



Towards Generalized Speech Enhancement with Generative Adversarial Networks

Santiago Pascual¹, Joan Serrà², Antonio Bonafonte^{1*}

¹Universitat Politècnica de Catalunya

²Telefónica Research

santi.pascual@upc.edu

Abstract

The speech enhancement task usually consists of removing additive noise or reverberation that partially mask spoken utterances, affecting their intelligibility. However, little attention is drawn to other, perhaps more aggressive signal distortions like clipping, chunk elimination, or frequency-band removal. Such distortions can have a large impact not only on intelligibility, but also on naturalness or even speaker identity, and require of careful signal reconstruction. In this work, we give full consideration to this generalized speech enhancement task, and show it can be tackled with a time-domain generative adversarial network (GAN). In particular, we extend a previous GAN-based speech enhancement system to deal with mixtures of four types of aggressive distortions. Firstly, we propose the addition of an adversarial acoustic regression loss that promotes a richer feature extraction at the discriminator. Secondly, we also make use of a two-step adversarial training schedule, acting as a warm up-and-fine-tune sequence. Both objective and subjective evaluations show that these two additions bring improved speech reconstructions that better match the original speaker identity and naturalness.

Index Terms: speech enhancement, generative adversarial networks, acoustic regression loss.

1. Introduction

Speech enhancement comprises the improvement of intelligibility and quality of speech contaminated by noise, typically in an additive form [1]. It is thus broadly applicable to scenarios that need to palliate masking artifacts over the speech, like communication systems, hearing aids, or cochlear implants, where enhancing the signal prior to amplification significantly reduces discomfort and increases intelligibility [2]. Deep learning has been broadly applied to this field, either in the form of regression of waveform samples or clean spectral components, as well as prediction of masking elements over spectral bins [3–9]. A recent trend in speech enhancement is also the use of generative adversarial networks (GANs) [10], with its first application being the speech enhancement GAN (SEGAN) [11]. Other GAN variations have been used to avoid using aligned clean/noisy corpora [12] or using other adversarial losses and/or domains [13, 14]. SEGAN has also been applied to speech regeneration as a whispered-to-voiced alaryngeal speech conversion [15], thus extending the previous denoising approach.

Although GANs were designed as an unsupervised learning strategy, they are proven to be more stable and effective when coupled with labels that are either injected as input conditionings [16, 17] or that help classify some traits about the data being generated [18, 19]. Also the successful speech synthesis

model parallel WaveNet [20] proved that a multi-task aggregation on top of the speech generative model is beneficial to improve speaker identity, prosodic, and content traits.

In this work, we first extend SEGAN towards a more generalized speech enhancement such that we recover cleaner speech out of severely distorted utterances. This generalization towards different speech distortions is interesting and directly applicable to modern communication technologies, where connections are interrupted (i.e. losing voice packets), and voice processing pipelines can distort our signals through amplifiers, compression of transmission encoders, etc. In this work we emulate these effects by introducing four applied distortions in this work, with different severity levels per distortion. Then, we propose the use of SEGAN to palliate them with three different schemes. First, a plain adversarially trained SEGAN version serves us as a baseline for our recovery experiments. Then, we propose a new regression component on top of the discriminator that serves as an additional self-supervised task, boosting the generated recovered speech quality. This new task is specially effective when we use the second modification, a two-stage adversarial training schedule as a warm up and fine-tuning sequence. Objective and subjective results show the effectiveness of both proposed mechanisms when they are jointly applied.

