

STATISTICAL PARAMETRIC SPEECH SYNTHESIS USING GENERATIVE ADVERSARIAL NETWORKS UNDER A MULTI-TASK LEARNING FRAMEWORK

Shan Yang¹, Lei Xie^{1,†}, Xiao Chen², Xiaoyan Lou², Xuan Zhu², Dongyan Huang³, Haizhou Li^{3,4}

¹School of Computer Science, Northwestern Polytechnical University, Xi'an, China

²Language Computing Lab, Samsung R&D Institute of China, Beijing, China

³Institute for Infocomm Research, A*STAR, Singapore

⁴Department of Electrical and Computer Engineering, National University of Singapore

ABSTRACT

In this paper, we aim at improving the performance of synthesized speech in statistical parametric speech synthesis (SPSS) based on a generative adversarial network (GAN). In particular, we propose a novel architecture combining the traditional acoustic loss function and the GAN's discriminative loss under a multi-task learning (MTL) framework. The mean squared error (MSE) is usually used to estimate the parameters of deep neural networks, which only considers the numerical difference between the raw audio and the synthesized one. To mitigate this problem, we introduce the GAN as a second task to determine if the input is a natural speech with specific conditions. In this MTL framework, the MSE optimization improves the stability of GAN, and at the same time GAN produces samples with a distribution closer to natural speech. Listening tests show that the multi-task architecture can generate more natural speech that satisfies human perception than the conventional methods.

Index Terms— Statistical parametric speech synthesis, deep neural network, generative adversarial network, conditional generative adversarial network, multi-task learning

1. INTRODUCTION

Statistical parametric speech synthesis (SPSS) has attracted significant attentions since the successful use of hidden Markov models (HMMs) [1, 2, 3]. In HMM based systems, Gaussian mixture model (GMM) was used to model the hidden states of observations. Considering the limitations of the decision tree clustering procedure in modeling the complex context dependencies in HMM-based statistical parametric speech synthesis [4, 5], deep neural networks (DNNs) have been proposed for acoustic modeling, which can produce more natural synthesized speech [4, 6]. More recently, advanced estimation criteria and novel network architectures

have been introduced to further improve the performance of SPSS [7, 8, 9, 10, 11].

Since the purpose of training in the statistical method is to maximize the likelihood or specifically to minimize the mean square error (MSE) between the synthesized (i.e., network outputs) and the original speech parameters in neural network based Text-to-Speech (TTS), the synthesized speech may achieve suboptimal human perceptual level. Hence there is an underlying reasonable-but-not-necessarily-optimal hypothesis that the most natural synthesized speech has the minimal value in the numerical loss, which may fall into the *perceptual deficiency* problem. In other words, the reduction in numerical errors may not necessarily lead to better perceived speech [12]. In this paper, we propose to use generative adversarial networks (GANs) [13] to remedy this deficiency.

Significant efforts have been made to remedy the perceptual deficiency problem by improving the training criteria [14, 15, 16, 17]. In [14], by incorporating the whole sequence parameters into training, the sequence generation error (SGE) minimization was proposed to eliminate the mismatch between training and testing. Considering the independence of frames in DNNs, the minimum trajectory error training was adopted to take into account the dynamic constraints from a wide acoustic context during training [15]. In SPSS, the speech features must be invertible for reconstruction through a vocoder, and this rules out the use of many perceptual representations of speech that can not be reconstructed to speech waveform. Hence one solution to the *perceptual suboptimality* issue is to bring more representative perceptual features into acoustic modeling [12]. In [12], under a multi-task learning (MTL) framework, along with the invertible spectral feature used in the vocoder, extra perceptual representations of speech, e.g., spectro-temporal excitation pattern, were included as a second prediction target in DNN-based SPSS.

