# Improving Voice Conversion for Dissimilar Speakers Using Perceptual Losses

Suhita Ghosh<sup>1</sup>, Yamini Sinha<sup>2</sup>, Ingo Siegert<sup>2</sup> and Sebastian Stober<sup>1</sup>

 $^{1}\ Artificial\ Intelligence\ Lab,\ Otto-von-Guericke-University\ Magdeburg,\ Germany$ 

 $^2$  Mobile Dialog Systems, Otto-von-Guericke-University Magdeburg, Germany

## Introduction

The rising trend of using voice as a means of interacting with smart devices has sparked worries over the protection of users' privacy and data security [1]. These concerns have become more pressing, especially after the European Union's adoption of the General Data Protection Regulation (GDPR). The information contained in an utterance encompasses critical personal details about the speaker, such as their age, gender, socio-cultural origins and more. If there is a security breach and the data is compromised, attackers may utilise the speech data to circumvent the speaker verification systems or imitate authorised users [2]. Therefore, it is pertinent to anonymise the speech data before being shared across devices, such that the source speaker of the utterance cannot be traced. Voice conversion (VC) can be used to achieve speech anonymisation, which involves altering the speaker's characteristics while preserving the linguistic content.

Many voice conversion approaches have been proposed over the years, where the deep learning-based methods outperform the conventional ones [3]. Further, the generative adversarial network (GAN) based approaches produce natural-sounding conversions [3]. However, the quality is dependant on the selection of the target speaker. This is because GAN-based VC methods typically use non-parallel data, which prevents the computation of loss between the source utterance and the conversion conditioned on a speaker other than the source. The quality of conversion degrades when the acoustic properties between the source and target speakers are diverse. However, to achieve a successful anonymisation, the source and target speakers should not have very similar acoustic properties, such as pitch.

In this work, we propose perceptual losses which are computed between the source and converted utterances. The losses facilitate the model to capture representations which are pertinent with respect to how humans perceive speech quality. The models trained with the proposed losses produce less robotic voices compared to the baseline, and improves the overall quality for all target speakers.

### **Related Work**

In earlier VC approaches, parallel data was utilised, where utterances of the same linguistic content are available from both the source and target speakers. Many traditional statistical modelling-based parametric [4, 5] and non-parametric [6, 7] parallel VC methods were proposed. Compared to the conventional methods, the sequence-to-sequence deep neural network (DNN) based method [8] using parallel data produces less robotic voices. However, it does not preserve prosody and produces mispronuncia-



Abbildung 1: StarGAN based Voice Conversion architecture

tions [9]. Further, the model learns one-to-one mapping, limiting its usage.

The recent works focus more on non-parallel data [3], as it is much easier and cheaper to collect. A few VC methods [10] use phonetic posteriorgrams (PPGs) as input to the encoder-decoder framework, which produces the translated acoustic features, consequently used by a vocoder to generate the converted speech. The PPGbased conversions are generally not smooth resulting in degraded voice quality and naturalness [3]. Many nonparallel variational auto encoder (VAE) [3, 11] methods were proposed, which typically disentangle the content and speaker embeddings using a reconstruction loss. The VAE-based approaches are prone to spectrum smoothing, which leads to a buzzy-sounding voice [3]. A plethora of GAN-based VC approaches [12, 13, 14] have been proposed to overcome this over-smoothing effect. GAN-based VC approaches use cycle-consistency loss [15], which enable them to use non-parallel data.

#### Method

Our architecture is based on the GAN-based method StarGANv2-VC [14]. We describe the architecture and then the perceptual losses used in the work.

#### StarGANv2-VC

StarGANv2-VC [14] is a non-parallel many-to-many voice conversion GAN-based approach. The architecture is shown in Figure 1. In StarGANv2-VC, only one generator is required for conversion among many pairs. We describe the pertinent components of the framework, which are portrayed in Figure 1:

- Generator: The generator G produces the converted mel-spectrogram  $\overline{X}$  using three inputs: log mel-spectrogram generated from the source utterance  $X_s$ , fundamental frequency (F0) embeddings  $h_f$  from the source utterance and target speaker's style-code  $h_s$ . The F0 embeddings are the convolutional outputs from a pre-trained joint detection and classification (JDC) network [16], which has convolutional layers followed by BLSTM units. The converted melspectrogram  $\overline{X}$  bears the style/timbre of the target speaker and the linguistic content of the source.
- Speaker Style Encoder: The speaker style encoder S captures representations of the speaker's style. The style may represent accent, mannerism and other attributes which can be associated with the speaker independent of the content spoken. Provided a mel-spectrogram  $X_r$ , which is generated from a reference utterance different from the source melspectrogram  $X_s$  and a speaker-code r, the S generates the speaker style embeddings  $h_s$ . The speakercode is a one-hot encoding of the speaker labels. The embedding  $h_s$  acts as one of the inputs to the generator G, which contributes to the style of the converted mel-spectrogram  $\overline{X}$ . S initially processes the input mel-spectrogram through multiple convolutional layers which are shared for all speakers, followed by a speaker-specific linear layer which maps the shared features into a speaker-specific style embedding.
- Discriminator and Speaker Classifier: The architecture has a discriminator D, as present in any GAN model, which performs the quality check of the conversions by capturing the representations for the real and fake samples. The additional adversarial speaker classifier (C) has the same architecture as D. When the D is trained, keeping the weights of G fixed, the C classifies the source speaker, which encourages G to produce conversions having no trace of source speaker's attributes. When G is trained, the weights of D are fixed, the C classifies the target speaker, which facilitates providing feedback to G, such that it produces conversions sounding like the target speaker.

