



Neural Vocoder is All You Need for Speech Super-resolution

Haohe Liu¹, Woosung Choi², Xubo Liu¹, Qiuqiang Kong³, Qiao Tian³, DeLiang Wang⁴

¹Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey, UK

²Centre for Digital Music, Queen Mary University of London, UK

³Speech, Audio, and Music Intelligence (SAMI) Group, ByteDance, China

⁴Department of Computer Science and Engineering, The Ohio State University, USA

hl01486@surrey.ac.uk, wschoi.cs@gmail.com, dwang@cse.ohio-state.edu

Abstract

Speech super-resolution (SR) is a task to increase speech sampling rate by generating high-frequency components. Existing speech SR methods are trained in constrained experimental settings, such as a fixed upsampling ratio. These strong constraints can potentially lead to poor generalization ability in mismatched real-world cases. In this paper, we propose a neural vocoder based speech super-resolution method (NVSR) that can handle a variety of input resolution and upsampling ratios. NVSR consists of a mel-bandwidth extension module, a neural vocoder module, and a post-processing module. Our proposed system achieves state-of-the-art results on the VCTK multi-speaker benchmark. On 44.1 kHz target resolution, NVSR outperforms WSRGlow and Nu-wave by 8% and 37% respectively on log spectral distance and achieves a significantly better perceptual quality. We also demonstrate that prior knowledge in the pre-trained vocoder is crucial for speech SR by performing mel-bandwidth extension with a simple replication-padding method. Samples can be found at <https://haoheliu.github.io/nvsr>.

Index Terms: neural vocoder, speech super resolution, bandwidth extension, deep learning, flexible resolution

1. Introduction

Speech super-resolution (SR) aims to improve speech fidelity by predicting the high-resolution (HR) speech given the low-resolution (LR) speech. Low-resolution speech is common for reasons such as compression and low sampling rate. From a spectrogram point of view, speech SR is also equally referred to as bandwidth extension (BWE) [1]. SR is an important technique for real-world applications, including speech quality enhancement [2] and text-to-speech synthesis [3].

A lot of early studies [4, 5] break SR into spectral envelop estimation and excitation generation from LR. At that time, the direct mapping from LR to HR is not widely explored since the dimension of HR is relatively high. Later, SR methods based on deep neural network [6, 7, 8, 9] show better subjective quality than traditional methods. Most neural network based methods study the 8 kHz to 16 kHz or 4 kHz to 16 kHz sampling rate upsample problems, such as AECNN [10] and TFNet [11]. Recently, Nu-wave [12] and WSRGlow [13] have explored the higher 48 kHz target resolution, where the input sampling rate is usually 12 kHz, 16 kHz, or 24 kHz.

Existing speech SR studies are usually performed under controlled experimental settings. For example, the input resolution and bandwidth in data simulations is always fixed value during training and evaluation. However, in real-world scenarios, the speech SR applications generally require the capability of handling diverse settings such as different input resolutions and bandwidths. These mismatches potentially lead to degraded

speech SR performance. Also, the demand for high upsampling-ratio (UPR) SR is common on low-quality historical recordings, where UPR stands for target bandwidth divided by input bandwidth. Nevertheless, less attention has been given to the high UPR cases. High UPR speech SR is challenging because it need the model to predict more low-frequency harmonic structures. To our knowledge, most speech SR methods only experiment on the UPR around 3, 4, and 6 [12, 13, 14].

To alleviate the difficulties in addressing flexible input resolution and high UPR mentioned above, we propose a neural-vocoder-based approach (NVSR) for speech SR. We train NVSR in two separate stages: (1) an HR mel spectrogram prediction stage; (2) a vocoder waveform synthesis and post-processing stage. Different from predicting HR in the linear frequency scale as in previous studies [11, 15], HR prediction on a low-dimensional mel frequency scale is more tractable, especially in high UPR cases. Neural Vocoder [16, 17] is the model that upsample a low dimensional feature (e.g., mel spectrogram) to high-resolution waveforms. This coarse to fine generation process is similar to speech SR. Thus the prior knowledge on speech provided by the neural vocoder, such as phoneme structures, would be useful for speech SR.

