



Investigation into Target Speaking Rate Adaptation for Voice Conversion

Michael Kuhlmann¹, Fritz Seebauer², Janek Ebberts¹, Petra Wagner², Reinhold Haeb-Umbach¹

¹Paderborn University, Germany ²Bielefeld University, Germany

{kuhlmann,ebbers,haeb}@nt.uni-paderborn.de,
{fritz.seebauer,petra.wagner}@uni-bielefeld.de

Abstract

Disentangling speaker and content attributes of a speech signal into separate latent representations followed by decoding the content with an exchanged speaker representation is a popular approach for voice conversion, which can be trained with non-parallel and unlabeled speech data. However, previous approaches perform disentanglement only implicitly via some sort of information bottleneck or normalization, where it is usually hard to find a good trade-off between voice conversion and content reconstruction. Further, previous works usually do not consider an adaptation of the speaking rate to the target speaker or they put some major restrictions to the data or use case. Therefore, the contribution of this work is two-fold. First, we employ an explicit and fully unsupervised disentanglement approach, which has previously only been used for representation learning, and show that it allows to obtain both superior voice conversion and content reconstruction. Second, we investigate simple and generic approaches to linearly adapt the length of a speech signal, and hence the speaking rate, to a target speaker and show that the proposed adaptation allows to increase the speaking rate similarity with respect to the target speaker.

Index Terms: voice conversion, any-to-any, speaking rate adaptation

1. Introduction

Voice Conversion (VC) deals with the transfer of speaker characteristics from one utterance to another while keeping the content information untouched. While some works [1] learn the mapping from one speaker to another using parallel data, i.e., data which contains same sentences uttered by different speakers, another course of research focuses on learning VC from non-parallel data, i.e., where each sentence is only uttered by a single speaker. Different approaches exist for non-parallel VC learning, e.g., based on Generative Adversarial Networks (GANs) [2] or based on information perturbation [3]. Another particularly popular approach to non-parallel VC is to perform disentanglement via autoencoding. The idea is to encode content and speaker information into two disjoint representations by, e.g., using a tight bottleneck [4, 5], normalization layers [5, 6, 7] or an adversarial approach [8]. Then, at test-time, VC can be performed by decoding the content representation of an utterance together with an exchanged speaker representation. So-called *zero-* [4] or *one-shot* [5, 6, 7] approaches can perform any-to-any voice conversion, where neither source nor target speakers need to be in the training data.

Most of above works mainly tackle timbre conversion, where the length of the source utterance, and hence the speaking rate, is kept unchanged. However, speaking rate is also an important characteristic of individuals [9] or situation specific speaking styles [10]. There are some works which jointly learn to also convert the speaking rate using sequence-to-sequence [11, 12, 13] or disentanglement [14, 15] approaches.

These works, however, require parallel data during training or known speaker identities during test time. In [16] it was shown that speech recognition for dysarthric speech can be enhanced by a simple linear interpolation of an utterance's spectrogram and outperforms cycle-based VC approaches on parallel data.

The contribution of this paper is two-fold. First, we propose to perform VC with our fully unsupervised Factorized Variational Autoencoder (FVAE) model [17]. The FVAE model, which achieves disentanglement by an adversarial Contrastive Predictive Coding (CPC) loss, has, so far, only been considered for learning disentangled features. We here investigate its suitability for VC when decoding the disentangled content representation with an exchanged speaker representation. Second, we investigate ways to estimate an interpolation factor for linearly adapting the speaking rate to the target speaker. We consider two variants. The first is based on an estimation of the phoneme rate of an utterance from transcribed training data, which then runs fully unsupervised at test-time. The second variant is based on acoustic unit rates from an unsupervised segmentation model and neither needs text transcriptions at train nor at test time. With our proposed fully unsupervised VC model combined with the fully unsupervised speaking rate adaptation, our model can be trained on and applied to fully unsupervised raw speech signals. We compare our approach against AutoPST [14], a state-of-the-art model for voice and prosody conversion without text transcriptions or parallel data. We show that a simple post-processing for speaking rate adaptation can achieve the desired effect of increased target speaker similarity and outperforms AutoPST in voice and speaking rate similarity.