In this paper, we propose to use GANs to solve the perceptual deficiency problem in acoustic modeling. GAN is a powerful generative model that has been successfully used in image generation [13, 18, 19] and other tasks [20, 21]. It consists of a generator G , which is treated as an acoustic model in

This work was supported by the National Natural Science Foundation of China (Grant No. 61571363) and the China Scholarship Council (Grant No. 201706290118). † Corresponding author

our framework to generate speech, and a discriminator D for discriminating the generated speech and the genuine speech. Specifically, the objective of G is to capture the distribution of the natural speech, while D aids the training of G by examining the data generated by G in reference to real data, and hence helping G learn the distribution that underpins the real data [13]. In our framework, GAN naturally addresses the perceptual deficiency problem: the updating of the generator is not directly from the data samples, while it comes from the back propagation of the discriminator. This means D can capture the essential difference between the natural speech and the synthesized speech and this ‘perceptual’ difference is used to guide G , the generator. Considering the mode collapse problem of the generated samples in GAN [18], we take conditional linguistic features as a guidance to control the generation process. Moreover, since the gradients of GANs are not stable, we also use the conventional MSE loss function to stabilize the training process. More specifically, inspired by [16, 12, 22, 23], we combine the MSE loss with the GAN loss under an MTL framework. The objective experiments show that our framework has comparable performance in numerical loss compared to the baseline BLSTM-based TTS, while promisingly, the subjective listening experiments indicate that the proposed architecture achieves significant improvement. That is, the proposed GAN approach results in better perceptual speech quality.

2. RELATED WORKS

We notice that there are several recent attempts of using GANs to improve the quality of synthesized speech. In [24], GAN was treated as a post-filter for acoustic models to overcome the over-smoothing problem. Specifically, natural speech was used as a conditional guidance of GAN, which tries to reproduce the natural speech texture from the synthesized one. In [25], variational autoencoding Wasserstein GAN (VAW-GAN) was proposed to build a voice conversion system from unaligned data, in which the GAN objective was incorporated into the decoder to improve the conditional variational autoencoder (C-VAE).

Our approach shares a similar idea with [16]. In order to compensate the difference between the synthesized speech and the natural speech in acoustic modeling, an Anti-Spoofing Verification (ASV) module (like the discriminator in GAN) was introduced to distinguish between the natural and the synthetic speech. The speech generator has no difference with a typical neural network acoustic model [4], i.e., learning a non-linear mapping from linguistic features to speech parameters, but the ASV discrimination loss was combined with the minimal generation error (MGE) loss, under an MTL framework, to train the network.

It is noted that our approach is different from [16] in terms of motivation and implementation. Instead of addressing the over-smoothed problem with additional ASV constraint as

compensation, we propose to use GANs to directly produce speech samples with closer distribution to natural speech from a *uniformly random noise distribution*. In other words, the input of our speech generator is random noise, while linguistic features are introduced to both the generator and the discriminator as conditions. As such, the prior uniformly random noise distribution creates new samples that approximate the training data distribution, and it brings diversity with conditions to the synthesized speech from the generator while the linguistic conditions add direct linguistic-discriminative information to the discriminator. On the other hand, as the Nash equilibrium is hard to achieve in network estimation, the training process becomes unstable during the adversarial game. To tackle this problem, we take other optimization methods, such as variational auto-encoder [22] or MSE, to restrain the process. Finally, in our implementation, we use state-of-the-art BLSTM network as a benchmark, which can produce speech with much better quality than feed-forward networks used in [16].

3. GAN-BASED MULTI-TASK LEARNING

Fig. 1 illustrates the architecture of the proposed GAN-based MTL framework, which consists of a *Generator* and a *Discriminator*. In the training process, different from [16], we use random noise as the input of *Generator*, and introduce the linguistic features to each hidden layer as the conditional information. Then the *Generator* can produce the synthesized speech, with which the *Discriminator* can distinguish between the synthesized speech and the natural speech under the same conditions. The estimation of this process is made up of two aspects: 1) For the *Discriminator* estimation, the *OR* operator means that the synthesized samples and natural samples are alternately used to train a binary classifier with corresponding class labels, and then the *Discriminator* can judge whether the speech is synthesized or genuine. 2) As for the *Generator*, the *AND* operator means that the synthesized and natural samples are simultaneously used to calculate the MSE, meanwhile the discriminant error will also affect the estimation of *Generator*.

