



StyleS2ST: Zero-shot Style Transfer for Direct Speech-to-speech Translation

Kun Song^{1,†}, Yi Ren², Yi Lei¹, Chunfeng Wang², Kun Wei¹, Lei Xie¹, Xiang Yin², Zejun Ma²

¹Audio, Speech and Language Processing Group (ASLP@NPU)
School of Computer Science, Northwestern Polytechnical University, Xi'an, China
²ByteDance

kunsong.npu.se@gmail.com, lxie@nwpu.edu.cn

Abstract

Direct speech-to-speech translation (S2ST) has gradually become popular as it has many advantages compared with cascade S2ST. However, current research mainly focuses on the accuracy of semantic translation and ignores the speech style transfer from a source language to a target language. The lack of high-fidelity expressive parallel data makes such style transfer challenging, especially in more practical zero-shot scenarios. To solve this problem, we first build a parallel corpus using a multi-lingual multi-speaker text-to-speech synthesis (TTS) system and then propose the *StyleS2ST* model with cross-lingual speech style transfer ability based on a style adaptor on a direct S2ST system framework. Enabling continuous style space modeling of an acoustic model through parallel corpus training and non-parallel TTS data augmentation, StyleS2ST captures cross-lingual acoustic feature mapping from the source to the target language. Experiments show that StyleS2ST achieves good style similarity and naturalness in both in-set and out-of-set zero-shot scenarios.

Index Terms: direct speech-to-speech translation, style transfer, zero-shot

1. Introduction

In recent years, direct speech-to-speech translation (S2ST) has attracted more and more attention [1, 2, 3]. Conventional S2ST [4] is usually based on cascade systems, including automatic speech recognition (ASR), machine translation (MT), and text-to-speech synthesis (TTS) modules, while direct S2ST aims to integrate the above modules into a unified model for directly synthesizing target language speech translated from the source language. Compared with the cascade system, direct S2ST takes advantage of a more elegant pipeline, which avoids relying on intermediate text and obtains less error propagation. Recent direct S2ST works have achieved promising performance by introducing speech-to-unit translation (S2UT) models [3, 5], which leverage discrete self-supervised units, such as features from HuBERT [6], as an intermediate representation bridging different languages in S2ST. These approaches adopt self-supervised pre-trained modules and data augmentation strategies to improve semantic accuracy, leading to state-of-the-art (SOTA) performance.

Since the unit of self-supervised learning can well represent the semantic information in speech, it helps the model better learn the alignment relationship between source and target languages. Therefore, direct S2ST has achieved semantic translation performance close to or even better than conventional cascade systems. However, current S2ST studies mainly focus

on linguistic information, while ignoring *speaking styles* (e.g., speaker timbre, intonation, rhythm, and intensity) that accompany the source language speech. The stylistic appearance of speech is complex and contains rich information such as speaker identity and attitude, which significantly affects the expressiveness and naturalness of speech. With the increase of application scenarios for S2ST, it is often critical to transfer the source speaking style to the target language speech in many scenarios, such as movie dubbing translation. While S2ST with style transfer is challenging due to the lack of paired multi-lingual speech data from the same speaker. Moreover, in real-world applications, it is an absolute limitation if only transferring the style seen during training since there may exist diverse style expressions in source language speech. Therefore, the generalization ability is essential in S2ST for an arbitrary unseen style from a new speaker, which we call the *zero-shot* style transfer of S2ST in this paper.

In this work, firstly, we would like to define the basic specifications of *style* in S2ST scenarios. Specifically, style here includes low-level characteristics such as speaker timbre, pitch, intensity, and rhythm, while high-level characteristics are always corresponding to speaker identity. We consider different speakers usually have distinguished speaking styles. Given the above assumption, this paper mainly focuses on zero-shot style transfer for S2ST, which has several challenges as follows:

- The stylistic expression and semantic content should be *disentangled* for style transfer in S2ST, while current works lack such ability since they only concern with the semantic translation. Moreover, style and linguistic contents are heavily entangled in speech and they are not easily to be disentangled.
- To conduct style transfer in direct S2ST, high-fidelity expressive paired data of source and target language speech from the same speaker is necessary but extremely hard to obtain.
- In zero-shot scenarios, it usually requires enormous amounts of speech data of diverse speaking styles to build a complete continuous speaker space. However, the demand for paired speech heavily increases the difficulty of zero-shot style transfer which desires good out-of-set speaker generalization.

