

# Improving Pronunciation and Accent Conversion through Knowledge Distillation And Synthetic Ground-Truth from Native TTS

Tuan Nam Nguyen\*, Seymanur Akti†, Ngoc Quan Pham‡, Alexander Waibel§

\*†‡§ Karlsruhe Institute of Technology, Karlsruhe, Germany

§ Carnegie Mellon University, Pittsburgh PA, USA

\*tuan.nguyen@kit.edu, †seymanur.akti@kit.edu, ‡ngoc.pham@kit.edu, §alexander.waibel@kit.edu

**Abstract**—Previous approaches on accent conversion (AC) mainly aimed at making non-native speech sound more native while maintaining the original content and speaker identity. However, non-native speakers sometimes have pronunciation issues, which can make it difficult for listeners to understand them. Hence, we developed a new AC approach that not only focuses on accent conversion but also improves pronunciation of non-native accented speaker. By providing the non-native audio and the corresponding transcript, we generate the ideal ground-truth audio with native-like pronunciation with original duration and prosody. This ground-truth data aids the model in learning a direct mapping between accented and native speech. We utilize the end-to-end VITS framework to achieve high-quality waveform reconstruction for the AC task. As a result, our system not only produces audio that closely resembles native accents and while retaining the original speaker’s identity but also improve pronunciation, as demonstrated by evaluation results.

**Index Terms**—accent conversion, voice conversion, speech synthesis

## I. INTRODUCTION

Second-language (L2) English learners typically present accents and mispronunciations, which can greatly impact their communication proficiency. AC basically aims to change the accent of speech while retaining the information of content and speaker identity. However, this task remains particularly challenging due to the extreme lack of parallel corpora and the large variations in different speakers. This research focuses on investigating techniques to enhance AC models’ capacity to improve pronunciation made by L2 speakers and devising appropriate evaluation measures for this purpose.

Conventional AC [1], [2] methods require reference utterances in the target accent during synthesis, limiting their applicability due to the challenge of obtaining linguistically identical references in different accents. This research, like recent studies [3]–[7], focuses on reference-free AC, which converts accents without needing reference utterances during inference. Reference-free AC methods generally fall into two categories: autoregressive and non-autoregressive architectures.

Autoregressive AC models, typically based on sequence-to-sequence (seq2seq) architectures, rely mostly on parallel data for training. However, such data—utterances from the same speakers in different accents—is scarce, making it difficult to obtain. To address this, data augmentation techniques like text-to-speech (TTS) [8] or voice conversion [9] are used to generate ground-truth audio for non-native speakers. These

ground-truth audios, originating from different speakers, often differ in duration and prosody compared to the original non-native audio. While the encoder-decoder model with attention mechanisms can help aligning inputs and outputs, the differences in length make these models less suitable for applications requiring precise time synchronization, such as video conference dubbing [10]–[12]. Additionally, there are some autoregressive models [13], [14] that trained with non-parallel data, but still share similar drawbacks, including slow inference speeds and unstable attention mechanism.

In contrary, non-autoregressive AC leverages non-parallel data, which is abundant and diverse. To effectively utilize this data, these models are built on a disentangle-resynthesis framework. Here, speech is disentangled into separate features such as speaker identity, content, prosody, and accent, which are then recombined to generate the original waveform. Content features are often represented by bottleneck features (BNFs), extracted from self-supervised pre-trained models [15] (e.g., WavLM [16], Wav2vec [17], HuBERT [18]), ASR bottlenecks [3], or ASR output in logits form [19]. However, BNFs also capture accent-related feature and non-native pronunciation issues. While methods like adversarial training [19] or Pseudo-Siamese networks [3] attempt to disentangle accent information, they still struggle to improve pronunciations due to the absence of ground-truth.

Building on the strengths and limitations of existing AC models, we propose a novel framework for training a non-autoregressive AC model using generated parallel data. We hypothesize that a TTS system trained solely on native speech will produce accent-independent linguistic representations. Additionally, this native TTS system is able to generate ideal ground-truth data for non-native speakers, ensuring native pronunciation, same speaker identity, duration, prosody, and precise alignment with the original non-native audio. The accent-independent linguistic representations learned by the native TTS model not only facilitate ground-truth generation but also distill knowledge into our AC model, aiding in the learning of accent-independent features.

