



Phase-aware music super-resolution using generative adversarial networks

Shichao Hu*, Bin Zhang*, Beici Liang, Ethan Zhao, Simon Lui

Tencent Music Entertainment (TME), Shenzhen, 518057, China

{shichaohu, robingzhang, beiciliang, ethanzhao, nomislui}@tencent.com

Abstract

Audio super-resolution is a challenging task of recovering the missing high-resolution features from a low-resolution signal. To address this, generative adversarial networks (GAN) have been used to achieve promising results by training the mappings between magnitudes of the low and high-frequency components. However, phase information is not well-considered for waveform reconstruction in conventional methods. In this paper, we tackle the problem of music super-resolution and conduct a thorough investigation on the importance of phase for this task. We use GAN to predict the magnitudes of the high-frequency components. The corresponding phase information can be extracted using either a GAN-based waveform synthesis system or a modified Griffin-Lim algorithm. Experimental results show that phase information plays an important role in the improvement of the reconstructed music quality. Moreover, our proposed method significantly outperforms other state-of-the-art methods in terms of objective evaluations.

Index Terms: Music super-resolution, Bandwidth expansion, Generative adversarial network, Phase estimation

1. Introduction

Audio super-resolution (SR) deals with the problem of recovering high-resolution (HR) audio from a low-resolution (LR) input. This is also known as “bandwidth expansion”. Traditionally, this technique is applied to telephone speech transmission where speech resolution is limited due to the low-bandwidth channels [1]. There are also studies showing that artificially expanding the speech bandwidth improves the perception for hearing-impaired persons [2]. In this paper, our study is mainly focused on SR tasks for music and aimed at improving the audio quality of LR music by predicting both the magnitudes and phase of the high frequency components (denoted as HFC).

Conventional approaches for SR problems are based on signal processing methods, such as linear mapping from the LR to HR spectral characteristics using linear predictive coding (LPC) analysis [3]. The mapping method can be extended with more complex statistical techniques, for instance, Gaussian mixture models (GMMs) [4] and Hidden Markov Model (HMM) [5].

With the development of deep learning, recent audio SR approaches based on deep neural networks (DNNs) present better performance than the conventional methods. Li et al. [6] used fully connected DNNs to learn the LR to HR mapping of spectral magnitude and then reconstruct the HR audio with the flipped phase of low frequency components (denoted as LFC). Kuleshov et al. [7] employed an auto-encoder network to generate the HR waveform from the LR waveform input. This time domain approach can avoid the phase problem and expand bandwidth through temporal interpolation. More recently, Eskimez et al. [8, 9] introduced a novel method that estimates log-power spectrograms of the HFC using the generative adversarial network (GAN). This was inspired by the great success of GAN

on generating highly realistic images [10, 11]. Experimental results of such GAN method achieved the state-of-the-art performance in the frequency domain.

To the best of our knowledge, existing literature using frequency-domain approach with DNNs has not fully considered the HFC phase to reconstruct HR signals. Considering the difficulty of phase modeling, directly mapping the HFC phase from LFC remains unsolved. Given the fact that phase plays an important role in speech perception, enhancement and synthesis [12, 13, 14], we propose to use a recent MelGAN-based audio vocoder [15] to predict the phase in Section 3. HR waveform can be generated using the predicted phase and the magnitude which is estimated using the GAN-based method mentioned above. Our proposed method outperforms the other state-of-the-art methods according to the evaluation results in Section 4. We further investigate the importance of phase in the music SR tasks through evaluating the quality of audio generated by different methods. Examples of the reconstructed music are available online ¹.

2. Related Works

Recent studies have shown deep-learning-based methods are effective on audio SR tasks. This section reviews the recent advanced audio SR approaches and introduces phase estimation methods.

2.1. Audio Super-Resolution with Deep Learning

Audio SR with deep learning is usually approached from either frequency or time domain. One of the earliest frequency-domain approaches was proposed in [6] using restricted Boltzmann machines as DNNs for training the mapping between magnitudes of LFC and HFC. A recent study introduced GANs for the purpose and shows the state-of-the-art results [8, 9]. A typical time-domain approach was proposed recently [7] which applied an auto-encoder network to generate HR waveform from the LR temporal signals.

