

# VOICE-PRESERVING ZERO-SHOT MULTIPLE ACCENT CONVERSION

Mumin Jin<sup>1,2</sup>, Prashant Serai<sup>1</sup>, Jilong Wu<sup>1</sup>, Andros Tjandra<sup>1</sup>, Vimal Manohar<sup>1</sup>, Qing He<sup>1</sup>

<sup>1</sup>Meta AI, <sup>2</sup> MIT

## ABSTRACT

Most people who have tried to learn a foreign language would have experienced difficulties understanding or speaking with a native speaker’s accent. For native speakers, understanding or speaking a new accent is likewise a difficult task. An accent conversion system that changes a speaker’s accent but preserves that speaker’s voice identity, such as timbre and pitch, has the potential for a range of applications, such as communication, language learning, and entertainment. Existing accent conversion models tend to change the speaker identity and accent at the same time. Here, we use adversarial learning to disentangle accent dependent features while retaining other acoustic characteristics. What sets our work apart from existing accent conversion models is the capability to convert an unseen speaker’s utterance to multiple accents while preserving its original voice identity. Subjective evaluations show that our model generates audio that sound closer to the target accent and like the original speaker.

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

## 1. INTRODUCTION

For many people, their own or others’ accents present a severe obstacle to communication. For some other people, watching a movie set in America with British accented speakers creates a dissonance that takes them out of the immersive experience. Therefore, a system that converts a speaker’s accent yet still preserves the original voice identity can have a great impact in a wide range of situations including communication, language learning, and entertainment.

The main challenge in voice-preserving accent conversion is the need to disentangle features related to a speaker’s voice, accent, and linguistic contents. Usually, each speaker in a data set only speaks with one accent, and there is relative scarcity in the available number of speakers and audio recordings for non-native accents. Previous works have used adversarial learning to disentangle features [1, 2], where both apply a discriminator to wipe-out speaker dependent information from content embeddings. Other works, such as [3, 4],

achieve disentanglement via quantization of different features to obscure undesired information.

Conventional accent conversion approaches require the availability of reference utterances with the same text with target accents, during synthesis [5, 6, 7, 8]. The applications of these approaches are very limited, as we often do not have access to reference utterances with the same linguistic content in a different accent. Recently, [9, 10, 11] have proposed systems that convert accents without needing a reference utterance during inference. The systems proposed in both [9] and [10], however, are not zero-shot accent conversion systems because they require further training on the input utterances. In the case of [9], the ASR component needed to be fine-tuned with the input speaker’s utterances, and in the case of [10], a dedicated model for each new speaker must be trained on parallel speech. The authors of [9] also acknowledge that their converted utterances were perceived to have a different voice from that of the original utterance. Our proposed system differs from the zero-shot, reference-free accent conversion system from [11] in that our system allows for synchronous accent conversion, and to multiple native and non-native accents.

Our proposed system is most similar to the accent conversion models proposed in [9], [11], and [2]. Unlike these works, our model converts unseen utterances with arbitrary accent to utterances with multiple target accents. Listeners from our perceptual tests agree that our model is good at preserving the original speaker’s voice characteristics. In our model, we utilize a pre-trained model checkpoint to extract speaker and accent independent text predictions prior to training. We further disentangle the accent-dependent features from other features with an accent discriminator. Finally, the processed, disentangled features are re-combined and fed to a HiFiGAN decoder to reconstruct the audio waveform. Our work make contributions in three major ways: 1) To the best of our knowledge, our model is the first to convert arbitrary accented, unseen speech to multiple target accents while preserving non-accent related voice characteristics. 2) We do not require text labels associated with accented speech or speaker ID labels during training, although we use an existing ASR model checkpoint trained on native accented English speech to extract linguistic features. 3) We convert accent while keeping the output synchronized to input, allowing for applications such as dubbing a video with different accents.

