
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Palkama, Kasper; Juvela, Lauri; Ilin, Alexander
Conditional Spoken Digit Generation with StyleGAN

Published in:
Proceedings of Interspeech


DOI:
[10.21437/Interspeech.2020-1461](https://doi.org/10.21437/Interspeech.2020-1461)

Published: 01/01/2020

Document Version
Publisher's PDF, also known as Version of record

Please cite the original version:
Palkama, K., Juvela, L., & Ilin, A. (2020). Conditional Spoken Digit Generation with StyleGAN. In *Proceedings of Interspeech* (pp. 3166-3170). (Interspeech). International Speech Communication Association (ISCA).
<https://doi.org/10.21437/Interspeech.2020-1461>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Conditional Spoken Digit Generation with StyleGAN

Kasper Palkama, Lauri Juvela, Alexander Ilin

Aalto University, Espoo, Finland

kasper.palkama@gmail.com, lauri.juvela@aalto.fi, alexander.ilin@aalto.fi

Abstract

This paper adapts a StyleGAN model for speech generation with minimal or no conditioning on text. StyleGAN is a multi-scale convolutional GAN capable of hierarchically capturing data structure and latent variation on multiple spatial (or temporal) levels. The model has previously achieved impressive results on facial image generation, and it is appealing to audio applications due to similar multi-level structures present in the data. In this paper, we train a StyleGAN to generate mel-spectrograms on the Speech Commands dataset, which contains spoken digits uttered by multiple speakers in varying acoustic conditions. In a conditional setting our model is conditioned on the digit identity, while learning the remaining data variation remains an unsupervised task. We compare our model to the current unsupervised state-of-the-art speech synthesis GAN architecture, the WaveGAN, and show that the proposed model outperforms according to numerical measures and subjective evaluation by listening tests.

Index Terms: speech synthesis, generative adversarial networks, deep learning

1. Introduction

Speech synthesis using neural networks has seen rapid advancements in recent years, and deep learning is a fundamental component for building state-of-the-art text-to-speech (TTS) applications [1]. A major factor in these improvements has been the adoption of generative models, such as WaveNet [2]. Conversely, speech data is interesting for benchmarking the performance of generative models, as it contains both deterministic structure and stochastic variation at multiple levels, ranging from utterance, word and phoneme level all the way to short-time segmental signal characteristics.

Generative adversarial networks (GANs) [3] have attracted enormous research attention since their introduction and have since achieved high synthesis quality in the image domain [4, 5]. In the audio and speech domain, research on GANs for unconditional (or weakly conditioned) synthesis tasks has been fairly limited. Convolutional GAN has demonstrated a capability to synthesize various raw audio sounds (WaveGAN) [6], although at limited quality, while [7] applied GANs on pitch-conditioned instrument sound synthesis.

Meanwhile in text-to-speech synthesis, GANs have been applied to the two sub-problems that constitute the current state-of-the-art: first, an acoustic model learns a mapping from a text sequence to acoustic features (i.e., mel-spectrogram in [1]), and second, a waveform generator model maps the acoustic features to a speech waveform. For the acoustic modeling task, adversarial training has been used as an auxiliary objective alongside conventional regression [8, 9] or as a generative post-filter to add stochasticity to regression-based predictions [10]. GAN-based waveform synthesis models that are capable of fast parallel inference have recently been proposed as alternative to autoregressive WaveNets for the latter task [11, 12, 13]. However,

all these approaches use strong conditioning that is either trivially aligned to the generated data or relies on external systems for alignment. Furthermore, they typically combine adversarial training with a regression task, which makes it difficult to assess the GAN performance in isolation.

In this work, we propose a GAN-based approach for generating mel-spectrograms of spoken digits from the speech commands dataset [14] using only a global conditioning on the digit identity and a purely adversarial training objective. In addition to the labeled variation of digits, the data contains a large amount of unlabelled variation in terms of different speakers and acoustic environments. This makes the dataset interesting for evaluating GANs ability of capturing the variation in an unsupervised manner. On the other hand, the task is connected to TTS acoustic modeling, and the present research takes a first step toward building a purely GAN-based acoustic model. The appeal of GAN for acoustic modeling includes not only parallel inference, but also the potential to capture and recreate non-annotated variation in the data.

