



Semi-supervised voice conversion with amortized variational inference

Cory Stephenson, Gokce Keskin, Anil Thomas, Oguz H. Elibol

Intel AI Lab, Santa Clara, CA, USA

cory.stephenson@intel.com

Abstract

In this work we introduce a semi-supervised approach to the voice conversion problem, in which speech from a source speaker is converted into speech of a target speaker. The proposed method makes use of both parallel and non-parallel utterances from the source and target simultaneously during training. This approach can be used to extend existing parallel data voice conversion systems such that they can be trained with semi-supervision. We show that incorporating semi-supervision improves the voice conversion performance compared to fully supervised training when the number of parallel utterances is limited as in many practical applications. Additionally, we find that increasing the number non-parallel utterances used in training continues to improve performance when the amount of parallel training data is held constant.

Index Terms: voice conversion, semi-supervised learning, variational inference, deep learning

1. Introduction

The goal of voice conversion (VC) is to take in speech produced by one person (source) and process it such that it sounds like it was produced by a different (target) speaker. VC systems have a diverse set of potential applications including the construction of more natural synthetic voices, anonymous transmission of recorded speech, voice spoofing, and data normalization for further speech processing applications. Due the broad applicability and inherent difficulty of the problem, the design of VC systems has been an area of consistent interest for decades and continues to see active research [1, 2].

Of particular interest are the statistical approaches that do not require access to word or phonetic transcriptions. These include the early work using vector quantization [3], which inspired the popular Gaussian mixture model (GMM) conversion method [4] and its later improvements [5]. While these methods can produce good results, they require many parallel utterances for training, are sensitive to misalignment in the training data, and special care must be taken to avoid the 'buzzing' that arises due to the usual maximum-likelihood training objective [5].

The necessity of parallel data for training is quite limiting, as the collection of this kind of data is a slow and expensive process. The desire to avoid this requirement has led to the development of VC approaches which use only non-parallel data from the source and target speakers [6, 7, 8]. While removing the parallel data requirement eases the burden on data collection, it introduces extra difficulty in ensuring the converted speech is both high quality and unchanged in phonetic content [9].

More recently, advances in training deep nonlinear models have led to renewed interest in applying neural network techniques to the VC problem. Methods which use deep neural networks as feature extractors or to parameterize the conditional distributions in generative models have proven effective in doing VC both with [10, 11, 12, 13] and without [14] parallel

data. These models excel at producing natural sounding speech. However, the increased model complexity comes with an increased demand for larger quantities of data in both the parallel and non-parallel case.

It is the goal of this work to fill the gap between parallel and non-parallel data voice conversion by introducing a method that uses both types of data simultaneously during training. To do this, we frame voice conversion as a semi-supervised learning problem. This follows naturally from previous shallow generative approaches such as [4, 5], which make use of a shared set of latent variables that generate both the source and target speech. While coupling this type of model with nonlinear transformations of the latent variables yields an intractable inference problem, we find that amortized variational inference as applied to deep generative models [15, 16] makes both training and conversion efficient.

We demonstrate the effectiveness of semi-supervised training in multiple ways. First, we extend a well-known neural network VC algorithm [11] such that it can be trained with semi-supervision. Then we show that incorporating non-parallel data in training leads to higher quality voice conversion when parallel data is scarce. Additionally, we verify that the semi-supervised training gives equivalent or better results to training with only parallel data when non-parallel data is scarce. Finally, we confirm that conversion accuracy continues to increase with increasing amounts of non-parallel training data, albeit with diminishing returns. We also observe that the semi-supervised training results in audio of equal or higher perceptual quality than the parallel data conversion systems.

2. Related work

While there are many proposed voice conversion methods in the literature (for a more in depth review, see [17]), here we examine some of the popular statistical approaches which use parallel training data. A common way to think about this problem is to posit the existence of a latent variable z_t that describes the sound is produced at time t , but not the characteristics of the source that produced it. The job of the voice conversion system then is to *infer* z_t from a sound produced by the source, and *generate* the corresponding sound in the target's voice.

