



WenetSpeech4TTS: A 12,800-hour Mandarin TTS Corpus for Large Speech Generation Model Benchmark

Linhan Ma^{1,†}, Dake Guo^{1,†}, Kun Song¹, Yuepeng Jiang¹, Shuai Wang^{2,3,4}, Liumeng Xue³, Weiming Xu¹, Huan Zhao¹, Binbin Zhang⁴, Lei Xie^{1,*}

¹Audio, Speech and Language Processing Group (ASLP@NPU),
School of Computer Science, Northwestern Polytechnical University, Xi'an, China

²Shenzhen Research Institute of Big Data, ³School of Data Science,
The Chinese University of Hong Kong, Shenzhen (CUHK-Shenzhen), China

⁴WeNet Open Source Community, China

{mlh2023, guodake}@mail.nwpu.edu.cn, lxie@nwpu.edu.cn

Abstract

With the development of large text-to-speech (TTS) models and scale-up of the training data, state-of-the-art TTS systems have achieved impressive performance. In this paper, we present *WenetSpeech4TTS*, a multi-domain Mandarin corpus derived from the open-sourced WenetSpeech dataset. Tailored for the text-to-speech tasks, we refined WenetSpeech by adjusting segment boundaries, enhancing the audio quality, and eliminating speaker mixing within each segment. Following a more accurate transcription process and quality-based data filtering process, the obtained WenetSpeech4TTS corpus contains 12,800 hours of paired audio-text data. Furthermore, we have created subsets of varying sizes, categorized by segment quality scores to allow for TTS model training and fine-tuning. VALL-E and NaturalSpeech 2 systems are trained and fine-tuned on these subsets to validate the usability of WenetSpeech4TTS, establishing baselines on benchmark for fair comparison of TTS systems. The corpus and corresponding benchmarks are publicly available on huggingface¹.

Index Terms: text-to-speech, corpus, benchmark

1. Introduction

Recently, language model-based and diffusion-based [1] text-to-speech (TTS) approaches have demonstrated remarkable performance. In addition to the strong modeling capabilities of the models, these achievements are mainly attributed to the scale-up of the data. For example, VALL-E [2] scaled up its training data to 60,000 hours of English speech [3] and NaturalSpeech 2 [4] used 44,000 hours of transcribed English speech data [5]. However, these existing large-scale open-source datasets are in English or multilingual. In contrast, there is a noticeable lack of comparably extensive datasets for Chinese TTS applications. To the best of our knowledge, the largest open-source transcribed Chinese speech dataset for TTS currently is DIDISPEECH [6], which contains nearly 800 hours of reading-style speech. This is insufficient for training large TTS models due to its small data scale and low diversity. Furthermore, there is currently a lack of a publicly available benchmark of TTS models trained on large-scale Chinese corpora for fair comparison. Due to the difficulty and high cost of reproducing and training

other systems, some studies can only obtain samples for comparison and evaluation by downloading audio samples from the demo pages of other systems. It poses a certain obstacle to the development of the community.

To the best of our knowledge, WenetSpeech [7] is the currently largest widely-used open-source Mandarin speech corpus for automatic speech recognition (ASR), containing 12,483 hours of transcribed Mandarin speech data collected from *YouTube* and *Podcast*. The speech data in WenetSpeech is sourced from real recordings, covering a large number of speakers and a wide range of domains such as audiobooks, interviews, reading, etc. Despite that some TTS studies [8] have used WenetSpeech as part of their training data to train large-scale models, the dataset itself exhibits certain limitations: Most notably, the speech data, sourced directly from *YouTube* and *podcast*, contains various noises and distortions and lacks further post-processing and refinement. While such a complex acoustic environment might be beneficial for constructing robust ASR systems, it is definitely unsuitable for TTS systems. Secondly, the segmentation of the data from *YouTube* videos is performed by an optical character recognition (OCR) system. A subtitle phrase from the source video is often split into several pieces that appear across different frames by annotators. Despite efforts to merge phrases during the OCR segmentation process, this approach frequently leads to sentences being split into multiple speech segments, many of which are too short and semantically incomplete. Furthermore, since it is specifically designed for ASR purposes, WenetSpeech does not ensure *speaker homogeneity* within segments, leading to segments where multiple speakers are present. Additionally, the speech and subtitle annotations of the original video data are not completely synchronized in time, leading to inaccurate subtitle boundaries detected by the OCR system. Consequently, some segments may have truncated words at the beginning or end.

