Investigating Cross-Domain Losses for Speech Enhancement

Sherif Abdulatif, Karim Armanious, Jayasankar T. Sajeev, Karim Guirguis, Bin Yang University of Stuttgart, Institute of Signal Processing and System Theory, Stuttgart, Germany

Abstract—Recent years have seen a surge in the number of available frameworks for speech enhancement (SE) and recognition. Whether model-based or constructed via deep learning, these frameworks often rely in isolation on either time-domain signals or time-frequency (TF) representations of speech data. In this study, we investigate the advantages of each set of approaches by separately examining their impact on speech intelligibility and quality. Furthermore, we combine the fragmented benefits of time-domain and TF speech representations by introducing two new cross-domain SE frameworks. A quantitative comparative analysis against recent model-based and deep learning SE approaches is performed to illustrate the merit of the proposed frameworks.

Index Terms—Speech enhancement, generative adversarial networks, automatic speech recognition, deep learning.

I. INTRODUCTION

Speech enhancement (SE) can be defined as recovering the desired speech from various unwanted effects, such as background noise, interference and reverberation. Accordingly, SE is considered a fundamental building block in many commonplace tasks such as hearing aids, smartphones and speech recognition. However, extracting the desired speech can be challenging in realistic environments where the noise level and frequency components are similar to the desired speech.

Advances in deep neural networks (DNNs) have led to leaps in performance in diverse fields such as image processing, remote sensing and medical applications [1]–[9]. In an effort to overcome the intrinsic challenges in SE tasks, a large body of research similarly adopted DNNs. This has led to a surge in the number of available deep learning (DL) approaches for SE as indicated by recent studies [10], [11]. These approaches can be divided into two main categories. The first utilizes the original time-domain signals while the latter relies on timefrequency (TF) domain representations.

Chronologically, speech-enhancement generative adversarial networks (SEGAN) is among the first DL frameworks to adopt raw time-domain waveforms as input [12]-[14]. This framework utilizes two networks (generator and discriminator) trained adversarially with each other. According to the original comparative results presented in [12], this framework surpasses the performance of traditional modelbased SE approaches. Also, a non-causal adaptation of the autoregressive generative Wavenet was introduced in [15]. Rather than adversarial training, this architecture incorporates residual blocks together with a time-domain L1 loss. Not only did this framework improve upon the performance of the prior SEGAN, but also it is easily adaptable to variable-length input waveforms. On the whole, deeper architectures are necessary for the aforementioned approaches due to the utilized higher dimensional input space and the challenging patterns in timedomain signals.

Conversely, the second category of SE approaches convert the raw input speech data into their corresponding TF representations. This transformation serves to represent the speech data in a more perceptually oriented manner with pronounced visual features. Accordingly, the majority of model-based and DL approaches uphold this data-feeding strategy. Due to the sequential nature of speech data, usually recurrent neural networks (RNNs) are suitable for the extraction of temporal speech features [16]-[20]. Also, convolutional neural networks (CNNs) were studied for SE by binary masking of the corresponding TF maps [21]-[23]. Analogous to SEGAN, adversarial models were also utilized for TF-based SE. For instance, FSEGAN provided a two-dimensional adaptation of the SEGAN framework to suit TF inputs [24]. Building upon FSEGAN, additional adversarial approaches were introduced to provide further improvements in the resultant speech quality [25]–[28]. All of the aforementioned research necessitates the sole utilization of the magnitude component for SE while ignoring the phase. This will affect the quality of the reconstructed speech, i.e. how comfortable is the listening experience. However, this deterioration in the resultant speech quality does not adversely affect the speech intelligibility [29, p. 94], which represents the degree of clarity of the speech content. To this end, a relatively smaller number of studies attempted dual magnitude and phase correction [30]-[32]. Moreover, the current TF-domain approaches are confined to input tracks of fixed duration. This hinders the training and inference efficiency as longer tracks must be divided into short segments which correspond to higher computational budget.

A new paradigm was recently proposed for SE, which explores the utilization of both time-domain and TF input components. In [33], authors introduced two parallel architectures operating on the time and TF representations with a joint auxiliary decision at the end. Both architectures employ a time-domain loss operating on the final combined output. Few studies attempted penalizing a time-domain framework with a TF-domain loss [13], [34], [35]. Moreover, some recent frameworks utilized a binary masking operation to enhance input TF representations which are then transformed back to the time-domain [36], [37]. Similarly, a single loss function is used to penalize the whole framework in the time-domain. Thus, despite utilizing both domains as input components, the above frameworks are trained solely using a single-domain loss function. To the best of our knowledge, no explicit investigation has been conducted regarding the impact of simultaneous utilization of cross-domain loss functions for SE.

