



# SEANet: A Multi-modal Speech Enhancement Network

Marco Tagliasacchi, Yunpeng Li, Karolis Misiunas, Dominik Roblek

Google Research

{mtagliasacchi, yunpeng, kmisiunas, droblek}@google.com

## Abstract

We explore the possibility of leveraging accelerometer data to perform speech enhancement in very noisy conditions. Although it is possible to only partially reconstruct user's speech from the accelerometer, the latter provides a strong conditioning signal that is not influenced from noise sources in the environment. Based on this observation, we feed a multi-modal input to SEANet (Sound Enhancement Network), a wave-to-wave fully convolutional model, which adopts a combination of feature losses and adversarial losses to reconstruct an enhanced version of user's speech. We trained our model with data collected by sensors mounted on an earbud and synthetically corrupted by adding different kinds of noise sources to the audio signal. Our experimental results demonstrate that it is possible to achieve very high quality results, even in the case of interfering speech at the same level of loudness. A sample of the output produced by our model is available at <https://google-research.github.io/seanet/multimodal/speech>.

**Index Terms:** speech denoising, multimodal, accelerometers.

## 1. Introduction

Enhancing the quality of speech is of paramount importance in digital communications. Speech degradation can occur for various reasons, e.g., from the interference of background noise, which can also contain overlapping speakers, to the effect of reverberations caused by room acoustics, to the artifacts introduced by compression and network impairments. This has motivated a very rich literature on speech enhancement and denoising. Traditional signal processing methods adopt spectral noise subtraction [1, 2], spectral masking [3, 4], statistical methods based on Wiener filtering [5] and Bayesian estimators [6, 7]. These methods make different assumptions about the underlying noise model (e.g., known signal-to-noise ratio (SNR), stationary noise, limited noise types, etc.), therefore they are unable to cope with challenging noisy conditions emerging when systems are deployed “in-the-wild”.

In recent years, data-driven methods have emerged, based on deep model architectures. Early works include methods based on denoising auto-encoders [8] and recurrent models [9]. More recently, deep architectures have been adopted to improve speech enhancement based on spectral masking [10]. Alternatively, generative models based on GANs [11, 12] and WaveNet [13] have been proposed. Speech denoising can also be seen as a special case of source separation, in which one of the sources represents the speech signal of interest [14, 15, 16, 17]. Our work belongs to the family of multi-modal models, which leverage additional conditioning signals to enhance the target speech. For example, the work in [18] uses tight crops of mouth images to denoise speech. This approach was later extended by Looking2Listen [19], which uses visual information from facial crops segmented from videos to disentangle different speakers that talk simultaneously. A similar ap-

proach is presented in [20], which adopts an attention mechanism to weight the contribution of the audio and visual modalities. Multi-modal cues can also be exploited for voice activity detection [21].

In this paper we consider the problem of multi-modal speech denoising. Instead of leveraging video as an additional modality, we consider data collected with a bone-conductance accelerometer mounted in an earbud, which operates synchronously with the microphone but at a lower sampling frequency. The sensor captures the local vibrations induced by the voice of the speaker, while being relatively insensitive to external sources. Hence, it can be used as a conditioning signal to enhance user's speech and suppress noise. The fact that inertial measurement sensors mounted in mobile devices can be sensitive to speech has been recognized in the past literature. For example, gyroscope signals were used to recognize speech in [22], while [23] reconstructs speech from accelerometer-sensed reverberations induced by smartphone loudspeakers. The work in [24] combines signals from a microphone and a bone sensor using a Gaussian mixture model on the high-resolution log spectra of each sensor. Similarly, multi-modal inputs are combined in [25] using deep denoising autoencoders that reconstruct Mel-scale features fed to an ASR system. An ad-hoc speech recovery stage is needed to reconstruct the time-domain denoised waveforms.

The proposed multi-modal SEANet (Speech Enhancement Network) model receives two waveforms, one acquired with a microphone and one with an accelerometer, and produces as output a denoised speech waveform. The model is fully convolutional and maps waveforms to waveforms, without resorting to explicit time-frequency representations like short-time Fourier Transform (STFT) or mel spectrograms. To train the model, we adopt a combination of adversarial and reconstruction losses inspired by the recent MelGAN model [26], which synthesizes waveforms from mel spectrograms. The adversarial losses induce the model to produce output waveforms that a discriminator cannot distinguish from clean speech. The reconstruction losses operate in the feature space defined by the discriminator and preserve speech content while suppressing noise.