To address the aforementioned issues of original WenetSpeech, in this paper, we introduce WenetSpeech4TTS, a 12,800-hour large-scale dataset tailored for the training of speech synthesis models. Based on the speech quality, WenetSpeech4TTS provides three subsets: *Basic*, *Standard*, and *Premium*, containing 7,226, 4,056, and 945 hours of effective data, respectively.

From WenetSpeech to WenetSpeech4TTS, we designed an automatic pipeline containing multiple processing steps, a similar approach of deriving TTS dataset from wild data have been

[†] Equal contribution. Random order.

^{*} Corresponding author.

¹Huggingface link: <https://huggingface.co/datasets/WenetSpeech4TTS/WenetSpeech4TTS>

proved effective in [9]². Specifically, we refined WenetSpeech for TTS by merging segments based on speaker similarity and pause duration, and by expanding segment boundaries to prevent truncated words. Audio quality was enhanced using the denoising model, followed by quality scoring. Furthermore, a speaker diarization system clustered segments from the same speaker, while a more advanced ASR system provided more accurate transcriptions.

To validate the usability of WenetSpeech4TTS, we trained the representative language model-based VALL-E [2] system and diffusion-based NaturalSpeech 2 [4] system on different subsets. Both subjective and objective evaluations were conducted, demonstrating that WenetSpeech4TTS is effective for training large TTS models, with higher-quality subsets yielding better performance. The WenetSpeech4TTS corpus with all metadata and the corresponding benchmark are publicly available on huggingface.

2. Processing pipeline

This section will detail the processing pipeline that customizes the WenetSpeech4TTS dataset through a series of operations, mainly including adjacent segments merging, boundary extension, speech enhancement, multi-speaker detection, speech recognition, and quality filtering.

2.1. Adjacent segments merging

To alleviate the issue that many segments are too short and semantically incomplete, we designed a merging strategy based on the interval time and speaker similarity. As illustrated in Figure 1, we first examine the interval time between adjacent segments based on the timestamp information in WenetSpeech. If the interval time is below 0.55 seconds, we consider them potentially belonging to the same sentence. We then apply speech enhancement to both segments to improve the speech quality, the specifics of the speech enhancement model will be elaborated in Section 2.3. For the extraction of speaker embeddings, Resemblyzer³ is employed. The cosine similarity between the two embeddings is then calculated as the speaker similarity. Segments with similarity scores above 0.65 are presumed to be from the same speaker and are merged, this process is repeated until the sentence duration reaches 20 seconds.

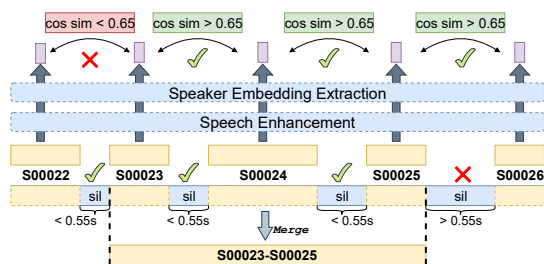


Figure 1: Demonstration of adjacent segments merging

2.2. Boundary extension

To address the issue of truncated words at the beginnings and ends of segments, we proposed to extend the speech segments after the merging operation described in Section 2.1. As illustrated in Figure 2, we extend the segment start backwards,

out surpassing the midpoint of the interval with the previous neighbor, and the maximum extension is limited to 0.5 seconds. The same operation is also applied to the segment ends.




Figure 2: Boundary extension

2.3. Speech enhancement

We refine the timestamp information of WenetSpeech based on the results of the previous two steps and then segment the source data to obtain new speech segments. The speech enhancement model was employed again to improve the speech quality of these segments. Specifically, we utilized MBTFNet [10], a multi-band time-frequency neural network designed specially to remove noise, background music, or other forms of interference, to obtain clean vocals. This speech enhancement operation removes interfering factors for subsequent operations such as quality evaluation, multi-speaker detection, speech recognition, and quality filtering. To validate the efficacy of our speech enhancement model, we conducted a quality assessment on a random sample of 10,000 segments after merging and extension, using the DNSMOS metric [11, 12]. We obtained and compared the DNSMOS P.808 scores for these segments before and after enhancement, illustrated in Figure 3, clearly demonstrate a significant improvement in segment quality following enhancement.