While the GANs-based spectra-mapping method shows promising results, the HFC phase for HR waveform reconstruction was not fully considered. In the existing works, it is artificially produced by flipping and repeating the LFC phase and adding a negative sign (denoted as FLIP). Since phase has been confirmed to effectively improve subjective perception on speech [12, 13], we attempt to predict HFC magnitudes complemented by the estimated phase in order to recover the HR waveform with improvements in music SR tasks.

¹<https://github.com/tencentmusic/TME-Audio-Super-Resolution-Samples>

*Indicating Equal Contribution.

2.2. Phase Estimation

Here two common methods for phase estimation are briefly introduced. One is Griffin-Lim algorithm (GLA), and the other one is based on speech vocoder. They are used and compared in our music SR tasks.

GLA has been commonly used as a phase recovery method based on the consistency of spectrogram[16, 17]. GLA method iteratively updates the complex-valued spectrogram while maintains the given magnitude (\mathbf{A}). The process of GLA can be explained by the following equation (here we largely followed the notation used in [18]):

$$\mathbf{X}^{[m+1]} = P_C(P_A(\mathbf{X}^{[m]})), \quad (1)$$

where $\mathbf{X}^{[m]}$ denotes the complex-valued spectrogram updated through the m -th iteration. P_A updates the phase and maintains the given magnitudes, i.e., \mathbf{A} . P_C calculates the corresponding consistent spectrograms. P_C and P_A are respectively given by:

$$P_C(\mathbf{X}) = \mathbb{F}\mathbb{F}^\dagger \mathbf{X}, \quad (2)$$

$$P_A(\mathbf{X}) = \mathbf{A} \odot \mathbf{X} \oslash |\mathbf{X}|, \quad (3)$$

where \mathbb{F} and \mathbb{F}^\dagger are the forward short-time Fourier transform (STFT) and inverse STFT, respectively. \odot is element-wise multiplication; \oslash represents element-wise division and division by zero is replaced by zero. Through iterations, complex-valued spectrograms are consistently optimized by minimizing the difference between the energy of \mathbf{X} and $P_C(\mathbf{X})$. To adapt GLA to our music SR tasks, we not only maintain the given LFC magnitude but also the LFC phase in Eq. 3. Although GLA has been widely used for the past decades due to its simplicity, it requires many iterations and resulted reconstruction quality is not good enough. In the experiments of this study, the GLA is implemented with 100 iterations, following the number of iterations used in [18].

Recently, neural-network-based speech vocoder has been extensively developed for waveform synthesis over the past few years. Although the neural-network-based vocoder has been generally developed for speech applications, it may have promising performance as well when applied to music audio modeling [19]. Here we decide to use a state-of-the-art method proposed in [15], which generates audio waveform from mel-spectrogram using GANs (denoted as MelGAN). The MelGAN model consists of a fully convolutional generator network that generate raw waveform from mel-spectrogram input, and a discriminator network with a multi-scale architecture. In our task, a waveform can be generated using the mel-spectrogram concatenating the LFC of the input LR signal and the estimated HFC. The phase information of HFC can be therefore extracted from the generated audio.

3. Proposed System

3.1. Joint Framework

We propose a joint music SR framework, mainly consisting of a ‘‘magnitude component’’ (MC) and a ‘‘phase component’’ (PC). Following the GANs architecture presented in [9], MC predicts the magnitudes of HFC from the LFC. Magnitude prediction is done by a generator network with a convolutional auto-encoder architecture. Besides, a discriminator network is designed to distinguish the estimated magnitudes from real ones. To train MC, we used Wasserstein distance [20] as GAN loss function to stabilize the training. As for PC, we used the the MelGAN

