

Enhancing Polyglot Voices by Leveraging Cross-Lingual Fine-Tuning in Any-to-One Voice Conversion

Giuseppe Ruggiero^{1,2}, Matteo Testa², Jurgen Van de Walle², Luigi Di Caro¹

¹Università degli studi di Torino, Turin, Italy

²Cerence Inc, Turin Italy

{giuseppe.ruggiero, luigi.dicaro}@unito.it

{matteo.testa, jurgen.vandewalle}@cerence.com

Abstract

The creation of artificial polyglot voices remains a challenging task, despite considerable progress in recent years. This paper investigates self-supervised learning for voice conversion to create native-sounding polyglot voices. We introduce a novel cross-lingual any-to-one voice conversion system that is able to preserve the source accent without the need for multi-lingual data from the target speaker. In addition, we show a novel cross-lingual fine-tuning strategy that further improves the accent and reduces the training data requirements. Objective and subjective evaluations with English, Spanish, French and Mandarin Chinese confirm that our approach improves on state-of-the-art methods, enhancing the speech intelligibility and overall quality of the converted speech, especially in cross-lingual scenarios.¹

1 Introduction

In recent years, advancements in speech synthesis (Zhang et al., 2023; Qin et al., 2023; van Niekerk et al., 2021a) have significantly improved the quality of artificial polyglot voices. Several noteworthy studies in cross-lingual voice conversion (XVC) merit attention, especially some any-to-one (A2O VC)² variants (van Niekerk et al., 2021a; Huang et al., 2021). The aim of XVC is to convert source voices into target voices without changing the linguistic content, even when source and target speak different languages (Abe et al., 1990). From an information perspective, VC can be accomplished by initially disentangling the source speech into representations describing the linguistic content and speaker characteristics and then synthesizing the converted speech from the extracted

linguistic content while incorporating the identity of the target speaker (Huang et al., 2021).

To disentangle content and speaker information, an ASR model trained on a labeled dataset is frequently employed to extract the supervised spoken content representation, such as text (Huang et al., 2020) or phonetic posteriorgram (PPG) (Sun et al., 2016). However, despite its promising results in VC, this approach encounters several limitations. First, the assembly of labeled datasets can be prohibitively expensive, especially in low-resource settings like the XVC scenario (Zhao et al., 2020). Second, the system’s sensitivity to ASR errors can lead to potential mispronunciations or inaccuracies in conversion. Finally, text and PPG do not capture all linguistic nuances such as accent, expressions and style resulting in neutral-sounding synthesized speech (Hussain et al., 2023).

In several recent works (van Niekerk et al., 2021a; Huang et al., 2021; Polyak et al., 2021), self-supervised learning (SSL) is used to train models for speech synthesis. This approach is primarily driven by the following advantages offered by SSL: 1) It requires only large-scale unlabeled data to learn rich and compact speech representations; 2) Numerous large-scale pre-trained SSL models are publicly available, including Wav2vec2 (Baevski et al., 2020), HuBERT (Hsu et al., 2021a), and WavLM (Chen et al., 2021b); 3) These models have been shown to generalize well across various languages, even though some of them have been trained on English data only. Since SSL representations contain a significant amount of information, including phonetic and speaker details, many studies (Huang et al., 2021; Polyak et al., 2021) have attempted to discretize them to extract disentangled linguistic content. However, while effective at removing speaker information, discretizing SSL representations results in lossy compression, which leads to mispronunciations in the converted speech (Qian et al., 2022).

¹Audio samples are available at: <https://giuseppe-ruggiero.github.io/a2o-vc-demo/>

²The aim of A2O VC is to convert speech from any speaker, including those unseen during training, to that of a fixed target speaker.

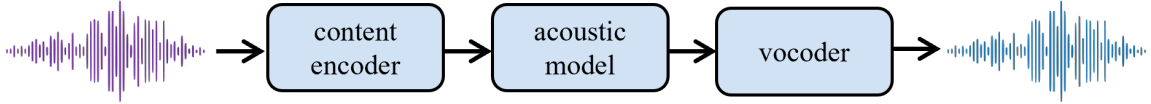


Figure 1: High-level overview of the voice conversion system.

To address this issue, alternative disentanglement strategies have been proposed. Soft-VC (van Niekerk et al., 2021a) creates soft speech units starting from discrete ones, which are obtained by applying k-means (Jin and Han, 2010) clustering over the output of a pre-trained HuBERT model. Specifically, a network is trained to predict a distribution over the discrete units to preserve more content information and thus correct mispronunciations. Neural Analysis and Synthesis (NANSY) (Choi et al., 2021) introduces information perturbation (prior to the SSL model) as a way to disentangle speaker-related information from other aspects of speech. In van Niekerk et al. (2021b), utterance-level standardization is used to remove speaker-specific information. However, while these methods are better at improving intelligibility than discretization, our findings indicate that the linguistic content is still compromised, especially in cross-lingual scenarios. Furthermore, other speaker-related features (e.g., the accent) are also distorted, as highlighted in Qian et al. (2022).

