

Reducing over-smoothness in speech synthesis using Generative Adversarial Networks

Leyuan Sheng

Department of Mathematics and Mechanics
Novosibirsk State University
Novosibirsk, Russia
l.sheng@ngs.nsu.ru

Evgeniy N. Pavlovskiy

Stream Data Analytics and Machine Learning laboratory
Novosibirsk State University
Novosibirsk, Russia
pavlovskiy@post.nsu.ru

Abstract—Speech synthesis is widely used in many practical applications. In recent years, speech synthesis technology has developed rapidly. However, one of the reasons why synthetic speech is unnatural is that it often has over-smoothness. In order to improve the naturalness of synthetic speech, we first extract the mel-spectrogram of speech and convert it into a real image, then take the over-smooth mel-spectrogram image as input, and use image-to-image translation Generative Adversarial Networks(GANs) framework to generate a more realistic mel-spectrogram. Finally, the results show that this method greatly reduces the over-smoothness of synthesized speech and is more close to the mel-spectrogram of real speech.

Keywords: *Speech synthesis, over-smoothness, mel-spectrogram, GANs*

I. INTRODUCTION

Speech synthesis is a technique that produces artificial speech. It is the core of text-to-speech(TTS) and has a wide range of applications: voice assistants, smart stereos, car voice navigation, etc.

Over time, different technologies have been applied to the field of speech synthesis. Concatenative TTS with HMM-based unit selection, the process of splicing together basic units from a large number of prerecorded voices was the most advanced technology for many years. Parameter TTS was used to generate speech parameters (including fundamental frequency, formant frequency, etc.) at every moment according to a statistical model, and then convert these parameters into waveforms, thus providing high intelligible and fluent speech [1], [2]. However, comparing with human speech, the audio produced by these systems usually sounds muffled and unnatural.

In recent years, TTS systems based on deep learning have been extensively studied, and some scientific works have produced surprisingly clear speech. Deep neural network based TTS system includes WaveNet [3], Char2Wav [4], DeepVoice1&2 [5], [6], and Tacotron1&2 [7], [8]. However, there is still a certain gap with real speech, where in the important cause is over-smoothness in synthesized speech produced form TTS.

Recently, we have seen emerging solutions for excessive smoothing problems. In order to close the gap, Takuhiro Kaneko et al. have incorporated a Generative Adversarial

Networks(GANs) as postfilter for short-term Fourier transform(STFT) spectrograms [9]. However, they used a F_0 information and linguistic properties as input features to predict STFT spectra, a very high-dimensional data, not easily trainable by GAN. Ori Ernst et al. proposed a speech enhancement method in [10] and D. Michelsanti et al. proposed a speech dereverberation method in [11], they are also using GAN treated spectrogram as an image, but in order to adapt the input shape of GAN, they removed the high frequency bin. Even for testing, they needed to zero-padd the spectrogram of each test sample.

In this paper, our main contribution is to propose a way that converts speech processing into image processing. Based on an end-to-end TTS model Tacotron 2, we use the texts as input to predict the mel-spectrograms. Then we convert the normalized mel-spectrograms into images, for training and testing to reshape the images into the same size. We then use a GAN framework to generate more realistic mel-spectrograms, in order to overcome the over-smoothness in synthesized speech.

II. RELATED WORKS

End-to-End TTS Tacotron 1 [7] and Tacotron 2 [8] are end-to-end models which are trained directly from text-audio pairs, greatly reducing dependence on hand-designed internal modules. One advantage of Tacotron 2 over Tacotron 1 is that the speech signal is generated from conditioning the Mel spectrogram which captures all content information, such as pronunciation, prosody, and speaker identity. Tacotron 2 also achieves state-of-the-art sound quality which uses a sequence-to-sequence model optimized for TTS to map a range of functions from sequential letters to encoded audio. These features are an 80-dimensional mel-spectrogram which calculates frames every 12.5 milliseconds, capturing not only the pronunciation of words, and the subtleties of human speech, including volume, speed, and intonation.

