



RefineGAN: Universally Generating Waveform Better than Ground Truth with Highly Accurate Pitch and Intensity Responses

Shengyuan Xu¹, Wenxiao Zhao², Jing Guo³

timedomain.Inc

¹xushengyuan@timedomain.ai, ²sean.z@timedomain.ai, ³joe.g@timedomain.ai

Abstract

Most GAN(Generative Adversarial Network)-based approaches towards high-fidelity waveform generation heavily rely on discriminators to improve their performance. However, GAN methods introduce much uncertainty into the generation process and often result in mismatches of pitch and intensity, which is fatal when it comes to sensitive use cases such as singing voice synthesis(SVS).

To address this problem, we propose RefineGAN, a high-fidelity neural vocoder focused on the robustness, pitch and intensity accuracy, and high-speed full-band audio generation. We applied a pitch-guided refine architecture with a multi-scale spectrogram-based loss function to help stabilize the training process and maintain the robustness of the neural vocoder while using the GAN-based training method.

Audio generated using this method shows a better performance in subjective tests when compared with the ground-truth audio. This result shows that the fidelity is even improved during the waveform reconstruction by eliminating defects produced by recording procedures. Moreover, it shows that models trained on a specified type of data can perform on totally unseen language and unseen speaker identically well. Generated sample pairs are provided on <https://timedomain-tech.github.io/refinegan/>.

Index Terms: speech synthesis, vocoder, singing voice synthesis

1. Introduction

1.1. Conventional neural vocoders

Vocoders have been employed in various fields in counter with human voice recreation, such as text-to-speech(TTS), singing voice synthesis(SVS), and voice conversion. Moreover, vocoders are often the bottleneck when it comes to the voice quality. Recently a series of GAN-based methods[1, 2] have been employed in vocoders to synthesize more realistic waveforms while also reducing synthetic artifacts in the voices. And these non-autoregressive CNN-based designs have increased the generation speed to an acceptable level.

While most neural vocoders are focused on generating band-limited waveforms, a majority of replay devices natively support full-band audio. The quality-reduced audio delivered by those conventional vocoders are becoming increasingly unsuitable to be consumed. And audio quality is crucial when deep learning-based technology is applied to the workflow of content creating. Moreover, in neural network singing voice synthesis[3], audio quality is considered as an obstacle constantly delays the technology from entering the existing business.

1.2. Towards finest waveform generation

We designed a way to use pitch guided methods to diminish the difficulty of generating voice waveform with higher sample rates. It makes stable, high-resolution voice synthesis affordable when using full-band Mel-spectrogram and the target pitch as network inputs. This approach is designed to make the task of conversion from the frequency domain to the time domain easier for the neural networks by using some simple external calculation.

Conventional neural or non-neural[4] vocoders center on eliminating most of the noticeable flaws and feature losses during the restoration. However, in our work, by using several methods to level down the learning difficulties, the amount of synthetic artifacts are highly reduced. Those defects are not audible to a large percentage of the content consumers or even listeners with professional experience in the audio industries.

1.3. Related Works

Hifi-GAN[2] introduced a unique multi-scale and multi-period discriminator design to produce high-fidelity voices. Evaluations on it had shown better audio quality and improved voice naturalness. Also, UnivNet[5] presented a new discriminator design based on the idea of deciding on a linear spectrogram rather than on the waveform. This dramatically enhance the discriminator's performance and improve the audio quality in higher frequencies.

2. RefineGAN

2.1. Generator

As we call it RefineGAN, we use an encoder-decoder-like architecture to refine a signal(as we call it, speech template) constructed from pitch information into the final speech waveform.

The refinement process includes an encoder module to encode the speech template into an intermediate form to apply Mel-spectrogram information into it, and a decoder module to generate the waveform from this hidden form. The encoding and decoding processes are done by first downsampling the waveform into a shape the same as the Mel-spectrogram and then upsampling it back into the final waveform. The downsampling process is done by using convolution layers and corresponding transposed convolution layers for the upsampling process.

To make pitch information fully accessible, we also add a cross-connection mechanism between each block in the encoder and decoder parts, using an UNet[6]-like strategy. We employ the design of using parallel ResBlocks[7] with differed kernel sizes and dilation sizes, which is similar to the architecture of Hifi-GAN. Each of these ResBlocks contains three sub-blocks connected in series. Furthermore, in each sub-block, inputs are passed through 2 weight normalized convolution layers with LeakyReLU[8] pre-applied.

