ni, A. Vyas, M. Fazel-Zarandi *et al.*, "Scaling speech technology to 1,000+ languages," *Journal of Machine Learning Research*, vol. 25, no. 97, pp. 1–52, 2024.
- [4] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [5] B. Zhang, D. Wu, Z. Peng, X. Song, Z. Yao, H. Lv, L. Xie, C. Yang, F. Pan, and J. Niu, "WeNet 2.0: More Productive End-to-End Speech Recognition Toolkit," in *Interspeech*, 2022.
- [6] Z. Yao, L. Guo, X. Yang, W. Kang, F. Kuang, Y. Yang, Z. Jin, L. Lin, and D. Povey, "Zipformer: A faster and better encoder for automatic speech recognition," in *The Twelfth International Conference on Learning Representations*, 2023.
- [7] Z. Gao, Z. Li, J. Wang, H. Luo, X. Shi, M. Chen, Y. Li, L. Zuo, Z. Du, and S. Zhang, "Funasr: A fundamental end-to-end speech recognition toolkit," in *Interspeech 2023*, 2023, pp. 1593–1597.
- [8] L. Floridi and M. Chiriatti, "Gpt-3: Its nature, scope, limits, and consequences," *Minds and Machines*, vol. 30, pp. 681–694, 2020.
- [9] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Roziere, N. Goyal, E. Hambro, F. Azhar *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [10] OpenAI, "GPT-4 Technical Report," *CoRR*, 2023.
- [11] G. Team, P. Georgiev, V. I. Lei, R. Burnell, L. Bai, A. Gulati, G. Tanzer, D. Vincent, Z. Pan, S. Wang *et al.*, "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," *arXiv preprint arXiv:2403.05530*, 2024.
- [12] A. Baeviski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [13] W. Hsu, B. Bolte, Y. H. Tsai, K. Lakhota, R. Salakhutdinov, and A. Mohamed, "HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units," *ACM*, 2021.
- [14] A. Baeviski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "Data2vec: A general framework for self-supervised learning in speech, vision and language," in *International Conference on Machine Learning*. PMLR, 2022, pp. 1298–1312.
- [15] R. Huang, M. Li, D. Yang, J. Shi, X. Chang, Z. Ye, Y. Wu, Z. Hong, J. Huang, J. Liu *et al.*, "Audiotape: Understanding and generating speech, music, sound, and talking head," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 21, 2024, pp. 23 802–23 804.
- [16] P. Dighe, Y. Su, S. Zheng, Y. Liu, V. Garg, X. Niu, and A. Tewfik, "Leveraging large language models for exploiting asr uncertainty," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 12 231–12 235.
- [17] R. Ma, M. Qian, P. Manakul, M. Gales, and K. Knill, "Can generative large language models perform asr error correction?" *arXiv preprint arXiv:2307.04172*, 2023.
- [18] B. Mu, Y. Li, Q. Shao, K. Wei, X. Wan, N. Zheng, H. Zhou, and L. Xie, "Mmger: Multi-modal and multi-granularity generative error correction with llm for joint accent and speech recognition," *arXiv preprint arXiv:2405.03152*, 2024.
- [19] Z. Tang, D. Wang, Y. Xu, J. Sun, X. Lei, S. Zhao, C. Wen, X. Tan, C. Xie, S. Zhou *et al.*, "Kespeech: An open source speech dataset of mandarin and its eight subdialects," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021.
- [20] 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," *CoRR*, 2023.
- [21] Z. Ma, G. Yang, Y. Yang, Z. Gao, J. Wang, Z. Du, F. Yu, Q. Chen, S. Zheng, S. Zhang, and X. Chen, "An Embarrassingly Simple Approach for LLM with Strong ASR Capacity," *CoRR*, 2024.
- [22] 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," *CoRR*, 2023.
- [23] S. Chen, Y. Wu, C. Wang, S. Liu, D. Tompkins, Z. Chen, W. Che, X. Yu, and F. Wei, "BEATs: Audio pre-training with acoustic tokenizers," in *Proceedings of the 40th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., vol. 202. PMLR, 23–29 Jul 2023, pp. 5178–5193. [Online]. Available: <https://proceedings.mlr.press/v202/chen23ag.html>
- [24] 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.
- [25] X. Geng, T. Xu, K. Wei, B. Mu, H. Xue, H. Wang, Y. Li, P. Guo, Y. Dai, L. Li *et al.*, "Unveiling the potential of llm-based asr on chinese open-source datasets," in *2024 IEEE 14th International Symposium on Chinese Spoken Language Processing (ISCSLP)*. IEEE, 2024, pp. 26–30.
- [26] B. Zhang, H. Lv, P. Guo, Q. Shao, C. Yang, L. Xie, X. Xu, H. Bu, X. Chen, C. Zeng, D. Wu, and Z. Peng, "WENETSPEECH: A 10000+ Hours Multi-Domain Mandarin Corpus for Speech Recognition," in *ICASSP*, 2022.
- [27] H. Bu, J. Du, X. Na, B. Wu, and H. Zheng, "AISHELL-1: An open-source Mandarin speech corpus and a speech recognition baseline," in *O-COCOSDA*, 2017.
- [28] A. Yang, B. Yang, B. Zhang, B. Hui, B. Zheng, B. Yu, C. Li, D. Liu, F. Huang, H. Wei *et al.*, "Qwen2.5 technical report," *arXiv preprint arXiv:2412.15115*, 2024.
- [29] H. Chen, Z. Li, G. Xia, B. Liu, Y. Yang, J. Kang, and J. Li, "Telespeech: Large-scale chinese multi-dialect and multi-accent speech pre-training," in *National Conference on Man-Machine Speech Communication*. Springer, 2024, pp. 183–190.
- [30] C.-C. Chiu, J. Qin, Y. Zhang, J. Yu, and Y. Wu, "Self-supervised learning with random-projection quantizer for speech recognition," in *International Conference on Machine Learning*. PMLR, 2022, pp. 3915–3924.
- [31] A. H. Liu, H.-J. Chang, M. Auli, W.-N. Hsu, and J. Glass, "Dinossr: Self-distillation and online clustering for self-supervised speech representation learning," *Advances in Neural Information Processing Systems*, vol. 36, pp. 58 346–58 362, 2023.
- [32] W. Yu, C. Tang, G. Sun, X. Chen, T. Tan, W. Li, L. Lu, Z. Ma, and C. Zhang, "Connecting speech encoder and large language model for asr," in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 12 637–12 641.
- [33] I. Loshchilov and F. Hutter, "Decoupled Weight Decay Regularization," in *ICLR*, 2019.
- [34] J. Zhang, T. He, S. Sra, and A. Jadbabaie, "Why Gradient Clipping Accelerates Training: A Theoretical Justification for Adaptivity," in *ICLR*, 2020.
- [35] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-Rank Adaptation of Large Language Models," in *ICLR*, 2022.