To our knowledge, we are the first attempt that tackles the flexible input resolution and high UPR problems for speech SR. Our proposed NVSR can achieve a UPR up to 22.05 (from 2 kHz to 44.1 kHz). NVSR achieves state-of-the-art result on the VCTK Multi-speaker benchmark under multiple experimental settings, outperforming previous methods by a large margin. We observe that even replacing neural network methods with a simple replication-padding method without learning, NVSR can still outperform existing speech SR methods. This demonstrates the superior performance offered by the neural vocoder for speech SR. Unlike most previous studies, we mainly focus on the multi-speaker setup, which is more challenging than the single-speaker setting studied in other literatures [7]. Our code and pre-trained models are open-sourced¹ to facilitate reproducibility.

The paper is organized as follows. Section 2 introduces the problem formulation of speech SR. Section 3 introduces the pipeline of NVSR. Section 4 discusses the experimental settings and analyses the results. Section 5 draw a final conclusion.

2. Problem Formulation

Given a discrete signal $\mathbf{x}_l = [x_i]_{i=1,2,\dots,T \cdot l}$ sampled at sampling rate l , speech SR system estimates signal $\mathbf{y}_h = [y_i]_{i=1,2,\dots,T \cdot h}$, where T is the length in seconds and $h > l$. According to the Nyquist theory, the highest bandwidth of \mathbf{x}_l and \mathbf{y}_h are $l/2$ (Hz) and $h/2$ (Hz), respectively. So, \mathbf{x}_l does not contain the frequency

¹https://github.com/haoheliu/ssr_eval

information between $h/2 - l/2$ (Hz), which can be considered as the generation target of the SR task.

Most frequency domain methods [7, 11] perform SR by predicting HR spectrogram from LR spectrogram, and transformed to waveform using the HR spectrogram prediction. This process can be formulated as

$$\mathbf{y}_h = F^{-1}(G(F(\mathbf{x}_h))), \quad (1)$$

where F and F^{-1} are the time to frequency spectrogram transformation and its reverse transformation respectively. \mathbf{x}_h is the upsampled version of \mathbf{x}_l , and function $G(\cdot)$ takes band-limited spectrogram as input and output full band spectrogram. Here \mathbf{x}_l is upsampled to ensure the system input and output have the same shape and the model design is not restricted by the ratio h/l . In most studies, F and F^{-1} are short-time-fourier-transform (STFT) and inverse STFT (iSTFT), respectively.

3. Methodology

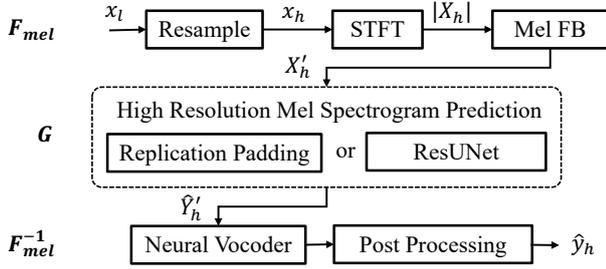


Figure 1: Overview of the NVSR pipeline.

Figure 1 shows the overall pipeline of NVSR including mel transform F_{mel} , high resolution mel spectrogram prediction with G , and inverse mel transform F_{mel}^{-1} . Similar to Equation 1, this process can be written as

$$\hat{\mathbf{y}}_h = F_{\text{mel}}^{-1}(G(F_{\text{mel}}(\mathbf{x}_h))) \approx V_{\phi}(G(F_{\text{mel}}(\mathbf{x}_h))). \quad (2)$$

Here G aims to predict the HR mel spectrogram with the LR input, as shown in Equation 3.

$$G: \mathbf{X}_{\text{mel}} \rightarrow \mathbf{Y}_{\text{mel}}, \quad (3)$$

where \mathbf{X}_{mel} and \mathbf{Y}_{mel} stand for the mel spectrogram of \mathbf{x} and \mathbf{y} . \mathbf{X}_{mel} is calculated by $F_{\text{mel}}(\mathbf{x}) = |\mathbf{X}| \mathbf{W}$, in which \mathbf{X} is the STFT of \mathbf{x} and \mathbf{W} is a set of mel filter banks.

Note that the inverse mel transform is not mathematically solvable. Thus, NVSR models the F_{mel}^{-1} with a neural vocoder V_{ϕ} . We show that V_{ϕ} achieves almost the same perceptual quality as the ideal inverse transform. For mel spectrogram prediction, we first introduce a ResUNet based method. Then we introduce a non-parametric replication-padding method, which mainly exploits the prior knowledge learned by the vocoder.

