



Non-parallel Emotion Conversion using a Deep-Generative Hybrid Network and an Adversarial Pair Discriminator

Ravi Shankar, Jacob Sager, Archana Venkataraman

Department of Electrical and Computer Engineering, Johns Hopkins University, USA

rshanka3@jhu.edu, jsager2@alumni.jhu.edu, archana.venkataraman@jhu.edu

Abstract

We introduce a novel method for emotion conversion in speech that does not require parallel training data. Our approach loosely relies on a cycle-GAN schema to minimize the reconstruction error from converting back and forth between emotion pairs. However, unlike the conventional cycle-GAN, our discriminator classifies whether a pair of input real and generated samples corresponds to the desired emotion conversion (e.g., $A \rightarrow B$) or to its inverse ($B \rightarrow A$). We will show that this setup, which we refer to as a variational cycle-GAN (VCGAN), is equivalent to minimizing the empirical KL divergence between the source features and their cyclic counterpart. In addition, our generator combines a trainable deep network with a fixed generative block to implement a smooth and invertible transformation on the input features, in our case, the fundamental frequency (F0) contour. This hybrid architecture regularizes our adversarial training procedure. We use crowd sourcing to evaluate both the emotional saliency and the quality of synthesized speech. Finally, we show that our model generalizes to new speakers by modifying speech produced by Wavenet.

Index Terms: Adversarial Networks, Unsupervised Learning, Emotion Conversion, Deformable Registration

1. Introduction

From automated customer support to hand-held devices, speech synthesis plays an important role in modern-day technology. While speech synthesis has undergone revolutionary advancements over the past few years, generating emotional cues remains an open challenge in the field. Emotional speech synthesis has the potential to facilitate more natural and meaningful human-computer interactions, and it provides a foundation for studying human intent, perception, and behavior [1].

The success of deep neural networks has brought about a swift change in how speech synthesis is approached. Deep neural networks can generate natural sounding speech given enough training examples [2, 3, 4]. However, these models have little control over the speaking style, including emotional inflection. One reason is the lack of training data to learn networks specific to each emotional class. Unsupervised models such as [5, 6] provide a promising middle ground by separating the speaking style from the content. However, it is difficult to tune the parameters of these models to synthesise speech in a predetermined emotion. Furthermore, the synthesis rate of these state-of-the-art models is slow due to their autoregressive scheme [7]. These limitations motivate the use of emotion conversion as an alternative to end-to-end synthesis. Broadly, the goal of emotion conversion is to modify the perceived affect of a speech utterance without changing its linguistic content or speaker identity. It allows a user greater control over the speaking style while being easy to train on limited data resources.

Emotional cues in speech are conveyed through vocal inflections known as prosody. Key attributes of prosody include the fundamental frequency (F0) contour, the relative energy of the signal, and the spectrum [8]. Many supervised and unsupervised algorithms have been proposed for emotion conversion. For example, the work of [9] proposed a Gaussian mixture model (GMM) to jointly model the source and target prosodic features. During inference, the target features are estimated from the source via a maximum likelihood optimization. A recent approach by [10] uses a Bidirectional LSTM (Bi-LSTM) to predict the spectrum and F0 contour. To overcome the data limitation, the authors pre-train their model on a voice conversion dataset and then fine-tune it for emotion conversion. The prosodic manipulation proposed by [11, 12] uses a highway neural network to predict the F0 and intensity for each frame of the input utterance. While these models have made significant contributions to the field, they require parallel emotional speech data for training, which limits their generalizability.

An unsupervised technique to disentangle style and content from speech has been proposed by [13]. This algorithm uses architecture based priors to separate style and content from spectrum while modifying the F0 using a linear Gaussian model. The authors of [14] offer a simpler cycle-GAN model for non-parallel emotion conversion, which independently modifies the spectrum and F0 contour. The latter is parameterized via a wavelet transform, which expands the input feature dimensionality. These approaches, however, are trained and evaluated on single speakers, with no validation on multispeaker conversion.

In this paper, we propose a novel method for non-parallel emotion conversion that blends the cycle-GAN architecture with implicit regularization from a generative curve registration method. Our novel loss function for the F0 conversion leads to an adversarial training where the discriminator classifies whether a pair of real and generated F0 contours represents a valid conversion. Another contribution of our model is that the generator combines a trainable deep neural network with a generative component to implement a smooth and invertible warping of the source F0 contour. The entire model is trained jointly by back-propagating through the generative block to optimize the parameters. We evaluate our model on the multi-speaker VESUS dataset [15] and use crowd sourcing to verify both the emotional saliency and speech quality of the converted utterances. We also demonstrate the generalizability of our model by converting speech produced by Google Wavenet [2].