2.1. Gaussian mixture model voice conversion

At a high level, this is how the well known GMM VC system [4] achieves conversion. In this approach, the input consists of a sequence of features $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T\}$ which are each assumed to be independently Gaussian mixture distributed. The means of the mixture components are parameterized with μ_i covariance matrices with Σ_i . In this model, z_t is a categorical latent variable that indicates which Gaussian is responsible for the observation. By virtue of the Gaussian mixture assumption, it is possible to compute the posterior distribution $P(z_t = i | \mathbf{x}_t)$ for z_t exactly.

As stated, this model can be trained entirely with non-parallel data with maximum likelihood. To achieve voice conversion, the target features are generated according to

$$\hat{\mathbf{y}}_t = \sum_{i=1}^K P(\mathbf{z}_t = i | \mathbf{x}_t) [\boldsymbol{\nu}_i + \boldsymbol{\Gamma}_i \boldsymbol{\Sigma}_i^{-1} (\mathbf{x}_t - \boldsymbol{\mu}_i)] \quad (1)$$

The parameters $\boldsymbol{\nu}_i$ and $\boldsymbol{\Gamma}_i$ must be learned using data consisting of $(\mathbf{x}_t, \mathbf{y}_t)$ pairs. This is done by minimizing the mean square error (MSE) $\|\mathbf{y}_t - \hat{\mathbf{y}}_t\|^2$ on the training data. After training, conversion proceeds by inferring $P(\mathbf{z}_t = i | \mathbf{x}_t)$ from the source features $\{\mathbf{x}_t\}$ and generating the converted features $\{\hat{\mathbf{y}}_t\}$ with Eq. 1.

This model treats conversion of each frame independently, and so the converted speech does not always mimic the dynamics of real speech. Later, this approach was extended to include temporal information relating to the trajectory of features [5, 18].

2.2. Neural network VC with parallel data

Motivated by the desire to move VC systems beyond the shallow conversion offered by the GMM based systems, there have been an increasing number of attempts at applying neural networks to the problem [10, 11, 12, 13]. Many methods treat VC as a purely supervised learning problem, in which some hidden features \mathbf{z}_t are inferred from the source input by the initial layers of the neural network, and then transformed by the final layers into the target output. Later methods [11, 13] build on this by exploiting the flexibility of recurrent and convolutional neural networks to model the temporal aspects of the input, resulting in better conversion and less sensitivity to alignment errors in the training data. A further improvement made in [13] replaces the usual cost function with an adversarially learned similarity metric that does away with the buzzing introduced by training with MSE.

While these methods have been successful when trained on datasets consisting of a hundred or more parallel utterances, performance degrades when less data is available. One way to avoid this problem is to consider a combined approach to increase the efficiency that the parallel data is used by augmenting the training with additional unpaired utterances. This combined approach has been relatively unexplored in the neural network VC literature. One of the few works [19] that addresses this possibility carries out unsupervised pre-training of a deep autoencoder with data from multiple speakers, followed by a fine tuning step using parallel data. This is reminiscent of the unsupervised GMM training followed by the supervised learning of the conversion parameters discussed above.

Further work on this autoencoding approach [20] showed that pre-training separate autoencoders for the source and target (a process that can be done without parallel data), and then fine tuning the source encoder and target decoder using parallel data improves the source to target voice conversion. Conversion using this method can be carried out by obtaining a latent \mathbf{z}_t from the source encoder and then decoding it into the target features by the target decoder. While this work clearly shows the benefits of semi-supervised training, to our knowledge the possibility of simultaneous training with a single objective remains unexplored.