3.1. Mel Spectrogram Prediction

3.1.1. ResUNet-based Method

We use ResUNet [18, 19] to model the $G(\cdot)$ in Equation 3 and estimate \mathbf{Y}_{mel} . This process can be written as

$$\hat{\mathbf{Y}}_{\text{mel}} = \text{ResUNet}(\mathbf{X}_{\text{mel}}) \odot (\mathbf{X}_{\text{mel}} + \epsilon), \quad (4)$$

where \odot means element-wise multiplication, and ϵ is a small constant to avoid the zero values in \mathbf{X}_{mel} . As shown in Figure 2,

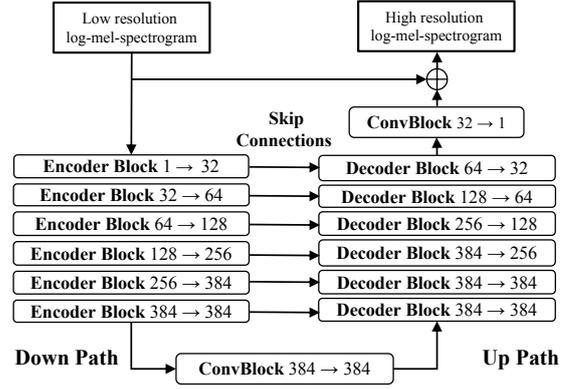


Figure 2: Predicting the higher frequencies of mel spectrogram with ResUNet. The numbers in each block means input and output channels.

ResUNet consists of six encoder and six decoder blocks. There are skip connections between encoder and decoder blocks at the same level. Both encoder and decoder blocks share the same structure, which is four convolutional blocks (ConvBlock). Each ConvBlock consists of two convolution operations with batch normalization [20], and leakyReLU [21] activation. The encoder blocks apply average pooling for downsampling. The decoder blocks apply transpose convolution for upsampling. After the last decoder block, we apply one more ConvBlock to estimate the output residual, which is added with the input to form the final estimation. Note that we use addition in Figure 2 instead of multiply in Equation 4 because we use log scale during our implementation. ResUNet is optimized using the mean absolute error (MAE) loss between the estimated mel spectrogram $\hat{\mathbf{Y}}_{\text{mel}}$ and the target mel spectrogram \mathbf{Y}_{mel} .

$$\text{loss} = \frac{1}{TF} \sum^T \sum^F |\mathbf{Y}_{\text{mel}} - \hat{\mathbf{Y}}_{\text{mel}}| \quad (5)$$

3.1.2. Replication padding-based Method

Replication padding-based methods do not use any training data for mel spectrogram prediction. Instead, it directly copies the energy of the cutoff frequency to the higher frequency bands. We use this method to demonstrate the importance of the neural vocoder. Specifically, we first search the cutoff frequency c based on the energy of \mathbf{X}_{mel} . Then we construct the frequency cutoff mask $\mathbf{M}_{T \times F}$ with binary values, where the frequency indexes greater than c are all zeros and others are all ones. $\mathbf{M}_{T \times F}$ is used later to select the higher/lower frequencies bands. Finally, we repeat the energy of the cutoff frequency in the high-frequency bands and add it with the LR mel spectrogram. The output is the HR mel spectrogram estimation $\hat{\mathbf{Y}}_{\text{mel}}$. This process can be written as Equation 6

$$\hat{\mathbf{Y}}_{\text{mel}} = \mathbf{M} \odot \mathbf{X}_{\text{mel}} + |1 - \mathbf{M}| \odot (\mathbf{X}_{\text{mel}}[:, c] \cdot \mathbf{1}_{1 \times F}), \quad (6)$$

where $\mathbf{1}_{1 \times F}$ is the all-one matrix, and $\mathbf{X}_{\text{mel}}[:, c]$ is the energy distribution across time at the cutoff frequency c .

3.2. Neural Vocoder

We choose TFGAN [17] as the neural vocoder V_{ϕ} . TFGAN can directly upsample mel spectrogram into waveform with transpose convolution and one-dimensional convolutional neural networks. By incorporating multi-resolution losses and discriminators in both the time and frequency domains, TFGAN can

achieve state-of-the-art performance on vocoding. In order to achieve speaker-independent, TFGAN is trained with a large corpus of more than one thousand speakers. As reported in [22], the mean opinion score (MOS) of the open-sourced TFGAN² is 3.74, with the ground truth MOS score at 3.95.