In our experiments we consider challenging scenarios in which the target speech signal is mixed with that of other speakers, or different kinds of background noise sampled from Freesound [27]. We demonstrate that by leveraging the conditioning signal collected by the accelerometer, it is possible to denoise speech even in very adverse conditions. We collected a dataset that contains speech and the corresponding accelerometer readings and observed an improvement in scale-invariant signal-to-distortion ratio (SI-SDRi) of 9.6dB when the interferer is mixed with a unit gain.

## 2. Method

The proposed SEANet model is trained in a fully supervised fashion using pairs  $\langle (x_m, x_a), y_m \rangle$ , where  $x_m$  denotes the input

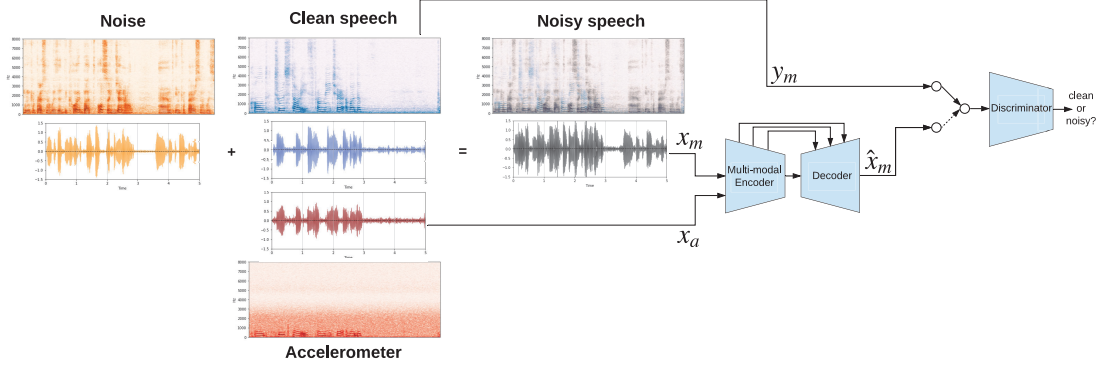


Figure 1: SEANet model overview. A noisy speech signal, obtained superimposing clean speech with a noise source, is fed to the multi-modal encoder together with the accelerometer signal. Spectrograms are shown only for illustration purposes, as they are not explicitly computed by the proposed wave-to-wave model.

noisy speech collected by the microphone,  $x_a$  the accelerometer signal used as conditioning, and  $y_m$  the target audio signal corresponding to clean speech. Note that  $x_a$  might have one or more channels, depending on the number of accelerometer axes used. We assume that  $x_m$ ,  $x_a$  and  $y_m$  are time-aligned and available at the same sampling rate. Since the sampling rate of accelerometers is typically smaller than that of the microphone, the former signal is interpolated before being fed to the model.

The model architecture consist of a UNet generator  $G(x_m, x_a)$ , which take as its input an audio  $x_m$  and one or more accelerometer readings  $x_a$  time-aligned with the audio. In Figure 1 we illustrate the case in which a single accelerometer axis is used. The generator produces as output a single-channel waveform  $\hat{x}_m$ , which represents the denoised speech. The discriminator is asked to determine whether its input comes from the distribution of clean speech, or from the output of the generator.

**Model architecture:** Our UNet generator is a symmetric encoder-decoder network with skip-connections. The decoder adopts the same architecture as the generator in [26], while the encoder mirrors the decoder in its layout. A skip-connection is added between each encoder block and its mirrored decoder block. The out-most skip connects only the speech channel needed by the output. The encoder and the decoder have each four blocks stacked together, which are sandwiched between two plain convolution layers. The encoder follows a down-sampling scheme of (2, 2, 8, 8) while the decoder up-samples in the reverse order. The number of channels is doubled whenever down-sampling and halved whenever up-sampling. Each decoder block consists of an up-sampling layer, in the form of a transposed 1D convolution, followed by three residual units each containing 1D convolutions with dilation rates of 1, 3, and 9, respectively. The encoder block again mirrors the decoder block, and consists of the same residual units followed by a strided 1D convolution for down-sampling. The overall structure of the generator is illustrated in Figure 2.