The proposed architecture is a conditional version of the style-based generator architecture for generative adversarial networks (StyleGAN) [15]. Experimental results show that the proposed method outperforms a DCGAN baseline in various objective metrics, as well as subjective naturalness evaluation by listening tests.

2. Model architecture

We train a StyleGAN model to generate mel-spectrograms [16], which are commonly used to represent audio signals (see Fig. 2 for illustration). Two variants of the model are presented: first, without any labeled conditioning, and second, a conditional model which receives the spoken utterance contents as an additional input. Our GAN model follows the design of the original StyleGAN with a few differences. The generation starts by sampling a random latent variable $\mathbf{z} \sim N(0, \mathbf{I})$ from the normal distribution. Optionally, \mathbf{z} is appended with a known conditioning vector \mathbf{c} and transformed with a multilayer perceptron network (a mapping network f) to produce a latent code

$$\mathbf{w} = f(\mathbf{z}, \mathbf{c}).$$

The additional input of the mapping network is the learned embedding of the word (digit class in this paper) that needs to be generated. The mapping network consists of an eight-layer fully connected network with the leaky ReLU activation functions and latent code normalization (normalization is done by dividing the latent code by the standard deviation computed from its elements). Vector \mathbf{c} is simply used as an extra input to each of the hidden layers of the mapping network [17] (see Fig. 1).

The latent code \mathbf{w} is used to modulate the generation process done by the synthesis network which is a convolutional network that transforms a constant 4×4 map with 128 channels into a 128×128 mel-spectrogram. Similarly to the orig-

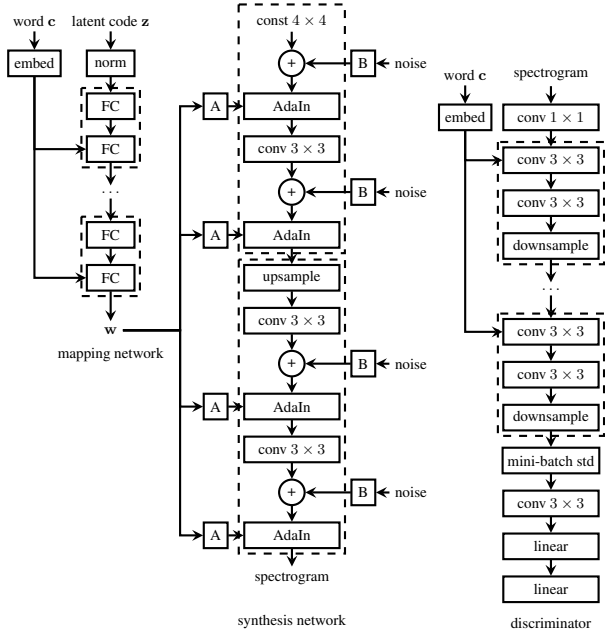


Figure 1: The proposed GAN architecture for limited-length audio generation. Like in the original StyleGAN, A is a learned affine transformation and B is a learned per-channel scaling factor, AdaIn is an adaptive instance normalization layer.

inal StyleGAN, the synthesis network consists of several convolutional blocks that now operate at different time-frequency scales. Each block contains an upsampling layer, two convolutional 3×3 layers, and two adaptive instance normalization layers:

$$\text{AdaIN}(\mathbf{x}_i, \mathbf{y}) = \mathbf{y}_{s,i} \frac{\mathbf{x}_i - \mu(\mathbf{x}_i)}{\sigma(\mathbf{x}_i)} + \mathbf{y}_{b,i}, \quad (1)$$

where \mathbf{x}_i is the i -th feature map of the input, $\mu(\mathbf{x}_i)$ and $\sigma(\mathbf{x}_i)$ are the mean and standard deviation computed from \mathbf{x}_i , and $\mathbf{y}_{s,i}$, $\mathbf{y}_{b,i}$ are the inputs controlling the normalization process. The style vectors $\mathbf{y} = (\mathbf{y}_s, \mathbf{y}_b)$ are computed using learned affine transformations of the latent codes \mathbf{w} (blocks A in the figure). To generate extra stochastic details, each block of the synthesis network contains an independent per-layer noise source, which is a single-channel Gaussian noise image broadcasted to match all the channels using learned scaling factors (blocks B in Fig. 1). In the generation process during training, we also use a style-mixing regularization proposed in [15]: blocks in the synthesis network derives their style vector \mathbf{y} from two different realization of \mathbf{z} .