3. Proposed semi-supervised method

In developing our semi-supervised approach, we wish to retain several desirable features from both the GMM VC systems and

the parallel data neural network systems discussed in Sec. 2. In the GMM systems, we note that the model parameters involved in computing the latent variables \mathbf{z}_t do not require parallel training data to be learned, which decreases the amount of information that must be obtained via supervision. Additionally, we note that there is an efficient (in this case exact) inference procedure for obtaining \mathbf{z}_t given the input features that is vital for computationally tractable training and inference.

The neural network systems also have an efficient procedure for computing \mathbf{z}_t , however it is qualitatively different from the probabilistic procedure in the GMMs owing to the highly nonlinear nature of the networks. By virtue of this more flexible inference procedure, the neural network models are better able to handle the sequential nature of the input features, and model more complex relationships between the latent variables and the source/target features.

Therefore, in constructing the semi-supervised approach, we seek a method that 1) can learn some model parameters at least partially from non-parallel data as in the GMM VC methods, 2) has an efficient and well defined probabilistic procedure for obtaining latent variables \mathbf{z}_t as in the GMM methods, 3) can model complex relationships between the latent variables and source/target features as can the prior neural network models, and 4) can flexibly model the long sequential nature of the input as in the neural network approaches.

To do this we assume a latent variable sequence $\mathbf{Z} = \{z_1, z_2, \dots, z_T\}$ generates both the source and target sequences, $\mathbf{X} = \{x_1, x_2, \dots, x_T\}$ and $\mathbf{Y} = \{y_1, y_2, \dots, y_T\}$ respectively. We do not treat each frame $t = \{1, 2, \dots, T\}$ independently, so each x_t, y_t depends on the entire sequence of latent variables. We model the conditional distributions with factorized Gaussians:

$$\begin{aligned} p(\mathbf{X} | \mathbf{Z}) &= \prod_{t=1}^T \mathcal{N}(x_t | f_{\theta_x}^{(x)}(\mathbf{Z}, t), \sigma^2 \cdot \mathbb{I}) \\ p(\mathbf{Y} | \mathbf{Z}) &= \prod_{t=1}^T \mathcal{N}(y_t | f_{\theta_y}^{(y)}(\mathbf{Z}, t), \sigma^2 \cdot \mathbb{I}) \\ p(\mathbf{Z}) &= \prod_{t=1}^T \mathcal{N}(z_t | 0, \mathbb{I}) \end{aligned} \quad (2)$$

Here, $\mathcal{N}(\cdot | f_{\theta}(\cdot), \sigma^2 \cdot \mathbb{I})$ is a multivariate normal distribution with mean given by the function $f_{\theta}(\cdot)$, which depends on parameters θ , with a diagonal covariance matrix with nonzero elements set to be σ^2 for simplicity. $f_{\theta_x}^{(x)}(\mathbf{Z}, t)$ and $f_{\theta_y}^{(y)}(\mathbf{Z}, t)$ are separate functions for the source and target speakers respectively that capture the dependence of the source and target features on the sequence of latent variables.

Exact inference is prohibitively expensive in this model, owing to both the nonlinearity of $f_{\theta_x}^{(x)}(\mathbf{Z}, t)$ and $f_{\theta_y}^{(y)}(\mathbf{Z}, t)$, and the large number of parent nodes for each x_t, y_t . However, approximate inference can be carried out in the variational autoencoder framework [15], which has shown success in semi-supervised classification problems [16]. This requires the use of an approximate inference model such that the problem of finding $p(\mathbf{Z} | \mathbf{X})$ or $p(\mathbf{Z} | \mathbf{Y})$ is replaced with the approximation $q(\mathbf{Z} | \cdot) \approx p(\mathbf{Z} | \cdot)$, where $q(\mathbf{Z} | \cdot)$ is defined as

$$q(\mathbf{Z} | \cdot) = \prod_{t=1}^T \mathcal{N}(z_t | g_{\phi_x}^{(\mu)}(\cdot, t), g_{\phi_y}^{(\sigma^2)}(\cdot, t) \cdot \mathbb{I}) \quad (3)$$