For the discriminator, we use the same multi-resolution convolutional architecture as [26]. Three structurally identical discriminators are applied to input audio at different resolutions: original, 2x down-sampled, and 4x down-sampled. Each discriminator consists an initial plain convolution followed by four grouped convolutions [28], each of which has a group size of 4, a down-sampling factor of 4, and a channel multiplier of 4 up

to a maximum of 1024 output channels. They are followed by two more plain convolution layers to produce the final output, i.e., the logits. Note that since the discriminator is fully convolutional, the number of logits in the output is more than one and proportional to the length of the input audio. Each logit judges the plausibility of a segment of the input that corresponds to its receptive field. We refer interested readers to [26] for more architectural details.

We use weight normalization [29] and ELU activation [30] in the generator, while layer normalization and Leaky ReLU activation [31] with  $\alpha = 0.3$  are used in the discriminator.

**Loss functions:** SEANet combines adversarial and reconstruction losses to train simultaneously the generator and the discriminators. The adversarial loss is a hinge loss averaged over multiple resolutions and over time. More formally, let  $k \in \{1, \dots, K\}$  index over the individual discriminators for different resolutions, and  $t$  index over the length of the output, i.e., the number of logits  $T_k$ , of discriminator  $k$ . The discriminator loss can be written as

$$\mathcal{L}_D = E_{y_m} \left[ \frac{1}{K} \sum_{k,t} \frac{1}{T_k} \max(0, 1 - D_{k,t}(y_m)) \right] + E_{(x_m, x_a)} \left[ \frac{1}{K} \sum_{k,t} \frac{1}{T_k} \max(0, 1 + D_{k,t}(G(x_m, x_a))) \right], \quad (1)$$

while the adversarial loss for the generator is

$$\mathcal{L}_G^{\text{adv}} = E_{(x_m, x_a)} \left[ \frac{1}{K} \sum_{k,t} \frac{1}{T_k} \max(0, 1 - D_{k,t}(G(x_m, x_a))) \right]. \quad (2)$$

For the reconstruction loss we use the “feature” loss proposed in [26], namely the normalized L1 distance between the discriminator internal layer output of the generated audio and that of the corresponding target audio:

$$\mathcal{L}_G^{\text{rec}} = E_x \left[ \frac{1}{K} \sum_{k,l} \frac{1}{L} \frac{\|D_k^{(l)}(y_m) - D_k^{(l)}(G(x_m, x_a))\|_1}{T_{k,l}} \right], \quad (3)$$

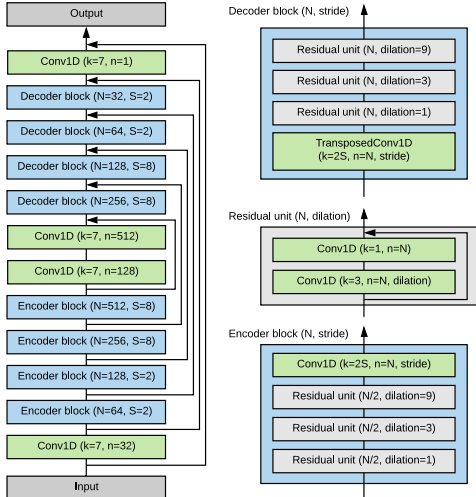


Figure 2: Generator architecture.

where  $x \triangleq \langle (x_m, x_a), y_m \rangle$  denotes a training example,  $L$  is the number of internal layers,  $D_k^{(l)}$  for  $l \in \{1, \dots, L\}$  is the output of layer  $l$  of discriminator  $k$ , and  $T_{k,l}$  is length of the feature layer  $D_k^{(l)}$ . Compared with per-sample losses, such as the average L1 distance between waveforms, the feature loss tends to be less sensitive to small misalignment. The overall generator loss is a weighted sum of the adversarial and the reconstruction loss, i.e.,

$$\mathcal{L}_G = \mathcal{L}_G^{\text{adv}} + \lambda \cdot \mathcal{L}_G^{\text{rec}}. \quad (4)$$

For all our experiments, we set the weight of the reconstruction loss  $\lambda$  to 100 and use a discriminator with  $K = 3$  scales. We train with the Adam optimizer, with a batch size of 16 and a constant learning rate of 0.0001 with  $\beta_1 = 0.5$  and  $\beta_2 = 0.9$ . We train for 200k iterations (2M iterations when training on Librispeech) on a single GPU. We evaluate results using the last checkpoint of each training run. No parameter tuning or early stopping were performed.