2.5.2. Loudness-based data augmentation

To make it possible for the model to generate a high-quality waveform no matter the loudness level of the utterance, we also applied a loudness-based data augmentation technique. This approach allows randomly adjusting the loudness level of each utterance within an acceptable range. We control the range by defining four parameters: p_{\min} , p_{\max} , which indicates the audio signal’s minimum and maximum peak value; r_{\min} , r_{\max} , which indicates the minimum and maximum rate of the loudness adjustments. We can randomly sample the target peak value :

$$p' \sim U[\max(p_{\min}, r_{\min}p), \min(p_{\max}, r_{\max}p)] \quad (3)$$

where p means the original peak value $\max(|y|)$.

And we can apply the gain to the audio signal by simply calculating $y' = \frac{yp'}{p}$. where y and y' represents the audio signal before and after the augmentation process.

In this case, we did a uniform random sampling on a linear scale. And we can train the model with a much greater loudness diversity and achieve a state where the model’s performance will not be affected by the loudness level.

2.6. Discriminator

We use two set of discriminators to bring a broader coverage of various features. One set of them is the same Multi-Period Discriminator we adopt from Hifi-GAN[2]. We use period parameters of 2,3,5,7,11 in the hope of being better suitable for 44100hz full-band generation. The other is the Multi-Resolution Discriminator we adopt from the UnviNet[5] and use the same three different sets of parameters as the author mentioned in the paper.

2.7. Training losses

2.7.1. Generator loss

For the generator, we add up 3 different loss elements we mentioned above to construct the final generator loss. And also add a weight λ to the \mathcal{L}_{mel} to help stable the training process:

$$\begin{aligned} \mathcal{L}_G = & \lambda \mathcal{L}_{\text{mel}}(\mathbf{y}, G(\mathbf{z}, \mathbf{c})) + \mathcal{L}_{\text{envelope}}(\mathbf{y}, G(\mathbf{z}, \mathbf{c})) \\ & + \frac{1}{n_{\text{MPD}}} \sum_{i=1}^{n_{\text{MPD}}} \mathbb{E}_{\mathbf{z}, \mathbf{c}} [\log(1 + \exp(-D_{\text{MPD}, i}(G(\mathbf{z}, \mathbf{c}))))] \\ & + \frac{1}{n_{\text{MRD}}} \sum_{i=1}^{n_{\text{MRD}}} \mathbb{E}_{\mathbf{z}, \mathbf{c}} [\log(1 + \exp(-D_{\text{MRD}, i}(M_i(G(\mathbf{z}, \mathbf{c})))))] \end{aligned} \quad (4)$$

2.7.2. Discriminator loss

Within both the Multi-Period Discriminator and the Multi-Resolution Discriminator, we use an average value of the output from modules using different parameters and a single Adam optimizer for both the discriminators. And as described below, the sum of the MPD’s and MRD’s output acts as the loss for back-propagation:

$$\begin{aligned} \mathcal{L}_{\text{MPD}} = & \frac{1}{n_{\text{MPD}}} \sum_{i=1}^{n_{\text{MPD}}} (\mathbb{E}_{\mathbf{y}} [\log(1 + \exp(-D_{\text{MPD}, i}(\mathbf{y})))]) \\ & + \mathbb{E}_{\mathbf{z}, \mathbf{c}} [\log(1 + \exp(D_{\text{MPD}, i}(G(\mathbf{z}, \mathbf{c})))))] \end{aligned} \quad (5)$$

$$\begin{aligned} \mathcal{L}_{\text{MRD}} = & \frac{1}{n_{\text{MRD}}} \sum_{i=1}^{n_{\text{MRD}}} (\mathbb{E}_{\mathbf{y}} [\log(1 + \exp(-D_{\text{MRD}, i}(M_i(\mathbf{y}))))]) \\ & + \mathbb{E}_{\mathbf{z}, \mathbf{c}} [\log(1 + \exp(D_{\text{MRD}, i}(M_i(G(\mathbf{z}, \mathbf{c})))))] \end{aligned} \quad (6)$$

$$\mathcal{L}_D = \mathcal{L}_{\text{MPD}} + \mathcal{L}_{\text{MRD}} \quad (7)$$

3. Experiments

3.1. Hyperparameters

We select 44100Hz as the sample rate to achieve the goal of full-band voice generation. Furthermore, to extract the Mel-spectrogram as the condition, we choose 128 as the number of the Mel channels. Then chose 2048 as both the FFT size and the windows size while 256 as the hop size.