Previous studies [13], [20] have used native TTS models to guide AC models, often by sharing the same decoder and distilling knowledge from the TTS text encoder to the AC encoder. A key challenge in these approaches is aligning the outputs of the TTS and AC encoders due to their differing

lengths. In [13], both models use the autoregressive Tacotron [21] architecture, necessitating an external unstable attention mechanism for alignment. Another approach [20] employs the non-autoregressive FastSpeech 2 [22] TTS model. In this approach, alignment is achieved through an external length regulator based on Montreal Force Alignment, which uses a GMM-HMM ASR model to upsample the TTS encoder’s output, enabling both encoders to produce outputs of similar length. This latter method closely aligns with our proposed approach. These approaches use a separate vocoder to convert mel-spectrograms into waveforms.

The non-autoregressive VITS framework has proven successful in both TTS [23] and Voice Conversion [15]. Inspired by this success, we adopt a modified VITS framework for our AC system due to several advantages. VITS uses internal monotonic alignment, simplifying audio-text alignment and improving training efficiency. Additionally, as an end-to-end architecture, VITS eliminates the need for a separate vocoder to convert mel-spectrograms into waveforms, setting it apart from previous AC methods.

Overall, to further enhance the AC model’s training process, we introduced several key improvements. In the first stage, we simultaneously pre-train the AC model while training the VITS-based native TTS, allowing the AC model to learn the distribution of native audio data comprehensively. After training the native VITS-based TTS, we leverage it to generate ideal ground-truth data for non-native audio, as described in the next section. In the second stage, we fine-tune the pre-trained AC model using the non-native input alongside the generated ground-truth output and also distillate knowledge from the native TTS. In summary, we train the VITS-based TTS model to assist in initializing the weights, generating ideal ground-truth data, and facilitating knowledge distillation for the AC model.

## II. METHODOLOGY

### A. Training native VITS and pretrained AC model

In the first stage, we pre-train the AC model and train the native TTS model using the VITS framework. VITS [23] is a conditional variational autoencoder architecture augmented with normalizing flow [24]. It consists of three main components: a posterior encoder, a prior encoder, and a waveform generator. These modules encode the distributions  $q_\phi(z|x)$ ,  $p_\theta(z|c)$ , and  $p_\psi(y|z)$ , respectively. Here,  $q_\phi(z|x)$  and  $p_\psi(y|z)$  represent the posterior and data distributions, parameterized by the posterior encoder  $\phi$  and HiFi-GAN [25] waveform generator  $\psi$ , where  $x$  is the speech input,  $z$  is the latent variable and  $y$  is the waveform output. The prior distribution of  $z$  is defined as  $p_\theta(z|c)$ , parameterized by the prior encoder  $\theta$  and refined through a normalizing flow  $f$ , where the latent variables are conditioned on the input  $c$ , which can be either audio or text.

Our AC and TTS models share the HiFi-GAN decoder, posterior encoder, speaker encoder, F0 encoder and normalizing flow, as shown in Fig. 1a. The posterior encoder takes linear spectrograms  $x_{lin}$  as input to sample a latent variable  $z$ .

This latent  $z$ , along with speaker embedding  $g$  from speaker encoder and F0 sequence embeddings  $F0$  from F0 encoder, are fed into the HiFi-GAN vocoder to generate the waveform.

The AC and TTS models each have separate prior encoders. The AC model’s prior encoder  $\theta_{audio}$  includes a content encoder which extracts speech features and a bottleneck extractor to generate a normal distribution out of BNFs. Previous works have utilized self-supervised pre-trained representations [15] or ASR features (bottleneck or logits) [19] as content encoder outputs. These features contain information related to content, speaker, and accent. To filter out unwanted information, the *BNFs* are processed by the bottleneck extractor. The TTS model’s prior encoder consists of a text encoder  $\theta_{text}$ , which takes the transcript as input. Monotonic Alignment Search (MAS) aligns the text encoder output with the normalizing flow output, to match the differing lengths for text representations and audio representations.