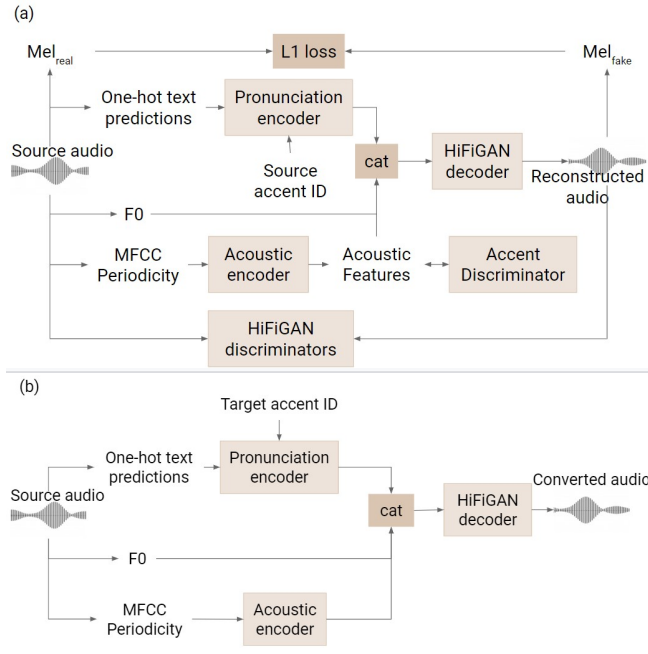
Work was done when Mumin was an intern at Meta AI  
Correspondence to Prashant Serai: pserai@meta.com

## 2. PROPOSED SYSTEM

Our proposed system during training and inference is shown in Fig. 1. Unlike [2] and [9], we do not train an accented ASR model with text labels corresponding to accented speech. Instead, we use an off-the-shelf wav2vec2.0 checkpoint<sup>1</sup> that has been pre-trained using self-supervised learning and fine-tuned on ASR task using 960 hours of LibriSpeech data[12].

Let  $\mathbf{x}$  denote the input audio waveform,  $\hat{\mathbf{x}}$  the output audio, and  $M$  the function that transforms a waveform into the corresponding 80-dimensional mel-spectrogram. We train our model to minimize the reconstruction loss in (1)

$$L_{mel} = \mathbb{E}_{\mathbf{x} \in \mathcal{X}} [\|M(\mathbf{x}) - M(\hat{\mathbf{x}})\|_1] \quad (1)$$




**Fig. 1.** (a) Training (b) During inference, we convert accent by feeding in the target accent ID to the pronunciation encoder.

### 2.1. Pronunciation Encoder

Inspired by many early works that have used phonetic posteriors (PPG) to capture accent-dependent features, we use a pronunciation encoder, as shown in Fig. 2, to synthesize accent-dependent pronunciation sequence given text predictions and accent ID [6, 13]. For each accent ID, the pronunciation encoder learns a unique embedding, which is concatenated with every frame of character-level wav2vec2.0 prediction. The four transformer layers with 8-head attention

mechanism accounts for how context affects pronunciations. Applying a dropout of 0.3 promotes the decoder to rely more on the acoustic encoder for non-accent related voice features.



**Fig. 2.** The pronunciation encoder.

### 2.2. Acoustic Encoder

The acoustic encoder maps the mel-frequency cepstrum coefficients (MFCC) and periodicity features to a single 256-dimensional vector. The acoustic encoder consists of four convolution layers with kernel sizes (5, 3, 3, 1) and dilations (1, 2, 1, 1), a self-attention layer, and finally an average pooling layer over time to output a single vector. Instead of using an absolute positional embedding, we use the convolution layers to act as relative positional embedding for the self-attention layer as in [12].

We use adversarial training to remove accent information from the output of the acoustic encoder. All the accents are labeled as either native or foreign. The accent discriminator, consisting of two fully connected layers, learns to predict 1 if the source audio accent is native and 0 if foreign, while the acoustic encoder tries to force the accent discriminator to predict 1 all of the time. Mathematically,  $\mathcal{N}$  the set of audios with native English accent,  $\mathcal{F}$  the set with foreign accents, and  $\mathcal{X} = \mathcal{N} \cup \mathcal{F}$  the entire set of training data. Let  $\mathbf{z}$  be the output of the acoustic encoder. The accent discriminator (AD) attempts to minimize  $L_{AD}$  from (2) while the acoustic encoder tries to minimize  $L_{AD,adv}$  from (3).

$$L_{AD} = -\mathbb{E}_{\mathcal{N}}[\log(AD(\mathbf{z}))] - \mathbb{E}_{\mathcal{F}}[\log(1 - AD(\mathbf{z}))] \quad (2)$$

$$L_{AD,adv} = -\mathbb{E}_{\mathcal{F}}[\log(AD(\mathbf{z}))] \quad (3)$$

### 2.3. HiFiGAN-based Voice Decoder

Finally, we re-combine the accent-dependent pronunciation encodings, the acoustic features, and F0 sequence. We use a modified HiFiGAN to invert the processed features back to audio waveform [14]. Our modifications to the HiFiGAN architecture include adding an additional convolution layer with kernel size 11, modifying the number of input channels to HiFiGAN, and modifying the upsampling rates so that the output length matches the source audio length.