To construct parallel bilingual speech data, an effective solution is to leverage a TTS model for pseudo-data pairs augmentation [1, 2]. Previous works usually synthesize monolingual speech from text extracted by speech-to-text translation, due to the high cost and limited ability of the pre-trained TTS models. In that way, the constructed pseudo-data pairs realize semantic parallelization, but not the speaker parallel parallelization. Existing S2UT-based S2ST works [3, 5] mainly use single-speaker unit decoder to achieve single-speaker speech-to-speech translation. Translatotron [1], proposed for S2ST without intermedi-

[†] Work done during an internship at ByteDance. Lei Xie is the corresponding author.

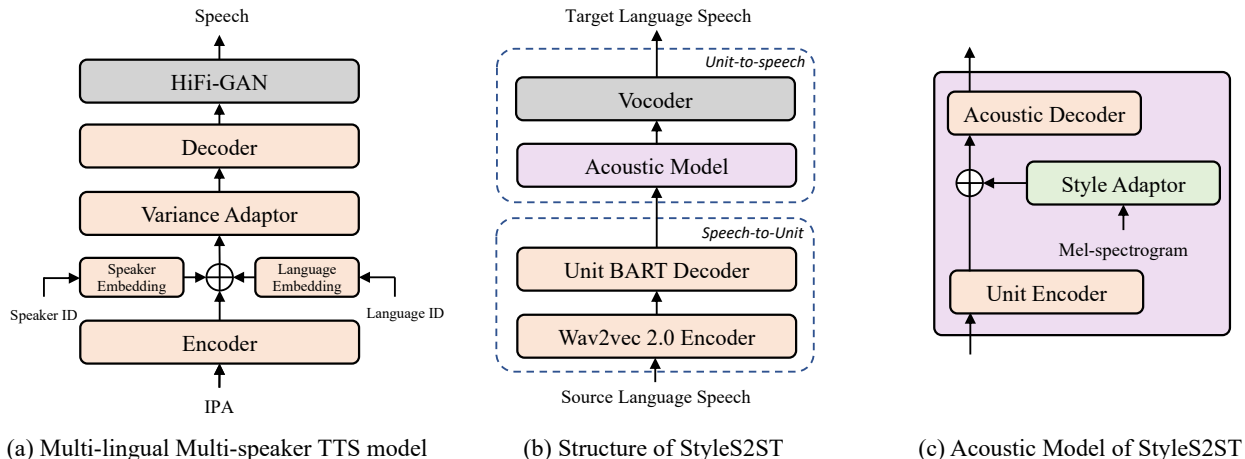


Figure 1: Architecture of our proposed methods, consists of a TTS model for building the bilingual parallel corpus and our proposed S2ST system – StyleS2ST.

ate representations, tries to achieve style transfer by a speaker encoder, but the performance is far from good enough for applications. Moreover, current direct S2ST systems can not reserve style from source to target speech, while zero-shot style transfer has even not been considered yet.

In this paper, we propose to conduct zero-shot style transfer for direct speech-to-speech translation. To construct a high-quality bilingual parallel corpus, we leverage a pre-trained multi-lingual multi-speaker TTS system taking the bilingual parallel texts as input. With the synthetic speech parallel in both semantics and speaking style, we build the proposed method based on a SOTA direct S2UT English-to-Chinese (En2Zh) system [5]. To achieve zero-shot style transfer, we adopt a pre-trained *style adaptor* based on a speaker embedding model in our S2UT system. The contributions of this paper can be summarized as follows. 1) To the best of our knowledge, this is the first attempt for zero-shot style transfer in direct S2ST. 2) We leverage a multi-speaker multi-style TTS system to construct parallel bilingual speech of the same semantic content and speaking style. 3) We propose to utilize a style adaptor to capture the speaking style of arbitrary unseen speakers for achieving zero-shot style transfer in direct S2ST. Experimental results show that the proposed approach, dubbed *StyleS2ST*, can successfully conduct zero-shot end-to-end speech-to-speech translation with the preservation of speaking style from the source to the target languages. Our audio samples are available on the demo website.¹

2. Methods

Our goal is to realize zero-shot style transfer in a direct S2ST system. In this section, we will introduce the proposed pseudo-data pairs construction process and the structure of our direct S2ST system based on S2UT. Next, we present our proposed style adaptor and the enhanced acoustic decoder to achieve style transfer and zero-shot, respectively.