Here, $g_{\phi_\mu}^{(\mu)}(\cdot, t)$ is a function parameterized by ϕ_μ that represents the mean of the multivariate normal, $g_{\phi_\sigma}^{(\sigma^2)}(\cdot, t)$ represents the diagonal elements of the covariance matrix. Note that the same function is used for both \mathbf{X} and \mathbf{Y} . This choice was made with the intent that \mathbf{Z} should be shared for both speakers. The functions $f_{\theta_x}^{(x)}(\mathbf{Z}, t)$, $f_{\theta_y}^{(y)}(\mathbf{Z}, t)$, $g_{\phi_\mu}^{(\mu)}(\cdot, t)$ and $g_{\phi_\sigma}^{(\sigma^2)}(\cdot, t)$ complete the specification of the model, and may be approximated with neural networks.

3.1. Training objective

Ideally, the parameters $\theta_x, \theta_y, \phi_\mu, \phi_\sigma$ would be learned via maximum likelihood. However, in this model exact likelihood calculations are prohibitive, and so we maximize a lower bound on the log-likelihood [15] $\log(p(\mathbf{X}, \mathbf{Y})) \geq \mathcal{L}_{X,Y}$ where

$$\begin{aligned} \mathcal{L}_{X,Y} = & \frac{1}{2\sigma^2} \sum_{t=1}^T \mathbb{E}_{z_t \sim q(\mathbf{Z}|\cdot)} \left(-\|x_t - f_{\theta_x}^{(x)}(\mathbf{Z}, t)\|^2 \right) + \\ & \frac{1}{2\sigma^2} \sum_{t=1}^T \mathbb{E}_{z_t \sim q(\mathbf{Z}|\cdot)} \left(-\|y_t - f_{\theta_y}^{(y)}(\mathbf{Z}, t)\|^2 \right) + \\ & - \mathcal{D}_{KL}(q(\mathbf{Z}|\cdot) \| p(\mathbf{Z})) \end{aligned} \quad (4)$$

For semi-supervised training, we must consider the case where both \mathbf{X}, \mathbf{Y} are known, the case where \mathbf{X} is known but \mathbf{Y} is not, and the case where \mathbf{Y} is known but \mathbf{X} is not. In the first case, the bound on the log-likelihood is given by Eq. 4. In the case where only \mathbf{X} is known, the expectation in Eq. 4 involving \mathbf{Y} is constant owing to the form of $p(\mathbf{Y}|\mathbf{Z})$, and so we want to maximize $\log(p(\mathbf{X})) \geq \mathcal{L}_{X,\cdot}$ where (after dropping the constant terms)

$$\begin{aligned} \mathcal{L}_{X,\cdot} = & \frac{1}{2\sigma^2} \sum_{t=1}^T \mathbb{E}_{z_t \sim q(\mathbf{Z}|\mathbf{X})} \left(-\|x_t - f_{\theta_x}^{(x)}(\mathbf{Z}, t)\|^2 \right) + \\ & - \mathcal{D}_{KL}(q(\mathbf{Z}|\mathbf{X}) \| p(\mathbf{Z})) \end{aligned} \quad (5)$$

Equivalently, when \mathbf{Y} is known but \mathbf{X} is not, we maximize $\log(p(\mathbf{Y})) \geq \mathcal{L}_{\cdot,Y}$, which has the same form as $\mathcal{L}_{X,\cdot}$, but has \mathbf{Y} in place of \mathbf{X} .