The training loss is split into Variational Auto-encoder (VAE) loss and Generative Adversarial Network (GAN) loss. GAN loss includes adversarial loss  $L_{adv}(D)$  and  $L_{adv}(G)$  for the discriminator and generator, along with feature matching loss  $L_{fm}(D)$  for the generator  $G$ . VAE loss consists of reconstruction loss  $L_{rec}$  (L1 distance between target and predicted mel-spectrogram) and  $L_{kl}$  (KL divergence between the prior and the posterior). To address training-inference mismatch of VITS, we integrate the full end-to-end inference during training to improve voice quality. Specifically, L1 distance between ground-truth  $x_{mel}$  and inference output  $x_{mel}^{e2e}$  generated by  $p_\psi(x|f^{-1}(z'), g, F0)$ , where the latent  $z'$  is sampled from  $p_{\theta_{audio}}(z'|c_{audio})$  or  $p_{\theta_{text}}(z'|c_{text})$ , and both  $g$  and  $F0$  are extracted from speaker encoder and F0 encoder respectively, is added to reconstruction loss. KL loss  $L_{kl}$ , which sums the KL divergence between the prior distributions  $p_{\theta_{text}}(z|c_{text})$  and  $p_{\theta_{audio}}(z|c_{audio})$ , and the posterior distribution  $q_\phi(z|x)$ . The overall training loss for pre-training phase can be expressed:

$$L_{kl} = KL(q_\phi(z|x)||p_{\theta_{text}}(z|c_{text})) + KL(q_\phi(z|x)||p_{\theta_{audio}}(z|c_{audio})) \quad (1)$$

$$L_{recon} = ||x_{mel} - \hat{x}_{mel}|| + ||x_{mel} - x_{mel}^{e2e}|| \quad (2)$$

$$L(G) = L_{rec} + L_{kl} + L_{adv}(G) + L_{fm}(G) \quad (3)$$

$$L(D) = L_{adv}(D) \quad (4)$$

### B. Generating ideal ground-truth

The native TTS model trained earlier is used to generate ground-truth audio for each non-native input. We start by sampling the latent variable  $z$  from the posterior distribution of the non-native audio  $x$ . We use MAS to find the alignment  $A$  to maximize the log-likelihood of  $f(z)$  on the distribution  $p_{\theta_{text}}(z'|c_{text}, A)$ , which is used to upsample the text input.

$$z \sim q_\phi(z|x_{non-native}) = N(z; \mu_\phi(x), \sigma_\phi(x)) \quad (5)$$

$$A = \operatorname{argmax} \log N(f(z); \mu_{\theta_{text}}(c_{text}, A), \sigma_{\theta_{text}}(c_{text}, A)) \quad (6)$$