In the synthesis stage, given a random noise and specific linguistic features, we can easily generate speech from *Generator* using the forward direction. We will describe the framework in details in the following.

3.1. Generative Adversarial Networks

GAN is a generative model that can learn a complex relationship between random noise input vector z and output parameters y by an adversarial process [13]. The estimation of GANs consist of two models: a generative model G that captures the data distribution from random noise z , and a discriminative model D that maximizes the probability of correctly discriminating between the *real* examples and *fake* samples generated

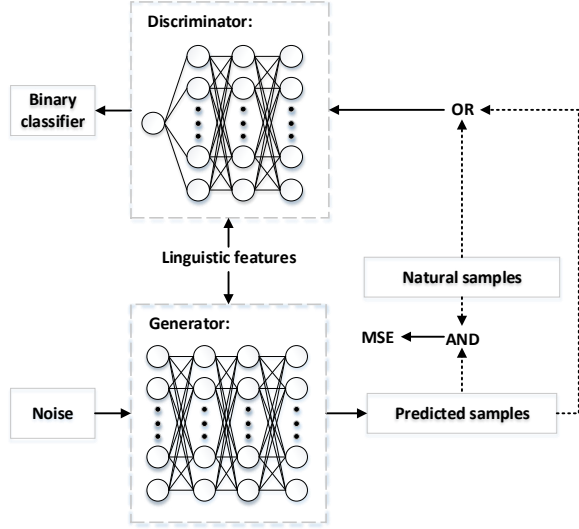


Fig. 1. System diagram of GAN-based multi-task learning framework.

from G .

In this adversarial process, the generator tends to learn a mapping function $G(z)$ to fit the real data distribution $p_{data}(x)$ from a uniformly random noise distribution $p_z(z)$, while the purpose of discriminator is to perfectly judge whether the sample is from $G(z)$ or $p_{data}(x)$. So the G and D are both trained simultaneously in the two-player min-max game with value function:

$$loss_{gan} = \min_G \max_D \mathbb{E}_{x \sim p_{data}(x)} [\log D(x)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z)))]. \quad (1)$$

In the above generative model, the modes of generated samples cannot be controlled because of the weak guidance. So the conditional generative adversarial network (CGAN) [18] is proposed to direct the generation by considering additional information y . Then the loss function can be expressed as

$$loss_{cgan} = \arg \min_G \max_D \mathbb{E}_{x \sim p_{data}(x)} [\log D(x|y)] + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z|y)|y))]. \quad (2)$$

3.2. Multi-Task Learning with GANs in SPSS

In the traditional acoustic model for SPSS, we usually minimize the MSE between the predicted parameters X_{model} and the natural speech X_{real} during the estimation. The objective can be written as

$$loss_{mse} = \arg \min \frac{\sum_{i=1}^n (X_{real,i} - X_{model,i})^2}{n}. \quad (3)$$

As Eq. (3) shows, the numerical difference (in terms of MSE) is only concerned in the estimation, and the numerical error reduction may not necessarily lead to perceptual improvement on the synthesized speech [12]. To solve this problem, we propose to use GANs to learn the essential differences between the synthesized speech and the natural speech through a discriminative process.