3.3. Post Processing

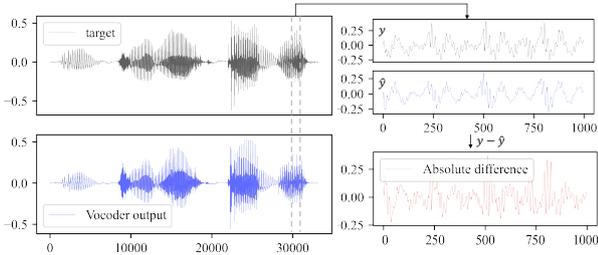


Figure 3: Visualization of the difference between the vocoder output and target waveforms.

Even though the output of TFGAN can achieve good perceptual quality, its output still has some trivial differences from the target, which will result in a poor results in the objective evaluation. For example, as shown in Figure 3, the vocoder output looks the same as the target in the waveform, but their absolute difference is still substantial. This is a common evaluation problem in generative model [23, 24, 25]. To alleviate the problem in evaluation, we propose to perform a lower-frequencies replacement (LFR) operation on the vocoder output as post-processing. Since the original lower-frequencies information in the input usually does not need to change, we replace the low-frequency bands in the vocoder output with the original input. Experimental result shows this can improve the metrics score.

4. Experiment

4.1. Dataset

We build the train and test sets using VCTK (version 0.92) [26], a multi-speaker English corpus that contains 110 speakers with different accents. We split it into a part for train (VCTK-Train) and a part for test (VCTK-Test). Following the data preparation strategy of [12], only the *mic1* microphone data is used for experiments, and *p280* and *p315* are omitted for the technical issues. For the remaining 108 speakers, the last eight speakers³ constitute the VCTK-Test. The remaining 100 speakers are defined as VCTK-Train. For the training of NVSR, all the utterances are resampled at the 44.1 kHz sample rate.

We follow the LR simulation process in [7, 10]. Given a target speech y_h in VCTK-Train, to obtain the low-resolution x_l , we first convolve y_h with an order eight *Chebyshev type I* lowpass filter with cutoff frequency $l/2$. Then we subsample the signal to l sample rate using polyphase filtering. We evaluate the performance of each system on different low sampling rates l .

4.2. Evaluation Metrics

We use Log-spectral distance (LSD) as the evaluation metrics following [10, 12, 23]. For target signal y_h and output estimation \hat{y}_h , LSD can be computed as Equation 7, where \mathbf{Y} and $\hat{\mathbf{Y}}$ stand

²<https://github.com/haoheliu/voicefixer>

³p360, p361, p362, p363, p364, p374, p376, s5

Table 1: LSD on VCTK-Test with 44.1 kHz target sampling rate. Model with * means it is trained with a fixed input resolution.

Input sampling rate (kHz)	4	8	12	24
*Nu-wave [12] (3.0M×4)	1.42	1.42	1.40	1.22
*WSRGlow [13] (229.9M×4)	1.12	0.98	0.87	0.79
NVSR-Pad (33.9M)	1.54	1.46	1.18	0.91
NVSR (99.0M)	0.98	0.91	0.85	0.70

Table 2: LSD on VCTK-Test with 16 kHz target sampling rate.

Input sampling rate (kHz)	2	4	8
*AudioUNet [7] (70.9M×3)	N/A	1.40	1.32
*TFNet [11] (58.8M×3)	N/A	N/A	1.36
*AECNN [10] (10.2M×3)	N/A	0.95	0.88
NVSR-Pad (33.9M)	2.51	2.24	1.89
NVSR (99.0M)	1.07	0.95	0.78

for the magnitude spectrogram of y_h and \hat{y}_h . A lower LSD value indicates better SR performance. We report the mean LSD of the VCTK-Test as the final score of a system.

$$\text{LSD}(\mathbf{Y}, \hat{\mathbf{Y}}) = \frac{1}{T} \sum_{t=1}^T \sqrt{\frac{1}{F} \sum_{f=1}^F \log_{10} \left(\frac{\mathbf{Y}(f, t)^2}{\hat{\mathbf{Y}}(f, t)^2} \right)^2} \quad (7)$$

4.3. Baselines

We include several state-of-the-art methods as baselines. Since NVSR uses a target sampling rate at 44.1 kHz, the baseline with 48 kHz output is downsampled to 44.1 kHz for a fair comparison. Similarly, the output of NVSR can also be downsampled to compare with other baselines with a smaller target sampling rate (e.g., 16 kHz). For the 48 kHz SR model, we reproduced Nu-wave⁴ and WSRGlow⁵ with their open-sourced code and default settings. We train each Nu-wave for 200 epochs and each WSRGlow for 100k iterations following [13]. For 16 kHz target sampling rate, we compared NVSR with the result reported in AECNN [10], TFNet [11], and AudioUNet [7].