$$c_{upsample} = \operatorname{upsampling}(c_{text}, A) \quad (7)$$

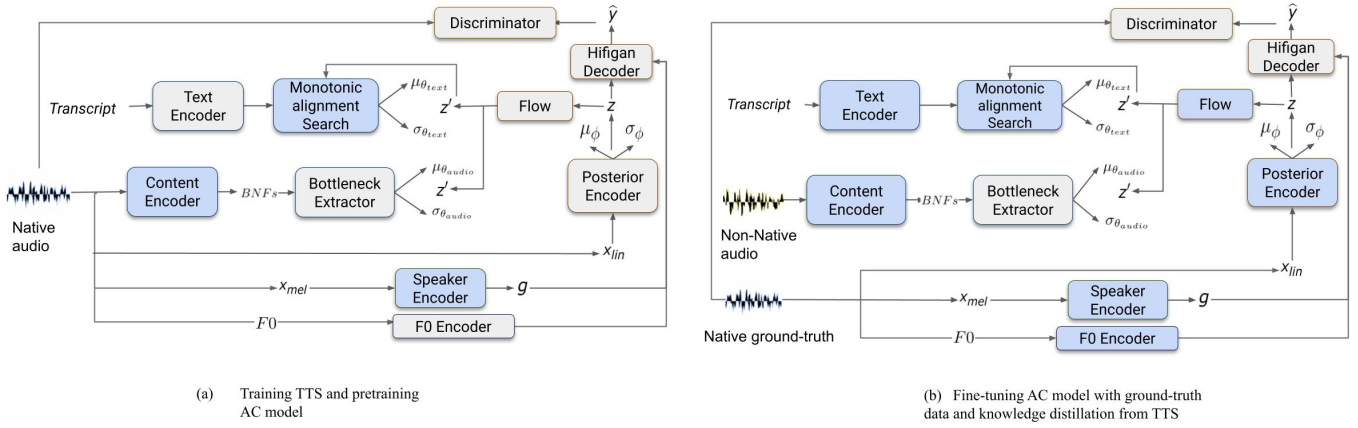


Fig. 1. Pre-training and fine-tuning procedure of our proposed model. The parameters of blue components are frozen.

Next, similar to the inference process of VITS, involves obtaining the latent variable  $z'$  by sampling from  $p_{\theta_{text}}(z' | c_{upsample})$ , which is then refined through an inverted normalizing flow  $f^{-1}$ . Finally, the ground-truth audio  $\hat{y}_{ground-truth}$  is generated from  $p_{\psi}(x | f^{-1}(z'), g, F0)$  using the HiFi-GAN decoder, where speaker embedding  $g$  and  $F0$  are extracted from non-native audio. In summary, the synthetic ground truth  $y_{ground-truth}$  is generated from the transcripts with perfect native pronunciation, while utilizing the MAS alignment,  $F0$  and  $g$  from the non-native audio to retain the original duration, prosody and speaker identity.

$$z' \sim p_{\theta_{text}}(z' | c_{upsample}) \quad (8)$$

$$y_{ground-truth} = HiFiGAN(f^{-1}(z'), g, F0) \quad (9)$$

### C. Finetuning AC model with ideal ground-truth and knowledge distillation from native TTS

In the second stage, we continue fine-tuning the AC model using synthetic ground-truth data. The primary objective is to enable the model to generate ideal native output from non-native input. Specifically, the content encoder processes the non-native input, while other components take the synthetic ground-truth, as illustrated in Figure 1b. During this process, we freeze the parameters of all components except the bottleneck extractor and HiFi-GAN decoder. The bottleneck extractor is fine-tuned to capture accent-independent content representations from non-native inputs. To facilitate this, we introduce a distillation loss into the training, which is the KL divergence between the two prior distributions  $p_{\theta_{text}}(z | c_{text})$  and  $p_{\theta_{audio}}(z | c_{audio})$ . The training loss during the fine-tuning stage is defined as:

$$L_{distill} = KL(p_{\theta_{audio}}(z | c_{audio}) || p_{\theta_{text}}(z | c_{text})) \quad (10)$$

$$L(G) = L_{rec} + L_{kl} + L_{distill} + L_{adv}(G) + L_{fm}(G) \quad (11)$$

### D. Overall model architecture detail

Our posterior encoder, HiFi-GAN decoder, normalizing flow, and text prior encoder follow the original VITS framework [23]. For the content encoder, we fine-tuned a pretrained Wav2vec 2.0 model on the Librispeech and L2Arctic dataset

using CMU-dict phoneme labels, utilizing the bottleneck of the final layer as the content representation. Afterward, the content encoder's parameters are frozen. To get better text-audio alignment, we input phoneme labels instead of text into the text encoder. Both bottleneck extractor and text encoder are the same as the VITS Transformer-based text encoder. We retain Wav2vec 2.0's downsampling rate of 320 for calculating mel-spectrograms, F0, and linear spectrograms, while the HiFi-GAN decoder uses a corresponding upsample rate of 320. F0 sequences are derived using the YAAPT algorithm with the same downsampling rate, while the F0 encoder is based on [26]. We employ the pre-trained speaker encoder from [27].