Recently, a variety of more versatile systems have emerged. One such example is kNN-VC (Baas et al., 2023), an any-to-any VC (A2A VC) approach that extracts frame-wise WavLM representations for both source and target speech. It then converts by directly replacing each frame of the source representation with its nearest neighbors in the target representation set. This strategy allows SSL features to be used without any further disentanglement step.

We propose an A2O VC system similar to that of van Niekerk et al. (2021a) except it learns to disentangle content information without the need for supplementary steps like discretization, and it can both enhance cross-lingual generation and reduce the training data requisites. Our main contributions are as follows:

- We show that a bottleneck information extractor similar to that of Li et al. (2022) is sufficient to improve intelligibility in cross-lingual scenarios without compromising on speech quality and target speaker identity.
- We propose a cross-lingual pre-training fine-

tuning strategy. This approach not only further enhances VC performance in cross-lingual scenarios but also allows low-quality data to be used for data intensive pre-training.

2 Model Architecture

The system consists of three components: a content encoder, an acoustic model, and a vocoder (Figure 1). The content encoder extracts speech representations from an audio of any speaker, the acoustic model translates these representations into a mel spectrogram of the target speaker, and the vocoder converts the resulting mel spectrogram into a time-domain waveform. For both training and inference, only audio data is required.

2.1 Problem Definition

Consider a dataset D of a target speaker containing M utterances in the time domain. Let us denote the i -th utterance of the target speaker as u_i , its features extracted by a pre-trained SSL model as $v_i \in \mathbb{R}^s$, and its mel spectrogram as x_i , where $i \in [1, M]$. The content encoder \mathcal{C} takes u_i as its input and generates v_i :

$$v_i = \mathcal{C}(u_i; w_{\mathcal{C}}) \quad (1)$$

where $w_{\mathcal{C}}$ represents the content encoder parameters. To remove speaker information in v_i while minimizing its impact on other components (e.g., linguistic content), the acoustic model \mathcal{A} uses a pre-net \mathcal{P} which projects v_i into a space \mathbb{R}^d , where $d \ll s$. This dimensionality gap imposes an information bottleneck, which allows the acoustic model to discard content-irrelevant information such as noise and speaker identity (Li et al., 2022). Thus, the acoustic model \mathcal{A} predicts \hat{x}_i :

$$\hat{x}_i = \mathcal{A}(\mathcal{P}(v_i; w_{\mathcal{P}}); w_{\mathcal{A}}) \quad (2)$$

where $w_{\mathcal{P}}$ represents the parameters of the pre-net, while $w_{\mathcal{A}}$ represents those of the acoustic model. Finally, the vocoder \mathcal{V} , given \hat{x}_i , generates \hat{u}_i :

$$\hat{u}_i = \mathcal{V}(\hat{x}_i; w_{\mathcal{V}}) \quad (3)$$

where $w_{\mathcal{V}}$ represents the vocoder parameters.

The three components of the VC system are trained independently. The content encoder \mathcal{C} is an off-the-shelf model pre-trained on a large set of multi-speaker data whereas the acoustic model \mathcal{A} and the vocoder \mathcal{V} are trained only on target speaker data. After generating pairs (v_i, x_i) for each u_i of the dataset D , the acoustic model \mathcal{A} is trained by optimizing the following objective function:

$$\min_{w_{\mathcal{P}}, w_{\mathcal{A}}} L_{\mathcal{A}}(x_i, \mathcal{A}(\mathcal{P}(v_i; w_{\mathcal{P}}); w_{\mathcal{A}})) \quad (4)$$

where $L_{\mathcal{A}}$ is a loss function in the time-frequency domain such as the L_1 distance between target and predicted mel spectrograms.

The vocoder \mathcal{V} is trained either directly on ground truth mel spectrograms or on the mel spectrograms predicted by the acoustic model:

$$\min_{w_{\mathcal{V}}} L_{\mathcal{V}}(u_i, \mathcal{V}(x_i; w_{\mathcal{V}})) \text{ or } \min_{w_{\mathcal{V}}} L_{\mathcal{V}}(u_i, \mathcal{V}(\hat{x}_i; w_{\mathcal{V}})) \quad (5)$$

where $L_{\mathcal{V}}$ is a loss function in the time domain.

By training only on target speaker data, the loss function in Eq. 4 makes the acoustic model target speaker specific: it generates spectrograms of that speaker. The primary objective is to convert the source speaker’s identity into that of the target speaker. Consequently, it is possible to perform XVC by converting a source utterance from a language different from that spoken by the target speaker, thereby enabling the target speaker to become polyglot.

2.2 Cross-Lingual Fine-Tuning

Although the system achieves intra- and cross-lingual A2O VC with good intelligibility and high quality overall, some phonemes are mispronounced or colored by the accent of the target speaker, especially in cross-lingual scenarios. To tackle this issue, we opted for a hybrid approach, leveraging the phonetic knowledge of a model trained on a specific language (e.g., French) to enhance the pronunciation and linguistic content of the target speaker (e.g., an English speaker) when it speaks that language. Let us consider a *lang1-lang2* cross-lingual scenario. Our objective is to enhance the acoustic model on *lang1* without compromising the intra-lingual and other cross-lingual tasks.