Generative Adversarial Networks Generative adversarial networks(GANs) is a generative model designed by Goodfellow et al. in 2014 [12]. In the GAN setup, the two differentiable functions represented by the neural network are locked in the game. These two participants, the generator G and the discriminator D, play different roles in this framework. The




Fig. 1. Normalized images of original mel-spectrogram, TTS synthesized mel-spectrogram and GAN-based mel-spectrogram)




Fig. 2. Original mel-spectrogram, TTS synthesized mel-spectrogram and GAN-based mel-spectrogram

generator G attempts to generate data from certain probability distribution. The discriminator D , like a judge, determines whether its input is from the generator G or from the real training set.

So far, GAN has been widely used in the field of computer vision and image processing, and has achieved many amazing results. [13]–[16]. Pix2pixHD [17] is the state-of-the-art for image-to-image translation.

III. PIX2PIXHD FRAMEWORK FOR MEL-SPECTROGRAM

Mel-spectrogram is a very low-dimensional representation of linear-frequency spectrogram containing spectral envelopes and harmonic information. The mel scaled acoustic features are the logarithmic sensitivity to the frequency perception of audio signals which are based on the system of human hearing, having an overwhelming advantage in emphasizing audio details [18]. The Mel-spectrogram features are calculated by dividing a long signal into frames, windowing, and then performing Fourier transform (FFT) on each frame, then stacking the results of each frame along another dimension, and finally passing the mel-scale filter banks like in Tacotron 2.

The pix2pixHD method is a Conditional Generative Adversarial Networks (CGANs) framework for image-to-image translation, which improves the pix2pix network by using coarse-to-fine generator G networks and multi-scale discriminator D networks. In our task, the goal is use the generator G reconstruct natural mel-spectrograms y from synthesized mel-spectrograms s , while the discriminator D needs to distinguish its input is synthesized or natural. The generator G and discriminator D are trained through the following a min-max problem:

$$\min_G \max_{D_1, D_2, D_3} \sum_{i=1,2,3} \mathcal{L}_{GAN}(G, D_i) \quad (1)$$

where the objective function $\mathcal{L}_{GAN}(G, D)$ ¹ is given by

$$\mathbb{E}_{(x,y)}[\log D(\mathbf{x}, \mathbf{y})] + \mathbb{E}_{\mathbf{x}}[\log(1 - D(\mathbf{x}, G(\mathbf{x})))] \quad (2)$$

We use an existing Tacotron 2 model which uses a neural text to mel-spectrogram conversion model as the baseline to predict the synthesized mel-spectrograms. We tried to draw a synthesized mel-spectrogram and saved it as an image. However, since the saved image was multi-colored and had four channels, it is difficult to convert the image into a synthesized mel-spectrogram. Therefore, we normalized the synthesized mel-spectrogram first and then saved it as a grayscale image. In fact, the grayscale image we saved can earlier convert into a synthesized mel-spectrogram. The advantage of this is that in addition to making some scaling changes to the image, we can also restore the synthesized mel-spectrogram features, which makes the speech processing and image processing unified.

IV. RESULTS

Database: we used Chinese speech data which contains spoken audio from a little girl, which contains approximately 17 hours of audio with a sample rate of 48 kHz, and consists of 15,000 utterances.

Evaluation: Figure 1 shows the comparison of normalized images of original mel-spectrogram(upper image), normalized TTS synthesized mel-spectrogram(in the middle) and normalized GAN-based mel-spectrogram(lower image). We fed the TTS synthesized mel-spectrogram as input to Pix2pixHD framework which generated the GAN-based mel-spectrogram. Figure 2 shows the comparison of the original mel-spectrogram(upper image), TTS synthesized mel-spectrogram(in the middle), and GAN-based mel-spectrogram(lower image). In both comparisons, the GAN-

¹we denote $\mathbb{E}_{(x,y)} \triangleq \mathbb{E}_{(x,y) \sim p_{data}(x,y)}$ and $\mathbb{E}_{\mathbf{x}} \triangleq \mathbb{E}_{\mathbf{x} \sim p_{data}(x)}$ for simplicity

based mel-spectrogram not only emphasized the harmonic structure, but also reproduced the detailed structure similar to the original mel-spectrogram than the over-smoothed synthetic mel-spectrogram.