2. Method

The foundation of our model is a cycle-GAN [16, 17], which optimizes the cycle consistency of converting back and forth between emotions. However, we adapt the traditional cycle-GAN framework to align the *distribution* of transformed source features to the distribution of target features. Accordingly, we refer to our model as a variational cycle-GAN (VCGAN).

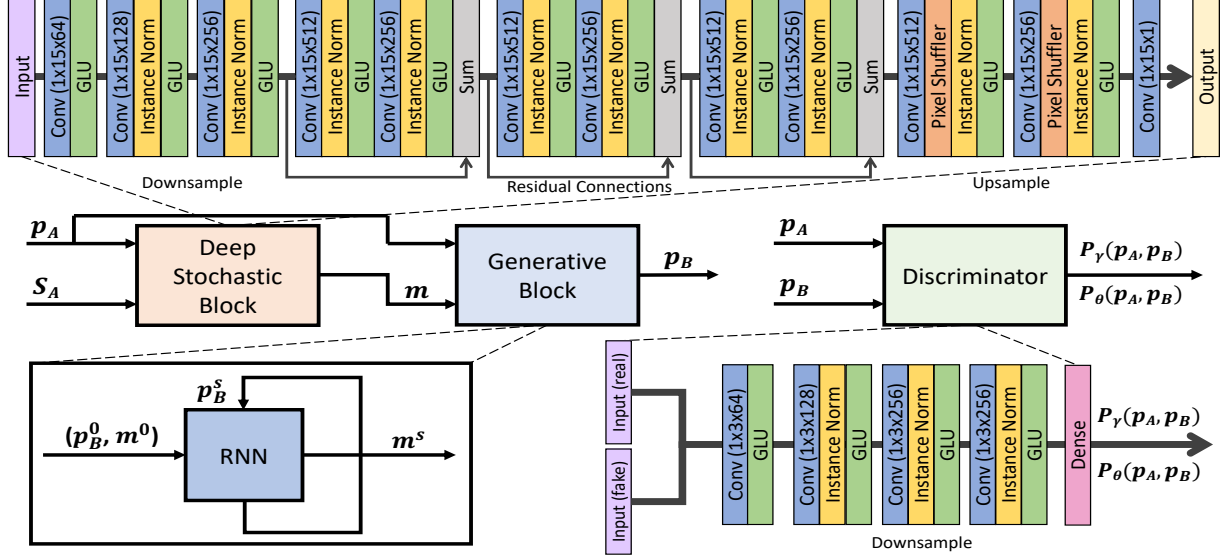


Figure 1: Block diagram and neural network architectures of VCGAN. The top row shows the neural network used as the stochastic component of the generators. The bottom left figure shows the deterministic component of generators as represented by an RNN. The bottom right figure illustrates the architecture of neural network used as discriminator for classifying the joint densities.

2.1. Generator Loss

We train the generators in our VCGAN using three different loss terms. The first is a cycle consistency loss, which ensures invertibility of the generator transforms. Formally, let A and B denote the source and target emotion classes, respectively. In the context of F0 conversion, the cycle consistency loss ensures that a given F0 contour, denoted by \mathbf{p}_A , is close to itself after undergoing the sequence of transformations $A \rightarrow B \rightarrow A$. Mathematically, this loss is expressed as $\mathcal{L}_C = E[\|\mathbf{p}_A - \mathbf{p}_A^c\|_1]$.