Figure 3: Distribution of DNSMOS P.808 scores for 10,000 random segments: enhanced (Red) vs. original (Blue).

2.4. Multi-speaker detection

For TTS models, speech segments containing multiple speakers are often not directly usable for training. To ensure such “speaker homogeneity”, we adopt a clustering based diarization process to obtain speaker assignments on segments. Specially, we train the ResNet293 [13] speaker embedding model on VoxCeleb [14] and VoxCeleb2 [15] datasets. Then we use the speech enhancement model in Section 2.3 to reconstruct the training data to fine-tune the speaker embedding model, eliminating the possible data mismatch introduced by the speech enhancement process. Each segment is divided into chunks with a 1.5s window and a 0.75s shift to compute speaker embeddings. Subsequently, a standard spectral clustering algorithm is applied to obtain the cluster centers and assignments. We label the 10% of embeddings in each cluster that are farthest from the cluster center as centrifugal embeddings. To avoid situations where the same person is assigned to different clusters,

²However, the processed dataset has not been released

³<https://github.com/resemble-ai/Resemblyzer>

we further merge the clusters if the cosine similarity between their cluster centers is greater than 0.75. Finally, if the speaker embeddings are assigned to different clusters or if more than half of the speaker embeddings are labeled as centrifugal in a segment, then the segment will be considered to contain more than one speaker.

2.5. Speech recognition

Compared to the original WenetSpeech, certain segments need to be transcribed since they have been restructured. Additionally, the transcripts obtained from the OCR system exhibit deletion errors in colloquial and repeated words, and the performance of the ASR system originally applied to the *Podcast* data is outdated. Therefore, we use the open-source industrial-level Paraformer-large [16] system⁴ to obtain new text transcriptions for the speech segments, this model achieved a 6.74% character error rate (CER) on the WenetSpeech “Test.Net” set. Meanwhile, we calculate the posterior probability of the Paraformer recognizing each segment as a confidence score.

2.6. Quality filtering

After the preceding steps, we conduct a quality-based data filtration. Segments with multiple speakers, as well as those with speech recognition confidence scores below 0.7, are excluded. The refined dataset, designated as the WenetSpeech4TTS corpus, comprises 12,800 hours of speech segments. *Some intervals between segments, which are not included in WenetSpeech, are merged into WenetSpeech4TTS through the merging operation. This enables us to obtain a larger dataset.*

3. The WenetSpeech4TTS corpus

We further divide our WenetSpeech4TTS corpus into subsets based on the DNSMOS P.808 scores. Figure 4 illustrates the score distribution of the dataset. Notably, segments with scores above 4.0 are labeled as *Premium*, those above 3.8 as *Standard*, and those above 3.6 as *Basic*. The remaining segments with a score lower than 3.6 are labeled as *Rest*. Specifics about these subsets are shown in Table 1.




Figure 4: The distribution of speech data quality. The horizontal axis represents DNSMOS P.808 scores, and the vertical axis represents the scale of data corresponding to the scores.

Figure 5 shows the distribution of segment lengths in WenetSpeech and the distribution of segment lengths in the WenetSpeech4TTS *Basic* subset. After the merging operation, the number of segments under 3 seconds greatly decreased, while segments over 5 seconds significantly increased. Importantly, many previously rare segments over 10 seconds are gen-

⁴https://www.modelscope.cn/models/iic/speech_paraformer-large-vad-punc_asr_nat-zh-cn-16k-common-vocab8404-pytorch/summary

Table 1: The training data subsets.

Training Subsets	DNSMOS Threshold	Hours	Average Segment Duration (s)
<i>Premium</i>	4.0	945	8.3
<i>Standard</i>	3.8	4,056	7.5
<i>Basic</i>	3.6	7,226	6.6
<i>Rest</i>	< 3.6	5,574	-
WenetSpeech4TTS (sum)	-	12,800	-
WenetSpeech (orig)	-	12,483	-

erated. This benefits training TTS models and demonstrates that the merging strategy is reasonable and necessary.




Figure 5: The distribution of segment lengths in WenetSpeech (a) and the WenetSpeech4TTS *Basic* subset (b).