2.1. Corpus Construction

To train the StyleS2ST, we construct a speaker-parallel En2Zh corpus through a multi-lingual multi-speaker TTS system. Figure 1(a) shows our TTS structure based on the acoustic model of Fastspeech 2 [7] and the vocoder of HiFi-GAN [8]. Firstly, we adopt the international phonetic alphabet (IPA), widely ap-

plied in cross-lingual TTS and voice conversion (VC) [9, 10], as input to model the shared speech space of different languages. Then, we use an encoder to encode the input phoneme sequence into a hidden representation, which is concatenated with injected speaker embedding and language embedding to synthesize speech in the target language and desired speaking style. Due to the lack of parallel corpus, each speaker has only English or Chinese data. Still, using IPA as input allows us to obtain bilingual high-fidelity synthesized speech for different speakers. To evaluate and filter the synthesized speech for training S2ST, we utilize word error rate (WER) and perceptual evaluation of speech quality (PESQ) [11] to check the semantic errors and speech quality. Any generated speech with a performance below the threshold is discarded from the corpus.

2.2. Structure of StyleS2ST System

Following the setup of [2], we extract the hidden representations of the pseudo-data pairs through the pre-trained HuBERT model, which is trained on a large amount of unlabeled speech data. The extracted hidden representations of speech are clustered by the k-means method to encode speech into discrete units as intermediate representations for the S2UT system. The direct StyleS2ST model composes of speech-to-unit and unit-to-speech modules. In addition, we keep the consecutive repeating units to preserve prosodic information in duration as much as possible.

As illustrated in Figure 1(b), the speech-to-unit module employs a pre-trained wav2vec 2.0 [12] speech encoder and a pre-trained unit BART [13] decoder based on the Transformer [14] architecture, following the approach presented in [5]. In contrast, the unit-to-speech module differs in that we do not directly use the HiFi-GAN vocoder. Instead, we divide the vocoder into two parts: the acoustic model and the vocoder itself. This pipeline is more common in text-to-speech (TTS) and facilitates easier style control. The Fastspeech [15] architecture is utilized as the acoustic model, while the DSPGAN [16] vocoder is used due to its superior robustness in universal scenarios. This is particularly significant for speech quality in zero-shot scenarios.

2.3. Acoustic Model

2.3.1. Style Adaptor

Retaining the duration features in the speech-to-unit module enables the inputs of the acoustic model to express local style

¹<https://StyleS2ST.github.io/StyleS2ST>

features such as prosody. To capture global style features, including speaker timbre, we adopt an utterance-level global style adaptor. However, constructing the speaker space and modeling speaker features for unseen speakers is a critical challenge when implementing zero-shot capability. A common method is to employ a reference encoder, such as the global style token (GST) [17] or variational autoencoder (VAE) [18], to generate a global style embedding from a reference speech. Nonetheless, the speaker space requires a significant amount of training data, and acquiring sufficient paired data is challenging. To address this issue, we utilize a pre-trained style adaptor based on the style embedding extracted from speaker embedding models. Specifically, we employ the ECAPA-TDNN [19] speaker embedding model as the style adaptor and pre-train it using a large amount of Chinese and English data. By training in two languages, we can avoid mismatch issues in speaker space between different languages, which is crucial for our S2UT system.

Since units are frame-level, there is no need for a duration predictor. Instead, we use the units as input to the unit encoder to capture contextual information. Figure 1(c) illustrates the implementation of the style adaptor module, which extracts a style embedding as an additional condition for the acoustic decoder. This enables style control ability and continuous acoustic space modeling.