In addition to the target speaker dataset D (*lang2*), let us also consider a dataset D' consisting of N speakers, each of whom has M' utterances in *lang1*. Let us denote the i -th utterance of

the j -th speaker as u'_{ji} , its SSL vector as v'_{ji} , and its mel spectrogram as x'_{ji} , where $j \in [1, N]$ and $i \in [1, M']$.

First, we pre-train the acoustic model \mathcal{A} using the loss in Eq. 4 but employing pairs (v'_{ji}, x'_{ji}) for each u'_{ji} from the dataset D' . Eq. 4 becomes:

$$\min_{w'_{\mathcal{P}}, w'_{\mathcal{A}}} L_{\mathcal{A}}(x'_{ji}, \mathcal{A}(\mathcal{P}(v'_{ji}; w'_{\mathcal{P}}); w'_{\mathcal{A}})) \quad (6)$$

Second, we fine-tune \mathcal{A} using the dataset D (or a subset thereof) from the target speaker, as explained in Section 2.1. This allows the acoustic model to acquire phonetic coverage of *lang1* during the pre-training, which can be leveraged during the fine-tuning on *lang2*. Consequently, the *lang2* target speaker will exhibit improved *lang1* linguistic characteristics after the conversion.

The cross-lingual fine-tuning technique offers several advantages. It enhances the linguistic content and pronunciation in XVC and reduces the amount of training data needed from the target speaker. While training the acoustic model without fine-tuning typically requires at least 10-12h of high-quality audio from the target speaker, the fine-tuning only requires 2h.

2.3 Content encoder

The content encoder transforms an input waveform into a more concise and information-dense representation known as SSL vector. We chose WavLM as pretrained SSL model. WavLM was pre-trained by Microsoft using a self-supervised strategy similar to that of Hsu et al. (2021a) on 60,000 hours of Libri-Light (Kahn et al., 2019), 10,000 hours of GigaSpeech (Chen et al., 2021a) and 24,000 hours of VoxPopuli (Wang et al., 2021), for a total of 94,000 hours of audio data. We used a publicly available WavLM-Large³ model employing the output of its 15th transformer layer as the representation for speech audio v_i in all of our experiments. This decision was motivated by the findings of Hsu et al. (2021b); Baevski et al. (2021), which demonstrated the efficacy of this layer in phone discrimination tests. Accordingly, s in Subsection 2.1 is configured to 1024, that is the dimension of the output vectors of WavLM-Large.

2.4 Acoustic Model and Vocoder

The acoustic model and vocoder are typical components in a text-to-speech (TTS) system.

³<https://huggingface.co/microsoft/wavlm-large>

The proposed acoustic model takes as input SSL vectors rather than graphemes or phonemes, and outputs target speaker mel spectrograms. Following van Niekerk et al. (2021a), the model’s architecture resembles that of Tacotron 2 (Shen et al., 2017) as it features an encoder and an autoregressive decoder. Both the encoder and decoder are preceded by a feed-forward pre-net, while a final linear layer with n -MELs units follows the decoder. No attention module is used (Huang et al., 2021). The encoder pre-net is a feed-forward neural net composed of a stack of two linear layers with 256 units, ReLU activations, and dropout (Srivastava et al., 2014). The encoder consists of a stack of three 1D-convolutional layers with 512 units, kernel size 5, stride 1, padding 2, and ReLU activations. Following Chou and Lee (2019), we added an Instance Normalization (IN) layer after each Conv1D layer, as this helps to further reduce speaker identity while preserving content information. The decoder predicts each spectrogram frame from the output of the encoder and the previously generated frames. First, a decoder pre-net, similar to the encoder one, is applied and then three LSTMs with 768 units each followed by a final linear layer with n -MELs units are used.

The vocoder transforms the predicted mel spectrograms into audio waveforms. We used a HiFi-GAN (Kong et al., 2020) model, a state-of-the-art real-time neural vocoder.

2.5 Length Regulator

As there is no attention mechanism, we use an optional duration adjustment strategy to mitigate potential length mismatches between the SSL input features and the target spectrogram sequence. Indeed, WavLM generates features at a 20 ms frame rate for audio sampled at 16 kHz (resulting in a hop length of 320), which might lead to a length mismatch when generating spectrograms with different time resolutions (e.g., different hop lengths and/or higher sample rates). This strategy involves integrating a length regulator module between the encoder and the decoder of the acoustic model which interpolates the SSL features extracted from the source speech. To do this, we leverage PyTorch’s `interpolate`⁴ function with its default settings.

⁴<https://pytorch.org/docs/stable/generated/torch.nn.functional.interpolate.html>

3 Experiments

To evaluate the effectiveness of our VC approach, experiments across a range of different tasks and datasets were conducted. All the experiments were run on a Linux machine with a single NVIDIA Titan RTX GPU with 24 GB of RAM.