2. Target speaker conversion

Given an utterance from an arbitrary speaker we want to match the voice and speaking rate to a target speaker while keeping the content unchanged. We want to impose as few restrictions as possible on the system. Specifically, the conversion process should work for arbitrary target speakers, needs only one utterance from the target speaker and does not rely on any text transcriptions. For this, we consider two successive steps:

1. Any-to-any VC: We use a neural-based disentanglement system that separates content from speaker style and swap in the target speaker style to perform a voice conversion. The output is a Mel spectrogram and we use a pretrained neural vocoder to obtain the converted audio signal.
2. Speech rate conversion: We use WSOLA [18] to globally adapt the tempo of the converted signal to match the speaking rate of the target utterance. The interpolation factor is the ratio of target phoneme rate to source phoneme rate. To obtain the phoneme rates from the utterances, we investigate two approaches.

With this two-step approach, we can decouple the speaking rate conversion from the VC model and the vocoder to separately evaluate the influence of adding the speaking rate conversion to the voice conversions. It is also possible to integrate the speaking rate conversion into the vocoder, e.g., non-trainable analysis-synthesis vocoders like STRAIGHT and WORLD or into non-autoregressive VC models [19]. We leave an analysis of the benefits and drawbacks of each of these approaches to future work. Figure 1 summarizes the proposed approach.

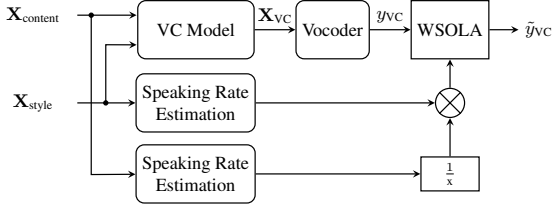


Figure 1: Overview of the proposed approach

2.1. Voice conversion model

To perform the voice conversion, we use an autoencoder to disentangle speaker information from the content and pair the content embedding with a neural speaker embedding from the target speaker. We use the Factorized Variational Autoencoder (FVAE) with an adversarial contrastive loss on the content embedding which achieves a well-balanced trade-off between informative linguistic content and speaker invariance [17].

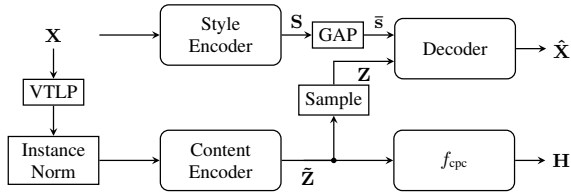


Figure 2: The architecture of the FVAE [17]

The FVAE consists of two fully-convolutional encoders and a fully-convolutional decoder (see Figure 2). The style encoder extracts a sequence of style embeddings \mathbf{S} from the input feature map \mathbf{X} . Static information is aggregated into a single embedding $\bar{\mathbf{s}} \in \mathbb{R}^{D_s}$ using global average pooling (GAP) over time on \mathbf{S} . Simultaneously, the content encoder extracts a content representation \mathbf{Z} from the input. To drive the decoder to infer the timbre from $\bar{\mathbf{s}}$ rather than from \mathbf{Z} during reconstruction, \mathbf{X} is first distorted by Vocal Tract Length Perturbation (VTLP) [20] and then transformed by Instance Normalization (IN) [21] before being passed to the content encoder. For each time frame, the content encoder outputs a tuple $\tilde{\mathbf{z}}_n = (\boldsymbol{\mu}_{z_n}, \boldsymbol{\xi}_{z_n})$, which are mean and log-variance vectors, respectively, and \mathbf{Z} is obtained through the reparameterization trick [22] during training. Then, $\bar{\mathbf{s}}$ is frame-wise concatenated with \mathbf{Z} and passed to the decoder, yielding the reconstructed speech feature $\hat{\mathbf{X}}$. Both encoders and the decoder are jointly trained to minimize the frame-wise Mean-Squared Error (MSE) between \mathbf{X} and $\hat{\mathbf{X}}$.