We use the same multi-scale discriminator (MSD) and multi-period discriminator (MPD) as in [14] to encourage the synthesis of natural sounding audio. Mathematically, if we

<sup>1</sup>Specifically, we downloaded the checkpoint corresponding to "Wav2Vec 2.0 Large (LV-60 + CV + SWBD + FSH)" from the link <https://github.com/facebookresearch/fairseq/tree/main/examples/wav2vec>.

consider the HiFiGAN discriminators, MSD and MPD, as one discriminator  $HD$ , the HiFiGAN discriminators try to minimize  $L_{HD}$  in (4).

$$L_{HD} = \mathbb{E}_{\mathcal{X}}[(HD(\mathbf{x}) - 1)^2 + (HD(\hat{\mathbf{x}}))^2] \quad (4)$$

As in [14], the adversarial loss  $L_{HD,adv}$ , and the feature mapping loss functions,  $L_{FM}$  applied to the rest of the model in (5) and (6).

$$L_{HD,adv} = \mathbb{E}_{\mathcal{X}}(HD(\hat{\mathbf{x}}) - 1)^2 \quad (5)$$

$$L_{FM} = \mathbb{E}_{\mathcal{X}}\left[\sum_{i=1}^T \frac{1}{N_i} \|HD^i(\mathbf{x}) - HD^i(\hat{\mathbf{x}})\|\right] \quad (6)$$

where  $HD^i$ ,  $N_i$  denote the features and the number of features in the  $i^{th}$  layer of the HiFiGAN discriminators.

### 3. EXPERIMENTS

#### 3.1. Data Sets

Our training data is summarized in Table 1. Our training data set includes 8 different accents: American (AM), Arabic (AR), British (BR), Hindi (HI), Korean (KO), Mandarin (MA), Spanish (SP), and Vietnamese (VI). We consider AM and BR accents as native and the rest foreign. Since our data set is highly unbalanced across different accents, we assign different weights to each subset. A weight of  $n$  assigned to a subset means that audio clips from that subset will appear around  $n$  times during one epoch of training. The weights can also be found in Table 1.

#### 3.2. Feature Processing

All audios are re-sampled at 16kHz and divided into 1.12s segments for training. We use the YAAPT algorithm to extract the F0 sequence at a frame shift of  $5ms$  and window length of  $20ms$  [20]. We up-sampled the outputs of the pronunciation encoder by repeating each time frame 4 times since text predictions from wav2vec2.0 are extracted at a  $20ms$  frame shift.

#### 3.3. Training Configurations

We trained our model with learning rate=0.0002 and decayed our learning rate by factor of 0.999 every 1000 iterations. We used the AdamW optimizer with  $\beta_1 = 0.8$  and  $\beta_2 = 0.99$ . For adversarial training, we optimized the accent discriminator for the first 50000 iterations before applying the adversarial loss to the acoustic encoder. The model was trained for a total of 3 million iterations using batch size 16.

#### 3.4. Evaluation Set-up

We conducted mean opinion score (MOS) tests to evaluate the performance of the proposed model on audio naturalness and speaker similarity and an XAB test to evaluate the results of accent conversion. A different set of 100 listeners participated in each listening test, and we recruited an additional 24 “experts” to evaluate accent conversion results. As in [9], all audios were shuffled before being presented to the listeners.

We trained an ablation model by disabling the accent discriminator. We used the same 20 prompt examples from the VCTK dataset used by [9] for evaluation. In our case, instead of using all examples from the Hindi accented female speaker  $p248$ , we replaced the first 10 examples with the corresponding ones from a male British accented speaker  $p226$ . We excluded audios from  $p248$  and  $p226$  from our training set so that they are unseen speakers. Using our proposed model (P) and the ablation model (AB), we converted the 20 audio clips to American (AM), Hindi (HI), and Korean (KO) accents.

### 4. RESULTS

#### 4.1. Audio Quality

The listeners were asked to rate the audio quality for the original (O) and the converted (AB, P) clips on a five-point scale (1-bad, 5-excellent). As seen in Table 2, the proposed model maintains a comparable quality to the original audio. We hypothesize it does better than the ablation model because accent invariance in the acoustic encoder improves generalization.