The proposed system was tested as two variants: *Proposed*, our VC method using the conventional training strategy, and *Proposed-F*, the alternative system employing the cross-lingual fine-tuning strategy. We selected two baseline models for comparison: Soft-VC⁵ (van Niekerk et al., 2021a) and kNN-VC⁶ (Baas et al., 2023). Both excel in cross-lingual voice conversion tasks, demonstrating strong content intelligibility, accent preservation, and speech quality.

3.1 Experimental Setup

We conducted experiments for both intra- and cross-lingual settings using LJSpeech (Ito and Johnson, 2017) as the target speaker dataset. We selected four languages for our evaluation: American English, European French, European Spanish, and Mandarin Chinese. For English, we used LibriSpeech (Panayotov et al., 2015); for French and Spanish, we used Multilingual LibriSpeech (MLS) (Pratap et al., 2020); for Mandarin Chinese, we used Aishell (Bu et al., 2017).

To ensure a fair comparison, when training the LJSpeech model we used the same configuration as van Niekerk et al. (2021a), including their HiFi-GAN vocoder pre-trained on LJSpeech, unless otherwise stated. More specifically, we used a 16 kHz sampling rate and we set n -MELs to 128, s to 1024 and d to 256.

For cross-lingual fine-tuning, we initially pre-trained three distinct acoustic models for each language using randomly selected 150h of audio data from the training sets of LibriSpeech, French MLS, Spanish MLS, and Aishell, respectively. The choice of 150h was based on Aishell’s training set size in order to ensure that each language gets the same amount of training hours. Each pre-trained acoustic model was fine-tuned with only 2h of LJSpeech data for 15k steps using a batch size of 8. We found this number of hours to be optimal for the fine-tuning step through several experiments conducted using a validation set composed of 100 utterances from LJSpeech. We also observed that

⁵<https://github.com/bshall/soft-vc>

⁶<https://github.com/bshall/knn-vc>

using less than 2h of data was insufficient to obtain a good overall quality.

3.2 Evaluation Metrics

We assessed intelligibility, speaker similarity, accent preservation and overall speech quality in subjective and objective evaluations.

We created a test set for each language to be used as the source speech for conversion into LJSpeech. Each set consists of 60 utterances obtained by randomly selecting 3 utterances from 20 speakers (10 male and 10 female) extracted from: the test-clean set of LibriSpeech for English; the test set of Spanish MLS for Spanish; a set composed of the dev and test sets of French MLS combined for French; and a set composed of the dev and test sets of Aishell combined for Chinese.

3.2.1 Objective Evaluations

To evaluate the *intelligibility* of the converted speech, we measured word error rate (WER) and phoneme error rate (PER) between the source and converted speech using the Whisper Medium ASR model⁷ (Radford et al., 2023) for English, French and Spanish and the FunASR model⁸ (Gao et al., 2023) for Chinese. Lower error rates indicate more intelligible speech, as the original words/phonemes are still recognizable after conversion. We measured *speaker similarity* (SSIM) by using a trained speaker verification model⁹. Specifically, we computed the cosine similarity between the d-vectors (Heigold et al., 2015) of each converted sample and that of the target speaker. The target speaker embedding is generated by averaging the d-vectors of 50 utterances from LJSpeech.

3.2.2 Subjective Evaluations

We conducted subjective evaluations to assess both the *overall quality* and the *native accent level* of the converted audio.

For each language, we selected all 60 utterances from the corresponding test set and converted them into LJSpeech with our two proposed methods and the two baselines, resulting in a total of 300 utterances per language (60 ground truth + 240 generated). Then, for both tests, 20 native-language participants were asked to listen to the samples, randomly mixed, and rate them. To evaluate the overall quality, we gathered a *mean opinion score*

(MOS) based on a 5-point scale, where 1 stands for “very poor” and 5 for “excellent”. The question asked was: “How do you rate the overall quality of this sample?”. For accent assessment, we conducted an *accent comparative mean opinion score* (ACMOS) by asking all the participants to rate the samples based on how convincing the accent was perceived to be on a 5-point scale, where 1 stood for “strong foreign accent” and 5 for “native”. The question asked was: “Does the voice sound like that of a native speaker?”. For this test, Proposed-F was set as the baseline. A negative Δ -ACMOS indicates that a system performs worse than the baseline, zero either refers to the baseline itself or indicates the same performance, and a positive value indicates an improvement over the baseline. The intra-lingual English-English scenario was not considered as the accent is consistent and correct after conversion.

3.3 Results

In this section, we demonstrate the effectiveness of our proposed approaches evaluating both intra- and cross-lingual scenarios. Additionally, we show that our systems are also capable of preserving the emotions of the source speakers.