## Perceptual Losses

Task specific perceptual losses facilitate models to capture pertinent representations, required to achieve the goal [17]. In our case, to improve the overall quality of voice conversions for all target speakers.

• Short Time Objective Intelligibility (STOI): STOI [18] is an intrusive metric that compares the degraded signal with the high quality ground truth to measure the intelligibility of the noisy signal. The STOI score ranges from 0 to 1, with higher values indicating better intelligibility. In our case, X and  $\bar{X}$  act as the ground truth and noisy signals respectively. To calculate STOI, firstly speech signals are divided into short frames where each frame overlaps with the adjacent frames to capture the temporal context. For each frame, the short-time power spectrum is calculated using a Fourier transform. The modulation spectrum of both the signals are calculated by applying a perceptual auditory filter-bank (onethird octave band) to the short-time power spectra. The correlation coefficient between the modulation spectra of the original and degraded speech signals are calculated, which provides a similarity measure between two spectra. The STOI score at time frame m is calculated by taking average over all one-third octave bands as shown in Equation 1, where j is the index of the one-third octave band,  $x_{i,m}$  and  $\bar{x}_{j,m}$  denote the vectors representing the short-term temporal envelopes for time frame m and one-third octave band j of the clean and noisy signals respectively.  $\mu(\cdot)$  denotes sample mean and J is the total number of the one-third octave bands.

$$f_{stoi}(X_m, \bar{X}_m) = \sum_{j=1}^J \frac{(x_{j,m} - \mu_{x_{j,m}})(\bar{x}_{j,m} - \mu_{\bar{x}_{j,m}})}{\|x_{j,m} - \mu_{x_{j,m}}\|_2 \|\bar{x}_{j,m} - \mu_{\bar{x}_{j,m}}\|_2}$$
(1)

The loss is calculated as shown in Equation 4, as done in [19], where mean squared error (Equation 3) is also considered along with STOI score (Equation 2), as STOI calculates the discrepancy only for frequencies below 4.3 KHz.  $\lambda_{stoi}$  and  $\lambda_{mse}$  are hyperparameters which weigh the contribution of  $L_{stois}$  and  $L_{stoim}$  respectively.

$$L_{stoi_s} = (1 - f_{stoi}(X_m, \bar{X}_m)) \tag{2}$$

$$L_{stoi_m} = (\|X_m^J - \bar{X}_m^J\|_1 / J)$$
 (3)

$$L_{stoi} = \frac{1}{m} (\lambda_{stoi} * L_{stoi_s} + \lambda_{mse} * L_{stoi_m})$$
(4)

• Predicted Mean Opinion Score (pMOS): MOS is a subjective measure which is used to assess the naturalness of the converted voice in voice conversion [14]. The measure correlates well with human perception of audio quality and naturalness. However, it is arduous and expensive as many participants' involvement is needed. Therefore, a measure similar to MOS is desirable, which captures the intrinsic naturalness of the conversions. MOSNet proposed in [20] can be used as a proxy MOS score generator. MOS-Net is a combination of a convolutional neural network (CNN) and bidirectional long short-term memory (BLSTM) architecture. The CNN layers extract the representations required to assess the quality of the frames. BLSTM can effectively incorporate prolonged temporal dependencies and sequential traits into representative features. At the end two fully-connected layers are used, which regresses the frame-wise features into a frame-level quality score, which is followed by a global averaging operation to obtain the utterance-level score. The loss is calculated as shown in Equation 5, where MOS(.) denotes MOSNet and  $\lambda_{mos}$  is a hyperparameter. The loss encourages G to produce conversions having naturalness similar to the original utterance.

$$L_{mos} = \lambda_{mos} * \|MOS(X) - MOS(\bar{X})\|_1$$
 (5)