Ultimately, the cycle-consistency loss \mathcal{L}_C , is a sample-specific loss and only provides a weak coupling between the two generators after each cyclic transformation. This problem is exaggerated by the discriminator, which is agnostic to the presence of a second generator in the cycle-GAN. The inference process can be improved by recognizing that the input distribution of one generator is the learnable output distribution of the other. To leverage this information, we add a KL divergence penalty on the input distribution of a generator and the target distribution of its complementary generator. It allows us to exploit the cyclic property at a global level. Specifically, let \mathbf{p}_A and \mathbf{p}_B denote the source and target F0 contours of utterances in emotion A and B, respectively. Let $G_\gamma : \mathbf{p}_A \rightarrow \mathbf{p}_B$ and $G_\theta : \mathbf{p}_B \rightarrow \mathbf{p}_A$ denote the two generators that transform the F0 contours between emotional classes. The corresponding learned data distributions are given by $P_\gamma(\mathbf{p}_B)$ and $P_\theta(\mathbf{p}_A)$. The KL divergence loss for generator G_γ can be expressed as:

$$\mathcal{L}_{KL} = KL(P(\mathbf{p}_A) \| P_\theta(\mathbf{p}_A)) \quad (1)$$

Optimizing this loss provides an additional coupling between the forward and backward transformations, one that entangles the two generators beyond cyclic consistency loss. Next, we show that this loss eliminates the need for the discriminator to classify samples from the marginal distributions i.e. real vs generated. By total probability law, we can write $P_\theta(\mathbf{p}_A)$ as:

$$P_\theta(\mathbf{p}_A) = \int P_\theta(\mathbf{p}_A | \mathbf{p}_B) P(\mathbf{p}_B) d\mathbf{p}_B \quad (2)$$

The integral form in Eq. (2) however, is intractable. To circumvent this, we use the variational trick and derive an upper bound

on the KL penalty. Combining Eq. (1) and Eq. (2), we get:

$$\begin{aligned} KL &= \int P(\mathbf{p}_A) \log \frac{P(\mathbf{p}_A)}{\int P_\theta(\mathbf{p}_A | \mathbf{p}_B) P(\mathbf{p}_B) d\mathbf{p}_B} d\mathbf{p}_A \\ &= - \int P(\mathbf{p}_A) \log \int \frac{P_\theta(\mathbf{p}_A | \mathbf{p}_B) P_\gamma(\mathbf{p}_B | \mathbf{p}_A) P(\mathbf{p}_B)}{P_\gamma(\mathbf{p}_B | \mathbf{p}_A)} d\mathbf{p}_B d\mathbf{p}_A \\ &\quad - H(\mathbf{p}_A) \\ &\leq - \int P(\mathbf{p}_A) \int P_\gamma(\mathbf{p}_B | \mathbf{p}_A) \log \frac{P_\theta(\mathbf{p}_A | \mathbf{p}_B) P(\mathbf{p}_B)}{P_\gamma(\mathbf{p}_B | \mathbf{p}_A)} d\mathbf{p}_B d\mathbf{p}_A \\ &\quad - H(\mathbf{p}_A) \\ &= \int P(\mathbf{p}_A) \int P_\gamma(\mathbf{p}_B | \mathbf{p}_A) \log \frac{P_\gamma(\mathbf{p}_B | \mathbf{p}_A) P(\mathbf{p}_A)}{P_\theta(\mathbf{p}_A | \mathbf{p}_B) P(\mathbf{p}_B)} d\mathbf{p}_B d\mathbf{p}_A \quad (3) \end{aligned}$$

where we have used Jensen's inequality between the second and third steps. The ratio of probabilities in Eq. (3) compares the joint densities of $(\mathbf{p}_A, \mathbf{p}_B)$. We estimate this ratio by a discriminator network denoted by D_γ . This discriminator acts as a global connector between the generators G_γ and G_θ by classifying the joint densities. Notice that the KL term in Eq. (1) appears only as a function of parameter θ but the variational trick allows us to introduce the parameters of G_γ into the picture. When training G_γ , optimizing the upper bound on KL divergence is equivalent to minimizing adversarial loss:

$$\mathcal{L}_{KL} \leq E_{\mathbf{p}_A \sim P(\mathbf{p}_A)} [E_{\mathbf{p}_B \sim P_\gamma} [\log(D_\gamma(\mathbf{p}_A, \mathbf{p}_B))]] \quad (4)$$

So far, we have derived the generator loss based on the F0 contour. In practice, we condition the generators on both spectrum and F0 contour i.e. $G_\gamma : \mathbf{S}_A \times \mathbf{p}_A \rightarrow \mathbf{p}_B$. Here, \mathbf{S}_A denotes the source emotion spectrum. Conditioning on spectrum allows VCGAN to learn the joint relationship between prosodic features. We can show that Eq. (3) is also an upper bound to:

$$E_{\mathbf{S}_A} [KL(P(\mathbf{p}_A | \mathbf{S}_A) \| P_\theta(\mathbf{p}_A | \mathbf{S}_A))] \quad (5)$$

Thus, we minimize the distance between conditional densities over F0 contours as opposed to the marginal densities in Eq. (1). The expectation in Eq. (5) averages over the spectral variations.