3.3.1 Speaker and Content Disentanglement

First, we show that any disentanglement step applied to the SSL representations, while effective in intra-lingual scenario, results in a significant performance loss in cross-lingual tasks. We compared the results obtained by directly using the output of the content encoder against some disentanglement strategies applied on top of the output of the 15th layer of WavLM-Large: discretization, perturbation, and per-utterance standardization. We omitted the soft strategy as it is included in Soft-VC.

For discretization, we followed the procedure outlined in van Niekerk et al. (2021a); for perturbation, we followed the procedure presented in Choi et al. (2021); and for per-utterance standardization, we followed the procedure described in van Niekerk et al. (2021b). Table 1 (middle) reports intelligibility in terms of WER and PER, and speaker similarity (SSIM) of each model. Compared to the other methods, the proposed systems significantly enhance the intelligibility across all languages in cross-lingual scenarios while also maintaining very good performance for English. Particularly notable is the significant improvement for Mandarin Chinese, with the WER decreasing from around

⁷<https://huggingface.co/openai/whisper-medium>

⁸<https://github.com/modelscope/FunASR>

⁹<https://github.com/resemble-ai/Resemblyzer>

	English					French				
	WER	PER	SSIM	MOS	Δ -ACMOS	WER	PER	SSIM	MOS	Δ -ACMOS
Ground truth	1.25	0.22	-	3.75 \pm 0.04	-	8.74	2.70	-	3.77 \pm 0.05	-
Discretization	18.81	10.28	90.86	-	-	70.00	60.54	90.54	-	-
Perturbation	2.82	0.88	92.42	-	-	27.17	12.08	90.84	-	-
Utterance std	4.76	2.56	88.62	-	-	15.75	5.58	84.03	-	-
Soft-VC	2.63	0.73	91.12	3.88 \pm 0.04	-	19.13	6.91	89.16	3.73 \pm 0.05	-0.13
kNN-VC	2.75	0.99	84.70	3.94 \pm 0.04	-	16.55	6.24	81.26	3.82 \pm 0.05	-0.10
Proposed	3.01	1.09	90.22	3.94 \pm 0.05	-	14.16	5.00	88.67	3.85 \pm 0.05	-0.09
Proposed-F	2.85	0.88	89.70	3.96 \pm 0.04	-	12.43	4.02	88.25	3.83 \pm 0.04	0.0

	Spanish					Mandarin Chinese				
	WER	PER	SSIM	MOS	Δ -ACMOS	WER	PER	SSIM	MOS	Δ -ACMOS
Ground truth	6.31	1.52	-	3.95 \pm 0.04	-	7.72	0.43	-	4.07 \pm 0.08	-
Discretization	75.04	66.74	92.33	-	-	90.01	77.35	87.26	-	-
Perturbation	17.95	6.52	89.64	-	-	47.73	20.03	87.14	-	-
Utterance std	12.51	3.66	85.43	-	-	28.34	5.47	82.27	-	-
Soft-VC	14.78	5.02	89.17	3.80 \pm 0.04	-0.08	51.82	21.04	86.15	3.01 \pm 0.08	-1.25
kNN-VC	11.91	3.47	80.31	3.83 \pm 0.04	-0.06	35.24	9.78	77.50	3.25 \pm 0.08	-0.47
Proposed	9.78	2.56	88.13	3.87 \pm 0.03	-0.05	19.95	3.64	84.78	3.67 \pm 0.08	-0.03
Proposed-F	8.61	2.21	86.53	3.86 \pm 0.04	0.0	18.39	3.33	83.02	3.73 \pm 0.08	0.0

Table 1: Objective and subjective evaluation. (middle) Results (%) in terms of intelligibility (W/PER \downarrow) and speaker similarity (SSMI \uparrow) considering three speaker and content disentanglement strategies: discretization, perturbation and per-utterance standardization. (bottom) W/PER, SSIM, overall quality (MOS \uparrow) with 95% confidence intervals, and the delta of the Accent Comparative MOS (Δ -ACMOS) for the proposed approaches and the baseline systems.

30% to 80%, and the PER decreasing from approximately 34% to 95%. Unlike perturbation, which fails to maintain the same performance observed for English, the overall effectiveness of the per-utterance standardization is significant. However, it does not adequately capture the target speaker’s identity like the alternative approaches do. On the other hand, the speaker identity of our systems is comparable to that of other techniques, indicating that in an A2O VC setup, the disentanglement steps are not strictly required.

3.3.2 Objective and Subjective Evaluation

After the initial assessment, we evaluated the quality of the two proposed methods by comparing our outcomes with those of the baselines. Table 1 (bottom) reports WER/PER and SSIM, along with MOS and ACMOS for each model. As observed previously, intelligibility remains clearly superior for the proposed systems for all the cross-lingual tasks and is comparable to that of the baselines for English. In addition, the MOS scores show that both proposed approaches outperform the baselines in terms of overall quality. While kNN-VC shows comparable MOS values, it fails to maintain good speaker similarity. Interestingly, the MOS values

of the ground truth are lower than those of all models for English and French. We hypothesize that VC systems can act as denoisers over noisy inputs and that most likely influenced listeners’ assignment of MOS scores. As for the accent assessment, Soft-VC and kNN-VC consistently exhibit negative delta values for the ACMOS, meaning that the generated samples are perceived as “less native-sounding” across all languages. In conclusion, while both proposed approaches outperform the baselines, Proposed-F leads to better intelligibility, higher MOS, and is also data efficient.