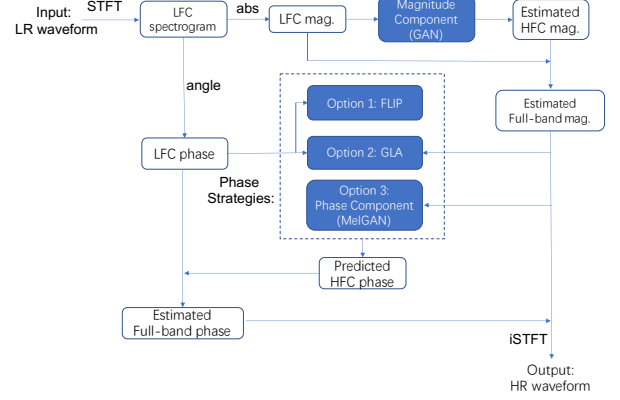


Figure 1: Schematic overview of the proposed system.

model presented in [15] to convert mel-spectrogram into waveform, from which the phase of HFC can be estimated.

Fig. 1 presents the whole joint framework. Both the magnitudes from MC and the phase from PC are used to obtain complex-valued spectrogram of high frequency bands. To estimate phase, PC with MelGAN is used in order to obtain a more accurate phase information for audio reconstruction. It is compared with another two strategies, namely FLIP and GLA. With this joint methods, quality of the reconstructed audio is expected to be improved. A detailed calculation process is illustrated as follows:

- Given the LR waveform input (\mathbf{x}_{LR}), calculate its LFC spectrogram (\mathbf{X}_{LFC}) with L time steps and K frequency bins;
- Use MC to predict magnitude of HFC ($\hat{\mathbf{M}}_{HFC}$) with N frequency bins from the magnitude of \mathbf{X}_{LFC} , i.e., \mathbf{M}_{LFC} ;
- Concatenate $\hat{\mathbf{M}}_{HFC}$ and \mathbf{M}_{LFC} as the input to PC for generating the corresponding waveform ($\hat{\mathbf{x}}_{MelGAN}$);
- HFC phase ($\hat{\mathbf{P}}_{HFC}$) can be extracted from $\hat{\mathbf{x}}_{MelGAN}$, or estimated using FLIP or GLA;
- Using $\hat{\mathbf{P}}_{HFC}$ and $\hat{\mathbf{M}}_{HFC}$, spectrogram of HFC ($\hat{\mathbf{X}}_{HFC}$) can be obtained;
- Concatenate $\hat{\mathbf{X}}_{HFC}$ and \mathbf{X}_{LFC} to generate HR spectrogram ($\hat{\mathbf{X}}_{HR}$);
- Finally, HR waveform ($\hat{\mathbf{x}}_{HR}$) can be obtained from $\hat{\mathbf{X}}_{HR}$ using iSTFT.

The proposed method is compared with the FLIP and GLA methods as different options for HFC phase estimation as presented in Fig. 1.

3.2. Implementation Details

For the magnitude components, we adjust the network architecture and configuration in [9] for this task. During the training process, generator of MC takes \mathbf{M}_{LFC} with $L = 32$ as input. The encoder block contains 4 convolution layers, which have 372,512,512 and 1024 channels with filter size of 7,5,3,3 and stride of 2,2,2,2. The decoder block contains 3 convolution layers, which have 1024, 1024 and 744 channels with filter size of 3,5,7 and stride 1,1,1. To upscale the temporal feature maps in the decoder block, the subpixel layer proposed in Shi et al. [21]

is used after each LeakyReLU activation. A skip connection is added between each encoder layer and the subpixel output of the corresponding decoder layer. Besides, batch normalization is used after each convolution layer and before the LeakyReLU activation. Following the decoder block, two convolution layers are used to generate the \hat{M}_{HFC} . They contain $2N$ channels with filter size of 7 and stride 1 and a subpixel layer, N channels with filter size of 9 and stride 1, respectively. As for the discriminator of MC, it consists of three convolution layers, each of which contain 1024 channels with stride of 7, 5 and 3, respectively. These layers are followed by two fully connected layers with 2048 and 1 neurons respectively, and a Sigmoid activation function for the final output of MC.

For the phase component, we use the network architecture described in [15]. The input of PC is 32×160 log-mel magnitudes transformed from the MC-predicted linear-scaled magnitudes. PC can then output a waveform of 8192 samples. We only make use of the HFC phase information extracted from the waveform. LFC phase is kept the same as the one calculated from LFC spectrogram of the LR waveform.