The WenetSpeech4TTS test set comprises 150 test sentences and 26 target speakers, of which 10 in-set speakers from the WenetSpeech4TTS corpus and 16 out-of-set speakers, with equal numbers of males and females. The out-of-set speakers include 6 from WenetSpeech Test.Net set, 5 from the open-source dataset Aishell3 [17], and the other 5 are amateur speakers collected.

The WenetSpeech4TTS corpus, including segments, transcripts, and DNSMOS scores, is open-sourced and audio samples can be available at our demo page⁵.

4. Experiments

To validate the usability of the WenetSpeech4TTS dataset and establish baselines on benchmark for large-scale TTS models, two popular large-scale TTS models, VALL-E [2] and NaturalSpeech 2 [4], are trained on the *Basic* set and finetuned on the other two sets with higher quality successively. We open-source these model weights along with our dataset.

4.1. VALL-E and NaturalSpeech 2

In VALL-E, both the Autoregressive (AR) and Non-Autoregressive (NAR) have the same transformer architecture with 12 layers, 16 heads, an embedding dimension of 1024, a feed-ward layer dimension of 4096, and a drop out of 0.1. To get better speech quality, we employ AudioDec⁶ as an acoustic codec rather than Encodec [18, 19]. We resample all recordings to 24k Hz and use the *libritts v1* model to extract acoustic tokens. For all transcripts, we use the front end of BERT-VITS2⁷

⁵<https://wenetspeech4tts.github.io/wenetspeech4tts/>

⁶<https://github.com/facebookresearch/AudioDec>

⁷<https://github.com/fishaudio/Bert-VITS2>

Table 2: The results of VALL-E and NaturalSpeech 2 with 95% confidence intervals. The training set *Basic* means training with the *Basic* subset. *Standard,FT* means fine-tuning with the *Standard* subset. *Premium,FT* means fine-tuning with the *Premium* subset after *Standard,FT*.

Method	Training set	10 seen speakers				16 unseen speakers			
		CER ↓	SECS ↑	NMOS ↑	SMOS ↑	CER ↓	SECS ↑	NMOS ↑	SMOS ↑
VALL-E	<i>Basic</i>	13.12	0.824	3.34 ± 0.15	3.68 ± 0.12	14.96	0.781	3.27 ± 0.12	3.37 ± 0.12
	<i>Standard, FT</i>	10.23	0.811	3.41 ± 0.14	3.59 ± 0.13	12.91	0.772	3.36 ± 0.10	3.46 ± 0.11
	<i>Premium, FT</i>	8.10	0.827	3.50 ± 0.15	3.61 ± 0.13	9.82	0.781	3.45 ± 0.09	3.58 ± 0.09
NS2	<i>Basic</i>	11.88	0.783	2.94 ± 0.13	3.15 ± 0.12	14.06	0.756	2.64 ± 0.11	2.95 ± 0.11
	<i>Standard, FT</i>	8.28	0.803	3.38 ± 0.14	3.50 ± 0.13	10.78	0.766	3.14 ± 0.10	3.29 ± 0.10
	<i>Premium, FT</i>	6.30	0.820	3.40 ± 0.13	3.48 ± 0.13	7.89	0.783	3.36 ± 0.09	3.45 ± 0.10

to convert text to Pinyin. Both the AR model and NAR model are trained using 4 NVIDIA A40 GPUs for a week, respectively. Initially, we use WenetSpeech4TTS *Basic* as the training set for the foundational VALL-E model. To improve audio quality and performance, we fine-tune the foundational VALL-E model with WenetSpeech4TTS *Standard* and then with WenetSpeech4TTS *Premium*.

The model configurations of NaturalSpeech 2 follow the default configurations in the Amphion toolkit⁸ [20]. The data pre-processing process is consistent with VALL-E, except that we use the default Encodec as the acoustic codec of NaturalSpeech 2. Additionally, we sample 2,000 hours of labeled data from WenetSpeech4TTS *Basic* to train the ASR model in Kaldi⁹ toolkit, and then obtain external duration information of all labeled data for NaturalSpeech 2 training. We train the model for a total of 600k steps on 8 NVIDIA A6000 GPUs. The training and fine-tuning strategy of NaturalSpeech 2 is similar to that of VALL-E.