### 3. Experiments

**Datasets:** We collected an in-house dataset with sensors mounted on an earbud, since a dataset with these characteristics is not available in the literature. The microphone sampled audio waveforms at 16kHz, while the 2-axis accelerometer operated at 4kHz. We selected one of the two axes and interpolated the accelerometer signal at 16kHz before feeding it to the model. We then applied high-pass filtering with a cut-off of 20Hz to all signals and normalized the amplitudes dividing all samples by a factor  $1.1 \cdot \text{quantile}(x; 0.9999)$  and clipping the result in the  $[-1, +1]$  range. This is necessary to deal with isolated spikes which were present in the raw output of the accelerometer.

We asked 25 subjects to speak while wearing one earbud in a relatively quiet office environment. In total we collected  $\sim 1.25$  hours of data, with each subject contributing  $\sim 3$  minutes. We organized the data in 5 folds, so that in each fold 20 speakers are used for training and 5 speakers for testing. Figure 3 shows the power spectral density of the signals acquired by the sensors. We observe that they share a similar response in the range of 100–300Hz, while the sensitivity of the accelerometer decreases rapidly above 300Hz.

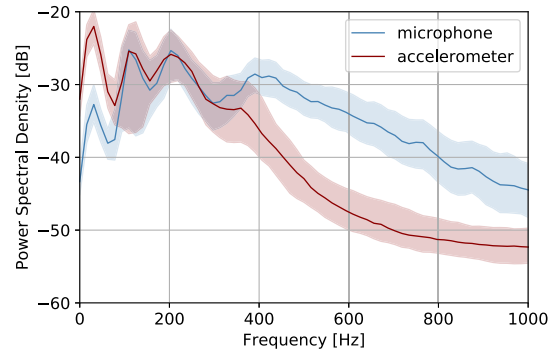


Figure 3: Power spectral density: microphone vs. accelerometer.

To explore the quality potentially achievable if we had access to more data, we created a synthetically generated multimodal dataset. First, we trained a variant of SEANet which learns to map audio waveforms to the corresponding accelerometer waveform, using the in-house dataset described above. This model uses the same architecture described in Section 2, with the only difference that it receives one input channel with clean audio and produces one output channel with the corresponding accelerometer signal. Note that learning this mapping is much easier than reconstructing audio samples from the accelerometer signal alone. Then, we fed audio samples from Librispeech [32] to synthesize the corresponding accelerometer signal. In this case, we followed the canonical split provided by Librispeech, using *train-clean-100* for training and *test-clean* for testing.

To generate the noisy input  $x_m$ , we mix the clean microphone recording  $y_m$  with other noise sources. We consider two scenarios: i) *mixed speech*, in which an utterance from a different speaker is mixed with the clean source; ii) *mixed noise*, in which we mix with samples taken at random from Freesound [27], to mimic a wide and diverse range of noise sources, with a unit mixing gain. In one of the experiments, we also limit the bandwidth of the accelerometer to simulate a sensor operating at lower sampling rates. In this case we use the following downsampling factors  $\{16, 20, 32, 40, 50, 64, 80, 100\}$ , corresponding to the sampling frequencies  $\{1000\text{Hz}, 800\text{Hz}, 500\text{Hz}, 400\text{Hz}, 320\text{Hz}, 250\text{Hz}, 200\text{Hz}, 160\text{Hz}\}$ . We also report results of an audio-only SEANet model, in which the accelerometer input is not used.

**Metrics and baselines:** In order to evaluate the quality of the enhanced speech, we measure the scale invariant signal-to-distortion ratio (SI-SDR), which accommodates for an amplitude gain mismatch between the estimated signal  $\hat{y}_m$  and the ground truth clean reference signal  $x_m$ . The SI-SDR is computed as described in [17].