4. Experimental Results

The dataset we used to train the model consists of 3000 music samples of lossless quality with a sampling rate of 44.1 kHz. These samples can generate the corresponding LR data using low-pass filtering. The spectrum width for LR samples was cut up to 4kHz, while the bandwidth for the target HR was set up to 8kHz. This is known as a $2x$ task. We used another 25 lossless-LR-paired music samples for evaluating the performance of our proposed system. The STFT is calculated with frame and hop length of 2048 and 256, respectively.

For an objective evaluation between different methods, we choose error metrics of the log-spectral distance (LSD) and the signal-to-noise ratio (SNR). LSD calculates the average spectral distance between the ground-truth and the generated signal in the frequency domain using:

$$LSD(M, \hat{M}) = \frac{1}{L} \sum_{l=1}^L \sqrt{\frac{1}{F} \sum_{f=1}^F |M(l, f) - \hat{M}(l, f)|^2}, \quad (4)$$

where $M(l, f)$ and $\hat{M}(l, f)$ are the ground truth and the estimated log-power spectra at l -th time step ($l = 1, \dots, L$) and f -th frequency bin ($k = 1, \dots, F$), respectively.

4.1. The Importance of Phase

To solely investigate the importance of phase in the audio SR task, we compare FLIP (i.e., phase flipping), GLA, and MelGAN methods with the assumption that we have knowledge of full-band magnitudes and LFC phase. Given the full-band magnitudes (0 – 8kHz) and LFC phase (0 – 4kHz), the task here is to generate phase information of HFC such that HR waveform can be reconstructed. For the FLIP method, the high-frequency phase is produced by flipping the phase of LFC and adding a negative sign. For the GLA method, as introduced in Sec. 2.2, a modified version of GLA is used to maintain both the magnitudes and phase of LFC through the iteration process. For the MelGAN method, HFC phase is extracted and complemented by the give magnitudes as well as LFC information in order to obtain the final HR output.

The comparison results of the three phase strategies are presented in Table 1. It shows that both the MelGAN and GLA method significantly outperforms the FLIP method used

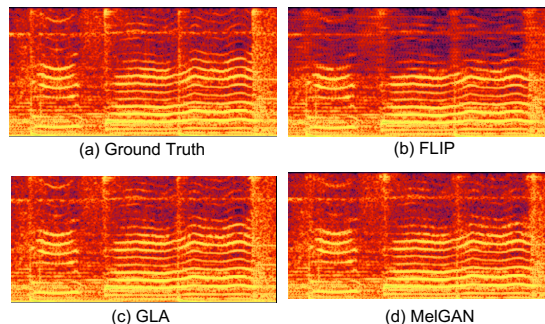


Figure 2: Spectrogram visualization with different phase strategies: Ground truth spectrogram (Top-left), Phase flipping (Top-right), GLA (Bottom-left), and MelGAN (Bottom-right).




Figure 3: Examples of waveform plotting: Ground truth HR waveform (Top), LR waveform (Middle), Waveform with GLA reconstruction (Bottom)

Table 1: Comparisons between different phase strategies.

Method	LSD HF (dB)	LSD Full (dB)	SNR (dB)
LR(input)	65.78	52.92	18.34
FLIP[6, 9]	9.56	7.63	17.95
GLA	5.62	4.46	15.39
MelGAN	6.61	5.27	15.86

in [6, 9] in terms of the LSD metric. An example of the generated spectrograms is shown in Fig. 2 which indicates that FLIP method weakens the energy of HFC due to the poor consistency between magnitudes and the produced phase, while GLA and MelGAN method generate magnitudes close to the ground truth.