### E. Inference

Our framework supports inference with or without a transcript. Without a transcript, the model extracts content information from non-native audio using content encoder and the bottleneck extractor in the prior encoder, then generates a waveform from the latent variable sampled from the prior distribution, conditioned on the speaker and prosody embedding. When a transcript is provided during inference, the system can utilize the prior text encoder instead of the audio encoder, as described in Section II-B, to generate output with native pronunciation. This dual inference approach enhances inference quality, particularly when a transcript is available.

## III. EXPERIMENT AND RESULT

### A. Data

We use the LJspeech dataset [28], which features consistent pronunciation from a single native speaker. Using FreeVC [15], a voice conversion model, we generate multispeaker native-accented utterances from the original voices. This augmented multi-speaker dataset is utilized, enabling the system to be trained in a multi-speaker setting during the first stage. In the second stage, we fine-tune the AC model on the L2ARCTIC dataset [29], featuring 24 accented speakers across 6 accents. For each accent, we select 3 speakers for training and the remaining speakers for testing. Each speaker's utterances are split into a training set of 1,032 non overlap utterances, a validation set of 50, and a test set of 50. The test set, chosen using our competitive ASR model [30], focuses on

utterances with the high average WER (larger than 10) across all speakers, with assumption that the higher WER rates mean stronger accents.

### B. Evaluation metrics

#### 1) Subjective tests: Nativeness and Speaker Similarity

**Tests:** Ten participants conducted two evaluations using a 5-point scale (1-bad, 2-poor, 3-fair, 4-good, 5-excellent). In the Nativeness test, they rated how closely the converted audios resembled native speech. In the speaker similarity test, participants assessed the similarity between the voice identity of the original input and the converted audio, providing a score out of 5 for each.

2) **Objective tests: Word error rates (WER), Accent classifier accuracy (ACC) and Speaker Embedding Cosine Similarity (SECS)** : To evaluate the improvement in pronunciation, we use the Word Error Rate (WER) from our competitive seq2seq Transformer ASR model [30]. A lower WER indicates better intelligibility in the conversion process. SECS is employed to measure speaker similarity by computing the cosine similarity between speaker embeddings of the original speech and the converted speech. We utilize the state-of-the-art speaker verification model from [16] to extract speaker embeddings, which suggests that a cosine similarity greater than 0.85 indicates that both audios likely originate from the same speaker.

Additionally, we train an accent classifier to determine whether an input audio is native or non-native. The classifier architecture and training setting is similar to [9]. We compute two accuracy scores, one for original non-native audio set and one for converted native audio set. A better AC is expected to show a larger gap between these two ACCs.

### C. Experimental setup

Trained on both native and non-native audio with ASR loss, the content encoder’s final layer yields representations closely aligned with linguistic content. This makes it more accent- and speaker-independent, facilitating accurate mapping from non-native pronunciation to the correct linguistic information. Therefore, we use the model which is pre-trained in the first stage as the baseline. The proposed model is the model fine-tuned in the second stage. To assess the effectiveness of fine-tuning on the synthetic ground-truth data, we train a variant that only uses knowledge distillation loss  $L_{distill}$  between the text and audio encoders in the second stage. Additionally, the audio quality of the Synthetic ground-truth is also evaluated, allowing us to assess the system’s performance when the correct transcript is provided. The training hyperparameters are set similarly to those in the original VITS, with the system trained for 600,000 steps in the first stage and 200,000 steps in the second stage. Sample evaluation audios are available at a github repository <sup>1</sup>

### D. Result

The objective metric evaluation, shown in Table II, indicates that the proposed method outperforms others in terms of WER.

<sup>1</sup><https://accentconversion.github.io/demo>

This demonstrates that the method successfully improves pronunciation, making it closer to native speech. Even without synthetic ground-truth, the model still benefits from knowledge distillation from the native TTS, although the pronunciation improvement is smaller. The synthetic ground-truth, generated from transcripts, has a WER of 5.1, confirming the high quality of these ground-truth audios. Additionally, SECS remains stable between 0.82 and 0.84 across all settings, verifying that the speaker identity is well-preserved in all methods. In terms of ACC, the results are not as strong as other metrics. Since we aim to preserve the prosody of the original audio, which can sometimes indicate the accent of the speaker, the model occasionally identifies the converted audio as non-native. This prosody retention may lead to less effective accent conversion in some cases. The subjective metric evaluation, also shown in Table I, indicates that our proposed method performs best in the nativeness test, while the Sim-MOS results, like in the objective evaluation, remain similar across all setups.