Our discriminator is formed of several repetitive blocks as well. Each block contains two convolutional layers with 3×3 kernels and a downsampling layer. The final block in the discriminator starts with a mini-batch standard deviation layer which is followed by a convolution layer and two linear layers. The first linear layer has the leaky ReLU activation function and the second one is plain linear.

We enhance the discriminator by providing it with the information about the desired class of the generated mel-spectrogram. We do this by concatenating the learned embedding of the desired class with intermediate feature maps produced at the beginning of each discriminator block (see Figure 1). The upsampling and downsampling layers are imple-




Figure 2: Mel-spectrogram of real utterances (left) and mel-spectrograms generated conditionally on the word (right). The first row contains examples of word “zero” and the second row contains examples of word “four”.

mented as it was done in the original StyleGAN model. All convolution layers use leaky ReLU activations with a leak factor 0.2 both in the discriminator and the generator.

3. Experimental setup

3.1. Speech commands dataset

To train the model, we used the Speech Commands dataset [14]. The dataset consists of 105,829 utterances of 35 short common words as a one-second or less WAVE format files. The sound samples have been uttered by a large variety of different speakers in different acoustic environments, and the data typically includes only a few samples per speaker. All samples in the dataset have been quality controlled by rejecting a sample if a human listener could not tell what word was being spoken, or it sounded like an incorrect word. We decided to use a subset of the dataset to limit the scope of the problem and focused on learning the digits from zero to nine (18,620 samples). The dataset uses a 16 kHz sample rate.

3.2. Conversion between audio and mel-spectrograms

To compute mel-spectrograms from raw audio, we first generate the linear-frequency spectrogram with the short-time Fourier transform (STFT), then apply a mel-filterbank transformation to map the magnitudes to a mel scale, and finally convert the resulting mel-spectrogram to a decibel scale via a logarithm. For the STFT, we use a 50 ms frame size, 12.5 ms frame hop, and the Hann window function. The mel filterbank consists of 128 filters spanning from 125 Hz to 7.6 kHz, and the filterbank output magnitudes are clipped to a minimum value of 0.01 in order to limit the dynamic range in the logarithmic domain and then compressed to log dynamic range.

For waveform generation from the mel-spectrogram, we first invert the logarithm and apply a mel filterbank pseudoinverse transform to obtain a linear scale magnitude spectrogram, followed by the Griffin-Lim [18] algorithm to generate the sound samples. Although neural waveform generation methods can result in higher synthesis quality, we opt to use this well-known signal processing method to focus on the mel-spectrogram synthesis performance and exclude potential interaction effects in a fully neural pipeline.

3.3. Details of model training

The audio files were first transformed into mel-spectrograms to create the training data for the StyleGAN. The StyleGAN was trained with many design choices borrowed from [15]. We used progressive growing to start generating mel-spectrograms at resolution 8×8 and gradually growing the size of the generated mel-spectrograms to 128×128 [4]. The lower-resolution mel-spectrograms were generated by bilinear interpolation of the 128×128 size melspectrograms. The mini-batch size was decreased during the progressive growing to maintain an approximately constant tensor volume in network activations throughout the training. We start with the mini-batch size of 256 for 8×8 resolution and halve the mini-batch size after we introduce a new higher-resolution layer until we reach the mini-batch size of 32 for the final 128×128 resolution. Due to the variable batch size, training progress is measured in samples-introduced-to-the-network instead of epochs or minibatch iterations. During progressive growing, we introduce 200,000 mel-spectrograms while fading in a new layer, and train the network for another 200,000 samples after the layer has been fully faded in. After progressive growing is finished, we train the network until 4.05M mel-spectrograms have been introduced to the network in total. This takes about three days on one NVIDIA Tesla V100 GPU.

We trained the StyleGAN using the WGAN-GP loss [19] with various modifications proposed in [4]. The discriminator and generator were optimized using minibatch updates at equal update schedules (i.e., the discriminator is updated once for each update of the generator). Additionally, we included a regularization loss term on the mean-squared discriminator activations on the real data to prevent the training from engaging in a “magnitude race” described in [4]. Specifically, we used an extra term $\epsilon E_{x \in \text{real}} [D(x)^2]$ in the discriminator loss which prevents a magnitude drift of the discriminator output $D(x)$ (ϵ was set to 0.001).