2. Speech enhancement GAN Review

GANs define a deep generative framework, involving two neural networks and an adversarial training mechanism [10]. The two networks are the generator (G) and the discriminator (D). Both G and D are trained with opposite objectives, hence the adversarial nature of the framework. On the one hand, D has to increasingly get better at classifying *real* or *fake* features. On the other hand, G has to make up samples realistic enough to make D fail in its prediction. This process is formulated as a minimax game between G and D , and we can use different loss functions at the output of D to express it. For instance a least squares loss makes it an LSGAN [21]. It is often interesting to add a conditioning factor to be used for a specific task, where we have some additional information $\tilde{\mathbf{x}}$ injected into G and D to perform the mapping and classification (see [22] and references therein). The formulation of G and D losses for a conditional LSGAN with these components is expressed as

$$\begin{aligned} \min_G V(G) &= \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z}), \tilde{\mathbf{x}} \sim p_{\text{data}}(\tilde{\mathbf{x}})} [(D(G(\mathbf{z}, \tilde{\mathbf{x}}), \tilde{\mathbf{x}}) - c)^2], \\ \max_D V(D) &= \frac{1}{3} \mathbb{E}_{\mathbf{x}, \tilde{\mathbf{x}} \sim p_{\text{data}}(\mathbf{x}, \tilde{\mathbf{x}})} [(D(\mathbf{x}, \tilde{\mathbf{x}}) - b)^2] \\ &+ \frac{1}{3} \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z}), \tilde{\mathbf{x}} \sim p_{\text{data}}(\tilde{\mathbf{x}})} [(D(G(\mathbf{z}, \tilde{\mathbf{x}}), \tilde{\mathbf{x}}) - a)^2] \\ &+ \frac{1}{3} \mathbb{E}_{\mathbf{x}, \tilde{\mathbf{x}}^r \sim p_{\text{data}}(\mathbf{x})} [(D(\mathbf{x}, \tilde{\mathbf{x}}^r) - a)^2], \end{aligned} \quad (1)$$

*A. Bonafonte is currently at Amazon Research, Cambridge, UK.

where $a = -1$, $b = 1$, and $c = 0$ fulfill the condition to minimize Pearson χ^2 divergence [21]. Also, $\tilde{\mathbf{x}}^r$ are conditioning samples unaligned with \mathbf{x} , thus creating a *fake* signal that captures that any misalignment between \mathbf{x} and $\tilde{\mathbf{x}}$ is a mistake [15].

In the speech enhancement setting, we have an input noisy signal $\tilde{\mathbf{x}}$ that has to be cleaned to obtain the enhanced signal $\hat{\mathbf{x}}$. SEGAN decimates and expands the input signal feature-wise, through a number of strided convolutional layers with multi-parametric rectified linear activations (PReLU) [23]. Decimation is applied until we obtain a compact representation of a few time samples \mathbf{c} . This result is concatenated with the generative noise component $\mathbf{z} \sim \mathcal{N}(0, I)$, which adds the generative stochastic behavior to the generator predictions $\hat{\mathbf{x}}$. Decimation discourages learning a trivial identity function in the reconstruction $\tilde{\mathbf{x}}$, and also accelerates the convolution operations when we go deeper in the decimation structure of the encoder (shorter sequence lengths involve faster processing). It also reduces the memory footprint by using shorter feature maps. The encoding process is then reversed in the decoding stage by means of transposed convolutions, followed again by PReLU until last layer which is \tanh to bound the output between -1 and 1 , stabilizing the adversarial process [16]. Additionally, skip connections linking the encoder and the decoder of G boosted stability and generated quality. The enhanced output signal is thus defined as the output of the decoder as $\hat{\mathbf{x}} = G(\mathbf{z}, \tilde{\mathbf{x}})$.

Successive SEGAN works have shown its applicability to new languages with little adaptation data [24] and its capacity to reconstruct damaged speech in the context of whispered-to-voiced alaryngeal speech conversion, also named dewhispering [15]. For reconstruction, however, some details of the original architecture were modified, which led to the whisper SEGAN (WSEGAN). The most critical ones were the addition of learnable skip connections and the removal of the L_1 regularization term at waveform level. This eliminated a strict temporal alignment between \mathbf{x} and $G(\mathbf{z}, \tilde{\mathbf{x}})$.