To foster disentanglement, a contrastive loss [23] is applied on the content embedding. An adversarial encoder f_{cpc} extracts an embedding $\mathbf{h}_n = f_{cpc}(\tilde{\mathbf{z}}_n)$ from the content encoder output for each time frame n . A prediction head $g_k(\cdot)$ is used to yield a prediction $\hat{\mathbf{h}}_n = g_k(\mathbf{h}_{n-k})$ for a future time step in the embedding dimension, where k is the lookahead shift. The contrastive

objective then tries to find the correct prediction out of a number of candidates. To remove static information from $\tilde{\mathbf{Z}}$, a large lookahead shift k is used (corresponding to 1s) and the content encoder is trained to maximize the contrastive loss. We use the same model architecture and training procedure as in [17] except that we set the size of the latent content embedding to $D_z = 16$ and use a downsampling factor of $S_{ds} = 2$.

Given the utterances $\mathbf{X}_{content}$ and \mathbf{X}_{style} , from which the content and style is to be extracted, respectively, we extract $\boldsymbol{\mu}_Z$ from $\mathbf{X}_{content}$ and $\bar{\mathbf{s}}$ from \mathbf{X}_{style} which are passed to the decoder to yield the converted speech feature \mathbf{X}_{vc} . We use Mel spectrograms to represent the speech signals. Waveform synthesis is performed with a pretrained ParallelWaveGAN [24].

2.2. Speaking rate estimation

To represent the speaking rate of an utterance, we opted for the utterance phoneme rate which is the number of phonemes in the utterance C_p over the total utterance duration N in seconds or samples, excluding silence segments. Another viable option could be the syllable rate [25, 26]. We use the phoneme rate due to the availability of corpora with hand-annotated phoneme labels or their likewise reliable estimation from transcriptions using forced alignment tools which allows us compare the results to an unsupervised speaking rate estimator (Section 2.2.2). The utterance duration can be simply measured after applying a Voice Activity Detection (VAD). To build a predictor for the number of phonemes, we consider two scenarios.

2.2.1. Scenario 1: Text-annotated training data

If transcribed audio data is available, we can perform a forced alignment to obtain the phoneme annotations. However, as mentioned above we do not want to be dependent of any transcriptions during inference time. Using the phoneme alignment, we train a local phoneme rate predictor which captures segmental variations in the speaking rate. We then compute the number of phonemes in the utterance by integrating over the local phoneme rate. This way, we can estimate the number of phonemes in the utterance without having the text transcription at inference time.

Following the definition from [25], the local phoneme rate $r_p(n)$ is the smoothed average of the instantaneous phoneme rate $I_p(n)$, $0 \leq n \leq N - 1$, with a window function $w(n)$:

$$r_p(n) = \frac{\sum_{m=0}^{K-1} I_p(n + m - \lfloor \frac{K-1}{2} \rfloor) w(m)}{\sum_{m=0}^{K-1} w(m)}, \quad (1)$$

where $w(n)$ is a Hann window of size K , and $I_p(n)$ is zero-padded at both sides. The instantaneous phoneme rate is a step-wise function where the step widths are given by the phoneme duration and the step heights are given by the inverse duration

$$I_p(n) = \frac{1}{S_{i+1} - S_i}, \quad S_i \leq n \leq S_{i+1}, \quad (2)$$

where S_i and S_{i+1} denote the start and stop times of the i -th phoneme. From this definition, it follows that $\int_{-\infty}^{\infty} I_p(n) dn = C_p$.

We train a local phoneme rate predictor consisting of stacked dilated 1d-convolution layers, two LSTM layers and a final convolution layer with softplus activation. The receptive field of the CNN is chosen to match the window size of $w(n)$ which is set to 625 ms [25]. We use Mel spectrograms with delta and delta-delta features as input and minimize the MSE between the prediction and true local phoneme rate. During inference, the number of phonemes can then be retrieved from the predicted local phoneme rate.