In this work, we attempt to combine the fragmented benefits of time and TF-domain SE frameworks. To this end, we empirically investigate the impact of single-domain loss functions



Fig. 1. Architecture details of all studied frameworks (SEGAN, Wavenet and AeGAN). The FSEGAN architecture is similar to the AeGAN with only one generator and no feature-based loss (\mathcal{L}_{FB}).

and their cross-domain counterparts. Namely, we incorporate the time-domain loss component into SE frameworks operating on TF representations and vice versa. Further, we separately examine the influence of each loss component on the resultant speech quality and speech intelligibility. To validate the advantages of the proposed cross-domain models, a quantitative comparative analysis between recent SE frameworks and the proposed cross-domain approaches is presented over different signal-to-noise-ratios (SNRs).

II. METHOD

In this section, the SE approaches investigated in the conducted comparative analysis will be briefly introduced. An overview of these frameworks is represented in Fig. 1. Moreover, the proposed cross-domain loss functions will be subsequently presented in details.

A. Time-domain Approaches

SEGAN [12] is a conditional adversarial framework composed of two networks, the generator G and the discriminator D, as depicted in Fig. 1a. The generator is a one-dimensional 22-layer U-net architecture which receives as input a fixedlength corrupted time-domain speech signal y(t). It is tasked with transforming it into a corresponding enhanced output $\hat{x}(t) = G(y(t))$. The discriminator network, with identical architecture as G's encoder, receives either the ground-truth clean speech signal x(t) or the generated speech output $\hat{x}(t)$ together with the corrupted input y(t). It acts as a binary classifier as it attempts to distinguish which input-pair is real and which is fake. This motivates G to improve its own performance by producing more realistic output which fools D, while conversely D attempts to improve its classification performance. This adversarial training component is represented by the following loss function:

$$\min_{G} \max_{D} \mathcal{L}_{adv} = \min_{G} \max_{D} \mathbb{E}_{x(t), y(t)} \left[\log D(x(t), y(t)) \right] + \mathbb{E}_{\hat{x}(t), y(t)} \left[\log \left(1 - D\left(\hat{x}(t), y(t) \right) \right) \right]$$
(1)

while an additional time-domain loss component is utilized to further improve the generator training:

$$\mathcal{L}_{\text{L1,time}} = \mathbb{E}_{x(t),\hat{x}(t)} \left[\| x(t) - \hat{x}(t) \|_1 \right]$$
(2)

Wavenet [15] is an audio domain adaptation of the autoregressive "PixelCNN" generative model for natural images [38]. It is composed of consecutive residual blocks preceded and followed by a series of one-dimensional convolutional layers. The residual blocks utilize dilated convolutional layers with exponentially increasing dilation factors resulting in a growing receptive field. Due to the sequential nature of time-series speech data, it is beneficial to capture long-term dependencies in order to improve the resultant speech quality. For this purpose, Wavenet employs gated activation units within the residual blocks, as illustrated in Fig. 1b. Finally, the resultant enhanced speech signal is penalized by an L1 loss, identical to that in Eq. 2. A particular advantage of this approach is the capability of accommodating variable-length inputs, whereas SEGANs are restricted to fixed-length data as they employ an encoder-decoder generator architecture.

B. Time-Frequency Approaches

FSEGAN [24] is a two-dimensional adaptation of the SEGAN framework. It follows the same design principles of conditional adversarial networks with a generator G (16layered U-net architecture) and a discriminator D trained simultaneously in competition with each other. In this approach, the inputs to the generator are TF-magnitude representations calculated by applying a short-time Fourier transform (STFT) on the raw speech data. Compared to prior time-domain approaches, the advantages of this framework are twofold. First, the TF-magnitude component is a visual interpretation of human speech phonetics. This assists in enhancing the performance of DL-related SE approaches, as reported in [24]. Further, this TF embedding reduces the required computational resources due to the efficient two-dimensional convolutional layers. For training this network, a similar adversarial loss function to that defined in Eq. 1 is utilized albeit with the TF-magnitude representations y_m , x_m and \hat{x}_m representing the corrupted inputs, clean targets and predicted outputs, respectively. To further enhance the quality of the resultant TF representations, an additional L1 loss is used to penalize the pixel-wise discrepancies between the outputs \hat{x}_m and targets x_m as expressed by:

$$\mathcal{L}_{L1,TF} = \mathbb{E}_{x_m, \hat{x}_m} \left[\| x_m - \hat{x}_m \|_1 \right]$$
(3)