3. Generalized Speech Enhancement

In this work, we depart from the classical denoising task and consider a more general class of enhancement problems. More specifically, we consider the problem of reconstructing speech that has been degraded by a number of (different) signal manipulations, each of them being potentially highly perceptually harmful. We propose to mix up such manipulations together with several speaker identities at training time, with the objective of being able to train generative algorithms to recover multiple (and diverse) speech components simultaneously. As a first approximation towards this generalized speech enhancement task, we consider the following signal manipulations:

- **Whispered speech:** This is the distortion is similar to the one introduced in the WSEGAN work [15]. However, it differs in the fact that instead of using magnetic sensors, we synthesize whispering speech by encoding the clean speech with a vocoder, removing the log-F0 (i.e. making all frames unvoiced), and recovering the signal into a version that whispers, hence artificially removing voicing. The vocoder used is Ahocoder with the default parameters [25].
- **Bandwidth reduction:** We downsample the audio signals with different severity factors, ranging from $\times 2$ to $\times 8$, thus reducing the bandwidth. Then the generalized enhancement model must reconstruct entire frequency bands, according to the clean signals seen in training.
- **Chunk removal:** For the parts of the waveform that con-

tain speech, a random number of chunks are subtracted by inserting silences. Thus, we cancel the signal at random temporal sections that contain speech. The length of a silence is sampled from one of two distributions $\mathcal{N}(0.05, 0.025)$, and $\mathcal{N}(0.1, 0.05)$ (numerical values in seconds).

- **Clipping:** The waveform is clipped globally by different severity factors relative to the maximum absolute peak of the whole utterance (e.g., 30%). The regenerated signal thus has to re-condition the signal inside a proper, non-distorted range of amplitudes.

4. Acoustic Mapping Discriminator

To better deal with the added difficulty that the previous signal manipulations can introduce to the speech signal, we propose two modifications to the existing SEGAN pipeline: the addition of two acoustic losses and the introduction of a two-stage training schedule for applying such losses.

4.1. Acoustic Losses

In general, D can be understood as a learnable loss function, where the realistic features we want to generate are implicit in the back-propagation, and G gets better because of D 's gradient flows [10]. This builds a need for D to be a competitive feature extractor, so that the better the features extracted by D , the better G can capture reality. The use of auxiliary classifying labels in D is generally helpful (Sec. 1). Additionally, multi-task setups can boost generative modeling performance, where multiple factors of the signal are predicted at the output of the generative model to enforce modeling better perceptual qualities of the signal [20]. This is applicable to the SEGAN setup where the discriminator only learns features that lead to a *real* or *fake* decision, but is not concerned about specific factors of the signal that are important to remark the reconstructed speaker identity, prosody, or contents. Note that we can also have some additional regularization loss in the output of G too, aggregated to the adversarial loss coming from D . For instance the first version of SEGAN used an L_1 regularization term that made G output a zero-centered signal [11], and the posterior WSEGAN implementation substituted it by a less restrictive power loss [20, 26], enforcing proper energy allocation in the frequency bands.

Following the aforementioned multi-task improvement criteria, we propose the use of an additional acoustic loss that opens a new branch in the end of D . This makes D necessarily aware of additional acoustic components that can inform G about important mistakes when mismatching identities, intonations, or contents in the reconstruction without necessarily injecting any additional conditioning in the input of G besides the distorted signal $\tilde{\mathbf{x}}$. This is advantageous in terms of avoiding restricting the generalization of G , as we do not inject any code limited by a set of possible classes to specify identities, distortion types, or any other kind of information. This new D output branch grows from a certain level in the convolutional structure, where the time decimation factor is 256, thus each time-step corresponds to 16 ms stride over the waveform (emulating a short-time Fourier transform sliding window). At each output frame, it predicts a concatenation of log-power spectral (LPS) bins, Mel-frequency cepstral coefficients (MFCC), and prosodic features. We design it using this collection of features as we hypothesize they properly convey different levels of acoustic cues, ranging from power allocation expressing identity, to content and intonation clues. This branch is just in addition to the




