



Leveraging Symmetrical Convolutional Transformer Networks for Speech to Singing Voice Style Transfer

Shrutina Agarwal[§], Sriram Ganapathy[§], Naoya Takahashi^{*}

[§]LEAP lab, Indian Institute of Science, Bangalore, India.

^{*}Sony Group Corporation, Tokyo, Japan.

sriramg@iisc.ac.in

Abstract

In this paper, we propose a model to perform style transfer of speech to singing voice. Contrary to the previous signal processing-based methods, which require high-quality singing templates or phoneme synchronization, we explore a data-driven approach for the problem of converting natural speech to singing voice. We develop a novel neural network architecture, called SymNet, which models the alignment of the input speech with the target melody while preserving the speaker identity and naturalness. The proposed SymNet model is comprised of symmetrical stack of three types of layers - convolutional, transformer, and self-attention layers. The paper also explores novel data augmentation and generative loss annealing methods to facilitate the model training. Experiments are performed on the NUS and NHSS datasets which consist of parallel data of speech and singing voice. In these experiments, we show that the proposed SymNet model improves the objective reconstruction quality significantly over the previously published methods and baseline architectures. Further, a subjective listening test confirms the improved quality of the audio obtained using the proposed approach (absolute improvement of 0.37 in mean opinion score measure over the baseline system).

Index Terms: Speech to singing style transfer, transformer networks, symmetrical neural networks.

1. Introduction

The task of speech to singing (STS) voice style transfer is the problem of converting the speech signal to a natural sounding singing voice. The inputs to the STS tasks are the speech waveform and a target melody contour. The key challenges involved in this problem relate to the ability of the STS system to generate natural voice which preserves the speaker and phonetic content while incorporating the melody.

The style transfer on other domains like images [1] has been more successful than in audio and speech domain. In the recent past, several works attempt the goal of transforming specific properties of audio like the speaker characteristics (referred to as “style”) while keeping linguistic characteristics intact (referred to as “content”). Haque et al. [2] proposed a system to transform the timbre of the speech. Similarly, Mor et al. [3] proposed a translation framework across various musical instruments, styles and genres. Wu et al. [4] explored a cycle-generative adversarial network (cycle-GAN) based framework to perform singing voice conversion. In this paper, we propose a novel neural architecture for the task of speech to singing voice conversion.

Typically, singing involves more vocal effort and tends to have a slower rate of change of syllabic content along with enhanced higher frequency formants [5, 6]. Due to the melodic

nature of singing, a singer produces a smoother fundamental frequency (F0) contour [7]. The previous works on STS conversion can be broadly categorized into two directions : (i) Template-based STS [8] and (ii) Model-based STS [9]. Both these methods require additional inputs like high quality singing templates or phoneme-score synchronization information, which are cumbersome in practice.

In this paper, inspired by the prior work by Parekh et. al. [10], we consider the STS task as a style transfer problem. The proposed model uses a template of melody/pitch contour and poses the style transfer as a data driven learning problem. We also propose a novel neural network architecture for the style transfer where a supervised singing voice target with a suitable objective function is introduced. Further, the adversarial loss function based on boundary equilibrium GAN [11] is used for training the STS model.

The model architecture is similar to the encoder-decoder model [12] popularly used in machine translation. The proposed architecture is realized with a symmetric network (called SymNet) consisting of convolutional, transformer, and attention layers. The alignment between the speech and singing voice is achieved with the attentional layer between the central transformer layers while local signal characteristics are captured with convolutional layers at the beginning of encoder network as well as the end of decoder network. We also explore a data augmentation procedure using unpaired singing/speech data to mitigate the necessity of large amounts of paired data. The experiments are performed on the NUS [13] and the NHSS datasets [14]. In these experiments, we compare the proposed approach with prior state-of-the-work for STS. Using the proposed SymNet architecture, data augmentation and loss function, we observe significant performance improvements over baseline models in terms of the objective and subjective reconstruction quality.

2. STS Framework

The block schematic of the proposed model for STS is shown in Fig 1. We represent the input speech with the log-magnitude spectrogram. A vocal melody extractor [15], is used to extract melody contour from the reference/target singing voice. We then time-stretch the spectrogram of speech to the same length as the target F0 contour. The speech spectrogram are concatenated with the melody contour features and are input to the deep neural model.