4.4. Training Details

The training high-resolution and low-resolution data pairs are built on the fly. we uniformly randomize the cutoff frequency $l/2$ of the training data in [1, 16] kHz. We use an Adam optimizer [27] with $\beta_1 = 0.5, \beta_2 = 0.999$ and a 3×10^{-4} learning rate to optimize the ResUNet. We apply the first 1000 steps as the warmup phase, during which the learning rate grows linearly from 0 to 3×10^{-4} . Then the learning rate is decayed by 0.85 every epoch. We stop the training after 60 epochs. For all the STFT and iSTFT, we use the Hanning window with a window length of 2048 and a hop length of 441. We use 128 mel filterbanks to calculate the mel spectrogram. We use eight Nvidia-V100-32GB GPUs to train the ResUNet, which takes about 3.9 hours.

4.5. Result and Discussion

Table 1 and Table 2 compare NVSR with other state-of-the-art models on both 44.1 kHz and 16 kHz target sampling rates. We use * to denote a model is trained with a fixed input resolution.

⁴<https://github.com/minds-lab-ai/nuwave>

⁵<https://github.com/zkx06111/WSRGlow>

⁶<https://www.youtube.com/watch?v=A7Q7fmIuKdM>

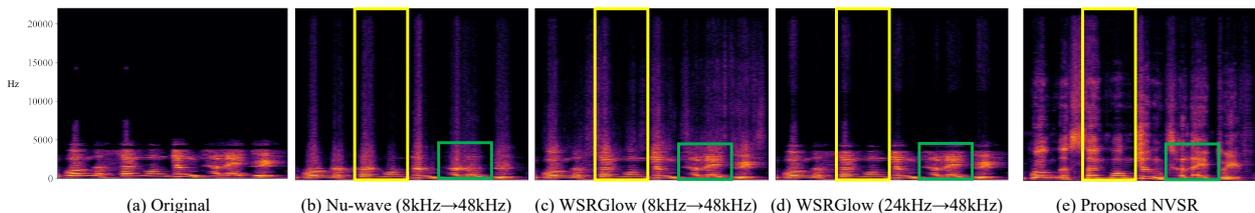


Figure 4: Robustness test with a low-resolution audio (a) from an old movie⁶. No ground truth is available.

Table 3: Evaluation results on VCTK-Test with different input sampling rate settings. Note that NVSR can also work on other input sampling-rates between 2 kHz to 32 kHz.

Input sampling rate (kHz)	2	4	8	12	16	24	32	AVG
Unprocessed	5.69	5.50	5.15	4.85	4.54	3.84	2.95	4.65
Groundtruth-mel	0.87	0.85	0.81	0.78	0.74	0.66	0.59	0.76
NVSR-Pad	1.55	1.54	1.46	1.18	1.11	0.91	0.76	1.21
NVSR	1.04	0.98	0.91	0.85	0.79	0.70	0.60	0.84
w/o post proc.	1.06	1.02	0.98	0.96	0.93	0.91	0.89	0.96
w/o mel pred.	2.37	2.37	2.16	1.95	1.77	1.48	1.13	1.89

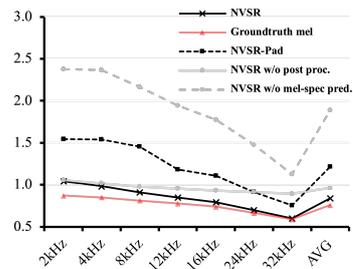


Figure 5: Visualization of Table 3

For example, we report four different sampling rate settings in Table 1 after training four different WSRGlow. Meanwhile, a single NVSR here can handle four different settings.