The bound on the entire dataset is therefore

$$\mathcal{L} = \sum_{\mathbf{X} \in \{\mathbf{X}, \cdot\}} \mathcal{L}_{X,\cdot} + \sum_{\mathbf{Y} \in \{\cdot, \mathbf{Y}\}} \mathcal{L}_{\cdot,Y} + \sum_{\mathbf{X}, \mathbf{Y} \in \{\mathbf{X}, \mathbf{Y}\}} \mathcal{L}_{X,Y} \quad (6)$$

In practice, we compute the expectations in $\mathcal{L}_{X,\cdot}, \mathcal{L}_{\cdot,Y}$ using a single sample. For the expectations in $\mathcal{L}_{X,Y}$ we would ideally use an approximation of $p(\mathbf{Z}|\mathbf{X}, \mathbf{Y})$. However, in the case where parallel training data is limited, directly learning such an approximation is prohibitively complex. Instead, we use two samples to compute this quantity, one using $q(\mathbf{Z}|\mathbf{X})$ and another using $q(\mathbf{Z}|\mathbf{Y})$.

3.2. Baseline systems and proposed modifications

As a baseline approach, we implemented the DBLSTM model from [11]. This model consists of four bidirectional long short-term memory (BLSTM) layers of sizes 128, 256, 256, 128 each. With the exception of using the WORLD vocoder instead of the STRAIGHT vocoder [21], we use this model as described in the original work.

We then extend the DBLSTM model so that it can be trained with semi-supervision. To do this, we interpret the

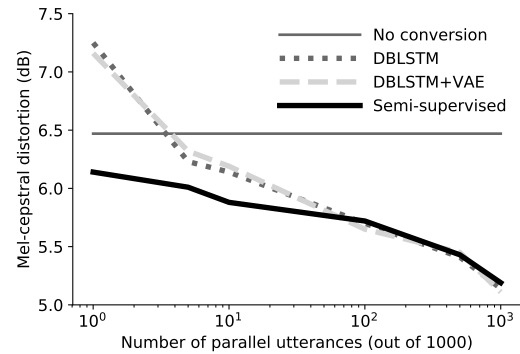


Figure 1: Performance of semi-supervised training vs. fully supervised training for a varying number of parallel training utterances out of a total of 1000 utterances.

first two BLSTM layers of the model as the encoder portion and apply two separate affine transformations to the 256 dimensional output for each time frame to compute $g_{\phi_\mu}^{(\mu)}(\cdot, t)$ and $\log(g_{\phi_\sigma}^{(\sigma^2)}(\cdot, t))$ respectively. In this way, we obtain the distribution $q(\mathbf{Z}|\cdot)$ describing a 256 dimensional latent variable \mathbf{Z} . This \mathbf{Z} acts as the inputs to the final two layers of the DBLSTM model which we interpret as the decoder portion. We have two separate decoder portions (each using $\sigma^2 = 10^{-3}$) with different parameters but an otherwise identical architecture such that one acts as $f_{\theta_x}^{(x)}(\mathbf{Z}, t)$ and the other as $f_{\theta_y}^{(y)}(\mathbf{Z}, t)$. The resulting model can be trained with the semi-supervised cost function described in Eq. 6.

To verify that this architectural modification by itself does not significantly impact performance in the absence of semi-supervision, we also carried out experiments on the modified model with purely supervised training. This corresponds to only optimizing the $\mathcal{L}_{X,Y}$ term of Eq. 6. For approximate inference in this model, we use \mathbf{X} . We call this model the DBLSTM+VAE, to distinguish it from the baseline DBLSTM and denote that it has been reinterpreted in the variational autoencoder (VAE) framework.

4. Experiments

To evaluate the performance of our semi-supervised method, we carried out experiments with varying amounts of parallel and non-parallel training data. To create the different datasets, we drew samples from the CLB→SLT (both females) pair in the CMU Arctic corpus [22]. To generate the parallel data corpus, we drew paired \mathbf{X}, \mathbf{Y} samples at random from both the A and B partitions of the CMU Arctic, and time aligned the target features to the source features using dynamic time warping (DTW). For the non-parallel \mathbf{X}, \cdot and \cdot, \mathbf{Y} datasets, we drew samples at random from the remaining unselected samples in the A and B partitions respectively, ensuring that the same prompt does not appear in both the \mathbf{X}, \cdot and \cdot, \mathbf{Y} datasets. We considered datasets consisting of at most 1000 utterances, and so the remaining 93 utterances from the A partition were used as a validation dataset, while testing was carried out on the remaining 39 utterances of the B partition.