Figure 1: Setup of new multi-task framework with acoustic loss, power loss, and adversarial outcomes (F and R).

binary output neuron that tells `real` and `fake`, and is only active when this binary activation is `real` with input \mathbf{x} or $v\tilde{x}$. The proposed acoustic loss in D and power regularization in G redefines the loss function as

$$\begin{aligned} \min_G V(G) &= \frac{1}{2} \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z}), \tilde{\mathbf{x}} \sim p_{data}(\tilde{\mathbf{x}})} [(D(G(\mathbf{z}, \tilde{\mathbf{x}}), \tilde{\mathbf{x}}) - c)^2] \\ &+ \frac{1}{2} \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z}), \mathbf{x}, \tilde{\mathbf{x}} \sim p_{data}(\mathbf{x}, \tilde{\mathbf{x}})} [(\delta(G(\mathbf{z}, \tilde{\mathbf{x}}), \tilde{\mathbf{x}}) - \Theta(\mathbf{x}))^2] \\ &+ \alpha \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z}), \mathbf{x}, \tilde{\mathbf{x}} \sim p_{data}(\mathbf{x}, \tilde{\mathbf{x}})} [\Phi(G(\mathbf{z}, \tilde{\mathbf{w}})) - \Phi(\mathbf{x})], \\ \max_D V(D) &= \frac{1}{4} \mathbb{E}_{\mathbf{x}, \tilde{\mathbf{x}} \sim p_{data}(\mathbf{x}, \tilde{\mathbf{x}})} [(D(\mathbf{x}, \tilde{\mathbf{x}}) - b)^2] \\ &+ \frac{1}{4} \mathbb{E}_{\mathbf{x}, \tilde{\mathbf{x}} \sim p_{data}(\mathbf{x}, \tilde{\mathbf{x}})} [(\delta(\mathbf{x}, \tilde{\mathbf{x}}) - \Theta(\mathbf{x}))^2] \\ &+ \frac{1}{4} \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z}), \tilde{\mathbf{x}} \sim p_{data}(\tilde{\mathbf{x}})} [(D(G(\mathbf{z}, \tilde{\mathbf{x}}), \tilde{\mathbf{x}}) - a)^2] \\ &+ \frac{1}{4} \mathbb{E}_{\mathbf{x}, \mathbf{x}^r \sim p_{data}(\mathbf{x})} [(D(\mathbf{x}, \mathbf{x}^r) - a)^2], \end{aligned}$$

where \mathbf{x} , $\tilde{\mathbf{x}}$, \mathbf{x}^r , \mathbf{z} , a , b , and c are the same as in Eq. 1, $\Phi(\mathbf{x})$ corresponds to the short-time Fourier transform magnitude in dBs (20 ms windows, 10 ms stride, and 2048 bias), and $\alpha = 10^{-3}$ is a weighting term [15]. The δ function is the discriminator output at some intermediate layer of our choice, and $\Theta(\mathbf{x})$ is the feature extractor of LPS, MFCC and prosodic features (see sec. 5.4). Hereafter, we will refer to the joint power loss regularization and acoustic loss as acoustic losses.

4.2. Adversarial Pre-Training

Whereas SEGAN with only an adversarial loss is stable and learns in a steady equilibrium (eq. 1), the addition of the acoustic losses induces a particular unbalance effect during training. The addition of these terms makes D learn quicker and converge faster in both losses, conversely making G converge slower. Importantly, both G and D should maintain an equilibrium learning from each other, so making one of them quickly better discourages the other to perform properly [10]. We hypothesize that scheduling the learning of the discriminator as a two-stage process makes the addition of these acoustic losses more effective for it first gets high-level representations classifying, and then focuses on specific speech properties by doing regressions. Hence we first do an adversarial warm up in D with Eq. 1 formulation, and then add the acoustic losses.