In Table 1, NVSR achieves an average LSD of 0.86, outperforming WSRGlow’s 0.94 and Nu-wave’s 1.37 by 0.08 and 0.51. The total parameter number of NVSR is 99.0 million (M), of which 33.9M comes from the vocoder. While WSRGlow is the largest model with 229.9M parameters on each setting. This demonstrates the state-of-the-art performance and the efficiency of NVSR. NVSR-Pad is the replication padding-based NVSR. NVSR-Pad can already achieve an average LSD of 1.27, outperforming Nu-wave on 12 kHz and 24 kHz. This proves the prior knowledge in vocoder can map the constant energy in the padded higher-frequency bands into meaningful energy distribution on the spectrogram. In Table 2, NVSR achieves the best performance on the 8 kHz input sampling rate. At 4 kHz, NVSR has a similar LSD score as AECNN. NVSR is the first to try 2 kHz to 16 kHz SR and achieve an LSD of 1.07. Note that the NVSR in Table 2 and Table 1 is the same model.

In Figure 4, we visualize the outputs of different methods on an old movie recording. The original speech is in low resolution, with the highest frequency around 4 kHz. As shown in the yellow boxes, the output of NVSR (Figure 4e) contains properly shaped harmonic components, while other methods (Figure 4b, and c) mainly fill in the high frequencies with stochastic components. WSRGlow trained with 24 kHz input resolution (Figure 4d) fails to predict because of the mismatch between the input resolution (8 kHz) and the resolution of its training data (24 kHz). Meanwhile, we found NVSR is the only model that can repair the distribution in lower-frequency bands (green boxes). Note that here we do not use post-processing in NVSR because the original input is of low quality.

In Table 3, Groundtruth-mel stands for the system using ground truth mel spectrogram directly as the input to the vocoder. This equals the model performance when the mel spectrogram prediction module works flawlessly. This experiment marks the ideally best performance of the NVSR system, with an average

LSD of 0.76. We also tried to remove the post-processing LFR operation in NVSR, which degrade the average performance from 0.84 to 0.96. If we do not perform mel spectrogram prediction on NVSR, the average LSD is 1.89, which is still better than the unprocessed audio with a 4.65 LSD. This result means even without the mel prediction stage, the vocoder can still improve the metrics performance on the evaluation set.

To better understand the result, we visualize Table 3 in Figure 5. The red line is the theoretical best performance of our proposed system using pre-trained TFGAN. The result of NVSR is close to red line, meaning our mel spectrogram prediction module works well. Figure 5 also shows the performance of NVSR will degrade without post-processing, especially in high input sampling rate cases. On 24 kHz input sampling rate, the performance of the NVSR-Pad is even on par with the NVSR without post-processing. Note that the proposed post-processing operation can also be applied to other methods like Nu-wave and is very likely to increase the model performance on the test set.

There are certain limitations of our proposed method. The performance of NVSR largely relies on the neural vocoder, which may become the bottleneck of NVSR. A possible future solution can be fine-tuning two stages in an end-to-end manner or designing an additional post-processing model for refinements. Besides, when extending NVSR to other kinds of sound, like music, vocoder may not as easy to train as the speech vocoder.

5. Conclusion

This paper presents a novel and powerful neural vocoder-based speech super-resolution model NVSR. It shows strong performance across a wide range of input sampling rates between 2 kHz to 32 kHz. On the VCTK Multi-Speaker SR benchmark, NVSR outperforms the state-of-the-art models trained with different input resolution settings both on 16 kHz and 44.1 kHz evaluation sets. We also demonstrate that prior knowledge in vocoder is crucial to speech SR using a simple replication padding-based mel spectrogram prediction method.