2.2.2. Scenario 2: Unlabeled training data

In the case that we do not have any annotated speech audio available, we can cast the estimation of the number of phonemes into an unsupervised phoneme segmentation task and count the number of predicted boundaries. Several works tackle phoneme segmentation without labels of which we chose one particularly simple approach that showed to generalize well to out-of-domain data [27].

A strided CNN takes the raw waveform as input and encodes it into a latent sequence vector. The encoder is trained to predict the next frame from the current one using a contrastive objective [23]. To predict the phoneme boundaries from the latent sequence vector, a score function is used which measures how dissimilar two adjacent frames are: High dissimilarity indicates a phoneme change. A peak detection algorithm is run over the dissimilarity curve and the estimated peaks are the predictions for the phoneme boundaries. We use the open-source implementation from [27] to train the model on the Timit+ training set.

3. Experiments

We evaluate the voice and speaking rate conversion approach on the CSTR VCTK corpus [28] which contains utterances from 110 English speakers. We split the VCTK dataset into three subsets for training, validation and testing with respectively 70, 20 and 20 speakers per subset. During training, we only use 90% of the utterances in the training and validation subsets and leave the remaining 10% for a closed speaker set evaluation. To train the local phoneme rate predictor, we trim the beginning and end silence and discard utterances which are shorter than 2 s after trimming. The forced alignments for scenario 1 training and speaking rate conversion evaluation are obtained with the Montreal Forced Aligner [29].

3.1. Accuracy of utterance phoneme rate prediction

We first evaluate how well we can predict the utterance phoneme rate from the speech signal under scenario 1 and 2. Starting off with scenario 1, we evaluate the local phoneme rate predictor in terms of the Pearson correlation coefficient r to the reference local phoneme rate. We compute the number of phonemes from the predicted local phoneme rate and evaluate the absolute difference e_{C_p} to the number of phonemes which were annotated from the forced alignment. We also show the relative error e_{C_p}/C_p which is equal to the relative difference in the utterance phoneme rates between prediction and target (the utterance duration N is the same for both). Table 1 shows the results for the local phoneme rate predictor. The predicted utterance phoneme rate is a reliable substitute for the true one because their relative difference is below the perception threshold of 5%, i.e., human listeners will likely not perceive any difference in tempo when listening to utterances with these phoneme rates given that the content is the same (see Section 3.2).

Table 2 shows the evaluation of the unsupervised segmentation on matched (Timit) and unmatched (VCTK unseen speakers) test sets. While the results for the matched case coincide with [27] there is a noticeable performance drop for the unmatched case caused by the domain mismatch. Still, the relative

Table 1: Performance of the local phoneme rate predictor evaluated against the ground truth local phoneme rate which is computed from the forced alignment.

Test set	$r \uparrow$	$e_{C_p} \downarrow$	$e_{C_p}/C_p \downarrow$
Seen speakers	.8319	2.29	4.14%
Unseen speakers	.8064	2.60	4.62%

Table 2: Quality of unsupervised phoneme segmentation [27]. Precision (P) and Recall (R) are calculated with a 20 ms tolerance window. All values are given in percent.

Test set	P \uparrow	R \uparrow	F1 \uparrow	R-Val \uparrow	$e_{C_p}/C_p \downarrow$
Timit	82.93	83.04	82.99	85.48	11.50
VCTK	69.70	72.34	71.00	74.92	11.43

difference in the utterance phoneme rate measured by e_{C_p}/C_p is similar for both test sets. While Timit has a higher ratio of predicted boundaries within the tolerance window (avg. 29.70 of 37.48) than VCTK (avg. 44.78 of 65.24) the total number of predicted boundaries is surprisingly accurate for VCTK (avg. 61.34 phonemes / sentence). Note also that the segmentation is applied to both the source and the target speaker’s utterance. Thus, the absolute segmentation quality may not be that important, as long as similar errors are made in both utterances.