5. Experimental Setup

5.1. Dataset

We employ the VCTK Corpus for our experiments [27]. We select 80 speakers for training and 14 for test. We trim large silence regions to 100 ms with the help of a voice activity detector.

Table 1: Frequency of number of training active transformations using $p = 0.4$.

Num. of transforms	0	1	2	3	4
Relative frequency	0.14	0.34	0.33	0.15	0.04

After this, we get roughly 20 h of training speech and 3 h of test speech. We incorporate the aforementioned distortions (Sec. 3) with an online process jointly working with training. When we construct a training minibatch, we (1) load the clean utterance, (2) get a random ≈ 1 s chunk (16384 samples) inside the utterance, and (3) apply a series of transforms that get activated independently under a probability $p = 0.4$. This way one or two distortions coincide often, and zero (thus purely autoencoder) or more than two coincide less. Table 1 shows this activation distribution with up to the 4 mentioned combinations. In addition to being active or not, each transform has a certain severity level or factor, as mentioned before (sec. 3). The only transform with just one level of severity is the whispering transform. Hence every time a transformation is activated in the pipeline, one of the possible factors is randomly selected. These possibilities are: clip factors of 30%, 40% and 50%, signal resampling factors of 2, 4, and 8 and, in the case of chunk removal, we allow the system to zero out up to 5 chunks from speech regions. Then, inside each region, the length of the chunk is sampled from one of the two mentioned Gaussian distributions (sec. 3).

For the test split, we have two different setups, one for an objective evaluation and another one for a subjective one. For the objective test set, we proceed as with the training set. For the subjective test, we generate a special split with two subsets in order to account for two different effects present in the system outcomes. The first subset is designed to focus on generated speaker identity (*SpkID* pool). To do so, we just use the most severe version of bandwidth reduction ($\times 8$) and the whispering manipulations, so the ones that affect the most this feature. Then, for each test speaker we degrade four utterances, two with bandwidth reduction and two with whispering. The second subset focuses on the generated speech naturalness (*Nat* pool). To evaluate this aspect, we pick 4 test speakers (two male and two female) and apply clipping, resampling, and whispering to 6 utterances in total (two utterances per distortion). Both clipping and resampling have severity factors 30% and $\times 8$, respectively.

5.2. Experiments

We experiment with three models applied over the aforementioned data. Firstly, we use SEGAN in a plain adversarial setup with the LSGAN loss as in Eq. 1, where the only signal comes from `real` or `fake` decisions. This is the baseline, which is trained for 400 epochs with two-timescale update rule (TTUR) learning rates [28] $\eta_D = 4 \cdot 10^{-4}$ and $\eta_G = 1 \cdot 10^{-4}$. Secondly, we apply the acoustic losses presented in Sec. 4.1. This system trains also for 400 epochs, but the learning rates are balanced equal $\eta_D = \eta_G = 5 \cdot 10^{-5}$, because the discriminator already has an advantage with the extra signals and TTUR involved a noisier result on G . We name this approach SEGAN-Aco. Finally, we include our improved proposal, for which we pre-train the system over 100 epochs the same way we do with the baseline, and then we activate the acoustic losses for the remaining 300 epochs, also lowering the learning rates to $\eta_D = \eta_G = 5 \cdot 10^{-5}$. We name this approach SEGAN-PTAco.

5.3. Evaluation

To assess each model result we perform two evaluations. First, we run a number of objective distortion metrics, often applied in speech synthesis problems: Mel cepstral distortion (MCD; in dB) [29], F0 root mean squared error (RMSE; in Hz), and the voiced/unvoiced frame prediction error (UV) [30]. These errors give us a first clue on how close is each system to the clean original signal in terms of content, identity, and intonation.