It should be noted that the original LR input yields a better SNR value than any other methods used in the comparison. To explain it, we plot an example of ground-truth HR waveform,

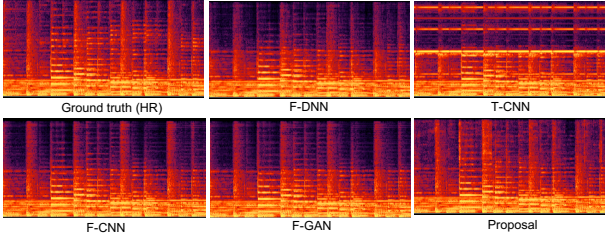


Figure 4: Visualization of predicted spectrograms by our audio SR approach and other existing methods.

along with its LR waveform and the reconstructed waveform by GLA as shown in Fig. 3. It presents that bandwidth expansion can effectively generate finer details that are missing in the LR waveform. These details correspond to audible improvements in the high-frequency range. However, the artifacts introduced by the expansion process may severely impact the SNR value. Therefore, the SNR metric may be not suitable for the frequency-domain music SR task. For this reason, we choose the LSD metric alone in the following experiment.

4.2. Comparison Results of SR Audio

To evaluate the effectiveness of the proposed joint framework, we implement our approach and compare it with the following state-of-the-art methods:

- F-DNN: Fully-connected network implementation in the frequency domain (Li et al. [6])
- T-CNN: CNN implementation in the time domain (Kuleshov et al. [7])
- F-CNN: CNN (generator network of GAN) implementation in the frequency domain (Eskimez et al. [9])
- F-GAN: GAN implementation in the frequency domain with phase flipping strategy (Eskimez et al. [9])
- Proposed: GAN implementation in the frequency domain for HB magnitude estimation and MelGAN implementation for HB phase estimation.

Table 2: Comparisons between our audio SR approach and other existing methods.

Method	LSD HF (dB)	LSD Full (dB)
LR(input)	65.78	52.92
F-DNN[6]	11.95	10.03
T-CNN[7]	17.08	14.34
F-CNN[9]	12.13	9.70
F-GAN[9]	11.84	9.46
Proposed	9.32	7.44

Fig. 4 plots the predicted spectrograms for each method. Table 2 presents the comparison results regarding the LSD on high-frequency and full range (denoted as HF and Full). It suggests that our proposed method obviously improves LSD. We achieve the best results among all other existing methods. It should be noted that T-CNN performs much worse than the F-DNN method. This conflicts with the conclusion in [7] which conducted the evaluation with speech samples. We believe this is because of the complex nature of polyphonic music, which

contains overlapping spectrograms of vocals and various instruments instead of monophonic speech. Thus it is difficult to indirectly expand the frequency bandwidth through temporal interpolation.

We also evaluate the three strategies of phase estimation. This is done by implementing the introduced frequency-domain approaches and replacing FLIP methods with GLA and MelGAN, respectively. In detail, the following methods are evaluated:

- F-DNN+GLA: Fully-connected network implementation for magnitude prediction (Li et al.[6]) and GLA phase method for phase estimation
- F-DNN+MelGAN: Fully-connected network implementation for magnitude prediction (Li et al.[6]) with MelGAN phase method for phase estimation
- F-CNN+GLA: CNN implementation[9] for magnitude prediction with GLA method for phase estimation
- F-CNN+MelGAN: CNN implementation[9] for magnitude prediction with MelGAN method for phase estimation
- F-GAN+GLA: GAN implementation[9] for magnitude prediction with GLA method for phase estimation
- F-GAN+MelGAN: GAN implementation[9] for magnitude prediction with MelGAN method for phase estimation.

Table 3: Comparisons between audio SR approaches with different phase strategies.

Method	LSD HF (dB)	LSD Full (dB)
F-DNN+FLIP[6]	11.95	10.03
F-DNN+GLA	10.11	8.63
F-DNN+MelGAN	9.63	8.26
F-CNN+FLIP[9]	12.13	9.70
F-CNN+GLA	9.67	7.72
F-CNN+MelGAN	9.58	7.65
F-GAN+FLIP[9]	11.84	9.46
F-GAN+GLA	9.41	7.51
F-GAN+MelGAN(proposed)	9.32	7.44

The LSD results shown in Table 3 suggests that the MelGAN method for phase estimation achieves the best result. Both the MelGAN and GLA methods consistently improve the LSD results over the FLIP method used in [6, 9]. Therefore, it is valid that a more accurate phase prediction can effectively solve the music SR tasks.