2.3.2. Enhanced by Unpaired Data

Given the limited number of speakers in the paired data, we employ a large number of unpaired Chinese data to pre-train the acoustic model. Fine-tuning during speaker adaptation can lead to catastrophic forgetting problems [20]. Therefore, we choose not to fine-tune the pre-trained model on our parallel corpus. Instead, we train the pre-trained model using both Chinese speeches in the parallel corpus and unpaired Chinese data, thereby increasing the amount of training data available.

3. Experiments

3.1. Datasets

For all the data used in this section, we downsample the audio to 16 kHz. The frame-level features utilize a 20 ms frame shift and a 50 ms frame length.

3.1.1. En2Zh Parallel Corpus

CoVoST 2 [21] is a commonly used speech-to-text translation dataset that contains 512K English speech and corresponding Chinese text data pairs. However, as the English speech data in CoVoST 2 is not high-fidelity and may result in a loss of speech quality when using speaker adaptation in TTS to construct corresponding Chinese speech, we use only the text data in our experiments. We train a multi-lingual multi-speaker TTS system using an internal dataset consisting of 20 speakers. Each speaker provides approximately 5,000 utterances and 5 hours of high-fidelity expressive Chinese or English data, with varying styles, including stories, broadcasts, etc., and covering speakers of different genders and ages. Using the TTS system, we distribute the Chinese texts in CoVoST 2 and corresponding English texts in Common Voice [22] datasets to these speakers and synthesize the corresponding English-Chinese speech pairs. We then set a threshold for WER and PESQ to filter the data. Specifically, we remove data with a WER higher than 6% or a PESQ below 20. After filtering, we obtain a total of 277K English-Chinese speech pairs, comprising 263 and 124 hours of speech, respectively.

3.1.2. Unpaired Chinese TTS Data

To improve the generalization performance of our model and enable zero-shot capabilities, we employ a large amount of Chinese speech data with diverse acoustic conditions. Specifically, we utilize an internal dataset along with two open-source datasets, AISHELL-3 [23] and DiDiSpeech [24], which contain 500 hours of speech from 1,000 speakers of varying ages and genders, including professional speakers who record high-quality speech in a studio environment and ordinary speakers who record speech in typical rooms with ambient noise and slight reverberation. Table 1 provides details on the datasets. From the total 500 hours of data, we randomly select 300 hours from 1K speakers for pre-training.

Table 1: *Detail of the unpaired Chinese TTS datasets, containing the duration of all samples and number of speakers.*

Dataset	Speakers	Hours
Internal	282	345
AISHELL-3 [23]	218	85
DiDiSpeech [24]	500	70

3.1.3. Test Sets

- **Seen speakers:** To evaluate the performance of our model, we randomly sample 200 utterances from the En2Zh parallel corpus to create separate validation and test sets. We use the test set to assess the semantic translation and style transfer performance for the seen speakers.
- **Unseen speakers:** Due to the limited number of bilingual speakers that can be synthesized by our TTS system, we utilize the English texts in the test set for seen speakers to synthesize speech spoken by English speakers with no overlap with the En2Zh parallel corpus. Specifically, we generate a test set of 100 utterances from 20 speakers, containing expressions in different styles.

3.2. Model Configurations

3.2.1. Evaluation Systems

Due to the lack of previous studies on zero-shot style transfer for direct S2ST, we mainly evaluate the performance of style transfer of our proposed methods compared to the original direct S2ST system. We conduct experiments on follows systems:

- **Direct S2UT:** Original direct S2UT system follows [5]. For consistency with our proposed model, we use acoustic model with vocoder structure instead of the unit decoder as the unit-to-speech module in this system.
- **StyleS2ST:** Our proposed model.
- **StyleS2ST-base:** Our proposed model without using unpaired Chinese data to pre-train the acoustic model. We use this system to validate the effectiveness of the pre-trained acoustic model.
- **GT-StyleS2ST:** System using ground-truth discrete units in the unit-to-speech module of our proposed model. We conduct this system to evaluate the topline of style transfer performance with no semantic error.

3.2.2. Implementation Details

Our direct S2ST system is based on Fairseq [25] and utilizes the Chinese HuBERT model² for feature extraction. We employ the large-scale unsupervised Chinese speech dataset WenetSpeech [26] to train a k-means model with a vocabulary of

²<https://huggingface.co/TencentGameMate/chinese-hubert-base>

Table 2: Experimental results of each system. MOS means naturalness MOS and SMOS means style similarity MOS. Higher is better for all indicators.