Given the importance of perceptual scores, we also conduct a subjective evaluation with 26 listeners. This has two stages, with two tasks to be evaluated by the listeners: speaker identification and naturalness rating. For speaker identification, listeners are asked to determine how close a reconstructed signal is towards the original speaker identity. They are presented with 4 randomly-selected utterances out of the *SpkID* pool (Sec. 5.1). For each utterance, the clean reference is shown, as well as the four systems to be rated: (1) the degraded input signal to G , (2) the SEGAN baseline, (3) SEGAN-Aco, and (4) SEGAN-PTAco. The rating is as simple as ordering them by preference, such that position 1 is for closest system towards the reference and position 4 is the furthest one. For naturalness rating, 6 utterances taken randomly out of the *Nat* pool (Sec. 5.1) are shown to each listener. For each utterance, the four different aforementioned systems are shown and asked to be ranked from most natural (1) to least natural (4). There is no reference shown for this case as there is no similarity trait like a speaker identity, it is only comparing the perceptual quality of the synthesized utterances. In all subjective evaluations, systems are shown in random order at every utterance.

5.4. SEGAN Setup

We use the same kernel widths, strides, and feature map configurations as in the WSEGAN work [15]. These are kernels of width 31 in all convolutional layers for both G and D . The feature maps are incremental in the encoder and decremental in the decoder, having {64, 128, 256, 512, 1024, 512, 256, 128, 64, 1} in G and {64, 128, 256, 512, 1024} in D . The discriminator is then the one diverging from its predecessors in earlier works [11, 15]. Firstly, it has a multi-layer perceptron (MLP) with 16384 inputs (1024×16 feature map, unrolled), 256 hidden PReLU units, and the single output for *real* and *fake* predictions. Secondly, we find the acoustic branch for SEGAN-Aco and SEGAN-PTAco models at the fourth convolutional layer, which outputs 512×64 feature maps for an input of approximately 1 s at 16 kHz sampling rate. These corresponds to 64 frames, each one injected into the acoustic prediction MLP of 128 hidden PReLU units and 277 linear outputs. These outputs predict 257 log-power spectral bins, 16 MFCCs, 1 log-F0 value of the frame, 1 voiced/unvoiced frame flag, 1 frame energy coefficient, and 1 frame zero crossing rate. Mini-batches of 150 samples are used for all models. Additionally, D implements spectral normalization to avoid sudden exploding gradients leading to training collapse [31, 32] and phase shift of 5 to reduce high-frequency artifacts in the output of G [33]. All the experiments were developed with the public SEGAN framework at https://github.com/santi-pdp/segan_pytorch.

6. Results

Objective results are shown in Table 2. We can see how, in all metrics, SEGAN-PTAco is the best system and the least noisy one. A lower MCD can be assumed to correlate with better con-

Table 2: *Objective metrics for the different systems (standard deviation into parenthesis). For each metric, lower is better.*

Model	MCD [dB]	RMSE [Hz]	UV [%]
Distorted	7.6 (5.6)	52.9 (75.9)	22.3 (25.5)
SEGAN	8.2 (3.5)	37.2 (43.5)	9.1 (11.4)
SEGAN-Aco	7.4 (2.7)	37.5 (53.5)	20.0 (23.4)
SEGAN-PTAco	6.5 (2.7)	22.7 (28.5)	5.8 (6.3)

Table 3: *Subjective mean ranking score (lower is better).*

Model	Speaker ID	Naturalness
Distorted	2.87 (1.07)	3.00 (1.04)
SEGAN	2.85 (1.03)	2.69 (1.04)
SEGAN-Aco	2.88 (0.83)	2.78 (0.92)
SEGAN-PTAco	1.41 (0.74)	1.52 (0.84)

tent preservation, identity reconstruction, and naturalness generation [34]. The lower values obtained for F0-RMSE and UV error also typically denote better intonation schemes matching the input identity with prosodic contents. The distorted signals have the noisiest behaviors, with large variances in all error metrics. Also expectedly, we observed that increasing the number of signal degradations yielded worse objective metrics.

