WGANSing: A Multi-Voice Singing Voice Synthesizer Based on the Wasserstein-GAN

Pritish Chandna*, Merlijn Blaauw*, Jordi Bonada* and Emilia Gómez*[†]

*Music Technology Group, Universitat Pompeu Fabra, Barcelona, Spain

[†]Joint Research Centre, European Commission, Seville, Spain

Email:pritchandna@upf.edu, merlijn.blaauw@upf.edu, jordi.bonada@upf.edu, emilia.gomez@upf.edu

Abstract—We present a deep neural network based singing voice synthesizer, inspired by the Deep Convolutions Generative Adversarial Networks (DCGAN) architecture and optimized using the Wasserstein-GAN algorithm. We use vocoder parameters for acoustic modelling, to separate the influence of pitch and timbre. This facilitates the modelling of the large variability of pitch in the singing voice. Our network takes a block of consecutive frame-wise linguistic and fundamental frequency features, along with global singer identity as input and outputs vocoder features, corresponding to the block of features. This block-wise approach, along with the training methodology allows us to model temporal dependencies within the features of the input block. For inference, sequential blocks are concatenated using an overlap-add procedure. We show that the performance of our model is competitive with regards to the state-of-the-art and the original sample using objective metrics and a subjective listening test. We also present examples of the synthesis on a supplementary website and the source code via GitHub.

Index Terms—Wasserstein-GAN, DCGAN, WORLD vocoder, Singing Voice Synthesis, Block-wise Predictions

I. INTRODUCTION

Singing voice synthesis and Text-To-Speech (TTS) synthesis are related but distinct research fields. While both fields try to generate signals mimicking the human voice, singing voice synthesis models a much higher range of pitches and vowel durations. In addition, while speech synthesis is controlled primarily by textual information such as words or syllables, singing voice synthesis is additionally guided by a score, which puts constraints on pitch and timing. These constraints and differences also cause singing voice synthesis models to deviate somewhat from their speech counterparts. Historically, both speech and singing voice synthesis have been based on concatenative methods, which involve transformation and concatenation of waveforms from a large corpus of specialized recordings. Recently, several machine learning based methods have been proposed in both fields, most of which also require a large amount of data for training. In terms of quality, the field of TTS has seen a revolution in the last few years, with the introduction of the WaveNet [1] autoregressive framework, capable of synthesizing speech virtually indistinguishable from a real voice recording. This architecture inspired the Neural Parametric Singing Synthesiser (NPSS) [2], a deep learning based singing voice synthesis method which is trained on a dataset of annotated natural singing and produces high quality synthesis.

The WaveNet [1] directly generates the waveform given local linguistic and global speaker identity conditions. While a high quality synthesis is generated, the drawback of this model is that it requires a large amount of annotated data. As such, some following works, like the Tacotron 2 [3], use the WaveNet as a vocoder for converting acoustic features to a waveform and use a separate architecture for modelling these acoustic features from the linguistic input. The WaveNet vocoder architecture, trained on unlabeled data, is also capable of synthesizing high-quality speech from an intermediate feature representation. The task that we focus on in this paper is generating acoustic features given an input of linguistic features.

Various acoustic feature representations have been proposed for speech synthesis, including the mel-spectrogram [3], which is a compressed version of the linear-spectrogram. However, for the singing voice, a good option is to use vocoder features, as they separate pitch from timbre of the signal. This is ideal for the singing voice as the pitch range of the voice while singing is much higher than that while speaking normally. Modelling the timbre independently of the pitch has been shown to be an effective methodology [2]. We note that the use of a vocoder for direct synthesis can lead to a degradation of sound quality, but this degradation can be mitigated by the use of a WaveNet vocoder trained to synthesize the waveform from the parametric vocoder features. As such, for the scope of this study, we limit the upper-bound of the performance of the model to that of the vocoder.