2.1. SymNet architecture

The proposed SymNet architecture is shown in Fig 2. The input speech spectrogram along with the melody features are passed

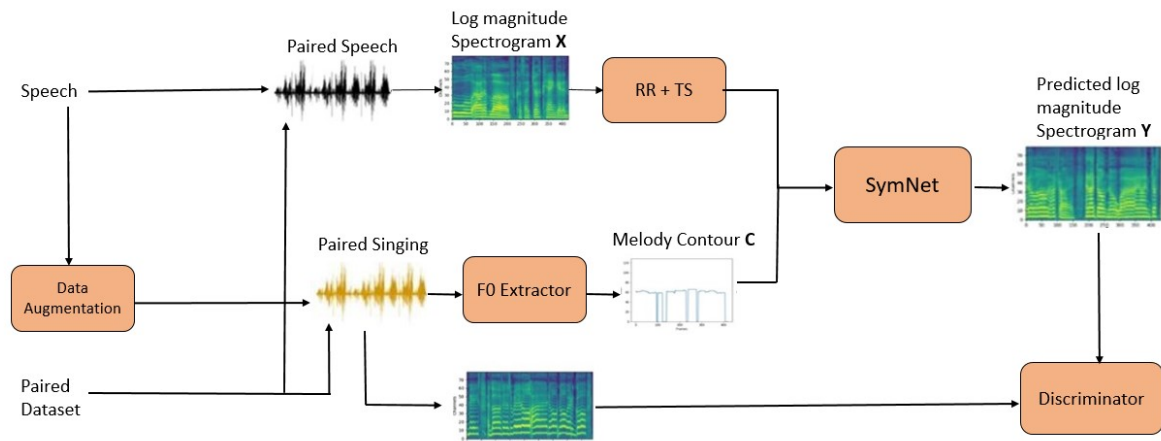


Figure 1: Block schematic of the proposed approach to STS. Here, RR is random sampling and TS is time-stretching.

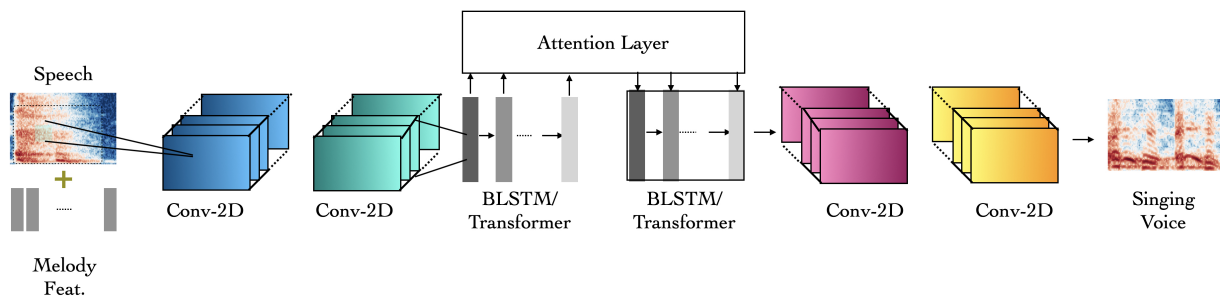


Figure 2: Proposed symmetrical convolutional recurrent network (SymNet) architecture with attention.

Table 1: SymNet configuration. K : kernels, D - dimensions

LAYER	SYMNET ₂₅₆	SYMNET ₅₁₂
CONV \times 2	64K, 3×3	64K, 3×3
BLSTM/TRANSFORMER	256 D	512 D
ATTEN.	256 D	512 D
BLSTM/TRANSFORMER	256 D	512 D
CONV \times 2	64K, 3×3	64K, 3×3

to two convolutional layers to model the local spectro-temporal patterns. The output of the convolutional layer is passed through a recurrent layer for modeling long-term dependencies. We consider bi-directional long-short term memory (BLSTM) or a transformer layers in this stage. The BLSTM model is a recurrent neural architecture implements using LSTM cells while the transformer models use self-attention layers to incorporate temporal dependency. The next layer of processing involves a separate attention network which performs the non-linear time alignment between the speech and singing voice spectrogram based on the target melody features. The rest of the network consisting of a BLSTM/transformer layer and the two convolutional layers allows the architecture to form a symmetrical model before and after the attention layer. We highlight the importance of the symmetrical architecture and the combination of three types of layers by comparing other possible architectures in our experiments. The detailed SymNet configuration is summarized in Table 1. Each transformer layer has 4 attention heads along with 2 encoder and 2 decoder layers. The dimension of the feed forward layer is set to 1024.