Algorithm 1: Warping to generate the target F0 contour given the momenta and source F0 contour

```

1 function GenerateF0( $\mathbf{m}$ ,  $\mathbf{p}_A$ );
   Input : momenta  $\mathbf{m}$  and source F0  $\mathbf{p}_A$ 
   Output: target F0  $\mathbf{p}_B$ 
2 Set  $s = 0$ ,  $[\mathbf{p}_B]^0 = \mathbf{p}_A$  and  $[\mathbf{m}]^0 = \mathbf{m}$ ;
3 if  $s < 3$  then
4    $d_{i,j} \leftarrow [\mathbf{p}_A]_i^s - [\mathbf{p}_A]_j^s$ ;
5    $K_{i,j} \leftarrow \exp\left(-\frac{(d_{i,j})^2}{\sigma^2}\right)$ ;
6    $[\mathbf{p}_B]_i^{s+1} \leftarrow [\mathbf{p}_B]_i^s + \sum_l K_{i,l} \cdot [\mathbf{m}]_l^s$ ;
7    $[\mathbf{m}]_i^{s+1} \leftarrow [\mathbf{m}]_i^s + 2 \cdot \sum_j \frac{-K}{\sigma^2} d_{i,j} \cdot [\mathbf{m}]_j^s$ ;
8    $s \leftarrow s + 1$ ;
9 else
10  return  $[\mathbf{p}_B]^s$ ;
11 end

```

2.2. Generative Hybrid Architecture

Adversarial training is susceptible to mode collapse, imbalance between generator-discriminator losses, and the architecture of neural networks. Keeping this in mind, we model the generated target F0 contour as a smooth and invertible warping of the source F0. Such warpings are also known as diffeomorphisms [18, 19] and can be parameterized by low dimensional embeddings called the momenta [20]. Therefore, our target F0 estimation is a two-step process: first, we estimate the momenta, and second, we modify the source F0 contour via a deterministic warping using momenta. As a result, our generators can be divided into two blocks, a stochastic component with trainable parameters and a deterministic component with static parameters. Specifically, let \mathbf{m} denote the latent momenta. The target F0 is generated by following Alg. 1. The dimensions of the momenta are the same as F0 contour. The kernel smoothness parameter, σ is empirically fixed at 50 to span the F0’s range. The warping function can be represented as a recurrent neural network (RNN) because of its iterative nature (Fig. 1). The advantage of this hybrid architecture is to stabilize the target F0 generation. In the absence of any such control mechanism, the F0 contours swing wildly, and causes mode collapse.

We constrain the generators to sample smoothly varying momenta by adding $\mathcal{L}_m = E[\|\nabla \mathbf{m}\|^2]$ to the loss. We approximate the gradient of momenta by its first-order difference. The final objective for generator G_γ is given by:

$$\begin{aligned} \mathcal{L}_{G_\gamma} = & \lambda_c E[\|\mathbf{p}_A - \mathbf{p}_A^s\|] + \lambda_m E[\|\nabla \mathbf{m}\|^2] \\ & + (1 - \lambda_c - \lambda_m) E_{(\mathbf{S}_A, \mathbf{P}_A)} \left[E_{\mathbf{P}_B \sim P_\gamma} \left[\log(D_\gamma(\mathbf{p}_A, \mathbf{p}_B)) \right] \right] \end{aligned} \quad (6)$$

To update the parameters of the stochastic part of generator network, the gradient back-propagates through the deterministic block, which is implemented as matrix-vector operations.