6. References

- [1] P. Ekstrand, "Bandwidth extension of audio signals by spectral band replication," in *Proceedings of the IEEE Benelux Workshop on Model Based Processing and Coding of Audio*. Citeseer, 2002.
- [2] S. Chennoukh, A. Gerrits, G. Miet, and R. Sluijter, "Speech enhancement via frequency bandwidth extension using line spectral frequencies," in *Proceedings of the IEEE Conference on Acoustics, Speech, and Signal Processing*, 2001, pp. 665–668.
- [3] K. Nakamura, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "A mel-cepstral analysis technique restoring high frequency components from low-sampling-rate speech," in *Conference of the International Speech Communication Association*, 2014.
- [4] Y. Nakatoh, M. Tsushima, and T. Norimatsu, "Generation of broadband speech from narrowband speech based on linear mapping," *Electronics and Communications in Japan*, vol. 85, no. 8, pp. 44–53, 2002.
- [5] J. Kontio, L. Laaksonen, and P. Alku, "Neural network-based artificial bandwidth expansion of speech," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 3, pp. 873–881, 2007.
- [6] K. Li and C.-H. Lee, "A deep neural network approach to speech bandwidth expansion," in *Proceedings of the IEEE Conference on Acoustics, Speech, and Signal Processing*, 2015, pp. 4395–4399.
- [7] V. Kuleshov, S. Z. Enam, and S. Ermon, "Audio super resolution using neural networks," *arXiv:1708.00853*, 2017.
- [8] A. Gupta, B. Shillingford, Y. Assael, and T. C. Walters, "Speech bandwidth extension with wavenet," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics*, 2019, pp. 205–208.
- [9] S. E. Eskimez and K. Koishida, "Speech super resolution generative adversarial network," in *Proceedings of the IEEE Conference on Acoustics, Speech, and Signal Processing*, 2019, pp. 3717–3721.
- [10] H. Wang and D. Wang, "Towards robust speech super-resolution," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 2058–2066, 2021.
- [11] T. Y. Lim, R. A. Yeh, Y. Xu, M. N. Do, and M. Hasegawa-Johnson, "Time-frequency networks for audio super-resolution," in *Proceedings of the IEEE Conference on Acoustics, Speech, and Signal Processing*, 2018, pp. 646–650.
- [12] J. Lee and S. Han, "NU-wave: A diffusion probabilistic model for neural audio upsampling," *arXiv:2104.02321*, 2021.
- [13] K. Zhang, Y. Ren, C. Xu, and Z. Zhao, "WSRglow: A glow-based waveform generative model for audio super-resolution," *arXiv:2106.08507*, 2021.
- [14] J. Kim, Y. Lee, S. Hong, and J. Ok, "Learning continuous representation of audio for arbitrary scale super resolution," *arXiv:2111.00195*, 2021.
- [15] S. Hu, B. Zhang, B. Liang, E. Zhao, and S. Lui, "Phase-aware music super-resolution using generative adversarial networks," *arXiv:2010.04506*, 2020.
- [16] J. Kong, J. Kim, and J. Bae, "HiFi-GAN: Generative adversarial networks for efficient and high fidelity speech synthesis," *arXiv:2010.05646*, 2020.
- [17] Q. Tian, Y. Chen, Z. Zhang, H. Lu, L. Chen, L. Xie, and S. Liu, "TFGAN: Time and frequency domain based generative adversarial network for high-fidelity speech synthesis," *arXiv:2011.12206*, 2020.
- [18] Q. Kong, Y. Cao, H. Liu, K. Choi, and Y. Wang, "Decoupling magnitude and phase estimation with deep resnet for music source separation," in *The International Society for Music Information Retrieval*, 2021.
- [19] H. Liu, Q. Kong, and J. Liu, "CWS-PResUNet: Music source separation with channel-wise subband phase-aware resnet," *arXiv:2112.04685*, 2021.
- [20] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Machine Learning*, 2015, pp. 448–456.
- [21] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.
- [22] H. Liu, Q. Kong, Q. Tian, Y. Zhao, D. Wang, C. Huang, and Y. Wang, "VoiceFixer: Toward general speech restoration with neural vocoder," *arXiv:2109.13731*, 2021.
- [23] R. Kumar, K. Kumar, V. Anand, Y. Bengio, and A. Courville, "NU-GAN: High resolution neural upsampling with GAN," *arXiv:2010.11362*, 2020.
- [24] Q. Liu, P. J. Jackson, and W. Wang, "A speech synthesis approach for high quality speech separation and generation," *IEEE Signal Processing Letters*, vol. 26, no. 12, pp. 1872–1876, 2019.
- [25] X. Liu, T. Iqbal, J. Zhao, Q. Huang, M. D. Plumbley, and W. Wang, "Conditional sound generation using neural discrete time-frequency representation learning," in *IEEE Workshop on Machine Learning for Signal Processing*, 2021, pp. 1–6.
- [26] J. Yamagishi, C. Veaux, K. MacDonald *et al.*, "CSTR VCTK corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2019.
- [27] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv:1412.6980*, 2014.