V. CONCLUSIONS

In this work we examined the process of reducing over-smoothness in speech synthesis by a generative adversarial networks. We achieved this, by adapting the Pix2pixHD framework to transform the speech feature processing into image-to-image translation. The image representation of speech features was directly processed, which was an extension of the previous method to the processing of speech features. We also showed the comparison of speech generated by our proposed method is comparable to original speech.

REFERENCES

- [1] H. Zen, A. Senior, and M. Schuster, Statistical parametric speech synthesis using deep neural networks, in Proc. ICASSP, 2013, pp. 79627966.
- [2] K. Tokuda, Y. Nankaku, T. Toda, H. Zen, J. Yamagishi, and K. Oura, Speech synthesis based on hidden Markov models, Proc. IEEE, vol. 101, no. 5, pp. 12341252, 2013.
- [3] A. van den Oord et al., WaveNet: A generative model for raw audio, arXiv:1609.03499, 2016.
- [4] J. Sotelo et al., Char2wav: End-to-end speech synthesis, in Proc. ICLR, 2017.
- [5] S. Arik et al., Deep voice: Real-time neural text-to-speech, in Proc. ICML, 2017, pp. 195204, arXiv:1702.07825.
- [6] S. Arik et al., Deep voice 2: Multi-speaker neural text-to-speech, in Proc. NIPS, 2017, arXiv:1705.08947.
- [7] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J. Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, Quoc Le, Yannis Agiomyrgiannakis, Rob Clark, and Rif A. Saurous. Tacotron: Towards end-to-end speech synthesis. In Proc. Interspeech, pages 40064010, August 2017.
- [8] Jonathan Shen, Ruoming Pang, Ron J. Weiss, Mike Schuster, Navdeep Jaitly, Zongheng Yang, Zhifeng Chen, Yu Zhang, Yuxuan Wang, RJ Skerry-Ryan, Rif A. Saurous, Yannis Agiomyrgiannakis, and Yonghui Wu. Natural TTS synthesis by conditioning WaveNet on mel spectrogram predictions. In Proc. IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP), 2018.
- [9] T. Kaneko, S. Takaki, H. Kameoka, J. Yamagishi, Generative Adversarial Network-based Postfilter for STFT Spectrograms, in INTERSPEECH, 2017.
- [10] O. Ernst, Shlomo E. Chazan, S. Gannot and J. Goldberger, Speech Dereverberation Using Fully Convolutional Networks arXiv preprint arXiv:1803.08243(2018)
- [11] D. Michelsanti and Z.-H. Tan, Conditional generative adversarial networks for speech enhancement and noise-robust speaker verification, in INTERSPEECH, 2017.
- [12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. WardeFarley, S. Ozair, A. Courville, and Y. Bengio, Generative adversarial nets, in Advances in neural information processing systems, 2014, pp. 26722680.
- [13] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, Image-to-image translation with conditional adversarial networks, arXiv preprint arXiv:1611.07004, 2016.
- [14] A. Radford, L. Metz, and S. Chintala, Unsupervised representation learning with deep convolutional generative adversarial networks, arXiv preprint arXiv:1511.06434, 2015.
- [15] S. Ioffe and C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, arXiv preprint arXiv:1502.03167, 2015.
- [16] J. T. Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, Striving for simplicity: The all convolutional net, arXiv preprint arXiv:1412.6806, 2014.
- [17] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2017. High-Resolution Image Synthesis and Semantic Manipulation with Conditional GANs. arXiv preprint arXiv:1711.11585 (2017).
- [18] Yi Zhao et al.: Wasserstein GAN and Waveform Loss-based Acoustic Model Training for Multi-speaker TTS Systems. arXiv preprint arXiv:1807.11679 (2018).