Like autoregressive networks, Generative adversarial networks (GANs) [4]–[6] are a family of generative frameworks for deep learning, which includes the Wasserstein-GAN [7] variant. While the methodology has provided exceptional results in fields related to computer vision, it has only a few adaptations in the audio domain and indeed in TTS, that we discuss in the following sections. We adapt the Wasserstein-GAN model for singing voice synthesis. In this paper, we present a novel block-wise generative model for singing voice synthesis, trained using the Wasserstein-GAN framework¹. The block-wise nature of the model allows us to model

¹The code for this model in the TensorFlow framework is available at <https://github.com/MTG/WGANSing>. Audio examples, including synthesis with voice change and without the reconstruction loss can be heard at https://pc2752.github.io/sing_synth_examples/

temporal dependencies among features, much like the inter-pixel dependencies are modelled by the GAN architecture in image generation. For this study, we use the original fundamental frequency for synthesis, leading to a performance driven synthesis. As a result, we only model the timbre of the singing voice and not the expression via the f_0 curve or the timing deviations of the singers.

The rest of the paper is organized as follows. Section II provides a brief overview of the GAN and Wasserstein-GAN generative frameworks. Section III discusses the state-of-the-art singing voice synthesis model that we use as a baseline in this paper and some of the recent applications of GANs in the field of TTS and in general, in the audio domain. The following sections, section IV and section V present our model for singing voice synthesis, followed by a brief discussion on the dataset used and the hyperparameters of the model in sections VI and VII respectively. We then present an evaluation of the model, compared to the baseline in section IX, before wrapping up with the conclusions of the paper and a discussion of our future direction in section X.

II. GANS AND WASSERSTEIN-GANS

Generative Adversarial Networks (GANs) have been extensively used for various applications in computer vision since their introduction. GANs can be viewed as a network optimization methodology based on a two-player non-cooperative training that tries to minimize the divergence between a parameterized generated distribution P_g and a real data distribution, P_r . It consists of two networks, a generator, G and a discriminator, D , which are trained simultaneously. The discriminator is trained to distinguish between a real input and a synthetic input output by the generator, while the generator is trained to fool the discriminator. The loss function for the network is shown in equation 1.

$$\mathcal{L}_{GAN} = \min_G \max_D \mathbb{E}_{y \sim P_r} [\log(D(y))] + \mathbb{E}_{x \sim P_x} [\log(1 - D(G(x)))] \quad (1)$$

Where y is a sample from the real distribution and x is the input to the generator, which may be noise or conditioning as in the Conditional GAN [5] and is taken from a distribution of such inputs, P_x .

While GANs have been shown to produce realistic images, there are difficulties in training including vanishing gradient, mode collapse and instability. To mitigate these difficulties, the Wasserstein-GAN [7] has been proposed, which optimizes an efficient approximation of the Earth-Mover (EM) distance between the generated and real distributions and has been shown to produce realistic images. The loss function for the WGAN is shown in equation 2. In this version of the GAN, the discriminator network is replaced by a network termed as critic, also represented by D , which can be trained to optimality and does not saturate, converging to a linear function.

$$\mathcal{L}_{WGAN} = \min_G \max_D \mathbb{E}_{y \sim P_r} [D(y)] - \mathbb{E}_{x \sim P_x} [D(G(x))] \quad (2)$$

We use a conditional version of the model, which generates a distribution, parametrized by the network and conditioned on a conditional vector, described in section V and follow the training algorithm proposed in the original paper [7], with the same hyperparameters for training.

III. RELATED WORK

GANs have been adapted for TTS in several variations over recent years. The work closest to ours was the one proposed by Zhao et al. [8], which uses a Wasserstein-GAN framework, followed by a WaveNet vocoder and a complimentary waveform based loss. Yang et al. [9] use the mean squared error (MSE) and a variational autoencoder (VAE) to enable the GAN optimization process in a multi-task learning framework. A BLSTM based GAN framework complemented with a style and reconstruction loss is used by Zhao et al. [10]. While these models use recurrent networks for frame-wise sequential prediction, we propose a convolutional network based system to directly predict a block of vocoder features, based on an input conditioning of the same size in the time dimension.

Other examples of the application of GANs for speech synthesis include work done by Kaneko et al., [11] and [12], which use GANs as a post-filter for acoustic models to overcome the oversmoothing related to the models used. GANs have also been adapted to synthesize waveforms directly; WaveGAN [13] is an example of the use of GANs to synthesize spoken instances of numerical digits, as well as other audio examples. GANSynth [14] has also been proposed to synthesize high quality musical audio using GANs.

For singing voice synthesis, Hono et al. [15] use a GAN-based architecture to model frame-wise vocoder features. This models the inter-feature dependencies within a frame of the output. In contrast, our model directly models a block of consecutive audio frames via the Wasserstein-GAN framework. This allows us to model temporal dependencies between features within that block. This temporal dependence is modelled via autoregression in the Neural Parametric Singing Synthesizer (NPSS) [2] model, which we use as a baseline in our study.