4.2. Evaluation metrics

We use the test set introduced in Section 3 to test these models’ performance. We synthesized all 150 utterances for each speaker for objective evaluation. We randomly selected 5 seen speakers and 5 unseen speakers, with 8 random samples for each speaker, for a total of 80 test cases for subjective evaluation. The objective evaluations included speaker embedding cosine similarity (SECS) and character error rate (CER) in ASR. The SECS metric was computed by extracting speaker embeddings with Resemblyzer¹⁰ and calculating the cosine similarity. The CER was measured between transcripts of real and synthesized utterances transcribed by Paraformer employed in Section 2.5. For subjective evaluation, We conduct Mean Opinion Score (MOS) experiments to evaluate speech naturalness (NMOS) and speaker similarity MOS (SMOS) reference to target speaker utterances. In each MOS test, a group of 20 native Chinese Mandarin listeners are asked to listen to synthetic speech and rate on a scale from 1 to 5 with a 0.5-point interval. The objective and subjective results are shown in Table 2.

4.3. Objective evaluation

For all models, the CER and SECS metrics are worse for unseen speakers compared to seen speakers, indicating that all models obtain better intelligibility and similarity on speakers they have seen during training. The CER for both VALL-E and NaturalSpeech 2 significantly decreases as the quality of the training data improves (from *Basic* to *Premium, FT*), which suggests

that higher-quality subsets lead to more stable speech synthesis. NaturalSpeech 2 outperforms VALL-E in terms of CER across all training sets as well as all speakers. This indicates that the NaturalSpeech 2 approach, which benefits from its external duration-aligned training, is more robust and achieves better intelligibility than VALL-E’s autoregressive modeling. Across both seen and unseen speakers, the SECS metric is relatively high for all models, indicating that their ability to model speaker timbre was generally good. As the quality of training data improves, the SECS metric of VALL-E remains relatively stable while that of NaturalSpeech 2 shows an upward trend. This suggests that NaturalSpeech 2 is more affected by data quality in speaker timbre modeling.

4.4. Subjective evaluation

The SMOS and NMOS scores of both VALL-E and NaturalSpeech 2 for seen speakers are higher than those for unseen speakers, which is the same situation as objective evaluation. As data quality improves, NMOS scores generally show an upward trend. Except that SMOS scores of VALL-E on in-set speakers are relatively stable, the other SMOS scores also gradually improve with the improvement of data quality. Overall, the naturalness of VALL-E is better than that of NaturalSpeech 2 under the same conditions. We attribute the reason to the error accumulation of NaturalSpeech 2’s external duration alignment model using automatically transcribed data. In the objective evaluation, the speaker similarity of VALL-E and NaturalSpeech 2 is basically the same, while in the subjective evaluation, the speaker similarity of NaturalSpeech 2 is lower. We speculate that it may be because the speech quality generated by Encodec is worse than that of AudioDec, which affects the listener’s perception of speaker timbre. Audio samples can be available at our demo page¹¹.

5. Conclusion

This paper proposed *WenetSpeech4TTS*, a multi-domain Mandarin corpus for large TTS model training which is derived from the open-source WenetSpeech dataset. We designed a series of processing operations to refine the WenetSpeech data and then divided it into subsets of different sizes according to the quality of the speech. To prove the usability of WenetSpeech4TTS and provide a benchmark, we trained and fine-tuned VALL-E and NaturalSpeech 2 on these subsets. Experimental results demonstrate that WenetSpeech4TTS can be used to train large TTS models and that higher-quality subsets achieve better performance.