3.2. Voice and speaking rate conversion

In order to measure whether the proposed approaches can effectively perform a speaking rate conversion, we rank the speakers in the test set by increasing average phoneme rate and perform a 3-fold split. We denote these splits as slow, normal and fast, respectively. The first 24 sentences in the VCTK corpus are parallel, i.e., they are read by all speakers. As an objective measure for the speaking rate conversion we evaluate the difference between the utterance duration after conversion and the duration of the parallel target sentence normalized to the target duration. The utterance phoneme rate for the target speaker is extracted from a different sentence from the target speaker. We compare the results under scenario 1 and 2 to the case when no speaking rate modification is performed. Figure 3 shows the results for the slow-to-fast (S2F), fast-to-slow (F2S) and normal-to-normal (N2N) conversion. The black dotted lines mark the Just Noticeable Difference (JND) which is the relative difference in duration to a reference stimulus that has to be exceeded such that the tempos of the two stimuli are perceived differently and was found to be around 5% relative duration difference [30]. Comparing the two variants to the unmodified case, we see an increased portion of the converted durations inside the JND which shows that the approach can indeed perform the desired speaking rate adaptation towards the target speaker.

We follow the methodology by [31] and perform speaker verification and automatic speech recognition to objectively assess the performance of the voice and joint voice and speaking rate conversion. For each source-target utterance conversion pair, we score both utterances against their respective speaker models (verification target) and against the other speaker (verification non-target) to get a topline Equal Error Rate (EER)

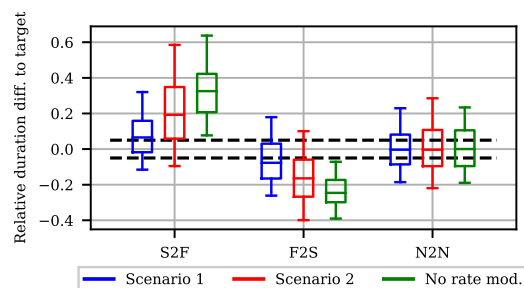


Figure 3: Normalized duration difference after speaking rate conversion on the first 24 parallel sentences in VCTK with unseen speakers. The black dotted lines mark the Just Noticeable Difference (JND) for speech tempo.

Table 3: *Objective evaluation metrics for voice and speaking rate conversion on seen speakers from VCTK. Copy-synthesis copies the source speaker vector during conversion, i.e., it reconstructs the source signal.*

Model	EER [%] ↓	WER [%] ↓
Natural	1.83	7.8
FVAE (copy-synthesis)	49.90	11.5
FVAE (no rate mod.)	6.81	19.5
+ Scenario 1	6.40	22.3
+ Scenario 2	7.01	22.0
AutoPST (stage 1)	8.94	49.0
AutoPST (stage 2)	7.42	74.1

for natural speech. For VC speech, we replace the target utterances with the conversions, i.e., lower EER is better. We use acoustic and language models from ESPNet [32] pretrained on LibriSpeech to report the Word Error Rate (WER). We use the open-source implementation [33] to train AutoPST where we slightly adapted the feature extraction to work with the Parallel-WaveGAN. Since AutoPST was trained with an one-hot speaker identity vector, we evaluate on the closed speaker test set. The results are summarized in Table 3.

The WSOLA algorithm introduces some artifacts which negatively impact the WER, which nevertheless is still much lower than for AutoPST. The WER for Scenario 1 is slightly higher despite the same voice conversion and time stretching approach as for Scenario 2 due to the stronger interpolations (see also Figure 3). Using an interpolation factor for WSOLA that significantly deviates from 1 increases the amount of artifacts. The bad WER for AutoPST stems from the autoregressive spectrogram generation with the Transformer decoder which can arbitrarily shift the content along the time and suffers from repetitions and elisions if not well regularized. We noticed that AutoPST was not able to reliably predict the utterance stop token in many cases which lead to a lot of abnormal speech patterns in the conversions. The first stage of AutoPST does not target to modify the tempo so we limited the duration of the conversion to its original duration which prevented the WER from rising too high. Adding the speaking rate conversion gives negligible improvements in EER which may stem from an increased utterance duration and therefore more accurate speaker statistics.