AeGAN [28] is an improved architecture which extends FSEGAN by introducing an additional non-adversarial featurebased loss function together with an enhanced generator



Fig. 2. SE cross-domain architectures with their respective losses. The weights given to cross-domain losses are chosen to reflect equal importance and thus penalization in both domains. All STFT blocks utilize dynamic time resolution technique [28], to embed all training tracks into 256×256 TF representations.

consisting of an end-to-end concatenation of three U-nets (CasNet), as shown in Fig. 1c. Moreover, the feature-based loss function (\mathcal{L}_{FB}) exploits the visual patterns in the input TF representations to produce globally consistent results. This is achieved by utilizing the discriminator as a trainable feature-extractor network and penalizing the discrepancies between the extracted feature-maps of the outputs \hat{x}_m and their ground-truth counterparts x_m . This loss is defined as:

$$\mathcal{L}_{\text{FB}} = \sum_{k=1}^{K} \lambda_k \| D_k (x_m) - D_k (\hat{x}_m) \|_1$$
(4)

with D_k representing the extracted feature-map from the k^{th} layer of the discriminator. K and λ_k are the number of layers and the individual weights given to each layer, respectively.

C. Cross-domain Loss Functions

In this study, we attempt to combine the benefits of both time-domain and TF-based SE approaches. To achieve this, we utilize from the prior two families of approaches the Wavenet and AeGAN as baselines and expand them with additional cross-domain loss functions, as illustrated in Fig. 2. Below are the two resultant frameworks:

• Time \rightarrow TF loss: The first framework, referred to as cross-domain Wavenet (CD-Wavenet), is an extension of the original Wavenet architecture. It implements an STFT on both the enhanced and target speech signals. The resultant TF-magnitude components are penalized using the cross-domain loss function $\mathcal{L}_{L1,TF}$, presented previously in Eq. 3. Since the TF-magnitude is a direct representation of speech phonetics, we hypothesize that additionally penalizing signals in the TF-domain will help to preserve the perceptual speech features, thus, enhancing the speech intelligibility with no adverse effect on the resultant quality.

• **TF** \rightarrow **Time loss**: Cross-domain AeGAN (CD-AeGAN) is the second investigated framework. It reverts back to the original time-domain waveforms via an inverse short-time Fourier transform (ISTFT) operation applied to the output and target TF-magnitude representations, together with the input noisy phase y_p . Discrepancies in the resultant time-series are then penalized using the aforementioned $\mathcal{L}_{L1,time}$ loss in Eq. 2. This guides the generator network to produce speech tracks that not only matches the perceptual features in the

TF-domain, but also the audible aspects of the original timedomain resulting in enhanced speech quality. The claim that time-domain loss inherits phase information, and thus directly impacts the speech quality, whereas speech intelligibility is rather affected by TF-magnitude is consistent with previous studies [39], [40].

III. DATASETS AND EXPERIMENTS

To investigate the impact of cross-domain loss functions in SE, we conduct a comparative study using the Voice Bank corpus dataset [41] as the ground-truth reference tracks. Additive noise from the DEMAND and QUT-TIMIT datasets [42], [43] were used to create the corrupted speech input tracks according to the following signal-to-noise ratios (SNRs): 0 and 5 dB. For training the different networks, data from 30 speakers (175 sentences) corrupted with 12 noise conditions were utilized (total of 49,510 tracks). As for the testing dataset, it consists of 20 new speakers reading 50 new sentences with 7 different noise conditions (1,000 tracks per SNR value). All input speech tracks are sampled at 16 kHz.

We quantitatively examine the performance of the abovedescribed DL approaches together with two model-based approaches: the Wiener filter [44] as a baseline and a recent Bayesian MMSE technique introduced in 2018 [45]. The SEGAN framework was trained on fixed one-second duration tracks, whereas the Wavenet utilized variable-length tracks. The TF-domain approaches were extended to incorporate the dynamic time resolution technique presented in [28]. This serves to embed tracks of variable duration into TF-magnitude embeddings of a fixed 256×256 dimensionality. To ensure a fair comparison, all trainable models were trained for 50 epochs using the architectures and hyper-parameter settings recommended in their respective original publications.