#### 4.2. Speaker Similarity

The raters first listen to a reference audio from each speaker with different text content. For each audio clip, the listeners were asked to rate how close did the voice sound to that of the reference audio’s on a five-point scale. The listeners were instructed to disregard the accent and recording conditions. Table 2 demonstrates that most listeners believed that the voice in the audios converted using the models sound almost as similar to the reference audio as other original audios from the same speaker.

#### 4.3. Accent Conversion

The listeners first listened to some reference audio clips with the target accent. Then, for each pair containing an original audio clip and its converted audio clip, the listeners were asked to choose the one that has a closer accent to the reference audio. As seen from the results in Fig. 3, for every target accent, the proposed model performs better than the ablation model. Significantly more listeners preferred the converted audios as sounding more American than the original audios. However, the listeners seemed unsure about whether


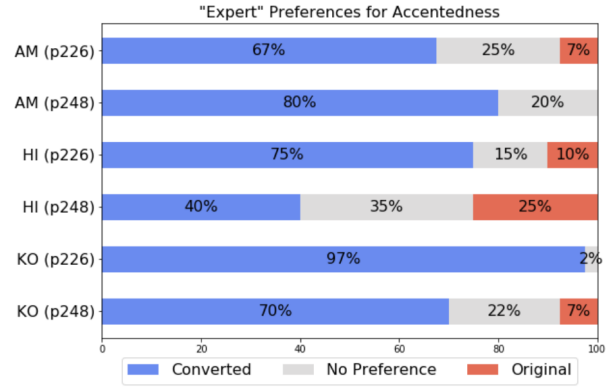
**Table 1.** Training Data Descriptions [15, 16, 17, 18, 19]

Data Set	Accents	Duration(hrs)	Speakers	Prompts	Weight
LibriTTS	AM	585	2456	LibriSpeech text	1
VCTK	BR	42	109	Newspaper clippings	6
SAA	BR	3.7	579	“Please call Stella...”	10
L2-Arctic	AR, HI, KO, MA, SP, VI	24	24 (2 male and 2 female per accent)	ARCTIC prompts	15
Indic TTS	HI	20.06	2 (1 male and 1 female)	ARCTIC prompts, Fairy tales	2

**Table 2.** MOS Study Results (with 95% confidence interval)

Samples	Audio Quality	Speaker Similarity
Original	$3.87 \pm 0.07$	$4.28 \pm 0.06$
Ablation	$3.27 \pm 0.12$	$3.89 \pm 0.10$
Proposed	$3.62 \pm 0.09$	$4.05 \pm 0.09$

the converted audios sounded more Korean or Hindi, likely because it is difficult to identify an accent after listening to only a few audio clips. The goal of conversion to other foreign accents is to help a non-native listener understand someone with a different foreign accent or to create a better entertainment value for the people knowledgeable of the desired accent. Therefore, we recruited 24 additional “expert” participants who were born and raised in USA, India, and Korea, to judge the corresponding accents. The results shown in Fig. 4 indicate that the listeners who are very familiar with the target accent were confident that the proposed model converted audios to sound like the target accent.

**Fig. 3.** Preference results from 100 random listeners. Inconclusiveness on accents other than American suggests that recognizing attributes of a specific accent may need more exposure than obtained from listening to a couple reference clips.**Fig. 4.** Preference results from raters having substantial exposure to respective accents (4 raters per accent). With the exception of HI-to-HI conversion for p248, the listeners strongly preferred the audios converted by our proposed model. Preference is stronger when converting from native to non-native accents or vice versa.

## 5. CONCLUSION

In this paper, we presented a novel accent conversion model that converts an unseen speaker’s utterance with an arbitrary accent to utterances with multiple different target accents<sup>2</sup>. The model is able to noticeably convert the accent of input audios while preserving speaker identity and audio quality. Future work may improve the perceptive accent accuracy of the converted audio by predicting pitch based on the accent.

## 6. REFERENCES

- [1] Yu Zhang, Ron J. Weiss, Heiga Zen, Yonghui Wu, Zhifeng Chen, R. J. Skerry-Ryan, Ye Jia, Andrew Rosenberg, and Bhuvana Ramabhadran, “Learning to speak fluently in a foreign language: Multilingual speech synthesis and cross-language voice cloning,” *Proceedings of the Annual Conference of the International Speech Communication Association, INTER-SPEECH*, vol. 2019-Sept, pp. 2080–2084, 2019.