In our model, we use each four layers in the encoder part and the decoder part. Between layers, we use downsampling rates of 2,2,8,8 on the encoding side and correspondingly 8,8,2,2 on the decoding side, and the kernel sizes of the these layers are all set to 2 times the rates. All the ResBlocks use the exact dilation sizes of 1,3 and 5, but parallel ResBlocks with different kernel sizes are only applied in the decoding side with kernel sizes of 3,7 and 11. In the encoding side, only one Res-Block with a kernel size of 7 are employed in each layer to reduce the model size.

3.2. Training

As we focused on Singing Voice Synthesis(SVS), we trained the model entirely using studio-grade quality singing recordings with a length of 59 hours from 12 different Chinese singers, the amounts of the data are shown in Table 1. In training, we randomly select slices from all the audio and calculate the Mel-spectrogram and the speech template after data augmentations.

Calculations of the conditions are all done simultaneously with the training process. With our data augmentation methods mentioned above, this approach formulates a great diversity of the training data and significantly strengthens the robustness of the model.

Table 1: Lengths and numbers of recordings of the singing voice dataset

	Singers	Length(Hours)	Recordings
Female Singers	8	41.68	2325
Male Singers	4	18.08	697
All	12	59.77	3022

4. Results

4.1. General Performance

We trained the model on 8 RTX3090 GPUs for 450k steps with a training time of approximately 130 hours using a batch size of 16. An on a single RTX3090 we can achieve a inference speed of approximately 110X real-time. And further experiments show that we can easily convert the whole model into FP16 (without using automatic mixed precision) in order to get benefits from optimized hardware with nearly no audible quality loss. Generated samples are available on our web page.

We also conducted a MOS test on 73 listeners. This test compares the subjective scores of ground-truth recordings, audio reconstructed by the model, and audio reconstructed by the

Table 2: MOS test result on speech and singing voice in three languages

	Griffin-Lim	Hifi-GAN	Univnet	RefineGAN (ours)	Ground Truth
Chinese speech	1.45 ± 0.05	3.19 ± 0.03	2.79 ± 0.04	4.33 ± 0.02	4.42 ± 0.02
Chinese singing	1.40 ± 0.05	2.71 ± 0.04	2.53 ± 0.04	4.34 ± 0.02	4.33 ± 0.03
Japanese speech	1.45 ± 0.05	3.14 ± 0.04	2.95 ± 0.04	4.29 ± 0.02	4.19 ± 0.03
Japanese singing	1.42 ± 0.05	2.76 ± 0.04	2.51 ± 0.04	4.25 ± 0.02	4.29 ± 0.03
English speech	1.61 ± 0.04	2.72 ± 0.04	2.46 ± 0.05	3.42 ± 0.04	4.06 ± 0.03

Griffin-Lim[10] algorithm. Each of the listeners scored 30 different 23 second-long utterances, and the result is shown in the Table 3.

Table 3: MOS test result on seen speakers

	Griffin-Lim	Predicted	Ground Truth
Male	1.93 ± 0.09	3.68 ± 0.09	3.67 ± 0.09
Female	2.15 ± 0.10	4.02 ± 0.09	3.97 ± 0.09
All	2.04 ± 0.07	3.85 ± 0.07	3.82 ± 0.07

4.2. Performance on unseen speakers, unseen languages, and speech signals

One of our goals is to achieve universal waveform regeneration, so we studied using this trained model to predict singing voices in unseen languages and with unseen speakers. Supplemental samples are also provided on the web page. After analysis on more samples, we can conclude that our model can perform on unseen data nearly identically well as on the seen data.

We also studied using the model trained with only singing voice recordings to reconstruct speech signals. We experimented on reconstructing speech in seen and unseen languages, and apart from Chinese[11], English[12], French[13] and Japanese[14], we also evaluated the performance on some under-resourced languages mentioned in [15]. In these provided pairs we can see that even when merely trained with Chinese singing data, the model can still perform well on speech signals in even unseen language.

5. Further Study

5.1. Training with a comprehensive dataset

Interested in the outperforming generalization capabilities shown in the experiments above, we further studied the model on a highly diverse dataset. Table 4 shows the composition of the dataset which includes studio-grade quality voices of different types(singing and speech) in several languages from over 150 speakers. And a mixed dataset like this usually presents enormous difficulties for conventional neural vocoders' convergence. We use the same hyperparameters as above and trained the model on the same setup for 450k steps with a batch size of 16.

Also we selected Hifi-GAN[2] and UnivNet[5] for comparison and use the official implements to train the models on the same dataset. We use the same hyperparameters provided in the implements and trained to convergence with the same 8-card setup, which takes nearly 150 hours each.