2.2. Data augmentation

We utilize a data augmentation method that generates singing voice data from natural speech data. Our prior work [16] used a data generation approach for automatic lyric recognition by leveraging the WORLD vocoder [17]. Specifically, the approach consisted of decomposing the natural speech and singing voice (opera vocals) into three constituent components - spectral envelope, fundamental frequency (F0) and aperiodicity parameter. The F0 contour from the singing voice is then used along with the spectral envelope and the aperiodicity parameter from the natural speech in the synthesizer. The synthesized output is the singing voice version of the natural speech. The western opera vocal dataset consists of both male and female opera singers and during the synthesis it was made sure that the speech and the opera vocals are gender and F0 matched [16]. We use the synthesized output and the input speech data as the paired data for augmentation. While this approach already generates a singing voice version of the natural speech, we find that the reconstruction quality is not sufficient for commercial applications. More details about the data augmentation can be found in Basak et. al. [16],

2.3. BEGAN loss

The boundary equilibrium generative adversarial networks (BEGAN) [18] is a energy based GAN, which matches the distribution of loss using an auto-encoder architecture, unlike the original GAN which matches the distributions between the real and generated samples directly. Moreover, the BEGAN relaxes the

equilibrium of the auto-encoder loss using a hyper-parameter $\gamma \in [0, 1]$. In order to implement the BEGAN loss (L_G) [11], the STS model also has an additional discriminator network.

2.4. Training

We experimented with two different training settings, (i) without the discriminator (and therefore without the BEGAN loss), and (ii) with the BEGAN discriminator. For the first one, our learning objective is based on mean square error (L_{MSE}) between predicted (\hat{Y}) and the true log-magnitude spectrograms of singing voice ($Y \in R^{F \times T}$). The combined loss then is, $L = L_{MSE} + \zeta L_G$. Here, ζ controls the trade-off between the MSE and the GAN loss.

2.5. Loss function annealing

We explore a loss function annealing, where the regularization term ζ in the joint loss function is annealed as a function of the iteration to focus on the reconstruction MSE in the latter part of the training. In our experiments (Sec 3.6), we explored various annealing schedules, such as a linear decay and a step decay.

3. Experiments

3.1. Dataset

We use the NUS data [13] as the paired training dataset containing 115 mins of singing data and the corresponding 54 mins of speech data. The 48 recordings of 20 unique English songs are sung/read by 12 subjects, and each song-speech paired audio can be time-aligned with their phone annotations. The phone annotations are in accordance with the CMU dictionary (39 phonemes). Out of the 20 unique songs in the dataset, one song (with two recordings) is used as the test set for NUS (similar to the dataset division followed in [10]). We also use NHSS [14] for training/testing the models. It consist of 100 songs with their respective speech recorded by 10 singers, resulting in a total of 7 hours of audio data. For data augmentation, the LibriSpeech [19] is converted to synthetic singing voice and about 10 hours of augmented data is added for model training (Sec. 2.2). We use the log spectral distance (LSD) as the performance metric in our objective evaluation. Further, a subjective test is performed to affirm the quality improvements seen in the objective metric.

3.2. Input pre-processing

Given an input time-domain speech signal and a target F0 contour, the pre-processing consists of the following steps:

Silence frame removal : All the silence frames from the speech signal are removed. This is achieved via a short-time energy threshold set at 40dB below the maximum energy frame. Any set of three or more consecutive silent frames (longer than approximately 50ms) are removed.

Time stretching : The input spectrogram is linearly time-warped to the length of the F0 contour.

Singing melody contour : The melody contour is extracted using the CREPE algorithm [20] applied on the singing voice. We also convert the continuous-valued F0 to one of the 128 MIDI levels. For training with paired data, we extract the melody contour from the singing counterpart of the input speech.