Systems	Seen Speakers					Unseen Speakers			
	Units-BLEU	Speech-BLEU	MOS	SMOS	SECS	Speech-BLEU	MOS	SMOS	SECS
Direct S2UT	58.21	37.07	3.66±0.06	3.51±0.05	0.774	18.21	2.58±0.05	2.15±0.06	0.646
StyleS2ST	58.21	39.12	3.76±0.05	3.83±0.06	0.844	30.15	3.43±0.06	3.56±0.05	0.820
StyleS2ST-base	58.21	38.85	3.72±0.05	3.85±0.04	0.838	20.14	2.57±0.05	2.28±0.07	0.688
GT-StyleS2ST	100.00	58.42	4.02±0.04	3.89±0.03	0.840	-	-	-	-
Records	100.00	65.50	4.23±0.03	4.09±0.03	0.885	-	-	-	-

1000 units, using the 11-layer HuBERT feature. Subsequently, we encode all Chinese speech into discrete units using the k-means model. For the speech-to-unit module, we use the large wav2vec 2.0 speech encoder³ for English. Additionally, we utilize the discrete units extracted by the k-means model from the WenetSpeech dataset to train the unit BART decoder with the configuration in [5].

In the unit-to-speech module, we follow the configurations in FastSpeech [15] for acoustic model and DSPGAN-mm in DSPGAN [16] for vocoder. For the style adaptor, we adopt ECAPA-TDNN [19] model configurations and train the model on speaker verification corpus, including the English dataset VoxCeleb 1&2 [27], the Chinese dataset CN-Celeb 1&2 [28], and other English and Chinese datasets⁴, with a total of more than 10K hours and 30K speakers.

3.3. Experimental Results

We evaluate each system for semantic translation and style transfer using various measures. To assess the quality of semantic translation, we first compute the BLEU scores of translated discrete units, referred to as units-BLEU. Next, we use an internal Chinese ASR model to transcribe the synthesized speech to text and calculate the BLEU scores, which we refer to as speech-BLEU. To evaluate the naturalness and style similarity of the synthesized speech, we conduct a mean opinion score (MOS) test across all the test sets for each system, with ten listeners rating each sample on a scale of 1 (worst) to 5 (best) for naturalness and style similarity. Additionally, we use the speaker encoder cosine similarity (SECS) between the speaker embeddings extracted from the synthesized Chinese and the ground-truth English speech as an objective evaluation, with SECS scores ranging from 0 to 1. A higher score indicates higher style similarity.

As shown in Table 2, our direct S2ST systems achieve good semantic translation accuracy. Notably, speech-BLEU is significantly lower than units-BLEU, and comparing the results of records and GT-StyleS2ST, we observe an adverse impact of the ASR system and synthetic speech quality to BLEU scores. While the direct S2ST systems demonstrate generalization in semantic translation on unseen speakers, there is also a decrease in performance. In the direct S2UT system and StyleS2ST-base system, the poor out-of-set generalization ability of the unit-to-speech module results in decreased intelligibility of synthesized speech, which is reflected in the lower speech-BLEU score. Regarding style transfer performance, we observe that the direct S2UT system achieves good style similarity within seen speakers, indicating that discrete units still contain some speaker-related information without the need for additional con-

ditions. Meanwhile, in other systems with a speaker adaptor, speaker similarity, and naturalness are significantly improved, indicating the ability of the speaker adaptor to represent speakers. In the case of unseen speakers, both the direct S2UT and StyleS2ST-base models exhibit a significant decrease in naturalness and style similarity compared to the seen speakers, while our proposed StyleS2ST model maintains a better degree of naturalness and style similarity. These results demonstrate the effectiveness of our pre-trained acoustic model and the direct S2UT and StyleS2ST-base models lacks the generalization performance for out-of-set timbre information.

3.4. Analysis on Acoustic Model Pre-training

To investigate the influence of data size and the speaker number used for pre-training the acoustic model in the zero-shot scenario, we conduct an analysis by varying the data size and speaker number and evaluating the naturalness and style similarity of the synthesized speech.