5.2. Results

We conducted a listening test on 67 listeners with the samples produced by these methods. and MOS results is showed be-

Table 4: Recording lengths of the extended dataset. *Including under-resourced languages and accents provided in [15]

Type	Length	Type	Length
Chinese singing	59.77	Chinese speech	47.38
Japanese singing	2.66	Japanese speech	9.06
English speech	5.83	French speech	10.06
Mixed speech*	27.69		
All	162.45		

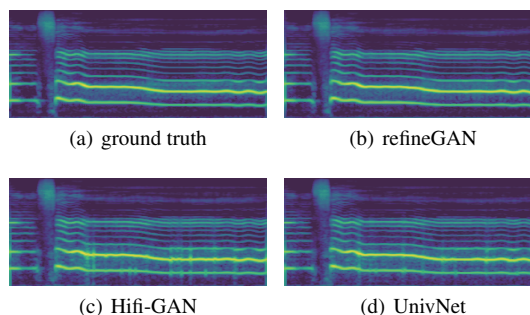


Figure 5: Comparison of the waveform reconstructed by three methods. (Using a slice from a Chinese singing recording for demonstration)

low in Table 2. The result shows our approach has a much advanced ability of universally modeling voice signals regardless of the language, the speaker or even whether it is singing voice or speech voice. And as showed in Figure 5, our method dramatically improved the pitch stability and solved the problem of generating pitch glitches when using conventional neural vocoders.

6. Conclusions

In this paper, we proposed RefineGAN, which consists of a series of methods to boost the performance of GAN-based neural vocoders. Parts of the methods mentioned above focused on improvements to the generator's architecture, while parts focused on improving training data's diversity.

The employment of these methods leads to model's outstanding robustness and stability, which enables the model trained on one particular type of data to be used to generate voices universally. Because the speaker or the language does not influence the voice quality much, this approach enables high-quality waveform synthesis even when the specific type of voice does not have sufficient high-fidelity recordings available for research.

7. References

- [1] 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.
- [2] J. Kong, J. Kim, and J. Bae, “Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis,” *arXiv preprint arXiv:2010.05646*, 2020.
- [3] J. Wu and J. Luan, “Adversarially trained multi-singer sequence-to-sequence singing synthesizer,” *arXiv preprint arXiv:2006.10317*, 2020.
- [4] 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.
- [5] W. Jang, D. Lim, J. Yoon, B. Kim, and J. Kim, “Univnet: A neural vocoder with multi-resolution spectrogram discriminators for high-fidelity waveform generation,” *arXiv preprint arXiv:2106.07889*, 2021.
- [6] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [7] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [8] B. Xu, N. Wang, T. Chen, and M. Li, “Empirical evaluation of rectified activations in convolutional network,” *arXiv preprint arXiv:1505.00853*, 2015.
- [9] M. Morise *et al.*, “Harvest: A high-performance fundamental frequency estimator from speech signals,” in *INTERSPEECH*, 2017, pp. 2321–2325.
- [10] 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.
- [11] X. X. S. Z. M. L. Yao Shi, Hui Bu, “Aishell-3: A multi-speaker mandarin tts corpus and the baselines,” 2015. [Online]. Available: <https://arxiv.org/abs/2010.11567>
- [12] E. Bakhturina, V. Lavrukhin, B. Ginsburg, and Y. Zhang, “Hi-Fi Multi-Speaker English TTS Dataset,” *arXiv preprint arXiv:2104.01497*, 2021.
- [13] C. Le Moine and N. Obin, “Att-HACK: An Expressive Speech Database with Social Attitudes,” in *Speech Prosody*, Tokyo, Japan, May 2020. [Online]. Available: <https://hal.archives-ouvertes.fr/hal-02508362>
- [14] R. Sonobe, S. Takamichi, and H. Saruwatari, “Jsut corpus: free large-scale japanese speech corpus for end-to-end speech synthesis,” *arXiv preprint arXiv:1711.00354*, 2017.
- [15] K. Sodimana, K. Pipatsrisawat, L. Ha, M. Jansche, O. Kjartansson, P. D. Silva, and S. Sarin, “A Step-by-Step Process for Building TTS Voices Using Open Source Data and Framework for Bangla, Javanese, Khmer, Nepali, Sinhala, and Sundanese,” in *Proc. The 6th Intl. Workshop on Spoken Language Technologies for Under-Resourced Languages (SLTU)*, Gurugram, India, Aug. 2018, pp. 66–70. [Online]. Available: <http://dx.doi.org/10.21437/SLTU.2018-14>