Random re-sampling (RR) : We have re-sampled the input speech in random fashion to change the rhythm and to disentangle the content and rhythm [11].

Table 2: Performance (LSD -dB) for the ablation study on the network architecture. Here, L - Layers, C - Cells, A - Attention, and K - convolutional kernels. All models here are trained with MSE loss.

ARCH.	DIM.	NUS	NHSS
BLSTM(3L)	256	10.45	11.54
BLSTM(3L)	512	10.42	11.47
CONV(2L)-BLSTM(2L)	256	10.38	11.35
CONV(2L)-BLSTM(2L)	512	10.33	11.31
BLSTM(2L)-A-BLSTM(2L)	256	10.34	11.02
BLSTM(2L)-A-BLSTM(2L)	512	10.29	10.96
CONV(4L)-BLSTM(2L)-A	256	10.16	10.88
CONV(4L)-BLSTM(2L)-A	512	10.13	10.85
CONV(4L)-TRANS.(2L)-A	256	9.99	10.63
CONV(4L)-TRANS.(2L)-A	512	9.89	10.56
SYMNET ₂₅₆ (BLSTM)	256	10.02	10.70
SYMNET ₅₁₂ (BLSTM)	512	9.95	10.61
SYMNET ₂₅₆ (TRANS.)	256	9.93	10.55
SYMNET ₅₁₂ (TRANS.)	512	9.87	10.51

Table 3: Comparison of BEGAN loss annealing schedule for ζ . The model architecture is SymNet₂₅₆ (BLSTM) trained on NUS dataset.

SCHEDULE FOR ζ	NUS	NHSS
CONSTANT $\zeta = 0$	10.02	10.70
CONSTANT $\zeta = 0.3$	9.52	10.41
$\zeta = 0.3$ DECAY BY 0.001 AT EACH EPOCH	9.45	10.40
FIRST 15 EPOCHS $\zeta = 0.3$, $\zeta = 0$ AFTER	9.33	10.27

3.3. Audio reconstruction

The magnitude spectrogram is recovered by applying $g(\hat{Y}) = e^{\hat{Y}-1}$ element-wise (the target spectrograms are $\hat{Y} = \log(1 + \hat{y})$, where y is the magnitude spectrogram of the singing voice). The reconstructed magnitude spectrogram $g(\hat{Y})$ is then converted to log-mel spectrogram with 80 frequency bins. Then, the MelGAN vocoder [21] is used to generate audio waveforms from the spectrogram.

3.4. Implementation details

The time domain signals are re-sampled from 44kHz to 16kHz. We compute short-time Fourier transform (STFT) with 1024 frequency samples using 64 ms window size and 16ms hop size. For the first setting (the one with just the MSE loss), we use Adam optimizer [22] with initial learning rate of 0.003 and an exponential decay factor of 0.9. All the networks were trained for 30 epochs, (with 1000 mini-batches in each epoch). For the second setting, the one with BEGAN loss, we use Adam optimizer with initial learning rate of 0.001 and exponential decay factor of 0.99. All the neural network implementations are performed using PyTorch [23] and Librosa [24].

3.5. Ablation study on architecture design

The first set of results, reported in Table 2, show the results for the ablation study with different model architecture choices. All the models reported here use the first setting of the loss function (only MSE). The results reported here are for models with various architectures - three BLSTM layers (with 256 and 512 cells), convolutional and BLSTM layers, BLSTM layers with

Table 4: Comparison of the prior work with the proposed approach in terms of log spectral distance (LSD) in dB. Also reported are the incremental performance gains observed using various components. Here, Symnet corresponds to symmetric network with either the transformer or BLSTM layers.