Table 3: Analysis on acoustic model pre-training in terms of naturalness & style similarity.

System	Naturalness	Style similarity	
	MOS	MOS	SECS
500hr+1000spk	3.40±0.04	3.47±0.05	0.815
300hr+1000spk	3.43±0.05	3.56±0.06	0.820
100hr+1000spk	2.96±0.05	3.41±0.06	0.808
300hr+500spk	3.27±0.06	3.53±0.05	0.826
300hr+250spk	3.23±0.04	3.46±0.04	0.812
StyleS2ST-base	2.57±0.06	2.28±0.07	0.688

Results in Table 3 indicate that both factors have a significant influence on the performance of naturalness and style similarity of the model. It's interesting to observe that there brings a drop in performance when the data size increases from 300 to 500 hours with 1000 speakers. We deduce that this may result from model capacity limitations, which prevent the model from fully converging on the larger dataset.

4. Conclusions

In this paper, we propose StyleS2ST, aiming to achieve zero-shot style transfer in a direct S2ST system. We first build an English-to-Chinese direct S2ST system using a parallel corpus generated by a multi-lingual multi-speaker TTS model. Then we introduce a style adaptor to achieve style transfer and use a large amount of Chinese speech to pre-train the acoustic model to improve the zero-shot ability. Experimental results demonstrate that StyleS2ST can translate English speech to Chinese speech with good preservation of source speaking style to the target, even in zero-shot scenarios. In future work, to further improve speech expressiveness in the translated speech, we would like to explore fine-grained style transfer [29].

³https://dl.fbaipublicfiles.com/fairseq/speech_to_speech/s2st_finetuning/w2v2/en/conformer_L.pt