2.3. Discriminator Loss

Similar to [21], we model the ratio term in Eq. (3) by a discriminator denoted by D_γ that distinguishes between the joint distributions of \mathbf{p}_A and \mathbf{p}_B learned by G_γ and G_θ , respectively. During training of the discriminator D_γ , we minimize:

$$\begin{aligned} \mathcal{L}_{D_\gamma} = & -E_{(\mathbf{S}_A, \mathbf{P}_A)} \left[E_{\mathbf{P}_B \sim P_\gamma} \left[\log(D_\gamma(\mathbf{p}_A, \mathbf{p}_B)) \right] \right] \\ & - E_{(\mathbf{S}_B, \mathbf{P}_B)} \left[E_{\mathbf{P}_A \sim P_\theta} \left[\log(1 - D_\gamma(\mathbf{p}_A, \mathbf{p}_B)) \right] \right] \end{aligned} \quad (7)$$

Similar discriminators have been proposed in [22, 23] to train autoencoders in adversarial setting. We use this discriminator

in VCGAN to establish a macro connection between the two generators. In fact, the optimal discriminators train the corresponding generators to minimize the Jensen-Shanon divergence between $P_\gamma(\mathbf{p}_A, \mathbf{p}_B)$ and $P_\theta(\mathbf{p}_A, \mathbf{p}_B)$ [23].

We use the 23 dimensional MFCC features for spectrum representation over a context of 128 frames extracted using a 5ms long window. The dimensionality of F0 contour is 128x1 while that of spectrum is 128x23. The momenta variable is of the same dimensionality as F0 which is 128x1. The hyperparameters were set to $\lambda_c = 1e-3$ and $\lambda_m = 1e-5$. The generator and discriminator networks are optimized alternately for one epoch each. We fix the mini-batch size to 1 and the learning rates are fixed at 1e-4 and 1e-7 for the generators and discriminators, respectively. We use Adam optimizer [24] with an exponential decay of 0.5 for the first moment. We implement the sampling process in the generators via dropout during training and testing. We convert the spectrum separately using a cycle-GAN proposed by [25]. Code can be downloaded from: <https://engineering.jhu.edu/nsa/links/>.

3. Experiments and Results

We evaluate our VCGAN model against three baselines via crowd-sourcing on Amazon mechanical Turk (AMT). Here, we play both the neutral speech and the converted speech. The listener is asked to classify the emotion in the converted speech and rate its quality on a scale from 1 to 5. We randomize the samples to weed out any non-diligent worker and identify bots.

3.1. Dataset

We use VESUS dataset [15] to carry out the experiments in this paper. VESUS contains a set of 250 utterances spoken by 10 actors in multiple emotions. We train one model for each pair of emotions i.e, neutral to angry, neutral to happy and neutral to sad. The dataset, also comes with human emotional ratings by 10 AMT workers. These ratings denote the ratio of AMT workers who correctly identify the intended emotion in a recorded utterance. For robustness, we only use utterances that are correctly rated as emotional by at least 50% of the total workers. The number of utterances for each emotion pair are:

- **Neutral to Angry conversion:** 1534 for training, 72 for validation and, 61 for testing.
- **Neutral to Happy conversion:** 790 for training, 43 for validation and, 43 for testing.
- **Neutral to Sad conversion:** 1449 for training, 75 for validation and, 63 for testing.

3.2. Baseline Algorithms

The first baseline is the GMM based joint modeling approach of [9]. This algorithm learns a mixture model in the joint space of source and target F0 and spectral features. During inference, a global variance constraint generates non-smooth target features using maximum likelihood. One caveat is that the GMM fails to generate intelligible speech when trained across multiple speakers. As a result, our GMM results are based on training single-speaker models and averaging the results across them. All other methods are trained on the full multi-speaker data.

The second baseline is the Bi-LSTM approach of [10]. This method parameterizes the F0 and the energy contours using a Wavelet transform. Following the authors’ strategy, we pre-train the model on a voice conversion dataset [26]. It is then fine-tuned for emotion conversion on the VESUS dataset.




Figure 2: Confidence of emotion conversion (top) and quality of reconstructed speech (bottom) on VESUS dataset. Marker * above the bars denote $p < 10^{-2}$ for a two sample t-test.




Figure 3: Confidence of emotion conversion (top) and quality of reconstructed speech (bottom) on wavenet speech. Marker * above the bars denote $p < 10^{-2}$ for a two sample t-test.

The third baseline is the unsupervised cycle-GAN proposed by [14]. It modifies the spectrum and F0 contour using two separate cycle-GANs. As described in [14], wavelet transform is applied to the F0 contour for expanding dimensionality.