GAN is able to generate data rather than estimate the density function. Due to the model collapse problem in the generative model in GAN [18], we propose the following generator loss function in order to guide GAN to converge to optimal solution such that the generative model produces desired data:

$$loss_G = \mathbb{E}_{z \sim p_z(z)} [G(z|y) - X_{real}]^2 + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z|y)|y))], \quad (4)$$

where $X_{real} \sim p_{data}(x)$, and X_{model} is generated by the generator G using uniformly random noise z under condition y . Combining Eq. (2) and Eq. (4), the final objective of our MTL framework is:

$$loss_{multi} = \arg \min_G \max_D \mathbb{E}_{x \sim p_{data}(x)} [\log D(x|y)] + \mathbb{E}_{z \sim p_z(z)} [G(z|y) - X_{real}]^2 + \mathbb{E}_{z \sim p_z(z)} [\log(1 - D(G(z|y)|y))]. \quad (5)$$

We treat the linguistic features as additional vector y , and make the input noise z obey a uniform distribution in the interval $[-1,1]$. Then our framework can generate the speech X_{model} by $G(z|y)$, and the $loss_{mse}$ and $loss_{cgan}$ are estimated at the same time during training. Note that the input of our speech generator is uniformly random noise and linguistic features are used as conditions for both the generator and the discriminator, which is different from [16].

Since the effective likelihood of GAN is unknown and intractable [22], several auto-encoder GAN variants use zero-mean Laplace distribution $\exp(-\lambda \|x - G(z)\|_1)$ [26, 27] to solve the problems. In order to directly show the likelihood of these variants, we can simply set $\lambda = 1$ and replace the $L1$ reconstruction loss with $L2$ norm, and then we can get the MSE format as traditional methods. That is to say, we can take other explicit likelihood (e.g., MSE) to solve the intractable inference of GANs. The $L1$ reconstruction loss will be investigated in the near future.

3.3. Phoneme Discrimination for GANs

In Section 3.2, the discriminator is a binary classifier to judge whether the data x is from $G(z)$ or $p_{data}(x)$ under the condition y . We also try to use phoneme information to guide the discrimination process in our multi-task framework, as shown in Fig.2.

Assume $label$ is a one-hot encoded vector representing the phoneme class, which is the category of both *fake* and *real*




Fig. 2. The discriminator with phoneme information.

samples for D . Then our goal is to minimize the cross entropy (CE) for the *real* and to maximize this loss for the *fake*, and the latter one means that we do not know which phoneme the *fake* belongs to. So the target function of GANs in $loss_{multi}$ can be updated with

$$\arg \min \mathbb{E}_{x \sim p_{data}(x)} [D_{CE}(x|y, label)] - \mathbb{E}_{z \sim p_z(z)} [D_{CE}(G(z|y)|y, label)]. \quad (6)$$

We obtain the new loss function considering the phoneme classification as follows.

$$loss_{multi-pc} = \arg \min \mathbb{E}_{x \sim p_{data}(x)} [\log D_{CE}(x|y, label)] + \mathbb{E}_{z \sim p_z(z)} [G(z|y) - X_{real}]^2 - \mathbb{E}_{z \sim p_z(z)} [D_{CE}(G(z|y)|y, label)]. \quad (7)$$

4. EXPERIMENTS

4.1. Experimental Setup

In the experiments, a Chinese speech corpus was used to evaluate the performance of our approach. The corpus consists of about 10,000 utterances from a single female speaker. We randomly selected around 8,000 sentences for network training, 1,000 utterances for model validation and another 1,000 for testing. Each speech waveform was sampled at 16 kHz, and we used WORLD [28] (D4C edition [29]) to extract 60-dimensional Mel-Cepstral Coefficients (MCCs), 1-dimensional band aperiodicities (BAP) and F_0 in log-scale in 5-ms step. So the final acoustic features were 63-dimensions including one extra binary voiced/unvoiced flag. As for the text, we made a complex text analysis module to get 138-dimensional linguistic features, including phoneme information, prosody boundary labeling, part of speech tagging, state information and corresponding position index.

To benchmark the performance of the GAN-based MTL framework, we compared four systems, listed as follows.