⁴<https://www.openslr.org/>

5. References

- [1] Y. Jia, R. J. Weiss, F. Biadys, W. Macherey, M. Johnson, Z. Chen, and Y. Wu, "Direct speech-to-speech translation with a sequence-to-sequence model," in *Interspeech 2019*. ISCA, 2019, pp. 1123–1127.
- [2] Y. Jia, M. T. Ramanovich, T. Remez, and R. Pomerantz, "Translatotron 2: High-quality direct speech-to-speech translation with voice preservation," in *ICML 2022*, vol. 162. PMLR, 2022, pp. 10 120–10 134.
- [3] A. Lee, P. Chen, C. Wang, J. Gu, S. Popuri, X. Ma, A. Polyak, Y. Adi, Q. He, Y. Tang, J. Pino, and W. Hsu, "Direct speech-to-speech translation with discrete units," in *ACL 2022*. Association for Computational Linguistics, 2022, pp. 3327–3339.
- [4] S. Nakamura, K. Markov, H. Nakaiwa, G. Kikui, H. Kawai, T. Jitsuhiro, J. Zhang, H. Yamamoto, E. Sumita, and S. Yamamoto, "The ATR multilingual speech-to-speech translation system," *IEEE Trans. Speech Audio Process.*, vol. 14, no. 2, pp. 365–376, 2006.
- [5] S. Popuri, P. Chen, C. Wang, J. Pino, Y. Adi, J. Gu, W. Hsu, and A. Lee, "Enhanced direct speech-to-speech translation using self-supervised pre-training and data augmentation," in *Interspeech 2022*. ISCA, 2022, pp. 5195–5199.
- [6] W. Hsu, B. Bolte, Y. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE ACM Trans. Audio Speech Lang. Process.*, vol. 29, pp. 3451–3460, 2021.
- [7] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *ICLR 2021*. OpenReview.net, 2021.
- [8] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," in *NeurIPS 2020*, vol. 33, 2020, pp. 17 022–17 033.
- [9] Y. Zhang, R. J. Weiss, H. Zen, Y. Wu, Z. Chen, R. J. Skerry-Ryan, Y. Jia, A. Rosenberg, and B. Ramabhadran, "Learning to speak fluently in a foreign language: Multilingual speech synthesis and cross-language voice cloning," in *Interspeech 2019*. ISCA, 2019, pp. 2080–2084.
- [10] H. Hemati and D. Borth, "Using ipa-based tacotron for data efficient cross-lingual speaker adaptation and pronunciation enhancement," *CoRR*, vol. abs/2011.06392, 2020.
- [11] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs," in *ICASSP 2001*. IEEE, 2001, pp. 749–752.
- [12] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "Wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NeurIPS 2020*, vol. 33, 2020, pp. 12 449–12 460.
- [13] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer, "BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension," in *ACL 2020*. Association for Computational Linguistics, 2020, pp. 7871–7880.
- [14] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *NeurIPS 2017*, 2017, pp. 5998–6008.
- [15] Y. Ren, Y. Ruan, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T. Liu, "Fastspeech: Fast, robust and controllable text to speech," in *NeurIPS 2019*, 2019, pp. 3165–3174.
- [16] K. Song, Y. Zhang, Y. Lei, J. Cong, H. Li, L. Xie, G. He, and J. Bai, "DSPGAN: a gan-based universal vocoder for high-fidelity TTS by time-frequency domain supervision from DSP," *CoRR*, vol. abs/2211.01087, 2022.
- [17] Y. Wang, D. Stanton, Y. Zhang, R. J. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, Y. Jia, F. Ren, and R. A. Saurous, "Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis," in *ICML 2018*, vol. 80. PMLR, 2018, pp. 5167–5176.
- [18] W. Hsu, Y. Zhang, R. J. Weiss, H. Zen, Y. Wu, Y. Wang, Y. Cao, Y. Jia, Z. Chen, J. Shen, P. Nguyen, and R. Pang, "Hierarchical generative modeling for controllable speech synthesis," in *ICLR 2019*. OpenReview.net, 2019.
- [19] B. Desplanques, J. Thienpondt, and K. Demuynck, "ECAPA-TDNN: emphasized channel attention, propagation and aggregation in TDNN based speaker verification," in *Interspeech 2020*. ISCA, 2020, pp. 3830–3834.
- [20] H. Hemati and D. Borth, "Continual speaker adaptation for text-to-speech synthesis," *CoRR*, vol. abs/2103.14512, 2021. [Online]. Available: <https://arxiv.org/abs/2103.14512>
- [21] C. Wang, A. Wu, J. Gu, and J. Pino, "Covost 2 and massively multilingual speech translation," in *Interspeech 2021*. ISCA, 2021, pp. 2247–2251.
- [22] R. Ardila, M. Branson, K. Davis, M. Kohler, J. Meyer, M. Henretty, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," in *LREC 2020*. European Language Resources Association, 2020, pp. 4218–4222.
- [23] Y. Shi, H. Bu, X. Xu, S. Zhang, and M. Li, "AISHELL-3: A multi-speaker mandarin TTS corpus," in *Interspeech 2021*. ISCA, 2021, pp. 2756–2760.
- [24] T. Guo, C. Wen, D. Jiang, N. Luo, R. Zhang, S. Zhao, W. Li, C. Gong, W. Zou, K. Han, and X. Li, "Didispeech: A large scale mandarin speech corpus," in *ICASSP 2021*. IEEE, 2021, pp. 6968–6972.
- [25] M. Ott, S. Edunov, A. Baevski, A. Fan, S. Gross, N. Ng, D. Grangier, and M. Auli, "Fairseq: A fast, extensible toolkit for sequence modeling," in *NAACL-HLT 2019*. Association for Computational Linguistics, 2019, pp. 48–53.
- [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*. IEEE, 2022, pp. 6182–6186.
- [27] A. Brown, J. Huh, J. S. Chung, A. Nagrani, and A. Zisserman, "Voxsrc 2021: The third voxceleb speaker recognition challenge," *CoRR*, vol. abs/2201.04583, 2022.
- [28] Y. Fan, J. W. Kang, L. T. Li, K. C. Li, H. L. Chen, S. T. Cheng, P. Y. Zhang, Z. Y. Zhou, Y. Q. Cai, and D. Wang, "Cn-celeb: A challenging chinese speaker recognition dataset," in *ICASSP 2020*. IEEE, 2020, pp. 7604–7608.
- [29] Y. Lei, S. Yang, X. Wang, and L. Xie, "MsEmoTTS: Multi-scale emotion transfer, prediction, and control for emotional speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 853–864, 2022.