We evaluated models recently proposed in the speech enhancement and separation literature, which receive as input only the audio signal. It is worth noting that a direct comparison with these methods is not meaningful, as SEANet receives as input an additional conditioning signal. However, this evaluation is useful to gauge the level of complexity of the dataset, highlighting the added value of leveraging the accelerometer signal. Namely, we include in our evaluation iTDCN++ [17] and Wavesplit [33]. The iTDCN++ model is inspired by ConvTasNet and predicts a mask with a sigmoid activation that is

Table 1: Mean SI-SDRi for the In-house dataset.

scenario	split	SEANet audio + accel	SEANet audio only
Mixed noise	1	9.9 ± 0.2	8.4 ± 0.2
	2	8.0 ± 0.2	7.9 ± 0.1
	3	8.3 ± 0.1	7.2 ± 0.2
	4	8.8 ± 0.1	8.1 ± 0.1
	5	9.9 ± 0.1	8.4 ± 0.1
	avg.	8.9	8.0
Mixed speech	1	10.1 ± 0.1	-0.9 ± 0.1
	2	8.6 ± 0.1	-0.9 ± 0.1
	3	9.2 ± 0.1	-0.7 ± 0.0
	4	9.0 ± 0.2	-1.0 ± 0.1
	5	11.1 ± 0.2	-0.9 ± 0.1
	avg.	9.6	-0.9

Table 2: Mean SI-SDRi for Librispeech.

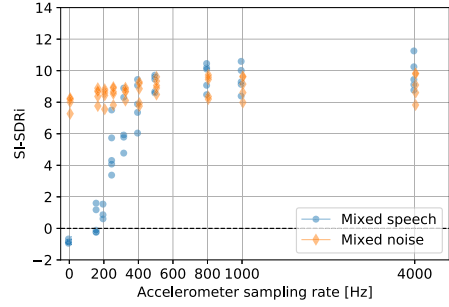
scenario	split	SEANet audio + accel	SEANet audio only
Mixed noise	test	12.4 ± 0.3	9.8 ± 0.2
Mixed speech	test	12.4 ± 0.3	-1.0 ± 0.2

applied to the mixture STFT coefficients. Wavesplit infers and clusters representations of each speaker and then estimates each source signal conditioned on the inferred representations.

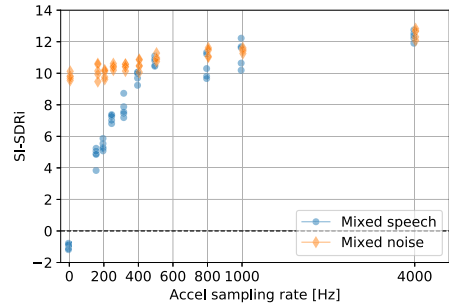
**Results:** Table 1 reports the results obtained repeating five replicas, on each of the five splits for the two scenarios. The average SI-SDRi is 8.9dB when mixing with background noise from Freesound and 9.6dB when mixing with speech. Note that the variability across replicas is small (standard deviation ±0.1 – 0.2dB), while there is a more significant variability across splits. We repeated the experiment by changing the gain used during mixing and observed that the SI-SDRi varies between 3.7dB (6.2dB), at 10dB mixing gain, and 15.0dB (15.1dB), at -10dB mixing gain for mixing noise (mixing speech). Table 1 includes results when SEANet is trained using audio only. In the mixed noise scenario, the model is still able to enhance speech, although attaining a lower SI-SDRi (7.9dB vs. 8.9dB). Conversely, in the mixed speech scenario the audio-only variant of SEANet is unable to separate the speakers. This is not surprising, since the model as described in this paper does not include a permutation invariant loss, which is needed to separate sources of the same kind. Using audio-only, iTDCN++ attains 7.5dB on mixed noise (trained on synthetically reverberated Libri-Light speech + synthetically-reverberated Freesound) and 4.2dB on mixed speech (trained on synthetically reverberated Libri-Light speech mixtures), while Wavesplit attains 8.8dB on mixed speech (trained on Librispeech mixtures, with no reverberation). This demonstrates the inherent difficulty of the in-house dataset and the fact that the availability of the conditioning signal makes the denoising problem significantly easier, especially in the scenario with mixed speech.

We also evaluate a model trained on Librispeech with synthetically generated accelerometer signals. Table 2 shows that this model achieves an SI-SDRi of 12.4dB on both mixed noise and mixed speech, thus hinting to the fact that better accuracy can be attained using a larger dataset during training. Examples of the denoised results produced by SEANet are publicly available at the following page: <https://google-research.github.io/seanet/multimodal/speech>.