- **BLSTM**: We used bidirectional long short-term memory (BLSTM) based acoustic model as the baseline, which contained three feed-forward layers with 512 nodes/layer, followed by two BLSTM layers with 512 cells and a fully-connected output layer.

Methods	MCD (dB)	F_0 RMSE (Hz)	V/UV (%)
BLSTM	4.624	18.544	6.447
ASV [16]	4.670	18.871	6.562
GAN	4.633	18.678	6.492
GAN-PC	4.628	18.616	6.464

Table 1. Objective evaluation results.

- **GAN-based MTL (GAN)**: The proposed framework shown in Fig. 1. For the generator G , we also used three feed-forward and two BLSTM layers corresponding to the baseline. But the input was replaced with 200-dimensional random noise under the $[-1,1]$ uniform distribution. The linguistic features were added to the output of each hidden layer in G as conditions. As for the discriminator, two convolutional layer were used with $5 * 5$ filter shape, and $LReLU$ was treated as the activation function followed by batch normalization [24, 19]. Besides, there was a fully-connected layer after the convolutional architecture and a binary classification layer in the end. The linguistic conditions were also introduced to all hidden layers in D .
- **ASV as a second task [16] (ASV)**: We realized the ASV approach with $\omega_D = 1$. The network architecture was the same as **BLSTM**.
- **GAN with phoneme classification (GAN-PC)**: The same GAN model architecture was used, except that the output layer of D became a 63-category phoneme classification task, as described in Section 3.3.

Note that, for fair comparison, we didn't use any post-processing methods (eg. global variance [30, 31] and modulation spectrum [32, 33]) to improve the quality of synthesized speech. All systems were optimized using the Adam optimizer [34, 35], and implemented with TensorFlow [36].

4.2. Objective Evaluation

We first conducted the objective measure to evaluate the performance of the four systems on the testing data. Specifically, Mel-cepstral distortion (MCD) was used to evaluate the distortion of spectrum, and RMSE was introduced to calculate the F_0 error. Besides, the V/UV error rate was also used to present the accuracy of the voice/unvoice flag judgements.

Table 1 shows the objective results. As shown, there seems to be no remarkable differences among these systems. Since one purpose of our framework was to compensate the deficiency of traditional numerical loss function by GANs, these numerical measures may not be suitable for evaluation because of the internal squared error [24]. But in another aspect, we can find that the proposed MTL framework can compensate the instability and the mode collapse issues of GANs and generate stable and diverse speech with the help

of traditional loss function. That is to say, GANs can utilize the numerical loss function to limit its adversarial process.

4.3. Subjective Evaluation

We conducted listening tests to assess the quality of the synthesized speech. we made four pairs of A/B preference test: BLSTM vs. GAN, GAN vs. ASV, GAN vs. GAN-PC and BLSTM vs. GAN-PC. For listening test, 20 sentences were randomly selected from the test data, and all listening pairs were presented in a shuffled order. There were 20 listeners participating in the test. In each test session, the listeners were asked to choose the better one considering the perceived speech quality, or choose the “neutral” option if there was no difference.

44.5% GAN	29.5% Neutral	27.0% BLSTM
40.8% GAN	30.5% Neutral	28.7% ASV
41.5% GAN	32.2% Neutral	26.3% GAN-PC
34.1% GAN-PC	36.8% Neutral	29.0% BLSTM

Fig. 3. The preference score (%) of A/B test.

Fig. 3 shows the preference bars of the four pairs. The first bar of **GAN** vs. **BLSTM** indicates that the GAN-based MTL framework can significantly improve the performance of the synthesized speech ($p < 0.0001$). The listeners pointed out that the **GAN** system could produce speech with less buzzy sounds in most cases and more natural prosody in some samples. The second bar of **GAN** vs. **ASV** shows that the proposed **GAN** approach is better than the **ASV** approach ($p < 0.0001$). As discussed in [16], the **ASV** optimization aims to make the distribution of the synthesized speech close to the natural speech. But this method theoretically lacks the linguistic conditional guidance in distinguishing between the distributions of natural and synthesized speech. As a result, as compared with **ASV**, we find that the proposed **GAN** approach can capture both subtle and rapid changes, leading to better brightness of the synthesized speech. Fig. 4 shows the distance of averaged global variance (GV) between natural speech and synthesized speech from different approaches. The smaller values mean that the GVs of synthesized speech are more similar to natural speech. The result indicates that the GVs of **GAN** are closer to the natural GVs than **ASV** especially in the first few coefficients.