5. Conclusions

In this paper, we investigated the importance of phase in music SR tasks and validated that phase plays an important role in the resulted audio quality. Moreover, we proposed a joint framework which combines a GAN-based mapping method for predicting magnitudes and a MelGAN-based method for predicting the phase information of the high frequency components. Our experimental results show that the proposed approach achieves significantly the best performance comparing with the other state-of-the-art methods in terms of the LSD metric.

6. References

- [1] M. Nilsson and W. B. Kleijn, "Avoiding over-estimation in bandwidth extension of telephony speech," in *2001 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No. 01CH37221)*, vol. 2. IEEE, 2001, pp. 869–872.
- [2] C. Liu, Q.-J. Fu, and S. S. Narayanan, "Effect of bandwidth extension to telephone speech recognition in cochlear implant users," *The Journal of the Acoustical Society of America*, vol. 125, no. 2, pp. EL77–EL83, 2009.
- [3] Y. Nakatoh, M. Tsushima, and T. Norimatsu, "Generation of broadband speech from narrowband speech based on linear mapping," *Electronics and Communications in Japan (Part II: Electronics)*, vol. 85, no. 8, pp. 44–53, 2002.
- [4] Kun-Youl Park and Hyung Soon Kim, "Narrowband to wideband conversion of speech using gmm based transformation," in *2000 IEEE International Conference on Acoustics, Speech, and Signal Processing. Proceedings (Cat. No.00CH37100)*, vol. 3, 2000, pp. 1843–1846 vol.3.
- [5] P. Jax and P. Vary, "Artificial bandwidth extension of speech signals using mmse estimation based on a hidden markov model," in *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings.(ICASSP'03)*, vol. 1. IEEE, 2003, pp. I–I.
- [6] K. Li and C.-H. Lee, "A deep neural network approach to speech bandwidth expansion," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 4395–4399.
- [7] V. Kuleshov, S. Z. Enam, and S. Ermon, "Audio super resolution using neural networks," *arXiv preprint arXiv:1708.00853*, 2017.
- [8] S. E. Eskimez and K. Koishida, "Speech super resolution generative adversarial network," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 3717–3721.
- [9] S. E. Eskimez, K. Koishida, and Z. Duan, "Adversarial training for speech super-resolution," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 347–358, 2019.
- [10] C. Dong, C. C. Loy, K. He, and X. Tang, "Image super-resolution using deep convolutional networks," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 2, pp. 295–307, 2015.
- [11] J. Kim, J. Kwon Lee, and K. Mu Lee, "Accurate image super-resolution using very deep convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1646–1654.
- [12] K. Paliwal, K. Wójcicki, and B. Shannon, "The importance of phase in speech enhancement," *speech communication*, vol. 53, no. 4, pp. 465–494, 2011.
- [13] P. Mowlae, R. Saeidi, and Y. Stylianou, "Advances in phase-aware signal processing in speech communication," *Speech communication*, vol. 81, pp. 1–29, 2016.
- [14] Y. Ai and Z.-H. Ling, "A neural vocoder with hierarchical generation of amplitude and phase spectra for statistical parametric speech synthesis," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 839–851, 2020.
- [15] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, "Melgan: Generative adversarial networks for conditional waveform synthesis," in *Advances in Neural Information Processing Systems*, 2019, pp. 14 881–14 892.
- [16] D. Griffin and J. Lim, "Signal estimation from modified short-time fourier transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 2, pp. 236–243, 1984.
- [17] N. Perraudin, P. Balazs, and P. L. Søndergaard, "A fast griffin-lim algorithm," in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*. IEEE, 2013, pp. 1–4.
- [18] Y. Masuyama, K. Yatabe, Y. Koizumi, Y. Oikawa, and N. Harada, "Deep griffin-lim iteration," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 61–65.
- [19] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *9th ISCA Speech Synthesis Workshop*, 2016, pp. 125–125.
- [20] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International Conference on Machine Learning*, 2017, pp. 214–223.
- [21] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang, "Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 1874–1883.