In line with [11], we extracted 50 Mel-cepstral coefficients (MCEPs) from the spectral envelope obtained using the WORLD vocoder [23] as features along with fundamental fre-

Table 1: Mean Opinion Score (MOS) of models trained on 1, 10, and 1000 parallel utterances, out of a total of 1000 utterances. Error bars represent 95% confidence intervals.

	1	10	1000
DBLSTM[11]	1.64±0.13	2.49±0.17	3.39±0.16
DBLSTM + VAE	1.73±0.14	2.56±0.14	3.40±0.14
Semi-Supervised	2.93±0.16	2.99±0.16	3.63±0.16

quency contours (F_0) and aperiodicities (APs). We used 16kHz audio, an FFT size of 1024 with a hop length of 5ms. The zeroth cepstral coefficient was left unmodified, and the remaining 49 coefficients were Gaussian normalized and used as the features for VC. The F_0 contours were converted via the usual log Gaussian normalization [24], and the APs were used directly from the input without modification.

To obtain an objective performance measure, we evaluated each model using mel-cepstral distortion [25, 26] (MCD). We measured the average MCD between paired source and target utterances before conversion to be 6.47 dB. While the proposed method treats the source and target symmetrically, we only carried out the evaluation for source to target conversion to directly compare to the supervised approaches which treat the problem asymmetrically.

4.1. Increasing amounts of parallel data

To verify the effectiveness of the semi-supervised training under realistic constraints on the amount of parallel training data, we considered a fixed training data budget of $N = 1000$ total utterances (roughly 1 hour from each speaker), and varied the number of parallel training utterances. The non-parallel utterances were evenly split between the source and target.

Results for varying amounts of parallel training data are shown in Fig. 1. When only a small fraction of the training data consists of parallel utterances, we find that training with semi-supervision gives far better performance than purely supervised training. We also see that for datasets that contain mostly parallel data, the proposed semi-supervised method gives equivalent performance to the fully supervised approach as expected. For intermediate amounts of parallel and non-parallel data, the semi-supervised training smoothly interpolates between these two cases, consistently performing better than or equal to the purely supervised training.

To verify that the audio quality remains high as the amount of parallel training data is varied, we also carried out a subjective evaluation of the quality of the converted audio from algorithms trained on 1, 10, and 1000 parallel utterances. We evaluated Mean Opinion Score (MOS) using 40 listeners on Amazon Mechanical Turk, over 90% of which were native English speakers. Each listener was asked to rate the quality of 3 utterances from each model on a 5-point scale (1=Bad, 2=Poor, 3=Fair, 4=Good, 5=Excellent). Results of this evaluation are shown in Table 1. We see that the semi-supervised training gives much higher quality audio when only a small amount of parallel data is available. Measuring both conversion (via MCD) and quality (with MOS), we find that when only a small amount of parallel data is available the semi-supervised approach achieves voice conversion of quality comparable to the supervised approaches trained with a much larger parallel dataset.

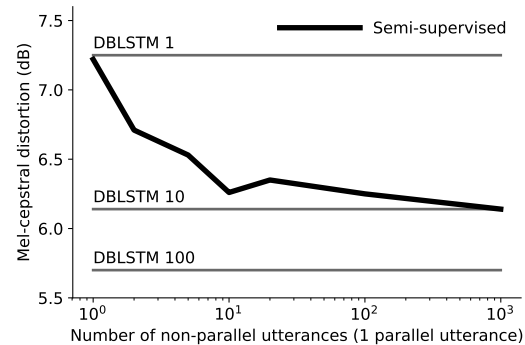


Figure 2: Performance of semi-supervised method with a single parallel utterance and an increasing number of non-parallel utterances. Horizontal lines show performance of supervised approach with a varying number of parallel utterances.