3.3.3 Ablation: Emotions Maintenance

In this section we show our system’s ability to preserve source emotions during VC. To achieve this, we extended the LJSpeech dataset by incorporating the four emotions in the Emotional Speech Dataset (ESD) (Zhou et al., 2022): angry, surprise, sad, and happy. As we exclusively focus on English for this experiment, we converted all emotional utterances from the 10 English speakers of ESD into LJSpeech using our highest quality setup. The resulting corpus, referred to as Emo-LJSpeech, consists of roughly 35h of data, 1h of which are created through VC.

	WER	PER	SSIM	MOS	Δ -ECMOS Angry	Δ -ECMOS Happy	Δ -ECMOS Sad	Δ -ECMOS Surprise
Ground truth	10.08	4.89	-	3.95 ± 0.04	-	-	-	-
Proposed	11.01	5.53	89.54	3.90 ± 0.05	+0.39	+0.63	+0.41	+0.92
Proposed-F	11.55	5.80	88.37	3.80 ± 0.04	+0.86	+0.74	+0.24	+0.80

Table 2: Results (%) measuring intelligibility (W/PER \downarrow), speaker similarity (SSIM \uparrow), overall quality (MOS \uparrow) with 95% confidence intervals, and the delta of the Emotion Comparative Mean Opinion Score (Δ -ECMOS) vs neutral samples for the emotional experiment.

To ensure compatibility with the original sample rate of LJSpeech (22050 Hz), we implemented the approach outlined in Section 2.5 to address length disparities between source and target features. We set n -MELs to 80, n -fft to 1024, window length to 1024, and hop length to 256. The acoustic model was trained on LJSpeech with a batch size of 32 using AdamW (Loshchilov and Hutter, 2019) with an initial learning rate of 0.0001 and a weight decay of 0.01 for 200k steps. Throughout the training, the learning rate was linearly decreased after a warm-up period in which it was linearly increased for the first 4k training steps. The above hyper-parameters are the results of an extensive hyper-parameters search phase conducted using a validation set of 100 utterances from LJSpeech. For cross-lingual fine-tuning, we pre-trained the acoustic model on 13h of data from the English subset of ESD (including neutral data). We then fine-tuned it with 2h of LJSpeech for 15k steps using a batch size of 8. For the vocoder, we used an internal implementation of HiFi-GAN working at 22050 Hz trained on a proprietary multi-speaker dataset.

To assess the quality of the newly generated emotional corpus, we created a test set by randomly selecting, for each of the 4 emotions, 3 utterances from 5 speakers extracted from the ESD English subset, resulting in a total of 60 emotional utterances. We then conducted objective and subjective evaluations by comparing the ground truth recordings with the 60 utterances converted into LJSpeech using the *Proposed* and *Proposed-F* models, trained as explained earlier. In addition to assessing WER/PER, SSIM, and MOS, we also analysed the proposed systems’ capacity to preserve source emotions through an *emotion comparative mean opinion score* (ECMOS). The test was conducted over each of the four emotions (referred to as e). 20 native participants rated the samples on a 5-point scale, indicating the extent to which the corresponding emotion (e.g., happy) was perceived, with 1 denoting “totally not e ” and 5 representing

“very e ”. The question was: “Does the voice sound e ?”. For this test, alongside the 60 emotional utterances, 60 corresponding neutral-style utterances were also generated using both models. These neutral samples served as reference data. We report the Δ -ECMOS which reflects the extent to which a certain emotion was perceived after VC compared to its neutral counterpart. The results are shown in Table 2. The intelligibility and overall quality are very similar to those of the ground truth recordings. Unlike previous experiments, Proposed outperformed Proposed-F. This could be attributed to the limited amount of data used for the pre-training (13h compared to the 150h used in previous experiments). We defer more detailed analysis of this to future work. However, both systems effectively preserve source emotions, as indicated by consistently positive Δ -ECMOS scores.

4 Conclusion

With this work we have developed an A2O VC system capable of generating human-level native-sounding polyglot voices without relying on multilingual data from the target speaker. We showed how several content and speaker disentanglement strategies applied to SSL features do not perform well in cross-lingual scenarios. To overcome this, we have demonstrated that by directly leveraging SSL features and performing cross-lingual fine-tuning, we are able to enhance the intelligibility and accent of the converted voices. Furthermore, we have illustrated that our models can preserve source emotions throughout the conversion process. Our experiments showcased improvements over strong baseline methods across all the examined metrics. Moreover, MOS tests conducted with multiple native participants revealed that our outcomes closely approximate human recordings in terms of overall quality. For future research, our aim is to explore the creation of high-quality polyglot voices in an A2A setup to further reduce data requirements.