We investigated the contribution of the conditioning provided by the accelerometer. To this end, we progressively dec-



(a) In-house dataset.



(b) Librispeech.

Figure 4: Improvement in SI-SDR for different accelerometer sampling rates (each point represents one replica).

imated the accelerometer signal before feeding it to our model during both training and evaluation. Figure 4a shows an interesting result. In the scenario with two overlapping speakers, a rapid decrease in SI-SDRi is observed when the sampling rate drops below 400Hz, and our model is unable to separate the speakers when the sampling rate is smaller than 200Hz. Conversely, for the scenario with background noise, only a small decrease in SI-SDRi is observed, also when the sampling rate of the accelerometer is drastically reduced. The average SI-SDRi across the splits drops from 8.9dB to 8.0dB. We can argue that this is a simpler scenario, giving the distinct acoustic characteristics of the background noise. These results are confirmed when training and evaluating on the multi-modal dataset generated from Librispeech, as illustrated in Figure 4b. In this case the average SI-SDRi drops from 12.4dB to 9.8dB.

## 4. Conclusions

In this paper we show that the accelerometer data collected from sensors mounted on earbuds provides a strong conditioning signal for speech denoising. This is especially useful in the challenging scenario with overlapping speakers. In our future work we plan to expand the multi-modal aspect of SEANet by exploring how to combine multiple microphone signals, accelerometer axes and visual cues.

## 5. Acknowledgements

We would like to thank Kevin Wilson, Scott Wisdom, John Hershey, Dick Lyon and Neil Zeghidour for their help with and feedback on this work. We also thank Alina Mihaela Stan for the help collecting the in-house dataset.