MODEL	MSE LOSS	BEGAN LOSS	DATA AUG.	W/ NHSS	LSD-NUS	LSD-NHSS
PAREKH ET. AL. [10]	✓				11.22	11.91
WU ET AL. [11]		✓			9.6	10.57
SYMNET ₅₁₂ (BLSTM)	✓				9.95	10.61
SYMNET ₅₁₂ (TRANSFORMER)	✓	✓			9.22	10.13
SYMNET ₅₁₂ (BLSTM)	✓	✓	✓		9.07	9.98
SYMNET ₅₁₂ (TRANSFORMER)	✓	✓	✓		8.99	9.62
SYMNET ₅₁₂ (BLSTM)	✓	✓		✓	9.01	8.79
SYMNET ₅₁₂ (TRANSFORMER)	✓	✓		✓	8.81	8.74
SYMNET ₅₁₂ (BLSTM)	✓	✓	✓	✓	8.80	8.76
SYMNET ₅₁₂ (TRANSFORMER)	✓	✓	✓	✓	8.63	8.67

attention, convolutional layers followed by BLSTM layers, convolutional and transformer layers and the proposed SymNet architectures. The key observations from these experiments are,

- The mixture of convolutional and transformer/LSTM layers is superior to the use of convolutional layers alone [11]. The use of 512D is slightly better than 256D.
- The use of attention in the model allows the non-linear time alignment between the speech and singing voice counterparts with only a minor increase in the number of parameters.
- The transformer improves over the BLSTM counterparts.
- Combining all three types of layers (Conv-Trans-A) demonstrates the effectiveness of the three types of layers, namely, local modeling with convolution, long-term modeling with the transformer, and dynamic alignment with attention.
- SymNet improves the LSD over the models with the same set of layers and parameters, but are asymmetrical. To the best of our knowledge, this is the first attempt to highlight the importance of the symmetric architecture for audio style transfer tasks.

3.6. BEGAN loss annealing

In Table 3, we consider three schedules for ζ which controls relative weighting of the BEGAN loss and the MSE loss - (i) ζ is fixed at 0.3, (ii) ζ is initialized to 0.3 and gradually reduced by 0.001 for every epoch, (iii) ζ is initialized to 0.3 and changed to 0 after 15 epochs. Table 3 shows the results for the scheduling choices along with the system trained without BEGAN loss ($\zeta = 0$). Even though BEGAN loss improves the LSD for fixed ζ , annealing ζ further improves the performance.

3.7. Analysis of different components

The results comparing different components of the proposed model are given in Table 4. The encoder-decoder model reported in [11] serves as the baseline system. The results reported here show that the incorporation of the SymNet architecture with BEGAN loss improves over the baseline system. Further, the incorporation of the data augmentation method improves the reconstruction quality. The effect of additional training data is explored by training the *SymNet*₅₁₂ model with additional data from NHSS dataset [14]. Further the combination of the data-augmentation approach along with the addition of the NHSS dataset improves the reconstruction quality with an absolute improvement of 0.97 dB over the baseline system on the NUS dataset and 1.46 dB on the NHSS dataset.

Table 5: Subjective evaluation (MOS) using hidden reference (ground truth) with 30 subjects listening to 20 audio snippets synthesized from the *SymNet*₅₁₂(Transformer) and the baseline system (We et al. [11]). Both the systems used the same training data as well.

	HIDDEN REF.	WU [11]	SYMNET
MOS	4.86 ±0.01	2.71 ±0.08	3.08 ±0.09

3.8. Subjective Evaluation

We conducted a mean opinion score (MOS) listening experiment¹ with 20 audio clips, each with three versions - hidden reference (ground truth), style transferred from speech using the proposed SymNet approach and using the baseline system [11]. The three audio files for each snippet were presented in a random order. The participants belonged to the age group of 21 – 39 with normal hearing. The subjects were asked to listen to the audio stimuli using high quality headphones and rate the files according to naturalness, reconstruction quality, intelligibility of the lyrics, and the preservation of speaker identity. The participants gave the rating on a 5 point scale, with 1 being extremely poor quality while 5 represents the natural singing voice. In total, 30 participants took part in the subjective listening test with 20 audio snippets. The summary of the subjective listening tests results is given in Table 5. The high rating for the hidden reference makes the results credible. The subjective results confirm the improved quality of the proposed SymNet model over the baseline system. The SymNet model improves the baseline system by an absolute MOS value of 0.37.