Limitations

To achieve high-quality polyglot voices, we focused on A2O and thus target-specific acoustic model. This approach requires a substantial amount of training data. While fine-tuning has helped reduce the required data, 2h can still be considered a significant amount. We leave to future research the development of approaches aimed at further reducing the amount of data needed to train models for specific target speakers.

We conducted our experiments on a single female speaker dataset (LJSpeech). Having limited testing resources, we decided to prioritize the number of experiments on the polyglot aspect over having multiple speakers. This allowed us to show that our model can work correctly with significantly different languages.

Based on the experiments conducted, the SSIM of the proposed methods is slightly lower, though comparable, to that of other approaches. This discrepancy likely stems from our direct use of SSL features without additional refinements. However, sacrificing a small degree of speaker similarity enhances other aspects such as intelligibility and accent, particularly in cross-lingual tasks. We believe that implementing a more efficient disentanglement scheme will enable us to achieve optimal quality across all components. This is a direction we plan to pursue in future research.

Acknowledgments

(Portions of) the research in this paper used the ESD Database made available by the HLT lab, National University of Singapore, Singapore.

References

- Masanobu Abe, Kiyohiro Shikano, and Hisao Kuwabara. 1990. [Cross-language voice conversion](#). *International Conference on Acoustics, Speech, and Signal Processing*, pages 345–348 vol.1.
- Matthew Baas, Benjamin van Niekerk, and Herman Kamper. 2023. [Voice Conversion With Just Nearest Neighbors](#). In *Proc. INTERSPEECH 2023*, pages 2053–2057.
- Alexei Baevski, Wei-Ning Hsu, Alexis Conneau, and Michael Auli. 2021. Unsupervised speech recognition. *Advances in Neural Information Processing Systems*, 34:27826–27839.
- Alexei Baevski, Henry Zhou, Abdel rahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A framework](#)

[for self-supervised learning of speech representations](#). *ArXiv*, abs/2006.11477.

- Hui Bu, Jiayu Du, Xingyu Na, Bengu Wu, and Hao Zheng. 2017. [Aishell-1: An open-source mandarin speech corpus and a speech recognition baseline](#). *2017 20th Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)*, pages 1–5.
- Guoguo Chen, Shuzhou Chai, Guan-Bo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, Mingjie Jin, Sanjeev Khudanpur, Shinji Watanabe, Shuaijiang Zhao, Wei Zou, Xiangang Li, Xuchen Yao, Yongqing Wang, Zhao You, and Zhiyong Yan. 2021a. [Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio](#). In *Proc. Interspeech 2021*.
- Sanyuan Chen, Chengyi Wang, Zhengyang Chen, Yu Wu, Shujie Liu, Zhuo Chen, Jinyu Li, Naoyuki Kanda, Takuya Yoshioka, Xiong Xiao, Jian Wu, Long Zhou, Shuo Ren, Yanmin Qian, Yao Qian, Micheal Zeng, and Furu Wei. 2021b. [Wavlm: Large-scale self-supervised pre-training for full stack speech processing](#). *IEEE Journal of Selected Topics in Signal Processing*, 16:1505–1518.
- Hyeong-Seok Choi, Juheon Lee, Wan Soo Kim, Jie Hwan Lee, Hoon Heo, and Kyogu Lee. 2021. [Neural analysis and synthesis: Reconstructing speech from self-supervised representations](#). *ArXiv*, abs/2110.14513.
- Ju-chieh Chou and Hung-yi Lee. 2019. [One-shot voice conversion by separating speaker and content representations with instance normalization](#). In *INTER-SPEECH*, pages 664–668.
- Zhifu Gao, Zerui Li, Jiaming Wang, Haoneng Luo, Xian Shi, Mengzhe Chen, Yabin Li, Lingyun Zuo, Zhihao Du, Zhangyu Xiao, and Shiliang Zhang. 2023. [Funasr: A fundamental end-to-end speech recognition toolkit](#). In *INTERSPEECH*.
- Georg Heigold, Ignacio Moreno, Samy Bengio, and Noam M. Shazeer. 2015. [End-to-end text-dependent speaker verification](#). *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5115–5119.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdel rahman Mohamed. 2021a. [Hubert: Self-supervised speech representation learning by masked prediction of hidden units](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Wei-Ning Hsu, Yao-Hung Hubert Tsai, Benjamin Bolte, Ruslan Salakhutdinov, and Abdel rahman Mohamed. 2021b. [Hubert: How much can a bad teacher benefit asr pre-training?](#) *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6533–6537.