## 6. References

- [1] M. Berouti, R. Schwartz, and J. Makhoul, "Enhancement of speech corrupted by acoustic noise," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 4, 1979, pp. 208–211.
- [2] S. Kamath and P. Loizou, "A multi-band spectral subtraction method for enhancing speech corrupted by colored noise," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 4, 05 2002.
- [3] A. M. Reddy and B. Raj, "Soft mask methods for single-channel speaker separation," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 6, pp. 1766–1776, 2007.
- [4] E. M. Grais and H. Erdogan, "Single channel speech music separation using nonnegative matrix factorization and spectral masks," in *International Conference on Digital Signal Processing (DSP)*, 2011, pp. 1–6.
- [5] P. Scalart and J. V. Filho, "Speech enhancement based on a priori signal to noise estimation," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, vol. 2, 1996, pp. 629–632 vol. 2.
- [6] H. Attias, J. C. Platt, A. Acero, and L. Deng, "Speech denoising and dereverberation using probabilistic models," in *Advances in Neural Information Processing Systems 13*, 2001, pp. 758–764.
- [7] P. C. Loizou, "Speech enhancement based on perceptually motivated bayesian estimators of the magnitude spectrum," *IEEE Transactions on Speech and Audio Processing*, vol. 13, no. 5, pp. 857–869, 2005.
- [8] X. Feng, Y. Zhang, and J. Glass, "Speech feature denoising and dereverberation via deep autoencoders for noisy reverberant speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2014, pp. 1759–1763.
- [9] 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 *Latent Variable Analysis and Signal Separation*, E. Vincent, A. Yeredor, Z. Koldovský, and P. Tichavský, Eds. Cham: Springer International Publishing, 2015, pp. 91–99.
- [10] S. Wisdom, J. R. Hershey, K. Wilson, J. Thorpe, M. Chinen, B. Patton, and R. A. Saurous, "Differentiable consistency constraints for improved deep speech enhancement," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2019, pp. 900–904.
- [11] S. Pascual, A. Bonafonte, and J. Serr, "SEGAN: Speech enhancement generative adversarial network," in *INTERSPEECH*, 2017, pp. 3642–3646.
- [12] C. Donahue, B. Li, and R. Prabhavalkar, "Exploring speech enhancement with generative adversarial networks for robust speech recognition," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5024–5028.
- [13] D. Rethage, J. Pons, and X. Serra, "A wavenet for speech denoising," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5069–5073.
- [14] Y. Luo and N. Mesgarani, "TaSNet: Time-domain audio separation network for real-time, single-channel speech separation," in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 696–700.
- [15] M. Kolbck, D. Yu, Z. Tan, and J. Jensen, "Multitalker speech separation with utterance-level permutation invariant training of deep recurrent neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 10, pp. 1901–1913, 2017.
- [16] J. L. Roux, G. Wichern, S. Watanabe, A. M. Sarroff, and J. R. Hershey, "The Phasebook: Building complex masks via discrete representations for source separation," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, IEEE, 2019, pp. 66–70.
- [17] I. Kavalerov, S. Wisdom, H. Erdogan, B. Patton, K. W. Wilson, J. L. Roux, and J. R. Hershey, "Universal sound separation," in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, WASPAA 2019, New Paltz, NY, USA, October 20-23, 2019*. IEEE, 2019, pp. 175–179.
- [18] J. Hou, S. Wang, Y. Lai, Y. Tsao, H. Chang, and H. Wang, "Audio-visual speech enhancement using multimodal deep convolutional neural networks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 2, pp. 117–128, 2018.
- [19] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," *ACM Trans. Graph.*, vol. 37, no. 4, Jul. 2018. [Online]. Available: <https://doi.org/10.1145/3197517.3201357>
- [20] T. Ochiai, M. Delcroix, K. Kinoshita, A. Ogawa, and T. Nakatani, "Multimodal speakerbeam: Single channel target speech extraction with audio-visual speaker clues," in *INTERSPEECH*, 2019, pp. 2718–2722.
- [21] I. Ariav and I. Cohen, "An end-to-end multimodal voice activity detection using wavenet encoder and residual networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 265–274, 2019.
- [22] Y. Michalevsky, D. Boneh, and G. Nakibly, "Gyrophone: recognizing speech from gyroscope signals," in *Proceedings of the 23rd USENIX conference on Security Symposium*, 2014, pp. 1053–1067.
- [23] S. A. Anand, C. Wang, J. Liu, N. Saxena, and Y. Chen, "Spearphone: A speech privacy exploit via accelerometer-sensed reverberations from smartphone loudspeakers," *CoRR*, vol. abs/1907.05972, 2019. [Online]. Available: <http://arxiv.org/abs/1907.05972>
- [24] J. R. Hershey, T. T. Kristjansson, and Z. Zhang, "Model-based fusion of bone and air sensors for speech enhancement and robust speech recognition," in *ISCA Tutorial and Research Workshop on Statistical and Perceptual Audio Processing, ICC, Jeju, Korea, October 3, 2004*. ISCA, 2004, p. 139.
- [25] H. Liu, Y. Tsao, and C. Fuh, "Bone-conducted speech enhancement using deep denoising autoencoder," *Speech Commun.*, vol. 104, pp. 106–112, 2018. [Online]. Available: <https://doi.org/10.1016/j.specom.2018.06.002>
- [26] K. Kumar, R. Kumar, T. de Boissiere, L. Gestin, W. Z. Teoh, J. Sotelo, A. de Brebisson, Y. Bengio, and A. Courville, "MelGAN: Generative adversarial networks for conditional waveform synthesis," in *Advances in Neural Information Processing Systems*, 2019.
- [27] F. Font, G. Roma, and X. Serra, "Freesound technical demo," in *ACM International Conference on Multimedia (MM'13)*, ACM, Barcelona, Spain: ACM, 21/10/2013 2013, pp. 411–412.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems*, 2012, pp. 1097–1105.
- [29] T. Salimans and D. P. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 901–909.
- [30] D.-A. Clevert, T. Unterthiner, and S. Hochreiter, "Fast and accurate deep network learning by exponential linear units (ELUs)," in *International Conference on Learning Representations*, 2016.
- [31] A. L. Maas, A. Y. Hannun, and A. Y. Ng, "Rectifier nonlinearities improve neural network acoustic models," in *ICML Workshop on Deep Learning for Audio, Speech and Language Processing*, 2013.
- [32] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [33] N. Zeghidour and D. Grangier, "Wavesplit: End-to-end speech separation by speaker clustering," *CoRR*, vol. abs/2002.08933, 2020. [Online]. Available: <https://arxiv.org/abs/2002.08933>