4. Summary

This paper proposes a novel architecture for the task of audio style transfer where the objective is to convert the read speech to singing voice (STS). The proposed architecture consists of a mixture of convolutional and transformer neural networks along with an attention layer. We investigate additional novel components for data augmentation and loss function annealing. The proposed system achieves significant improvements in reconstruction quality over the state-of-art approach using objective metrics. Several experiments have been performed to analyze the incremental performance benefits for various components in the proposed system. Finally, subjective evaluations also confirm the superior performance of the proposed method.

¹The study is hosted at <https://symnettesting.herokuapp.com/>

5. References

- [1] L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2414–2423.
- [2] A. Haque, M. Guo, and P. Verma, "Conditional end-to-end audio transforms," *10.21437/Interspeech.2018-38*, pp. 2295–2299, 09 2018.
- [3] N. Mor, L. Wolf, A. Polyak, and Y. Taigman, "A universal music translation network," *arXiv preprint arXiv:1805.07848*, 2018.
- [4] C.-W. Wu, J.-Y. Liu, Y.-H. Yang, and R. J.-S. Jang, "Singing style transfer using cycle-consistent boundary equilibrium generative adversarial networks," *arXiv: arXiv:1807.02254*, 2018.
- [5] J. Sundberg, "The acoustics of the singing voice," *Scientific American*, vol. 236, no. 3, pp. 82–91, 1977.
- [6] I. R. Titze and J. Sundberg, "Vocal intensity in speakers and singers," *the Journal of the Acoustical Society of America*, vol. 91, no. 5, pp. 2936–2946, 1992.
- [7] M. Umbert, J. Bonada, M. Goto, T. Nakano, and J. Sundberg, "Expression control in singing voice synthesis: Features, approaches, evaluation, and challenges," *IEEE Signal Processing Magazine*, vol. 32, no. 6, pp. 55–73, 2015.
- [8] L. Cen, M. Dong, and P. Chan, "Template-based personalized singing voice synthesis," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2012, pp. 4509–4512.
- [9] T. Saitou, M. Goto, M. Unoki, and M. Akagi, "Speech-to-singing synthesis: Converting speaking voices to singing voices by controlling acoustic features unique to singing voices," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 2007, pp. 215–218.
- [10] J. Parekh, P. Rao, and Y. Yang, "Speech-to-singing conversion in an encoder-decoder framework," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 261–265.
- [11] D.-Y. Wu and Y.-H. Yang, "Speech-to-Singing Conversion Based on Boundary Equilibrium GAN," in *Proc. Interspeech*, 2020, pp. 1316–1320.
- [12] K. Cho, B. Van Merriënboer, D. Bahdanau, and Y. Bengio, "On the properties of neural machine translation: Encoder-decoder approaches," *arXiv preprint arXiv:1409.1259*, 2014.
- [13] Z. Duan, H. Fang, B. Li, K. C. Sim, and Y. Wang, "The NUS sung and spoken lyrics corpus: A quantitative comparison of singing and speech," in *Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, 2013, pp. 1–9.
- [14] B. Sharma, X. Gao, K. Vijayan, X. Tian, and H. Li, "NHSS: A speech and singing parallel database," *arXiv preprint arXiv:2012.00337*, 2020.
- [15] L. Su, "Vocal melody extraction using patch-based cnn," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 371–375.
- [16] S. Basak, S. Agarwal, S. Ganapathy, and N. Takahashi, "End-to-end lyrics recognition with voice to singing style transfer," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.
- [17] 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.
- [18] D. Berthelot, T. Schumm, and L. Metz, "Began: Boundary equilibrium generative adversarial networks," *arXiv preprint arXiv:1703.10717*, 2017.
- [19] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [20] J. Kim, J. Salamon, P. Li, and J. Bello, "CREPE: A convolutional representation for pitch estimation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018.
- [21] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. Courville, "Melgan: Generative adversarial networks for conditional waveform synthesis," *arXiv preprint arXiv:1910.06711*, 2019.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [23] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, et al., "Pytorch: An imperative style, high-performance deep learning library," *arXiv preprint arXiv:1912.01703*, 2019.
- [24] B. McFee, C. Raffel, D. Liang, D. P. Ellis, M. McVicar, E. Battenberg, and O. Nieto, "Librosa: Audio and music signal analysis in Python," in *Proceedings of the 14th python in science conference*, 2015, vol. 8, pp. 18–25.