• Pitch correlation coefficient (PCC): Pitch is the perceptual measure of F0. The pitch contour contributes to the intonation or prosody of an utterance [21]. PCC is the Pearson correlation between two normalised F0 contours, which provides the similarity between two utterances with respect to prosody [21]. The F0 contours for two utterances having same content and intonation will vary for two groups (age, gender, etc). However, there should not be a large difference between the normalised F0 contours, i.e. the change in F0 over time should not vary much. Therefore, a higher PCC is desirable. PCC Loss is represented in Equation 6, where Pearson(.) is the Pearson correlation operator.

$$L_{pcc} = 1 - Pearson\left(\frac{F(X)}{\|F(X)\|_{1}}, \frac{\bar{F}(X)}{\|\bar{F}(X)\|_{1}}\right) \quad (6)$$

### **Objective Function**

The generator G is trained with loss  $L_G$ , where  $L_{adv}$  is the typical GAN adversarial loss,  $L_{aspk}$  is the adversarial speaker classification loss,  $L_{sty}$  is the style reconstruction loss and  $L_{cyc}$  is the cyclic consistency loss, as proposed in [14].  $L_p$  denotes one of the proposed perceptual losses.

$$L_{G} = \min_{G,S} L_{adv} + \lambda_{aspk} L_{aspk} + \lambda_{sty} L_{sty} + \lambda_{cyc} L_{cyc} + \lambda_{p} L_{p}$$
(7)

The discriminator D and classifier C are trained using the objective function shown in Equation 8, where  $L_{spk}$ is the speaker classification loss [14].

$$L_D = \min_{D,C} - L_{adv} + \lambda_{spk} L_{spk} + \lambda_p L_p \tag{8}$$

The  $\lambda$  for the corresponding loss denotes the hyperparameter which weighs the loss's contribution.

#### **Experiment Details**

We train all the models using English utterances of the 20 speakers from VCTK [22] dataset, as done in [14]. The utterances are resampled to 24 kHz and randomly split as 80%/10%/10% (train/val/test). The models are trained for 150 epochs and with batch size of 64. The log-melspectrograms are derived from 2 second long utterances. AdamW optimizer is used with initial learning rate of 0.0001. The hyperparameters are set as:  $\lambda_{spk} =$  $0.1, \lambda_{aspk} = 0.5, \lambda_{sty} = 1, \lambda_{cyc} = 1, \lambda_{stoi} = 1, \lambda_{mse} = 0.1.$ STOI is computed using hyperparameters same as in [18]. The naturalness of the conversions is evaluated using pMOS. The intelligibility of the conversions is measured using character error rate (CER), using the transcriptions from Whisper [23] medium-English model. We use automatic speaker verification (ASV) to measure speaker similarity as done in [14]. We trained an AlexNet as done in [14] for speaker classification for the selected twenty speakers. The classification accuracy (Speaker CLS) serves as the objective metric to assess how close the conversions sound to the target speaker.

#### **Results and Discussion**

We randomly selected 5 male and 5 female speakers as the target speakers. For each source speaker, randomly



**Abbildung 2:** Change in F0 contour (blue) for the original, conversions using baseline and PCC. The pictures are produced using Praat software, which also shows the F0 contour and intensity contour (yellow).

50 utterances are selected, which leads to 1000 conversions. The model trained using  $L_{pcc}$  produces the best results with respect to naturalness and intelligibility, followed by the model trained using pMOS loss. It is also observed the standard deviation for the baseline is much higher than the ones trained using target perceptual losses. Therefore, the proposed losses produce better quality conversions overall, and not just for specific target speakers. With respect to speaker similarity, all the models perform similarly, where PCC outperforms. It can also

**Tabelle 1:** Mean and standard deviation with objective metrics.

 Standard deviation in brackets.

Model	Objective Metrics		
model	MOS	CER	Speaker
			CLS
Baseline	3.20(1.23)	8.42(9.39)	78.7%
STOI	3.72(0.32)	3.01 (4.50)	80.8%
pMOS	3.73(0.23)	3.10(3.31)	83.1%
PCC	<b>3.80</b> (0.25)	<b>2.96</b> (3.78)	85.8%

be seen that the models trained with perceptual losses maintain the intonation better than the baseline, as seen in Figure 2. Further, it can been seen that in case of a large change in F0 contour (between L1 and L2), the baseline fails to maintain the structure.

#### Conclusion

In this work we propose perceptual losses which are calculated between the original and converted samples conditioned on the target speaker. The losses facilitate the model to disentangle the content and speaker representations, which leads to improved conversions not dependant on the target speaker selection. In this work, we focus only on the naturalness and intelligibility aspect of voice quality. As future work, we will perform listening tests to validate the results obtained through objective measures. Further, we intend to incorporate perceptual losses which capture the emotional content as well. This would be useful for the intelligent speech devices, whose response is driven by the emotion of the end-user.