The third bar of **GAN** vs. **GAN-PC** in Fig. 3 shows that there is a decline in performance with phoneme classification in GANs ($p < 0.001$). In order to explain this phenomenon, we compared the **BLSTM** system with **GAN-PC**,




Fig. 4. The difference of averaged GVs per mel-cepstral coefficients compared to natural speech.

as shown in the last bar. In this preference test, **GAN-PC** slightly outperforms **BLSTM**, but the difference is not significant between the two systems ($p = 0.0253$). As we know, the phoneme information is related to the contents of speech, which highly correlates to the intelligibility of the synthesized samples. The **BLSTM** based acoustic model can already produce speech with high intelligibility. However, the purpose of treating the phoneme label as a guidance for the discriminator in **GAN-PC** is to improve the intelligibility, not to make the distribution of synthesized examples closer to the natural samples. So simply letting the discriminator distinguish whether x is a natural sample in **GAN** can make the synthesized speech be more related to human perception, resulting in better subjective listening performance.

5. CONCLUSION

This paper proposed to use GANs to improve the quality of synthesized speech. We use a multi-task learning architecture with GANs, where the GANs can compensate the deficiency of traditional MSE loss function while the MSE can also help to solve the instability of GANs. Evaluation results show that the proposed method can compensate the weakness of numerical loss function and improve the performance of SPSS. The proposed framework has a little increase in the computation cost, compared to traditional acoustic models during the generation process, as the extra computation only comes from noise generation and feature concatenation.

In our future work, we will focus on improving the performance of the generator in GAN and try to use GAN in end-to-end speech synthesis [37, 38, 39, 40, 41]. Since the MSE loss is still used to stabilize the adversarial process in our framework, we attempt to find a self-stabilizing architecture to directly estimate the distribution of synthe-

sized speech, such as using Wasserstein GAN [35] and VAE-GAN [22].