3.2.1. Mixed Speaker Evaluation

Fig. 2 shows the result of our multispeaker emotion conversion based on the VESUS dataset. Our proposed VCGAN outperforms the baselines on two emotion pairs, namely, neutral to angry and neutral to sad. The cycle-GAN comes a close second ahead of the GMM and Bi-LSTM models. This shows that generative models contain the needed flexibility for this task. Note that GMM’s emotion saliency is close to or better than Bi-LSTM, largely because we train a separate model for each speaker. The poor saliency ratings for the Bi-LSTM likely reflect the difficulty of training recurrent architectures on small datasets. VCGAN performs slightly worse than the cycle-GAN for neutral to happy conversion due to the smaller number of samples for training. For the other two emotion pairs, our hybrid generative approach outperforms the baselines by learning the complex relationship between spectrum and F0 contour.

VCGAN does extremely well in retaining the quality of speech after conversion which is evident from the mean opinion scores (MOS) shown in Fig. 2. This is mainly because the prediction of F0 contour is conditioned on spectrum, which allows




Figure 4: Comparison of F0 estimation by the proposed model and the Cycle-GAN on VESUS parallel utterances. Marker * above the bars denote $p < 10^{-2}$ for a two sample t-test.

the generator to exploit the harmonicity present in the spectrum. Bi-LSTM method has the best MOS among the baseline algorithms because empirically it does not change the utterance but merely copies the source features as output.

3.2.2. Wavenet Evaluation

To simulate an unseen speaker, we generate 100 neutral utterances using Wavenet [2]. We then apply the models learned on the VESUS dataset without any fine-tuning. We have omitted the GMM, since it can only be trained on single speakers, and we do not have access to “emotional” Wavenet utterances. Fig. 3 illustrates the results of this experiment. As seen, the Bi-LSTM does just as poorly on an unseen speaker as on the VESUS dataset. Empirically, we observe that the Bi-LSTM output resembles a distorted identity mapping. While the cycle-GAN largely retains its performance, it achieves a lower emotional saliency than our model in all cases. This is because the smooth warping between source and target F0 automatically adjusts to the frequency range of a new speaker.

Both the cycle-GAN and VCGAN exhibit a decrease in MOS for the Wavenet utterances. Here, the minimal conversion allows the Bi-LSTM to produce more natural sounding speech for neutral to angry conversion. Nonetheless, our method comes in a close second. Taken together, we can conclude that there is a trade-off between the emotion saliency and the speech quality. VCGAN balances it better in comparison to the baselines.

3.2.3. Quantitative Comparison

The parallel utterances in VESUS allow us to objectively measure the difference between the converted and real F0 contours. We compare our approach to the Cycle-GAN model which is also a non-parallel technique. As seen in Fig. 4, our approach has lower mean absolute error for all three emotion pairs. This indicates that intonations might have a unique trend for each emotion, which VCGAN can exploit better than cycle-GAN.

4. Conclusions

We proposed a novel approach to train a pair of GANs in a cyclic schema by comparing their induced joint densities. The GAN generators were composed of a trainable and a static component. The trainable component generated a latent embedding called momenta, which was then used by the fixed component to warp the source F0 contour. Our model is both objectively and subjectively superior to the existing state-of-the-art methods. It achieves a good balance between the emotion saliency and reconstruction quality. The novel loss function used for training the generators helps unfold the complex relationship between spectrum and F0. Further, the deformation based modeling of target F0 contour makes it robust for new unseen speakers.