3.3. Subjective listening tests

In a subjective evaluation¹, we compared the proposed approaches and AutoPST against the same recordings of the converted target speakers. Participants rated the samples on overall quality, and compared their similarity to the reference sample in the dimensions of voice and speaking rate. The reference recording was modified by the same vocoder used across models [34] to allow a more fine grained comparison between the systems. The reference recording for similarity was the utterance from which the target speaking rate was estimated to ensure a meaningful comparison of speaking rate. 24 different samples were selected randomly and stratified by gender of source speaker, target speaker, as well as the direction of speaking rate conversion (S2F and F2S). To keep the experiment reasonably short, the samples were split into two lists across the source speaker’s gender, with half of the participants listening to each. 26 (13 female, 13 male) participants (native speakers of English) were recruited and paid via Prolific. The collected Mean Opinion Scores (MOS) and standard deviations can be found in Table 4.

¹Audio examples: go.upb.de/interspeech2022

Table 4: *Subjective MOS for overall quality, voice- and speaking rate conversion. Control denotes the optimal topline of the same recording, mixed in with the conversion samples.*

Model	Quality	Simil. Voice	Simil. Speed
Control	4.24 ± 0.99	4.85 ± 0.51	4.76 ± 0.67
FVAE (no rate mod.)	2.85 ± 1.33	3.29 ± 1.35	3.00 ± 1.40
+ Scenario 1	2.65 ± 1.22	3.35 ± 1.33	3.29 ± 1.47
+ Scenario 2	2.79 ± 1.29	3.33 ± 1.33	3.10 ± 1.45
AutoPST (stage 1)	2.56 ± 1.36	2.82 ± 1.39	3.01 ± 1.47
AutoPST (stage 2)	1.81 ± 1.15	2.88 ± 1.44	2.41 ± 1.48

To test the statistical significance of the MOS differences, we computed Linear Mixed Effects Regression models [35] for each quality dimension, with participants added as random effects. We found a significant improvement in overall quality of FVAE over AutoPST (stage 1) ($\beta=0.29^*$). The voice dimension also showed a significant improvement of FVAE over AutoPST (stage 1) ($\beta=0.47^{**}$). This improvement is also significant for the speaking rate adaptation models Scenario 1 ($\beta=0.53^{***}$) and Scenario 2 ($\beta=0.51^{***}$). In the judgements of perceived speaking rate similarity we found a significant interaction between the VC model and the direction of speed conversion, so the analysis was carried out independently for both speed conditions. In the S2F case, there were no significant differences in perceived speaking rate scores between the models. More notably, in the opposing F2S case we found improvements of Scenario 1 over the baseline FVAE ($\beta=0.65^{**}$) and AutoPST (stage 1) ($\beta=0.71^{***}$). As observed in the objective evaluations, the failing duration conversion of AutoPST (stage 2) leads to inferior quality and speaking rate distortions as judged by the participants. This showed to be significant for all quality dimensions, i.e., overall quality (FVAE: $\beta=-1.03^{***}$; Scenario 1: $\beta=-0.84^{***}$; Scenario 2: $\beta=-0.98^{***}$), voice similarity (FVAE: $\beta=-0.40^*$; Scenario 1: $\beta=-0.47^{**}$; Scenario 2: $\beta=-0.45^{**}$) and speaking rate similarity (FVAE: $\beta=-0.87^{***}$; Scenario 1: $\beta=-1.52^{***}$; Scenario 2: $\beta=-2.08^{***}$).

4. Conclusion

This paper investigated the joint conversion of voice and speaking rate. For voice conversion, we evaluated the performance of a neural disentanglement system that runs fully unsupervised during training and testing. Speaking rate adaptation was carried out by linear interpolation where the interpolation factor was estimated as the ratio between phoneme rates of source and target utterance. Adding the speaking rate conversion helped to improve the similarity to the target speaker while only introducing minor artifacts to quality and intelligibility. Our proposed approach showed to consistently outperform the reference baseline in all metrics and is applicable to unknown speakers. Future works may investigate the combined use of phoneme and syllable rates for speaking rate conversion which showed to better correlate with perceived tempo [25]. Also, alternative approaches using neural networks [36, 37] may replace the WSOLA algorithm to reduce the amount of artifacts.