Subjective results are shown in Table 3. In general, they also highlight the importance of both the acoustic losses and the two-stage adversarial training schedule. We can see that the baseline and SEGAN-Aco are clustered together with the distorted signals for the Speaker ID test. This implies that, for both bandwidth extension and dewhispering problems, neither systems reconstruct a proper speaker identity from the input signal. A qualitative listening in this case tells us that the baseline imposes an arbitrary identity to the reconstruction, and SEGAN-Aco sounds robotic and muffled. In contrast, SEGAN-PTAco is consistently ranked above the distorted signal, and we can clearly hear a competitive identity consistency with the conditioning. In the naturalness task, the baseline and SEGAN-Aco systems do better than the distorted signals. This was expected, as they involve an enhancement process that removes very noticeable artifacts and recover speech nuances. Nevertheless, SEGAN-PTAco is still the best system by a large margin. Some audio samples are available at <http://veu.talp.cat/gsegan/>.

7. Conclusion

In this work, we propose the application of speech enhancement GANs in a more generalized signal recovery framework. We first introduce four aggressive distortions applied in our experiments emulating aggressive distortions to the speech signals. Next, we introduce a new acoustic regression component in the SEGAN D and, in addition, we propose a two-stage adversarial training method to make the new acoustic regression component reach an appropriate performance. Objective and subjective results are provided, showing how both the regression losses and the two-stage schedule outperform the regression losses alone and the purely adversarial approach.

8. Acknowledgements

This research was supported by the project TEC2015-69266-P (MINECO/FEDER, UE). We deeply thank the participants of the subjective evaluation.