All weights in convolutional, fully-connected, and affine transform layers were initialized with values drawn randomly from the standard normal distribution. All bias terms and the constant 4×4 feature map in the synthesis network were initialized with zeros. We used the Adam optimizer [20] with $\alpha = 0.001$, $\beta_1 = 0.0$, $\beta_2 = 0.99$, and $\epsilon = 10^{-8}$. However, for the final 128×128 resolution, we increased the learning rate to 0.0015. Following [15], we also reduced the learning rate for the mapping net by two orders of magnitude. We did not use an exponential running average for the weights of the generator.

4. Results

We evaluate the quality of the generated audio samples using two objective metrics: the first method builds upon the Fréchet inception distance [21] and the second method evaluates the error rates of the output of the Deep Speech automatic speech recognition system [22, 23] run on the generated audio samples.

Table 1: Fréchet distance for various generator designs (the smaller the better).

Model	Label cond.	Progressive growing	Style-mix regul.	FD
StyleGAN-U1	no	no	no	49.0
StyleGAN-U2	no	yes	no	27.1
StyleGAN-C1	yes	yes	no	41.6
StyleGAN-C2	yes	yes	yes	31.3

Furthermore, we conducted a MOS listening test to evaluate a the subjective quality of the generated samples.

4.1. Evaluation with Fréchet distance

Fréchet distance (FD) [21] is commonly used in GAN research to measure the quality and variation of generated samples by examining the activation statistics of a pre-trained classifier model. In image generation applications, an Inception net classifier is typically used [24], but we found that simply viewing spectrograms as monochrome images provided inconsistent results. Instead, we use two domain-specific classifiers to score the generated samples in terms of content and speaker variation.

To score the generated samples in terms of content, we trained a classifier on the spoken digits dataset to recognize the ten different digit classes from mel-spectrograms. The architecture of the classifier was similar to the one used in [6] to classify melspectrograms to compute the inception score. We modified that architecture by adding an extra average-pooling layer to produce the activations used for computing the Fréchet distance (FD) score. The classifier achieved a 97% accuracy on the test set of 2552 samples after training with 150,000 mel-spectrograms. The FD score was computed as

$$FD = \|\mathbf{m}_r - \mathbf{m}_g\|_2^2 + \text{Tr}(\mathbf{C}_r + \mathbf{C}_g - 2(\mathbf{C}_r \mathbf{C}_g)^{1/2}), \quad (2)$$

where $\text{Tr}()$ denotes the trace of a matrix, \mathbf{m}_r , \mathbf{C}_r are the mean and covariance matrix computed from the classifier activations on real data samples, and \mathbf{m}_g , \mathbf{C}_g are the same statistics computed from the classifier activations on the generated data.

Table 1 presents FD scores computed using the described digit-classifier for different designs of the StyleGAN generator. The scores were calculated multiple times during the course of training using the training set (18,620 mel-spectrograms). The lowest values of the FD scores are reported in Table 1. We can clearly see that progressive growing and style-mixing regularization improve the FD score. Label conditioning slightly decreases the quality of the generated mel-spectrograms but gives a way to control the generation process.

To score the generated samples in terms of speaker variation, we computed the FD score (2) using the activations of a pre-trained speaker embedding model from [25]. The motivation for using this kind of embeddings was to provide means for measuring variation of speaker identity in the generated samples regardless of the contents (conditioning input \mathbf{c}) of the utterances. This speaker information was unlabeled in the present experiments, but the generative model should ideally learn to embed such variability in the latent code \mathbf{z} . The results are presented in Table 2. We can see that the StyleGAN model outperforms the comparison WaveGAN [6] method.

4.2. Evaluation with the Deep Speech recognizer

FD is a measure which assesses the quality of unconditional generation of mel-spectrograms. To evaluate the quality of

Table 2: Fréchet distance in speaker embedding space (the smaller the better).

Method	FD
Griffin-Lim reconstructed	0.11
StyleGAN (proposed)	0.24
WaveGAN	0.33




Figure 3: The violin plot of the total CER scores of the (generated) audio samples converted into text with the Deep Speech recognizer (the smaller the better). The thick black line represents the median CER scores.

the conditional generation, we convert the generated mel-spectrograms into audio with the Griffin-Lim transform and then attempt to decode the audio to text using a pre-trained Deep Speech end-to-end speech recognition system [22, 23]. Thus, the Deep Speech recognizer performance acts as a proxy for intelligibility evaluation. In this case, the metric for conditional generation quality is the character error rate (CER) [26] between the desired sequence of characters and the sequence produced by the Deep Speech recognizer.