5. Acknowledgements

This research has been funded by (Deutsche Forschungsgemeinschaft DFG, German Research Foundation), projects 446378607 and TRR 318/1 2021 – 438445824. Computational resources were provided by the Paderborn Center for Parallel Computing (PC²).

6. References

- [1] B. Sisman, J. Yamagishi, S. King, and H. Li, "An overview of voice conversion and its challenges: From statistical modeling to deep learning," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 132–157, 2020.
- [2] T. Kaneko, H. Kameoka, K. Tanaka, and N. Hojo, "Stargan-vc2: Rethinking conditional methods for stargan-based voice conversion," in *Proc. Interspeech*, 2019, pp. 679–683.
- [3] H.-S. Choi, J. Lee, W. Kim, J. Lee, H. Heo, and K. Lee, "Neural analysis and synthesis: Reconstructing speech from self-supervised representations," in *Proc. Advances in Neural Information Processing Systems*, 2021.
- [4] K. Qian, Y. Zhang, S. Chang, X. Yang, and M. Hasegawa-Johnson, "AutoVC: Zero-shot voice style transfer with only autoencoder loss," in *Proc. 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, 2019, pp. 5210–5219.
- [5] D.-Y. Wu, Y.-H. Chen, and H.-y. Lee, "VQVC+: One-shot voice conversion by vector quantization and u-net architecture," in *Proc. Interspeech*, 2020, pp. 4691–4695.
- [6] J.-C. Chou and H.-Y. Lee, "One-Shot Voice Conversion by Separating Speaker and Content Representations with Instance Normalization," in *Proc. Interspeech*, 2019, pp. 664–668.
- [7] Y.-H. Chen, D.-Y. Wu, T.-H. Wu, and H.-y. Lee, "Again-vc: A one-shot voice conversion using activation guidance and adaptive instance normalization," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021, pp. 5954–5958.
- [8] J.-C. Chou, C.-C. Yeh, H.-Y. Lee, and L.-S. Lee, "Multi-target Voice Conversion without Parallel Data by Adversarially Learning Disentangled Audio Representations," in *Proc. Interspeech*, 2018, pp. 501–505.
- [9] B. L. Smith, B. L. Brown, W. J. Strong, and A. C. Rencher, "Effects of speech rate on personality perception," *Language and speech*, vol. 18, no. 2, pp. 145–152, 1975.
- [10] B. Lindblom, "Explaining phonetic variation: A sketch of the h&h theory," in *Speech production and speech modelling*. Springer, 1990, pp. 403–439.
- [11] K. Tanaka, H. Kameoka, T. Kaneko, and N. Hojo, "Atts2s-vc: Sequence-to-sequence voice conversion with attention and context preservation mechanisms," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 6805–6809.
- [12] T. Hayashi, W.-C. Huang, K. Kobayashi, and T. Toda, "Non-autoregressive sequence-to-sequence voice conversion," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2021, pp. 7068–7072.
- [13] J.-X. Zhang, Z.-H. Ling, and L.-R. Dai, "Non-parallel sequence-to-sequence voice conversion with disentangled linguistic and speaker representations," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 540–552, 2019.
- [14] K. Qian, Y. Zhang, S. Chang, J. Xiong, C. Gan, D. Cox, and M. Hasegawa-Johnson, "Global prosody style transfer without text transcriptions," in *Proc. International Conference on Machine Learning*. PMLR, 2021, pp. 8650–8660.
- [15] J. Wang, J. Li, X. Zhao, Z. Wu, S. Kang, and H. Meng, "Adversarially learning disentangled speech representations for robust multi-factor voice conversion," in *Proc. Interspeech 2021*, 2021, pp. 846–850.
- [16] L. Prananta, B. M. Halpern, S. Feng, and O. Scharenborg, "The effectiveness of time stretching for enhancing dysarthric speech for improved dysarthric speech recognition," *arXiv preprint arXiv:2201.04908*, 2022.
- [17] J. Ebberts, M. Kuhlmann, T. Cord-Landwehr, and R. Haeb-Umbach, "Contrastive predictive coding supported factorized variational autoencoder for unsupervised learning of disentangled speech representations," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2021.