4.2. Increasing amounts of non-parallel data

As the results of Fig. 1 and Table 1 show, the inclusion of non-parallel data in the semi-supervised training leads to improved performance, and performance of all methods continues to improve for larger amounts of parallel training data. However, in the semi-supervised case, it may also be possible to improve performance by increasing the amount of non-parallel data used in training.

To test this effect we considered a fixed data budget of only one parallel utterance, and varied the amount of non-parallel utterances used in the semi-supervised training. Results of this experiment are shown in Fig. 2. We find that increasing the amount of non-parallel data used in training improves the VC performance as was the case for increasing amounts of parallel data. However the rate of improvement decreases for larger amounts of non-parallel data, while we did not observe this with larger amounts of parallel data.

This suggests that in creating a dataset for a semi-supervised VC system, a trade off must be made between gathering harder to obtain but more informative parallel training examples, and easier to obtain but less informative non-parallel examples. While we do not investigate this here, it is an interesting and promising avenue of future work to devise methods of improving the efficiency in which non-parallel training data is used as this will determine the overall cost and difficulty of creating a VC training dataset.

5. Conclusion

We have proposed a new semi-supervised method for achieving voice conversion using both parallel and non-parallel data. This method incorporates both types of data simultaneously during training by optimizing a variational objective defined for paired and unpaired utterances. When only a small number of parallel utterances are available, we show that incorporating this method into an existing neural network model improves the accuracy and perceptual quality of the converted speech compared to supervised training. We also find that increasing the amount of non-parallel data continues to improve voice conversion. This opens up the possibility of training VC systems with more flexible datasets consisting of mixed parallel and non-parallel data.