We generated 500 samples for each digit with each assessed generator design and label the generated samples with 10 digit classes utilizing the classifier introduced in Section 4.1. For conditionally-generated samples we also used labels produced by the classifier instead of the ones used in the generation process. As the baseline, we also computed the CER scores for random samples from the training data. In addition, we computed the scores for the samples generated by the current state-of-the-art GAN in waveform generation, the WaveGAN [6], using a pre-trained model provided by the authors [27].

The results are presented in Fig. 3. We can see that StyleGAN models achieved lower CER scores compared to the WaveGAN. In addition, the StyleGAN model with label conditioning outperformed the one without conditioning.

4.3. Listening test

For subjective quality evaluation, we conducted a listening test on the Amazon Mechanical Turk crowdsourcing platform (limiting workers by location to the English speaking countries). The listeners were presented with samples from each system under evaluation and asked to rate the naturalness of the sample on a five-point absolute category rating scale ranging from 1 (Bad) to 5 (Excellent). Four systems were included in the comparison: “Natural” samples are unprocessed utterances from the dataset, while “Re-synthesis” samples are synthesized from unmodified mel-spectrograms of natural samples. This represents the upper limit in quality using the present waveform synthesis

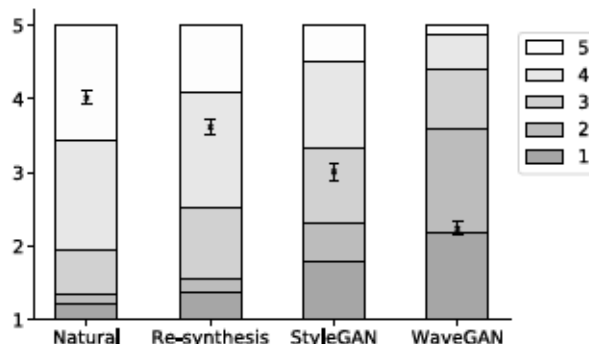


Figure 4: Naturalness mean opinion score (MOS) ratings with 95% confidence intervals. Stacked distribution histograms for the ratings are shown in the background.

method.

The tests were split into eight batches of 100 test cases, and each batch was evaluated by five individual workers. A total number of 3946 valid evaluations was collected in the listening test. Samples were drawn randomly for each system, but balanced between different digit classes. Figure 4 shows mean opinion scores (MOS) for naturalness with t-statistic based 95% confidence intervals Bonferroni corrected for multiple comparisons. Stacked histograms for the answer distributions are shown in the background. In the plot, the ratings have been averaged over listeners and digits. The results show that the proposed method StyleGAN-C2 outperforms WaveGAN.

5. Conclusion

In this paper, we adapted the recently developed StyleGAN model for speech generation with minimal or no conditioning on text. The proposed model produced higher-quality audio samples and captured better the data distribution compared to previous GAN-based speech generation models. There are two clear obstacles for the use of the presented model for generic text-to-speech mapping, but these are left as future work. First, the present model uses static conditioning, whereas a TTS system input is typically a sequence of characters (or phonemes). A natural extension would be to include a sequence encoder similar to Tacotron [1], perhaps combined with a convolutional self-attention mechanism [28]. Second, the synthesis network currently generates a fixed-length spectrogram output. Introducing a duration prediction model to generate the synthesis network input feature map could allow varying output lengths, as the network itself is otherwise fully convolutional. Another interesting direction is to move towards end-to-end training and eliminate the intermediate step of conversion between raw audio and mel-spectrograms, which may potentially further improve the quality of the generated samples.

Source code and audio samples are available at <https://github.com/kapalk/cStyleGAN>

6. Acknowledgements

We acknowledge the computational resources provided by the Aalto Science-IT project.