The NPSS uses an autoregressive architecture, inspired by the WaveNet [1], to make frame-wise predictions of vocoder features, using a mixture density output. The network models the frame-wise distribution as a sum of Gaussians, the parameters of which are estimated by the network. In contrast, the use of adversarial networks for estimation, imposes no explicit constraints on the output distribution. The NPSS model has been shown to generate high quality singing voice synthesis, comparable or exceeding state-of-the-art concatenative methods. A multi-singer variation of the NPSS model has also been proposed recently [16], and is used as the baseline for our study.

Fig. 1: The framework for the proposed model. A conditioning vector, consisting of frame-wise phoneme and f_0 annotations along with speaker identity is passed to the generator. The critic is trained to distinguish between the generated sample and a real sample.

IV. PROPOSED SYSTEM

We adopt an architecture similar to the DCGAN [17], which was used for the original WGAN. For the generator, we use an encoder-decoder schema, shown in figure 2 wherein both the encoder and decoder consist of 5 convolutional layers with filter size 3. Connections between the corresponding layers of the encoder and decoder, as in the U-Net [18] architecture are also implemented.

Fig. 2: The architecture for the generator of the proposed network. The generator consists of an encoder and a decoder, based on the U-Net architecture [18].

While convolutional networks are capable of modelling arbitrary length sequences, the critic in our model takes a block of a fixed length input, thereby modelling the dependencies within the block. Furthermore, the encoder-decoder schema leads to conditional dependence between the features of the generator output, within the predicted block of features. This approach implies implicit dependence between vocoder features of a single block but not within the blocks themselves. As a result, for inference, we use overlap-add of consecutive blocks of output vocoder features, as shown in figure 3. An overlap of 50% was used with a triangular window across

features. This approach is similar to that used by Chandna et al. [19]

Fig. 3: The overlap add process for the generated features. As shown, predicted features from time t to time $t + N$ are overlap-added with features from time $t + N/2$ to $t + 3N/2$, where t is the start time of the process in view. A triangular window is used for the adding process, applied across each of the features.

As proposed by Radford et al. [17], we use strided convolutions in the encoder instead of deterministic pooling functions for downsampling. For the decoder, we use linear interpolation followed by normal convolution for upsampling instead of transposed convolutions, as this has been shown to avoid the high frequency artifacts which can be introduced by the latter [20]. Blocks of size N consecutive frames are passed as input to the network and the output has the same size. Like the DCGAN, we use ReLU activations for all layers in the generator, except the final layer, which uses a tanh activation. We found that the use of batch normalization did not affect the performance much. We also found it helpful to guide the WGAN training by adding a reconstruction loss, as shown in equations 3 and 4. This reconstruction loss is often used in conditional image generation models like Lee et al. [21].

$$\mathcal{L}_{recon} = \min_G \mathbb{E}_{x,y} \|G(x) - y\| \quad (3)$$

$$\mathcal{L}_{total} = \mathcal{L}_{WGAN} + \lambda_{recon} \mathcal{L}_{recon} \quad (4)$$

Where λ_{recon} is the weight given to the reconstruction loss. The networks are optimized following the scheme described in Arjovsky et al. [7]. The critic for our system uses an architecture similar to the encoder part of the generator, but uses LeakyReLU activation instead of ReLU, as used by Radford et al. [17].

Convolutional neural networks offer translation invariance across the dimensions convolved, making them highly useful in image modelling. However, for audio signals, this invariance is useful only across the time-dimension but undesirable across the frequency dimension. As such, we follow the approach of NPSS [2], representing the features as a 1D signal with multiple channels.

V. LINGUISTIC AND VOCODER FEATURES

The input conditioning to our system consists of frame-wise phoneme annotations, represented as a one-hot vector and

continuous fundamental frequency extracted by the spectral autocorrelation (SAC) algorithm. This conditioning is similar to the one used in NPSS. In addition, we condition the system on the singer identity, as a one-hot vector, broadcast throughout the time dimension. This approach is similar to that used in the WaveNet [1]. The three conditioning vectors are then passed through a 1×1 convolution and concatenated together along with noise sampled from a uniform distribution and passed to the generator as input.