6. REFERENCES

- [1] Takayoshi Yoshimura, Keiichi Tokuda, Takashi Masuko, Takao Kobayashi, and Tadashi Kitamura, "Simultaneous modeling of spectrum, pitch and duration in HMM-based speech synthesis," in *Proceedings of Eurospeech*, 1999, pp. 2347–2350.
- [2] Keiichi Tokuda, Takayoshi Yoshimura, Takashi Masuko, Takao Kobayashi, and Tadashi Kitamura, "Speech parameter generation algorithms for HMM-based speech synthesis," in *Proceedings of ICASSP*, 2000, vol. 3, pp. 1315–1318.
- [3] Heiga Zen, Keiichi Tokuda, and Alan W Black, "Statistical parametric speech synthesis," *Speech Communication*, vol. 51, no. 11, pp. 1039–1064, 2009.
- [4] Heiga Zen, Andrew Senior, and Mike Schuster, "Statistical parametric speech synthesis using deep neural networks," in *Proceedings of ICASSP*, 2013, pp. 7962–7966.
- [5] Zhen-Hua Ling, Shi-Yin Kang, Heiga Zen, Andrew Senior, Mike Schuster, Xiao-Jun Qian, Helen M Meng, and Li Deng, "Deep learning for acoustic modeling in parametric speech generation: A systematic review of existing techniques and future trends," *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 35–52, 2015.
- [6] Oliver Watts, Gustav Eje Henter, Thomas Merritt, Zhizheng Wu, and Simon King, "From HMMs to DNNs: where do the improvements come from?," in *Proceedings of ICASSP*, 2016, pp. 5505–5509.
- [7] Zhen-Hua Ling, Li Deng, and Dong Yu, "Modeling spectral envelopes using restricted boltzmann machines for statistical parametric speech synthesis," in *Proceedings of ICASSP*, 2013, pp. 7825–7829.
- [8] Shiyin Kang, Xiaojun Qian, and Helen Meng, "Multi-distribution deep belief network for speech synthesis," in *Proceedings of ICASSP*, 2013, pp. 8012–8016.
- [9] Heiga Zen and Andrew Senior, "Deep mixture density networks for acoustic modeling in statistical parametric speech synthesis," in *Proceedings of ICASSP*, 2014, pp. 3844–3848.
- [10] Yuchen Fan, Yao Qian, Feng-Long Xie, and Frank K Soong, "TTS synthesis with bidirectional LSTM based recurrent neural networks," in *Proceedings of Interspeech*, 2014, pp. 1964–1968.
- [11] Xin Wang, Shinji Takaki, and Junichi Yamagishi, "An autoregressive recurrent mixture density network for parametric speech synthesis," in *Proceedings of ICASSP*, 2017.
- [12] Zhizheng Wu, Cassia Valentini-Botinhao, Oliver Watts, and Simon King, "Deep neural networks employing multi-task learning and stacked bottleneck features for speech synthesis," in *Proceedings of ICASSP*, 2015, pp. 4460–4464.
- [13] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *Proceedings of NIPS*, 2014, pp. 2672–2680.
- [14] Fan Yuchen, Qian Yao, Frank K Soong, and Lei He, "Sequence generation error (SGE) minimization based deep neural networks training for text-to-speech synthesis," in *Proceedings of INTERSPEECH*, 2015, pp. 864–868.
- [15] Zhizheng Wu and Simon King, "Minimum trajectory error training for deep neural networks, combined with stacked bottleneck features," in *Proceedings of INTERSPEECH*, 2015, pp. 309–313.
- [16] Yuki Saito, Shinnosuke Takamichi, and Hiroshi Saruwatari, "Training algorithm to deceive anti-spoofing verification for DNN-based speech synthesis," in *Proceedings of ICASSP*. IEEE, 2017, pp. 4900–4904.
- [17] Zhizheng Wu and Simon King, "Improving trajectory modelling for dnn-based speech synthesis by using stacked bottleneck features and minimum generation error training," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 7, pp. 1255–1265, 2016.
- [18] Mehdi Mirza and Simon Osindero, "Conditional generative adversarial nets," *arXiv preprint arXiv:1411.1784*, 2014.
- [19] Alec Radford, Luke Metz, and Soumith Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *Proceedings of ICLR*, 2016.
- [20] Santiago Pascual, Antonio Bonafonte, and Joan Serra, "SEGAN: Speech enhancement generative adversarial network," *arXiv preprint arXiv:1703.09452*, 2017.
- [21] Li-Chia Yang, Szu-Yu Chou, and Yi-Hsuan Yang, "MidiNet: A convolutional generative adversarial network for symbolic-domain music generation using 1D and 2D conditions," *arXiv preprint arXiv:1703.10847*, 2017.