7. References

- [1] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerrv-Ryan, R. A. Saurous, Y. Agiomvrgiannakis, and Y. Wu, "Natural TTS synthesis by conditioning Wavenet on mel spectrogram predictions," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4779–4783.
- [2] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, "WaveNet: A generative model for raw audio," *arXiv pre-print*, 2016. [Online]. Available: <http://arxiv.org/abs/1609.03499>
- [3] 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 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 2672–2680. [Online]. Available: <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- [4] T. Karras, T. Aila, S. Laine, and J. Lehtinen, "Progressive growing of GANs for improved quality, stability, and variation," in *International Conference on Learning Representations*, 2018. [Online]. Available: <https://openreview.net/forum?id=Hk99zCeAb>
- [5] A. Brock, J. Donahue, and K. Simonyan, "Large scale GAN training for high fidelity natural image synthesis," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=B1xsqj09Fm>
- [6] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=ByMVTsR5KQ>
- [7] 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*, 2019. [Online]. Available: <https://openreview.net/forum?id=H1xQVn09FX>
- [8] Y. Saito, S. Takamichi, and H. Saruwatari, "Statistical parametric speech synthesis incorporating generative adversarial networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 1, pp. 84–96, Jan 2018.
- [9] Y. Zhao, S. Takaki, H.-T. Luong, J. Yamagishi, D. Saito, and N. Minematsu, "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. 60 478–60 488, 2018.
- [10] 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*, 2017, pp. 4910–4914.
- [11] L. Juvela, B. Bollepalli, J. Yamagishi, and P. Alku, "GELP: GAN-excited linear prediction for speech synthesis from mel-spectrogram," in *Proc. Interspeech*, 2019, pp. 694–698.
- [12] K. Kumar, R. Kumar, T. de Boissiere, L. Gestein, W. Z. Teoh, J. Sotelo, A. de Brébisson, Y. Bengio, and A. C. Courville, "Mel-GAN: Generative adversarial networks for conditional waveform synthesis," in *Advances in Neural Information Processing Systems 32*. Curran Associates, Inc., 2019, pp. 14 881–14 892.
- [13] M. Bińkowski, J. Donahue, S. Dieleman, A. Clark, E. Elsen, N. Casagrande, L. C. Cobo, and K. Simonyan, "High fidelity speech synthesis with adversarial networks," in *Proc. ICLR*, 2020. [Online]. Available: <https://openreview.net/forum?id=r1gfQgSFDr>
- [14] P. Warden, "Speech Commands: A Dataset for Limited-Vocabulary Speech Recognition," *arXiv e-prints*, p. arXiv:1804.03209, Apr 2018.
- [15] T. Karras, S. Laine, and T. Aila, "A style-based generator architecture for generative adversarial networks," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [16] S. Davis and P. Mermelstein, "Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 28, no. 4, pp. 357–366, August 1980.
- [17] M. Mirza and S. Osindero, "Conditional Generative Adversarial Nets," *arXiv e-prints*, p. arXiv:1411.1784, Nov 2014.
- [18] D. Griffin and Jae Lim, "Signal estimation from modified short-time fourier transform," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 32, no. 2, pp. 236–243, April 1984.
- [19] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, "Improved training of Wasserstein GANs," in *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc., 2017, pp. 5767–5777.
- [20] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, Dec. 2014.
- [21] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, "GANs trained by a two time-scale update rule converge to a local Nash equilibrium," in *Advances in Neural Information Processing Systems 30*. Curran Associates, Inc., 2017, pp. 6626–6637.
- [22] A. Hannun, C. Case, J. Casper, B. Catanzaro, G. Diamos, E. Elsen, R. Prenger, S. Sathesh, S. Sengupta, A. Coates, and A. Y. Ng, "Deep Speech: Scaling up end-to-end speech recognition," *arXiv e-prints*, p. arXiv:1412.5567, Dec 2014.
- [23] Mozilla Corporation, "Deepspeech," <https://github.com/mozilla/DeepSpeech>, 2017.
- [24] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [25] T. Kaseva, A. Rouhe, and M. Kurimo, "SphereDiar – an effective speaker diarization system for meeting data," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2019.
- [26] R. W. Soukoreff and I. S. MacKenzie, "Measuring errors in text entry tasks: An application of the Levenshtein string distance statistic," in *CHI '01 Extended Abstracts on Human Factors in Computing Systems*, ser. CHI EA '01. New York, NY, USA: Association for Computing Machinery, 2001, p. 319–320. [Online]. Available: <https://doi.org/10.1145/634067.634256>
- [27] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *International Conference on Learning Representations*, 2019.
- [28] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, "Self-attention generative adversarial networks," in *International Conference on Machine Learning*, 2019, pp. 7354–7363.