We use the WORLD vocoder [22] for acoustic modelling of the singing voice. The system decomposes a speech signal into the harmonic spectral envelope and aperiodicity envelope, based on the fundamental frequency f_0 . We apply dimensionality reduction to the vocoder features, similar to that used in [2].

VI. DATASET

We use the NUS-48E corpus [23], which consists of 48 popular English songs, sung by 12 male and female singers. The singers are all non-professional and non-native English speakers. Each singer sings 4 different songs from a set of 20 songs, leading to a total of 169 min of recordings, with 25,474 phoneme annotations. We train the system using all but 2 of the song instances, which are used for evaluation.

VII. HYPERPARAMETERS

A hopsize of 5 ms was used for extracting the vocoder features and the conditioning. We tried different block-sizes, but empirically found that $N = 128$ frames, corresponding to 640 ms produced the best results.

We used a weight of $\lambda_{recon} = 0.0005$ for $\mathcal{L}_{recon}$ and trained the network for 3000 epochs. As suggested in [7], we used RMSProp for network optimization, with a learning rate of 0.0001. After dimension reduction, we used 60 harmonic and 4 aperiodic features per frame, leading to a total of 64 vocoder features.

VIII. EVALUATION METHODOLOGY

For objective evaluation, we use the Mel-Cepstral Distortion metric. This metric is presented in table I. For subjective evaluation, we used an online AB test wherein the participants were asked to choose between two presented 5 s to 7 s examples², representing phrases from the songs. The participant's choice was based on the criteria of Intelligibility and Audio Quality. We compare our system to the NPSS, trained on the same dataset. Along with the NPSS, we use a re-synthesis with the WORLD vocoder as the baseline as this is the upper limit of the performance of our system. We compared 3 pairs for this evaluation:

- WGANsing - Original song re-synthesized with WORLD vocoder.
- WGANsing - NPSS

²We found that WGANsing without the reconstruction loss as a guide did not produce very pleasant results and did not include this in the evaluation. However, examples for the same can be heard at https://pc2752.github.io/sing_synth_examples/

- WGANsing, original singer - WGANsing, sample with different singer.

Fig. 4: Subjective test results for the WGANsing-NPSS pair.

For the synthesis with a changed singer, we included samples with both singers of the same gender as the original singer and of a different gender. The input f_0 to the system was adjusted by an octave to account for the different ranges of the genders. For each criteria, the participants were presented with 5 questions for each of the pairs, leading to a total of 15 questions per criteria and 30 questions overall³.

IX. RESULTS

There were a total of 27 participants from over 10 nationalities, including native English speaking countries like the USA and England, and ages ranging from 18 to 37 in our study. The results of the tests are shown in figures 5, 4 and 6.

Fig. 5: Subjective test results for the WGANsing-Original pair.

From the first two figures, it can be seen that our model is qualitatively competitive with regards to both the original baseline and the NPSS, even though a preference is observed for the later. This result is supported by the objective measures, seen in table I, which show parity between WGANsing and the NPSS models. Figure 6 shows that the perceived intelligibility of the audio is preserved even after speaker change, even though there is a slight compromise on the audio quality.

Fig. 6: Subjective test results for the WGANsing-WGANsing Voice Change pair.

Variability in the observed results can be attributed to the subjective nature of the listening test, the diversity of

³The subjective listening test used in our study can be found at https://trompa-mtg.upf.edu/synth_eval/

participants and the dataset used, which comprises of non-native, non-professional singers. We note that there is room for improvement in the quality of the system, as discussed in next section.

Song	WGAN + $\mathcal{L}_{recon}$	WGAN	NPSS
Song 1 JLEE 05	5.36 dB	9.70 dB	5.62 dB
Song 2 MCUR 04	5.40 dB	9.63 dB	5.79 dB

TABLE I: The MCD metric for the two songs used for validation of the model. The three models compared are the NPSS [2] and the WGANsing model with and without the reconstruction loss.

X. CONCLUSIONS AND DISCUSSION

We have presented a multi-singer singing voice synthesizer based on a block-wise prediction topology. This block-wise methodology, inspired by the work done in image generation, allows the framework to model inter-block feature dependency, while long-term dependencies are handled via an overlap-add procedure. We believe that synthesis quality can further be improved by adding the previously predicted block of features as a further conditioning to current batch of features to be predicted. Further contextual conditioning, such as phoneme duration, which is currently implicitly modelled, can also be added to improve synthesis. Synthesis quality can also be improved through post-processing techniques such as the use of the WaveNet vocoder on the generated features.