## Acknowledgements

This research has been partly funded by the Federal Ministry of Education and Research of Germany in the project Emonymous (project number S21060A) and partly funded by the Volkswagen Foundation in the project AnonymPrevent (AI-based Improvement of Anonymity for Remote Assessment, Treatment and Prevention against Child Sexual Abuse).

## Literatur

- [1] C. Wienrich, C. Reitelbach, and A. Carolus, "The trustworthiness of voice assistants in the context of healthcare investigating the effect of perceived expertise on the trustworthiness of voice assistants, providers, data receivers, and automatic speech recognition," Frontiers in Computer Science, vol. 3, p. 685250, 2021.
- [2] N. Zhang, X. Mi, X. Feng, X. Wang, Y. Tian, and F. Qian, "Dangerous skills: Understanding and mitigating security risks of voice-controlled third-party functions on virtual personal assistant systems," in IEEE Symposium on Security and Privacy, 2019, pp. 1381–1396.
- [3] Sisman, Berrak, et al. "An overview of voice conversion and its challenges: From statistical modeling to deep learning." IEEE/ACM Transactions on Audio, Speech, and Language Processing 29 (2020): 132-157.
- [4] Kobayashi, Kazuhiro, et al. "The NU-NAIST Voice Conversion System for the Voice Conversion Challenge 2016." Interspeech. 2016.
- [5] Helander, Elina, et al. "Voice conversion using dynamic kernel partial least squares regression." IEEE transactions on audio, speech, and language processing 20.3 (2011): 806-817.
- [6] Takashima, Ryoichi, Tetsuya Takiguchi, and Yasuo Ariki. "Exemplar-based voice conversion using sparse representation in noisy environments." IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences 96.10 (2013): 1946-1953.
- [7] Jin, Zeyu, et al. "Cute: A concatenative method for voice conversion using exemplar-based unit selection." 2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2016.
- [8] Liu, Ruolan, Xiao Chen, and Xue Wen. "Voice conversion with transformer network." ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2020.
- [9] Lian, Zheng, et al. "ARVC: An Auto-Regressive Voice Conversion System Without Parallel Training Data." INTERSPEECH. 2020.
- [10] Sun, Lifa, et al. "Phonetic posteriorgrams for manyto-one voice conversion without parallel data trai-

ning." 2016 IEEE International Conference on Multimedia and Expo (ICME). IEEE, 2016.

- [11] Wu, Da-Yi, and Hung-yi Lee. "One-shot voice conversion by vector quantization." ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2020.
- [12] Kameoka, Hirokazu, et al. "StarGAN-VC: Nonparallel many-to-many voice conversion using star generative adversarial networks." 2018 IEEE Spoken Language Technology Workshop (SLT). IEEE, 2018.
- [13] Sakamoto, Shoki, et al. "StarGAN-VC+ ASR: Stargan-based non-parallel voice conversion regularized by automatic speech recognition" arXiv preprint arXiv:2108.04395 (2021).
- [14] Li, Yinghao Aaron, Ali Zare, and Nima Mesgarani. "StarGANv2-VC: A diverse, unsupervised, nonparallel framework for natural-sounding voice conversion." arXiv preprint arXiv:2107.10394 (2021).
- [15] Kaneko, Takuhiro, and Hirokazu Kameoka. "CycleGAN-VC: Non-parallel voice conversion using cycle-consistent adversarial networks." 2018 26th European Signal Processing Conference (EUSIPCO). IEEE, 2018.
- [16] Kum, Sangeun, and Juhan Nam. "Joint detection and classification of singing voice melody using convolutional recurrent neural networks. Äpplied Sciences 9.7 (2019): 1324.
- [17] Ghosh, Suhita, et al. "Perception-aware losses facilitate CT denoising and artifact removal." 2021 IEEE 2nd International Conference on Human-Machine Systems (ICHMS). IEEE, 2021.
- [18] Taal, Cees H., et al. "An algorithm for intelligibility prediction of time–frequency weighted noisy speech." IEEE Transactions on Audio, Speech, and Language Processing 19.7 (2011): 2125-2136.
- [19] Zhao, Yan, et al. "Perceptually guided speech enhancement using deep neural networks." 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE, 2018.
- [20] Lo, Chen-Chou, et al. Mosnet: Deep learning based objective assessment for voice conversion.ärXiv preprint arXiv:1904.08352 (2019).
- [21] D. Hirst, "A Praat plugin for momel and intsint with improved algorithms for modelling and coding intonation. icphs xvi, Saabrücken," 2007.
- [22] J. Yamagishi, C. Veaux, K. MacDonald et al., "CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92)," 2019.
- [23] Radford, Alec, et al. "Robust speech recognition via large-scale weak supervision." arXiv preprint ar-Xiv:2212.04356 (2022).