6. References

- [1] T. Toda, L.-H. Chen, D. Saito, F. Villavicencio, M. Wester, Z. Wu, and J. Yamagishi, "The voice conversion challenge 2016," in *Interspeech*, 2016, pp. 1632–1636.
- [2] J. Lorenzo-Trueba, J. Yamagishi, T. Toda, D. Saito, F. Villavicencio, T. Kinnunen, and Z. Ling, "The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods," in *Proc. Odyssey 2018 The Speaker and Language Recognition Workshop*, 2018, pp. 195–202.
- [3] M. Abe, S. Nakamura, K. Shikano, and H. Kuwabara, "Voice conversion through vector quantization," *Journal of the Acoustical Society of Japan (E)*, vol. 11, no. 2, pp. 71–76, 1990.
- [4] Y. Stylianou, O. Cappé, and E. Moulines, "Continuous probabilistic transform for voice conversion," *IEEE Transactions on speech and audio processing*, vol. 6, no. 2, pp. 131–142, 1998.
- [5] T. Toda, A. W. Black, and K. Tokuda, "Voice conversion based on maximum-likelihood estimation of spectral parameter trajectory," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 8, pp. 2222–2235, 2007.
- [6] D. Erro, A. Moreno, and A. Bonafonte, "Inca algorithm for training voice conversion systems from nonparallel corpora," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 18, no. 5, pp. 944–953, 2010.
- [7] A. Mouchtaris, J. Van der Spiegel, and P. Mueller, "Non-parallel training for voice conversion by maximum likelihood constrained adaptation," in *Acoustics, Speech, and Signal Processing (ICASSP), 2004. IEEE International Conference on*, vol. 1. IEEE, 2004, pp. 1–1.
- [8] T. Kinnunen, L. Juvela, P. Alku, and J. Yamagishi, "Non-parallel voice conversion using i-vector plda: Towards unifying speaker verification and transformation," in *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 5535–5539.
- [9] Y. Saito, Y. Ijima, K. Nishida, and S. Takamichi, "Non-parallel voice conversion using variational autoencoders conditioned by phonetic posteriorgrams and d-vectors," in *Acoustics, Speech and Signal Processing (ICASSP), 2018 IEEE International Conference on*, 2018, pp. 5274–5278.
- [10] S. Desai, E. V. Raghavendra, B. Yegnanarayana, A. W. Black, and K. Prahallad, "Voice conversion using artificial neural networks," in *Acoustics, Speech and Signal Processing (ICASSP), 2009 IEEE International Conference on*. IEEE, 2009, pp. 3893–3896.
- [11] L. Sun, S. Kang, K. Li, and H. Meng, "Voice conversion using deep bidirectional long short-term memory based recurrent neural networks," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4869–4873.
- [12] T. Nakashika, R. Takashima, T. Takiguchi, and Y. Ariki, "Voice conversion in high-order eigen space using deep belief nets," in *Interspeech*, 2013, pp. 369–372.
- [13] T. Kaneko, H. Kameoka, K. Hiramatsu, and K. Kashino, "Sequence-to-sequence voice conversion with similarity metric learned using generative adversarial networks," in *Interspeech*, 2017, pp. 1283–1287. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2017-970>
- [14] C.-C. Hsu, H.-T. Hwang, Y.-C. Wu, Y. Tsao, and H.-M. Wang, "Voice conversion from non-parallel corpora using variational auto-encoder," in *Signal and Information Processing Association Annual Summit and Conference (APSIPA), 2016 Asia-Pacific*. IEEE, 2016, pp. 1–6.
- [15] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *CoRR*, vol. abs/1312.6114, 2013. [Online]. Available: <http://arxiv.org/abs/1312.6114>
- [16] D. P. Kingma, S. Mohamed, D. J. Rezende, and M. Welling, "Semi-supervised learning with deep generative models," in *Advances in Neural Information Processing Systems*, 2014, pp. 3581–3589.
- [17] S. H. Mohammadi and A. Kain, "An overview of voice conversion systems," *Speech Communication*, vol. 88, pp. 65–82, 2017.
- [18] S. Takamichi, T. Toda, A. W. Black, and S. Nakamura, "Modulation spectrum-constrained trajectory training algorithm for gmm-based voice conversion," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 4859–4863.
- [19] S. H. Mohammadi and A. Kain, "Voice conversion using deep neural networks with speaker-independent pre-training," in *Spoken Language Technology Workshop (SLT), 2014 IEEE*. IEEE, 2014, pp. 19–23.
- [20] S. H. Mohammadi and A. Kain, "Semi-supervised training of a voice conversion mapping function using a joint-autoencoder," in *Sixteenth Annual Conference of the International Speech Communication Association*, 2015.
- [21] H. Kawahara, I. Masuda-Katsuse, and A. De Cheveigne, "Re-structuring speech representations using a pitch-adaptive time-frequency smoothing and an instantaneous-frequency-based f0 extraction: Possible role of a repetitive structure in sounds1," *Speech communication*, vol. 27, no. 3-4, pp. 187–207, 1999.
- [22] J. Kominek and A. W. Black, "CMU ARCTIC databases for speech synthesis," 2003.
- [23] M. Morise, F. Yokomori, and K. Ozawa, "WORLD: a vocoder-based high-quality speech synthesis system for real-time applications," *IEICE TRANSACTIONS on Information and Systems*, vol. 99, no. 7, pp. 1877–1884, 2016.
- [24] K. Liu, J. Zhang, and Y. Yan, "High quality voice conversion through phoneme-based linear mapping functions with straight for mandarin," in *Fuzzy Systems and Knowledge Discovery, 2007. FSKD 2007. Fourth International Conference on*, vol. 4. IEEE, 2007, pp. 410–414.
- [25] R. Kubichek, "Mel-cepstral distance measure for objective speech quality assessment," in *Communications, Computers and Signal Processing, 1993., IEEE Pacific Rim Conference on*, vol. 1. IEEE, 1993, pp. 125–128.
- [26] J. Kominek, T. Schultz, and A. W. Black, "Synthesizer voice quality of new languages calibrated with mean mel cepstral distortion," in *Spoken Languages Technologies for Under-Resourced Languages*, 2008.