We also note that the synthesis was greatly helped by the addition of the reconstruction loss, while the use of batch normalization did not affect performance either way. The synthesis quality of the model was evaluated to be comparable to that of state-of-the-art synthesis systems, however, we note that variability in subjective measures like intelligibility and quality is introduced during the listening test, owing to the diversity of the population participating. This variability can be reduced through a targeted listening test with expert participants. The generative methodology used allows for potential exploration in expressive singing synthesis. Furthermore, the fully convolutional nature of the model leads to faster inference than autoregressive or recurrent network based models.

ACKNOWLEDGMENTS

This work is partially supported by the European Commission under the TROMPA project (H2020 770376). The TITAN X used for this research was donated by the NVIDIA Corporation.

REFERENCES

- [1] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *SSW*, 2016.
- [2] M. Blaauw and J. Bonada, "A neural parametric singing synthesizer," in *INTERSPEECH*, 2017.
- [3] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan *et al.*, "Natural tts synthesis by conditioning wavenet on mel spectrogram predictions," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 4779–4783.
- [4] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014.
- [5] M. Mirza and S. Osindero, "Conditional generative adversarial nets," *CoRR*, vol. abs/1411.1784, 2014.
- [6] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen, "Improved techniques for training gans," in *Advances in neural information processing systems*, 2016, pp. 2234–2242.
- [7] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *International Conference on Machine Learning*, 2017, pp. 214–223.
- [8] Y. Zhao, S. Takaki, H.-T. Luong, J. Yamagishi, D. Saito, and N. Mine-matsu, "Wasserstein gan and waveform loss-based acoustic model training for multi-speaker text-to-speech synthesis systems using a wavenet vocoder," *IEEE Access*, vol. 6, pp. 60478–60488, 2018.
- [9] S. Yang, L. Xie, X. Chen, X. Lou, X. Zhu, D. Huang, and H. Li, "Statistical parametric speech synthesis using generative adversarial networks under a multi-task learning framework," in *2017 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2017, pp. 685–691.
- [10] S. Ma, D. McDuff, and Y. Song, "A generative adversarial network for style modeling in a text-to-speech system," in *International Conference on Learning Representations*, 2019.
- [11] T. Kaneko, H. Kameoka, N. Hojo, Y. Ijima, K. Hiramatsu, and K. Kashino, "Generative adversarial network-based postfilter for statistical parametric speech synthesis," in *Proc. ICASSP*, vol. 2017, 2017, pp. 4910–4914.
- [12] T. Kaneko, S. Takaki, H. Kameoka, and J. Yamagishi, "Generative adversarial network-based postfilter for stft spectrograms," in *Proceedings of Interspeech*, 2017.
- [13] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *International Conference on Learning Representations*, 2019.
- [14] J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts, "Gansynth: Adversarial neural audio synthesis," in *International Conference on Learning Representations*, 2018.
- [15] Y. Hono, K. Hashimoto, K. Oura, Y. Nankaku, and K. Tokuda, "Singing voice synthesis based on generative adversarial networks," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6955–6959.
- [16] M. Blaauw, J. Bonada, and R. Daido, "Data efficient voice cloning for neural singing synthesis," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019.
- [17] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *CoRR*, vol. abs/1511.06434, 2016.
- [18] 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.
- [19] P. Chandna, M. Miron, J. Janer, and E. Gómez, "Monoaural audio source separation using deep convolutional neural networks," in *International Conference on Latent Variable Analysis and Signal Separation*. Springer, 2017, pp. 258–266.
- [20] D. Stoller, S. Ewert, and S. Dixon, "Wave-u-net: A multi-scale neural network for end-to-end audio source separation," in *ISMIR*, 2018.
- [21] S. Lee, J. Ha, and G. Kim, "Harmonizing maximum likelihood with GANs for multimodal conditional generation," in *International Conference on Learning Representations*, 2019.
- [22] 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.
- [23] Z. Duan, H. Fang, B. Li, K. C. Sim, and Y. Wang, "The nus sung and spoken lyrics corpus: A quantitative comparison of singing and speech," in *2013 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*. IEEE, 2013, pp. 1–9.