- [22] Mihaela Rosca, Balaji Lakshminarayanan, David Warde-Farley, and Shakir Mohamed, “Variational approaches for auto-encoding generative adversarial networks,” *arXiv preprint arXiv:1706.04987*, 2017.
- [23] Heiga Zen, “Generative model-based text-to-speech synthesis,” 2017, Invited talk given at CBMM workshop on speech representation, perception and recognition.
- [24] Takuhiro Kaneko, Hirokazu Kameoka, Nobukatsu Hojo, Yusuke Ijima, Kaoru Hiramatsu, and Kunio Kashino, “Generative adversarial network-based postfilter for statistical parametric speech synthesis,” in *Proceedings of ICASSP*, 2017, pp. 4910–4914.
- [25] Chin-Cheng Hsu, Hsin-Te Hwang, Yi-Chiao Wu, Yu Tsao, and Hsin-Min Wang, “Voice conversion from unaligned corpora using variational autoencoding wasserstein generative adversarial networks,” in *Proceedings of INTERSPEECH*, 2017.
- [26] David Berthelot, Tom Schumm, and Luke Metz, “BEGAN: Boundary equilibrium generative adversarial networks,” *arXiv preprint arXiv:1703.10717*, 2017.
- [27] Anh Nguyen, Jason Yosinski, Yoshua Bengio, Alexey Dosovitskiy, and Jeff Clune, “Plug & play generative networks: Conditional iterative generation of images in latent space,” *arXiv preprint arXiv:1612.00005*, 2016.
- [28] Masanori Morise, Fumiya Yokomori, and Kenji 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.
- [29] Masanori Morise, “D4C, a band-a-periodicity estimator for high-quality speech synthesis,” *Speech Communication*, vol. 84, pp. 57–65, 2016.
- [30] Tomoki Toda and Keiichi Tokuda, “A speech parameter generation algorithm considering global variance for HMM-based speech synthesis,” *IEICE TRANSACTIONS on Information and Systems*, vol. 90, no. 5, pp. 816–824, 2007.
- [31] Tomoki Toda, Takashi Muramatsu, and Hideki Banno, “Implementation of computationally efficient real-time voice conversion,” in *Proceedings of Interspeech*, 2012, pp. 94–97.
- [32] Shinnosuke Takamichi, Tomoki Toda, Graham Neubig, Sakriani Sakti, and Satoshi Nakamura, “A postfilter to modify the modulation spectrum in HMM-based speech synthesis,” in *Proceedings of ICASSP*. IEEE, 2014, pp. 290–294.
- [33] Shinnosuke Takamichi, Tomoki Toda, Alan W Black, and Satoshi Nakamura, “Modulation spectrum-based post-filter for gmm-based voice conversion,” in *Proceedings of APSIPA*. IEEE, 2014, pp. 1–4.
- [34] Diederik Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” *arXiv preprint arXiv:1412.6980*, 2014.
- [35] Martin Arjovsky, Soumith Chintala, and Léon Bottou, “Wasserstein GAN,” *arXiv preprint arXiv:1701.07875*, 2017.
- [36] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, et al., “Tensorflow: Large-scale machine learning on heterogeneous distributed systems,” *arXiv preprint arXiv:1603.04467*, 2016.
- [37] Yuxuan Wang, RJ Skerry-Ryan, Daisy Stanton, Yonghui Wu, Ron J Weiss, Navdeep Jaitly, Zongheng Yang, Ying Xiao, Zhifeng Chen, Samy Bengio, et al., “Tacotron: A fully end-to-end text-to-speech synthesis model,” *arXiv preprint arXiv:1703.10135*, 2017.
- [38] Soroush Mehri, Kundan Kumar, Ishaan Gulrajani, Rithesh Kumar, Shubham Jain, Jose Sotelo, Aaron Courville, and Yoshua Bengio, “SampleRNN: An unconditional end-to-end neural audio generation model,” *arXiv preprint arXiv:1612.07837*, 2016.
- [39] Aaron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [40] Sercan Arik, Gregory Diamos, Andrew Gibiansky, John Miller, Kainan Peng, Wei Ping, Jonathan Raiman, and Yanqi Zhou, “Deep voice 2: Multi-speaker neural text-to-speech,” *arXiv preprint arXiv:1705.08947*, 2017.
- [41] Jose Sotelo, Soroush Mehri, Kundan Kumar, Joao Felipe Santos, Kyle Kastner, Aaron Courville, and Yoshua Bengio, “Char2wav: End-to-end speech synthesis,” in *Proceedings of ICLR workshop submission*, 2017.