9. References

- [1] P. C. Loizou, *Speech Enhancement: Theory and Practice*, 2nd ed. Boca Raton, FL, USA: CRC Press, Inc., 2013.
- [2] L.-P. Yang and Q.-J. Fu, "Spectral subtraction-based speech enhancement for cochlear implant patients in background noise," *Journal of the Acoustical Society of America*, vol. 117, no. 3, pp. 1001–1004, 2005.
- [3] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, "Speech enhancement based on deep denoising autoencoder," in *Proc. of the Conf. of the Int. Speech Communication Association (INTERSPEECH)*, 2013, pp. 436–440.
- [4] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Transactions on Audio, Speech and Language Processing*, vol. 23, no. 1, pp. 7–19, 2015.
- [5] D. S. Williamson and D. Wang, "Time-frequency masking in the complex domain for speech dereverberation and denoising," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 7, pp. 1492–1501, 2017.
- [6] F. Weninger, H. Erdogan, S. Watanabe, E. Vincent, J. Le Roux, J. R. Hershey, and B. Schuller, "Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR," in *Proc. of the Int. Conf. on Latent Variable Analysis and Signal Separation*, 2015, pp. 91–99.
- [7] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 708–712.
- [8] D. Rethage, J. Pons, and X. Serra, "A wavenet for speech denoising," in *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5069–5073.
- [9] S. R. Park and J. Lee, "A fully convolutional neural network for speech enhancement," in *Proc. of the Conf. of the Int. Speech Communication Association (INTERSPEECH)*, 2017.
- [10] 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 (NeurIPS)*, 2014, pp. 2672–2680.
- [11] S. Pascual, A. Bonafonte, and J. Serrà, "Segan: Speech enhancement generative adversarial network," in *Proc. of the Conf. of the Int. Speech Communication Association (INTERSPEECH)*, 2017, pp. 3642–3646.
- [12] T. Higuchi, K. Kinoshita, M. Delcroix, and T. Nakatani, "Adversarial training for data-driven speech enhancement without parallel corpus," in *IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2017, pp. 40–47.
- [13] S. Qin and T. Jiang, "Improved wasserstein conditional generative adversarial network speech enhancement," *EURASIP Journal on Wireless Communications and Networking*, vol. 2018, no. 1, p. 181, 2018.
- [14] C. Donahue, B. Li, and R. Prabhavalkar, "Exploring speech enhancement with generative adversarial networks for robust speech recognition," in *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5024–5028.
- [15] S. Pascual, A. Bonafonte, J. Serrà, and J. A.-G. López, "Whispered-to-voiced alaryngeal speech conversion with generative adversarial networks," *Proc. IberSPEECH 2018*, pp. 117–121, 2018.
- [16] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *ArXiv: 1511.06434*, 2015.
- [17] A. Brock, J. Donahue, and K. Simonyan, "Large scale gan training for high fidelity natural image synthesis," *arXiv preprint arXiv:1809.11096*, 2018.
- [18] A. Odena, C. Olah, and J. Shlens, "Conditional image synthesis with auxiliary classifier gans," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 2642–2651.
- [19] M. Lucic, M. Tschannen, M. Ritter, X. Zhai, O. Bachem, and S. Gelly, "High-fidelity image generation with fewer labels," *arXiv preprint arXiv:1903.02271*, 2019.
- [20] A. v. d. Oord, Y. Li, I. Babuschkin, K. Simonyan, O. Vinyals, K. Kavukcuoglu, G. v. d. Driessche, E. Lockhart, L. C. Cobo, F. Stimberg *et al.*, "Parallel wavenet: Fast high-fidelity speech synthesis," *arXiv preprint arXiv:1711.10433*, 2017.
- [21] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. P. Smolley, "Least Squares Generative Adversarial Networks," in *Proc. of the IEEE Int. Conf. on Computer Vision (ICCV)*. IEEE, 2017, pp. 2813–2821.
- [22] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proc. of the IEEE Int. Conf. on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017, pp. 5967–5976.
- [23] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," in *Proc. of the IEEE Int. Conf. on Computer Vision (ICCV)*, 2015, pp. 1026–1034.
- [24] S. Pascual, M. Park, J. Serrà, A. Bonafonte, and K.-H. Ahn, "Language and noise transfer in speech enhancement generative adversarial network," in *Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5019–5023.
- [25] D. Erro, I. Sainz, E. Navas, and I. Hernáez, "Improved hnm-based vocoder for statistical synthesizers," in *Twelfth Annual Conference of the International Speech Communication Association*, 2011.
- [26] W. Ping, K. Peng, and J. Chen, "Clarinet: Parallel wave generation in end-to-end text-to-speech," *arXiv preprint arXiv:1807.07281*, 2018.
- [27] C. Veaux, J. Yamagishi, K. MacDonald *et al.*, "Superseded-cstr vctk corpus: English multi-speaker corpus for cstr voice cloning toolkit," 2016.
- [28] 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*, 2017, pp. 6626–6637.
- [29] R. Kubichek, "Mel-cepstral distance measure for objective speech quality assessment," in *Proc. of the IEEE Pacific Rim Conf. on Communications Computers and Signal Processing (PACRIM)*, vol. 1. IEEE, 1993, pp. 125–128.
- [30] S. Pascual and A. Bonafonte, "Multi-output rnn-lstm for multiple speaker speech synthesis and adaptation," in *Proc. of the European Signal Processing Conference (EUSIPCO)*. IEEE, 2016, pp. 2325–2329.
- [31] H. Zhang, I. Goodfellow, D. Metaxas, and A. Odena, "Self-attention generative adversarial networks," *arXiv preprint arXiv:1805.08318*, 2018.
- [32] T. Miyato, T. Kataoka, M. Koyama, and Y. Yoshida, "Spectral normalization for generative adversarial networks," *arXiv preprint arXiv:1802.05957*, 2018.
- [33] C. Donahue, J. McAuley, and M. Puckette, "Synthesizing audio with generative adversarial networks," *ArXiv: 1802.04208*, 2018.
- [34] S. Pascual de la Puente, "Deep learning applied to speech synthesis," Master's thesis, Universitat Politècnica de Catalunya, 2016.