TABLE I  
SUBJECTIVE METRICS

Models	Nativeness	Sim-MOS
Original	1.67 ± 0.05	-
Synthetic ground-truth	3.93 ± 0.05	3.84 ± 0.15
Baseline	3.32 ± 0.08	3.94 ± 0.18
w/o Synthetic ground-truth	3.71 ± 0.08	3.9 ± 0.19
Proposed	<b>3.87</b> ± 0.08	3.85 ± 0.19

TABLE II  
OBJECTIVE METRICS

Models	WER	ACC	SECS
Original	18.3	98.3	1.0
Synthetic ground-truth	5.1	13.9	0.83
Baseline	17.1	18.9	0.82
w/o Synthetic ground-truth	14.3	17.0	0.84
Proposed	<b>12.4</b>	17.2	0.83

## IV. CONCLUSION

Our work presents a significant advancement in accent conversion and native-like pronunciation correction. Through the introduction of a two-stage training process that leverages VITS framework, we enable the AC model to better capture the accent-independent feature of speech while maintaining the speaker’s identity. The use of generated ground-truth data, alongside knowledge distillation from native TTS, further enhances the system’s ability to correct non-native pronunciations effectively. Objective and subjective evaluations demonstrate that our approach not only improves accent conversion but also addresses pronunciation issues, providing a more natural and comprehensible output.

## V. ACKNOWLEDGMENT

This research was supported in part by a grant from Zoom Video Communications, Inc. T; European Commission Project Meetween (101135798), the Federal Ministry of Education and Research (BMBF) of Germany under the number 01EF1803B (RELATER), and the pilot program Core-Informatics of the Helmholtz Association (HGF). The work is enabled by the Horeka supercomputer (BMBF) and the DeltaGPU at NCSA/ACCESS grant [31].