- Wen-Chin Huang, Tomoki Hayashi, Shinji Watanabe, and Tomoki Toda. 2020. [The sequence-to-sequence baseline for the voice conversion challenge 2020: Cascading asr and tts](#). *ArXiv*, abs/2010.02434.
- Wen-Chin Huang, Shu wen Yang, Tomoki Hayashi, Hung yi Lee, Shinji Watanabe, and Tomoki Toda. 2021. [S3prl-vc: Open-source voice conversion framework with self-supervised speech representations](#). *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6552–6556.
- Shehzeen Samarah Hussain, Paarth Neekhara, Jocelyn Huang, Jason Li, and Boris Ginsburg. 2023. [Acvc: Adaptive and controllable voice conversion using explicitly disentangled self-supervised speech representations](#). *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Keith Ito and Linda Johnson. 2017. The lj speech dataset. <https://keithito.com/LJ-Speech-Dataset/>.
- Xin Jin and Jiawei Han. 2010. *K-Means Clustering*, pages 563–564. Springer US, Boston, MA.
- Jacob Kahn, Morgane Rivière, Weiyi Zheng, Evgeny Kharitonov, Qiantong Xu, Pierre-Emmanuel Mazar’è, Julien Karadayi, Vitaliy Liptchinsky, Ronan Collobert, Christian Fuegen, Tatiana Likhomanenko, Gabriel Synnaeve, Armand Joulin, Abdel rahman Mohamed, and Emmanuel Dupoux. 2019. [Libri-light: A benchmark for asr with limited or no supervision](#). *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7669–7673.
- Jungil Kong, Jaehyeon Kim, and Jaekyoung Bae. 2020. [Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 17022–17033. Curran Associates, Inc.
- Jingyi Li, Weiping Tu, and Li Xiao. 2022. [Freevc: Towards high-quality text-free one-shot voice conversion](#). *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *International Conference on Learning Representations*.
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. [Librispeech: An asr corpus based on public domain audio books](#). In *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210.
- Adam Polyak, Yossi Adi, Jade Copet, Eugene Kharitonov, Kushal Lakhota, Wei-Ning Hsu, Abdel rahman Mohamed, and Emmanuel Dupoux. 2021. [Speech resynthesis from discrete disentangled self-supervised representations](#). In *Interspeech*.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. [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*, pages 2757–2761. ISCA.
- Kaizhi Qian, Yang Zhang, Heting Gao, Junrui Ni, Cheng-I Lai, David Cox, Mark A. Hasegawa-Johnson, and Shiyu Chang. 2022. [Contentvec: An improved self-supervised speech representation by disentangling speakers](#). In *International Conference on Machine Learning*.
- Zengyi Qin, Wenliang Zhao, Xumin Yu, and Xin Sun. 2023. [Openvoice: Versatile instant voice cloning](#). *ArXiv*, abs/2312.01479.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. [Robust speech recognition via large-scale weak supervision](#). In *International Conference on Machine Learning*, pages 28492–28518. PMLR.
- Jonathan Shen, Ruoming Pang, Ron J. Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Z. Chen, Yu Zhang, Yuxuan Wang, R. J. Skerry-Ryan, Rif A. Saurous, Yannis Agiomyrgiannakis, and Yonghui Wu. 2017. [Natural tts synthesis by conditioning wavenet on mel spectrogram predictions](#). *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4779–4783.
- Nitish Srivastava, Geoffrey E. Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. [Dropout: a simple way to prevent neural networks from overfitting](#). *J. Mach. Learn. Res.*, 15:1929–1958.
- Lifa Sun, Kun Li, Hao Wang, Shiyin Kang, and Helen M. Meng. 2016. [Phonetic posteriorgrams for many-to-one voice conversion without parallel data training](#). *2016 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6.
- Benjamin van Niekerk, Marc-André Carbonneau, Julian Zaïdi, Matthew Baas, Hugo Seuté, and Herman Kamper. 2021a. [A comparison of discrete and soft speech units for improved voice conversion](#). *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6562–6566.
- Benjamin van Niekerk, Leanne Nortje, Matthias Baas, and Herman Kamper. 2021b. [Analyzing speaker information in self-supervised models to improve zero-resource speech processing](#). *ArXiv*, abs/2108.00917.
- Chaghan Wang, Morgane Riviere, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux. 2021. [VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation](#). In *Proceedings of the 59th Annual*

Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 993–1003, Online. Association for Computational Linguistics.

Zi-Hua Zhang, Long Zhou, Chengyi Wang, Sanyuan Chen, Yu Wu, Shujie Liu, Zhuo Chen, Yanqing Liu, Huaming Wang, Jinyu Li, Lei He, Sheng Zhao, and Furu Wei. 2023. [Speak foreign languages with your own voice: Cross-lingual neural codec language modeling](#). *ArXiv*, abs/2303.03926.

Yi Zhao, Wen-Chin Huang, Xiaohai Tian, Junichi Yamagishi, Rohan Kumar Das, Tomi H. Kinnunen, Zhenhua Ling, and Tomoki Toda. 2020. [Voice conversion challenge 2020: Intra-lingual semi-parallel and cross-lingual voice conversion](#). *ArXiv*, abs/2008.12527.

Kun Zhou, Berrak Sisman, Rui Liu, and Haizhou Li. 2022. Emotional voice conversion: Theory, databases and esd. *Speech Communication*, 137:1–18.