⁸<https://github.com/open-mmlab/Amphion>

⁹<https://github.com/kaldi-asr/kaldi>

¹⁰<https://github.com/resemble-ai/Resemblyzer>

¹¹<https://wenetspeech4tts.github.io/wenetspeech4tts/>

6. References

- [1] J. Ho, A. Jain, and P. Abbeel, "Denosing diffusion probabilistic models," in *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [2] C. Wang, S. Chen, Y. Wu, Z. Zhang, L. Zhou, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, "Neural codec language models are zero-shot text to speech synthesizers," *CoRR*, vol. abs/2301.02111, 2023.
- [3] J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux, "Libri-light: A benchmark for ASR with limited or no supervision," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*. IEEE, 2020, pp. 7669–7673.
- [4] K. Shen, Z. Ju, X. Tan, E. Liu, Y. Leng, L. He, T. Qin, sheng zhao, and J. Bian, "Naturalspeech 2: Latent diffusion models are natural and zero-shot speech and singing synthesizers," in *The Twelfth International Conference on Learning Representations*, 2024.
- [5] V. Pratap, Q. Xu, A. Sriram, G. Synnaeve, and R. Collobert, "MLS: A large-scale multilingual dataset for speech research," in *Interspeech 2020, 21st Annual Conference of the International Speech Communication Association, Virtual Event, Shanghai, China, 25-29 October 2020*. ISCA, 2020, pp. 2757–2761.
- [6] T. Guo, C. Wen, D. Jiang, N. Luo, R. Zhang, S. Zhao, W. Li, C. Gong, W. Zou, K. Han *et al.*, "Didispeech: A large scale mandarin speech corpus," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6968–6972.
- [7] B. Zhang, H. Lv, P. Guo, Q. Shao, C. Yang, L. Xie, X. Xu, H. Bu, X. Chen, C. 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)*. IEEE, 2022, pp. 6182–6186.
- [8] Z. Zhang, L. Zhou, C. Wang, S. Chen, Y. Wu, S. Liu, Z. Chen, Y. Liu, H. Wang, J. Li, L. He, S. Zhao, and F. Wei, "Speak foreign languages with your own voice: Cross-lingual neural codec language modeling," *CoRR*, vol. abs/2303.03926, 2023.
- [9] J. Yu, H. Chen, Y. Bian, X. Li, Y. Luo, J. Tian, M. Liu, J. Jiang, and S. Wang, "Autoprep: An automatic preprocessing framework for in-the-wild speech data," *CoRR*, vol. abs/2309.13905, 2023.
- [10] W. Xu, Z. Chen, Z. Tan, S. Lv, R. Han, W. Zhou, W. Zhao, and L. Xie, "MBTFNET: multi-band temporal-frequency neural network for singing voice enhancement," in *IEEE Automatic Speech Recognition and Understanding Workshop, ASRU 2023, Taipei, Taiwan, December 16-20, 2023*. IEEE, 2023, pp. 1–8.
- [11] C. K. Reddy, V. Gopal, and R. Cutler, "Dnsmos: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2021, Toronto, ON, Canada, June 6-11, 2021*. IEEE, 2021, pp. 6493–6497.
- [12] C. K. A. Reddy, V. Gopal, and R. Cutler, "Dnsmos P.835: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2022, Virtual and Singapore, 23-27 May 2022*. IEEE, 2022, pp. 886–890.
- [13] H. Wang, C. Liang, S. Wang, Z. Chen, B. Zhang, X. Xiang, Y. Deng, and Y. Qian, "Wespeaker: A research and production oriented speaker embedding learning toolkit," in *IEEE International Conference on Acoustics, Speech and Signal Processing ICASSP 2023, Rhodes Island, Greece, June 4-10, 2023*. IEEE, 2023, pp. 1–5.
- [14] A. Nagrani, J. S. Chung, and A. Zisserman, "Voxceleb: A large-scale speaker identification dataset," in *Interspeech 2017, 18th Annual Conference of the International Speech Communication Association, Stockholm, Sweden, August 20-24, 2017*. ISCA, 2017, pp. 2616–2620.
- [15] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Interspeech 2018, 19th Annual Conference of the International Speech Communication Association, Hyderabad, India, 2-6 September 2018*. ISCA, 2018, pp. 1086–1090.
- [16] Z. Gao, S. Zhang, I. McLoughlin, and Z. Yan, "Paraformer: Fast and accurate parallel transformer for non-autoregressive end-to-end speech recognition," in *INTER_SPEECH*. ISCA, 2022, pp. 2063–2067.
- [17] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, "AISHELL-3: A Multi-Speaker Mandarin TTS Corpus," in *Proc. Interspeech 2021*, 2021, pp. 2756–2760.
- [18] Y.-C. Wu, I. D. Gebru, D. Marković, and A. Richard, "AudioDec: An open-source streaming high-fidelity neural audio codec," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [19] A. Défossez, J. Copet, G. Synnaeve, and Y. Adi, "High fidelity neural audio compression," *Transactions on Machine Learning Research*, 2023.
- [20] X. Zhang, L. Xue, Y. Gu, Y. Wang, H. He, C. Wang, X. Chen, Z. Fang, H. Chen, J. Zhang, T. Y. Tang, L. Zou, M. Wang, J. Han, K. Chen, H. Li, and Z. Wu, "Amphion: An open-source audio, music and speech generation toolkit," *arXiv*, vol. abs/2312.09911, 2024.