- [18] W. Verhelst and M. Roelands, "An overlap-add technique based on waveform similarity (WSOLA) for high quality time-scale modification of speech," in *Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1993, pp. 554–557.
- [19] S.-H. Lee, H.-R. Noh, W.-J. Nam, and S.-W. Lee, "Duration controllable voice conversion via phoneme-based information bottleneck," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 1173–1183, 2022.
- [20] N. Jaitly and G. E. Hinton, "Vocal tract length perturbation (VTLP) improves speech recognition," in *Proc. ICML Workshop on Deep Learning for Audio, Speech and Language*, 2013, p. 21.
- [21] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," *arXiv preprint arXiv:1607.08022*, 2016.
- [22] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. International Conference on Learning Representations*, 2014.
- [23] A. v. d. Oord, Y. Li, and O. Vinyals, "Representation learning with contrastive predictive coding," *arXiv preprint arXiv:1807.03748*, 2018.
- [24] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel WaveGAN: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2020, pp. 6199–6203.
- [25] H. R. Pfister, "Reducing segmental duration variation by local speech rate normalization of large spoken language resources," in *Proc. International Conference on Language Resources and Evaluation*, 2002.
- [26] Y. Jiao, M. Tu, V. Berisha, and J. Liss, "Online speaking rate estimation using recurrent neural networks," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing*, 2016, pp. 5245–5249.
- [27] F. Kreuk, J. Keshet, and Y. Adi, "Self-supervised contrastive learning for unsupervised phoneme segmentation," in *Proc. Interspeech*, 2020, pp. 3700–3704.
- [28] J. Yamagishi, C. Veaux, and K. MacDonald, "CSTR VCTK Corpus: English Multi-Speaker Corpus for CSTR Voice Cloning Toolkit (version 0.92)," University of Edinburgh. The Centre for Speech Technology Research (CSTR).
- [29] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kaldii," in *Proc. Interspeech*, 2017, pp. 498–502.
- [30] H. Quené, "On the just noticeable difference for tempo in speech," *Journal of Phonetics*, vol. 35, pp. 353–362, 2007.
- [31] R. K. Das, T. Kinnunen, W.-C. Huang, Z. Ling, J. Yamagishi, Y. Zhao, X. Tian, and T. Toda, "Predictions of subjective ratings and spoofing assessments of voice conversion challenge 2020 submissions," in *Proc. Joint workshop for the Blizzard Challenge and Voice Conversion Challenge 2020*, 2020, pp. 99–120.
- [32] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. E. Y. Soplin, J. Heymann, M. Wiesner, N. Chen *et al.*, "Espnet: End-to-end speech processing toolkit," in *Proc. Interspeech*, 2018, pp. 2207–2211.
- [33] K. Qian, "Global rhythm style transfer without text transcriptions," Official implementation. [Online]. Available: <https://github.com/auspicious3000/AutoPST>
- [34] T. Hayashi, "Unofficial ParallelWave GAN with PyTorch," version 0.5.4. [Online]. Available: <https://github.com/kan-bayashi/ParallelWaveGAN>
- [35] R. H. Baayen, D. J. Davidson, and D. M. Bates, "Mixed-effects modeling with crossed random effects for subjects and items," *Journal of memory and language*, vol. 59, pp. 390–412, 2008.
- [36] M. Morrison, Z. Jin, N. J. Bryan, J.-P. Caceres, and B. Pardo, "Neural pitch-shifting and time-stretching with controllable LPC-Net," *arXiv preprint arXiv:2110.02360*, 2021.
- [37] T. Okamoto, K. Matsubara, T. Toda, Y. Shiga, and H. Kawai, "Neural speech-rate conversion with multispeaker wavenet vocoder," *Speech Communication*, vol. 138, pp. 1–12, 2022.