5. References

- [1] D. Schacter, D. T. Gilbert, and D. M. Wegner, *Psychology (2nd Edition)*. Worth Publications, 2011.
- [2] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," *CoRR*, vol. abs/1609.03499, 2016.
- [3] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. J. Skerry-Ryan, R. A. Saurous, Y. Agiomyrgiannakis, and Y. Wu, "Natural TTS synthesis by conditioning wavenet on mel spectrogram predictions," *CoRR*, vol. abs/1712.05884, 2017.
- [4] N. Kalchbrenner, E. Elsen, K. Simonyan, S. Noury, N. Casagrande, E. Lockhart, F. Stimberg, A. van den Oord, S. Dieleman, and K. Kavukcuoglu, "Efficient neural audio synthesis," in *Proc. ICML 2018*, J. G. Dy and A. Krause, Eds., vol. 80. PMLR, 2018, pp. 2415–2424.
- [5] Y. Wang, R. J. Skerry-Ryan, Y. Xiao, D. Stanton, J. Shor, E. Battenberg, R. Clark, and R. A. Saurous, "Uncovering latent style factors for expressive speech synthesis," *CoRR*, vol. abs/1711.00520, 2017.
- [6] E. Battenberg, S. Mariooryad, D. Stanton, R. Skerry-Ryan, M. Shannon, D. Kao, and T. Bagby, "Effective use of variational embedding capacity in expressive end-to-end speech synthesis," *CoRR*, vol. abs/1906.03402, 2019.
- [7] Y. Ning, S. He, Z. Wu, C. Xing, and L.-J. Zhang, "A review of deep learning based speech synthesis," *Applied Sciences*, vol. 9, p. 4050, 09 2019.
- [8] K. R. Scherer, "Vocal communication of emotion: A review of research paradigms," *Speech Communication*, vol. 40, no. 1, pp. 227 – 256, 2003.
- [9] R. Aihara, R. Takashima, T. Takiguchi, and Y. Ariki, "Gmm-based emotional voice conversion using spectrum and prosody features," *American Journal of Signal Processing*, vol. 2, pp. 134–138, 12 2012.
- [10] H. Ming, D.-Y. Huang, L. Xie, J. Wu, M. Dong, and H. Li, "Deep bidirectional lstm modeling of timbre and prosody for emotional voice conversion," in *Proc. Interspeech 2016*, 09 2016, pp. 2453–2457.
- [11] R. Shankar, J. Sager, and A. Venkataraman, "A Multi-Speaker Emotion Morphing Model Using Highway Networks and Maximum Likelihood Objective," in *Proc. Interspeech 2019*, 2019, pp. 2848–2852.
- [12] R. Shankar, H.-W. Hsieh, N. Charon, and A. Venkataraman, "Automated Emotion Morphing in Speech Based on Diffeomorphic Curve Registration and Highway Networks," in *Proc. Interspeech 2019*, 2019, pp. 4499–4503.
- [13] J. Gao, D. Chakraborty, H. Tembine, and O. Olaleye, "Nonparallel Emotional Speech Conversion," in *Proc. Interspeech 2019*, 2019, pp. 2858–2862.
- [14] K. Zhou, B. Sisman, and H. Li, "Transforming spectrum and prosody for emotional voice conversion with non-parallel training data," *CoRR*, vol. abs/2002.00198, 2020.
- [15] J. Sager, R. Shankar, J. Reinhold, and A. Venkataraman, "VESUS: A Crowd-Annotated Database to Study Emotion Production and Perception in Spoken English," in *Proc. Interspeech 2019*, 2019, pp. 316–320.
- [16] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio, "Generative adversarial networks," *CoRR*, vol. abs/1406.2661, 2014.
- [17] J. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *Proc. ICCV 2017*. IEEE Computer Society, 2017, pp. 2242–2251.
- [18] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *International journal of computer vision*, vol. 61, no. 139–157, 2005.
- [19] S. C. Joshi and M. I. Miller, "Landmark matching via large deformation diffeomorphisms," *IEEE transactions on image processing*, vol. 9, no. 8, pp. 1357–1370, 2000.
- [20] H.-W. Hsieh and N. Charon, "Diffeomorphic registration of discrete geometric distributions," *CoRR*, vol. abs/1801.09778, 2018.
- [21] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel, "Infogan: Interpretable representation learning by information maximizing generative adversarial nets," in *Proc. NIPS 2016*, D. D. Lee, M. Sugiyama, U. von Luxburg, I. Guyon, and R. Garnett, Eds., 2016, pp. 2172–2180.
- [22] A. Srivastava, L. Valkov, C. Russell, M. U. Gutmann, and C. A. Sutton, "VEEGAN: reducing mode collapse in gans using implicit variational learning," in *Proc. NIPS, 2017*, I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, and R. Garnett, Eds., 2017, pp. 3308–3318.
- [23] V. Dumoulin, I. Belghazi, B. Poole, A. Lamb, M. Arjovsky, O. Mastropietro, and A. C. Courville, "Adversarially learned inference," *CoRR*, vol. abs/1606.00704, 2016.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *CoRR*, vol. abs/1412.6980, 2015.
- [25] T. Kaneko and H. Kameoka, "Parallel-data-free voice conversion using cycle-consistent adversarial networks," *CoRR*, vol. abs/1711.11293, 2017.
- [26] J. Kominek and A. W. Black, "The cmu arctic speech databases," *SSW5-2004*, 01 2004.