<sup>2</sup>Demo samples at <https://accent-conversion.github.io>

- [2] Zhichao Wang, Wenshuo Ge, Xiong Wang, Shan Yang, Wendong Gan, Haitao Chen, Hai Li, Lei Xie, and Xiulin Li, "Accent and speaker disentanglement in many-to-many voice conversion," *2021 12th International Symposium on Chinese Spoken Language Processing, ISCSLP 2021*, pp. 2–6, 2021.
- [3] Andros Tjandra, Ruoming Pang, Yu Zhang, and Shigeki Karita, "Unsupervised learning of disentangled speech content and style representation," *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 4, pp. 3191–3195, 2021.
- [4] Adam Polyak, Yossi Adi, Jade Copet, Eugene Kharonov, Kushal Lakhotia, Wei Ning Hsu, Abdelrahman Mohamed, and Emmanuel Dupoux, "Speech resynthesis from discrete disentangled self-supervised representations," *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, vol. 5, pp. 3531–3535, 2021.
- [5] Shaojin Ding, Guanlong Zhao, and Ricardo Gutierrez-Osuna, "Accentron: Foreign accent conversion to arbitrary non-native speakers using zero-shot learning," *Computer Speech & Language*, vol. 72, pp. 101302, 2022.
- [6] Guanlong Zhao, Sinem Sonsaat, John Levis, Evgeny Chukharev-Hudilainen, and Ricardo Gutierrez-Osuna, "Accent conversion using phonetic posteriorgrams," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5314–5318.
- [7] Guanlong Zhao and Ricardo Gutierrez-Osuna, "Using phonetic posteriorgram based frame pairing for segmental accent conversion," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 10, pp. 1649–1660, 2019.
- [8] Sandesh Aryal and Ricardo Gutierrez-Osuna, "Can voice conversion be used to reduce non-native accents?," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2014, pp. 7879–7883.
- [9] Songxiang Liu, Disong Wang, Yuewen Cao, Lifa Sun, Xixin Wu, Shiyin Kang, Zhiyong Wu, Xunying Liu, Dan Su, Dong Yu, et al., "End-to-end accent conversion without using native utterances," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6289–6293.
- [10] Guanlong Zhao, Shaojin Ding, and Ricardo Gutierrez-Osuna, "Converting foreign accent speech without a reference," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2367–2381, 2021.
- [11] Waris Quamer, Anurag Das, John Levis, Evgeny Chukharev-Hudilainen, and Ricardo Gutierrez-Osuna, "Zero-shot foreign accent conversion without a native reference," *Proc. Interspeech 2022*, pp. 4920–4924, 2022.
- [12] Alexei Baevski, Henry Zhou, Abdelrahman Mohamed, and Michael Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 2020-Decem, pp. 1–12, 2020.
- [13] Guanlong Zhao, Shaojin Ding, and Ricardo Gutierrez-Osuna, "Foreign accent conversion by synthesizing speech from phonetic posteriorgrams," in *INTER-SPEECH*, 2019, pp. 2843–2847.
- [14] Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae, "Hifigan: Generative adversarial networks for efficient and high fidelity speech synthesis," *Advances in Neural Information Processing Systems*, vol. 2020-Decem, 2020.
- [15] Heiga Zen, Viet Dang, Rob Clark, Yu Zhang, Ron J. Weiss, Ye Jia, Zhifeng Chen, and Yonghui Wu, "LibriTTS: A corpus derived from librispeech for text-to-speech," *CoRR*, vol. abs/1904.02882, 2019.
- [16] Christophe Veaux, Junichi Yamagishi, Kirsten MacDonal, et al., "Cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," *University of Edinburgh. The Centre for Speech Technology Research (CSTR)*, 2017.
- [17] Nishanthi N L Arun Baby, Anju Leela Thomas and Hema A Murthy, "The speech accent archive: towards a typology of english accents," in *Community-based Building of Language Resources (International Conference on Text Speech and Dialogue)*, pp. 37–43. 2016.
- [18] Guanlong Zhao, Sinem Sonsaat, Alif Silpachai, Ivana Lucic, Evgeny Chukharev-Hudilainen, John Levis, and Ricardo Gutierrez-Osuna, "L2-arctic: A non-native english speech corpus," in *Proc. Interspeech*, 2018, p. 2783–2787.
- [19] Steven H A. Baby and Stephen A Kunath, "Resources for indian languages," in *Corpus-based studies in language use, language learning, and language documentation*, pp. 265–281. Brill, 2011.
- [20] Kavita Kasi and Stephen A Zahorian, "Yet another algorithm for pitch tracking," in *2002 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2002, vol. 1, pp. 1–361.