## REFERENCES

- [1] G. Zhao, S. Sonsaat, J. Levis, E. Chukharev-Hudilainen, and R. Gutierrez-Osuna, "Accent conversion using phonetic posteriorgrams," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5314–5318.
- [2] G. Zhao, S. Ding, and R. Gutierrez-Osuna, "Foreign Accent Conversion by Synthesizing Speech from Phonetic Posteriorgrams," in *Proc. Interspeech 2019*, 2019, pp. 2843–2847.
- [3] D. Jia, Q. Tian, K. Peng, J. Li, Y. Chen, M. Ma, Y. Wang, and Y. Wang, "Zero-Shot Accent Conversion using Pseudo Siamese Disentanglement Network," in *Proc. INTERSPEECH 2023*, 2023, pp. 5476–5480.
- [4] M. Jin, P. Serai, J. Wu, A. Tjandra, V. Manohar, and Q. He, "Voice-preserving zero-shot multiple accent conversion," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [5] G. Zhao, S. Ding, and R. Gutierrez-Osuna, "Converting foreign accent speech without a reference," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2367–2381, 2021.
- [6] W.-C. Huang and T. Toda, "Evaluating methods for ground-truth-free foreign accent conversion," 2023. [Online]. Available: <https://arxiv.org/abs/2309.02133>
- [7] A. Waibel and C. Fuegen, "Simultaneous translation of open domain lectures and speeches," Jan. 3 2012, uS Patent 8,090,570.
- [8] T.-N. Nguyen, N.-Q. Pham, and A. Waibel, "Syntacc : Synthesizing multi-accent speech by weight factorization," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [9] T. N. Nguyen, N.-Q. Pham, and A. Waibel, "Accent Conversion using Pre-trained Model and Synthesized Data from Voice Conversion," in *Proc. Interspeech 2022*, 2022, pp. 2583–2587.
- [10] A. Waibel, H. Steusloff, R. Stiefelbogen *et al.*, "Chil: Computers in the human interaction loop," 2005.
- [11] A. Waibel, M. Bett, M. Finke, and R. Stiefelbogen, "Meeting browser: Tracking and summarizing meetings," in *Proceedings of the DARPA broadcast news workshop*, 1998, pp. 281–286.
- [12] F. Huang, S. Vogel, and A. Waibel, "Automatic extraction of named entity translanguing equivalence based on multi-feature cost minimization," in *Proceedings of the ACL 2003 workshop on Multilingual and mixed-language named entity recognition*, 2003, pp. 9–16.
- [13] Y. Zhou, Z. Wu, M. Zhang, X. Tian, and H. Li, "Tts-guided training for accent conversion without parallel data," *IEEE Signal Processing Letters*, vol. 30, pp. 533–537, 2023.
- [14] S. Liu, D. Wang, Y. Cao, L. Sun, X. Wu, S. Kang, Z. Wu, X. Liu, D. Su, D. Yu, and H. Meng, "End-to-end accent conversion without using native utterances," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 6289–6293.
- [15] J. Li, W. Tu, and L. Xiao, "Freevc: Towards hifigagh-quality text-free one-shot voice conversion," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [16] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, J. Wu, M. Zeng, X. Yu, and F. Wei, "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.
- [17] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 12 449–12 460.
- [18] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 29, p. 3451–3460, oct 2021. [Online]. Available: <https://doi.org/10.1109/TASLP.2021.3122291>
- [19] M. Jin, P. Serai, J. Wu, A. Tjandra, V. Manohar, and Q. He, "Voice-preserving zero-shot multiple accent conversion," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [20] X. Chen, J. Pei, L. Xue, and M. Zhang, "Transfer the linguistic representations from tts to accent conversion with non-parallel data," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 12 501–12 505.
- [21] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgianakis, and Y. Wu, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4779–4783.
- [22] Y. Ren, C. Hu, X. Tan, T. Qin, S. Zhao, Z. Zhao, and T.-Y. Liu, "Fastspeech 2: Fast and high-quality end-to-end text to speech," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=piLPYqxtWuA>
- [23] J. Kim, J. Kong, and J. Son, "Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech," in *International Conference on Machine Learning*, 2021.
- [24] D. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 1530–1538. [Online]. Available: <https://proceedings.mlr.press/v37/rezende15.html>
- [25] J. Kong, J. Kim, and J. Bae, "Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 17 022–17 033.
- [26] D. Wang, L. Deng, Y. T. Yeung, X. Chen, X. Liu, and H. Meng, "Vqmvic: Vector quantization and mutual information-based unsupervised speech representation disentanglement for one-shot voice conversion," in *Interspeech 2021*, 2021, pp. 1344–1348.
- [27] S. Liu, Y. Cao, D. Wang, X. Wu, X. Liu, and H. Meng, "Any-to-many voice conversion with location-relative sequence-to-sequence modeling," *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, vol. 29, p. 1717–1728, apr 2021. [Online]. Available: <https://doi.org/10.1109/TASLP.2021.3076867>
- [28] K. Ito and L. Johnson, "The lj speech dataset," <https://keithito.com/LJ-Speech-Dataset/>, 2017.
- [29] G. Zhao, S. Sonsaat, A. Silpachai, I. Lucic, E. Chukharev-Hudilainen, J. Levis, and R. Gutierrez-Osuna, "L2-arctic: A non-native english speech corpus," in *Interspeech 2018*, 2018, pp. 2783–2787.
- [30] N.-Q. Pham, T.-L. Ha, T.-N. Nguyen, T.-S. Nguyen, E. Salesky, S. Stüker, J. Niehues, and A. Waibel, "Relative positional encoding for speech recognition and direct translation," in *Interspeech 2020*, 2020, pp. 31–35.
- [31] T. J. Boerner, S. Deems, T. R. Furlani, S. L. Knuth, and J. Towns, "Access: Advancing innovation: Nsf's advanced cyberinfrastructure coordination ecosystem: Services & support," in *Practice and Experience in Advanced Research Computing*, 2023, pp. 173–176.