To investigate both the resultant speech quality and intelligibility, multiple metrics were used in the comparative analysis. For speech quality assessment, the segmental SNR (SSNR) [46], the perceptual evaluation of speech quality (PESQ) [47], the mean opinion score (MOS) predictions of the signal distortion (CSIG), background noise (CBAK) and the overall condition (COVL) [48] are utilized. With regards to human speech intelligibility, we evaluate the word error rate (WER) of a pre-trained Deep Speech model [49] as well as the shorttime objective intelligibility measure (STOI) [50].

TABLE I. Quantitative comparison of different approaches on the test set. Bold font indicates the best scores. The metric (1-WER) is utilized to represent improvements in the resultant intelligibility score (higher is better).

	(a) SNR 0 dB						
Model	PESQ	CSIG	CBAK	COVL	SSNR	STOI	1-WER
					[dB]	[%]	[%]
Wiener	1.83	1.35	2.04	1.49	-0.52	63.1	12.2
Bayesian	2.05	1.84	1.86	1.75	-0.37	64.2	20.4
SEGAN	2.16	2.47	2.18	2.20	0.24	67.8	28.5
Wavenet	2.34	3.33	2.59	2.78	2.44	69.2	51.2
CD-Wavenet	2.43	2.98	2.62	2.64	2.74	69.7	58.6
FSEGAN	2.33	3.26	2.50	2.75	0.59	68.2	46.9
AeGAN	2.48	3.57	2.68	2.99	1.59	69.9	54.2
CD-AeGAN	2.53	3.61	2.70	3.06	3.03	70.2	56.6
	(b) SNR 5 dB						
			(b)) SNR 5	dB		
Model	PESQ	CSIG	(b) CBAK) SNR 5 COVL	dB SSNR	STOI	1-WER
Model	PESQ	CSIG	(b) CBAK) SNR 5 COVL	dB SSNR [dB]	STOI [%]	1-WER [%]
Model Wiener	PESQ 2.20	CSIG 1.85	(b) CBAK 2.38) SNR 5 COVL 1.95	dB SSNR [dB] 1.23	STOI [%] 72.2	1-WER [%] 35.9
Model Wiener Bayesian	PESQ 2.20 2.38	CSIG 1.85 2.32	(b) CBAK 2.38 2.23) SNR 5 COVL 1.95 2.19	dB SSNR [dB] 1.23 1.82	STOI [%] 72.2 73.1	1-WER [%] 35.9 37.6
Model Wiener Bayesian SEGAN	PESQ 2.20 2.38 2.50	CSIG 1.85 2.32 2.83	(b) CBAK 2.38 2.23 2.52) SNR 5 COVL 1.95 2.19 2.56	dB SSNR [dB] 1.23 1.82 2.05	STOI [%] 72.2 73.1 74.5	1-WER [%] 35.9 37.6 47.5
Model Wiener Bayesian SEGAN Wavenet	PESQ 2.20 2.38 2.50 2.73	CSIG 1.85 2.32 2.83 3.74	(b) CBAK 2.38 2.23 2.52 2.91) SNR 5 COVL 1.95 2.19 2.56 3.19	dB SSNR [dB] 1.23 1.82 2.05 4.03	STOI [%] 72.2 73.1 74.5 75.9	1-WER [%] 35.9 37.6 47.5 62.8
Model Wiener Bayesian SEGAN Wavenet CD-Wavenet	PESQ 2.20 2.38 2.50 2.73 2.77	CSIG 1.85 2.32 2.83 3.74 3.49	(b) CBAK 2.38 2.23 2.52 2.91 2.96) SNR 5 COVL 1.95 2.19 2.56 3.19 3.09	dB SSNR [dB] 1.23 1.82 2.05 4.03 4.73	STOI [%] 72.2 73.1 74.5 75.9 76.7	1-WER [%] 35.9 37.6 47.5 62.8 70.9
Model Wiener Bayesian SEGAN Wavenet CD-Wavenet FSEGAN	PESQ 2.20 2.38 2.50 2.73 2.77 2.68	CSIG 1.85 2.32 2.83 3.74 3.49 3.68	(b) CBAK 2.38 2.23 2.52 2.91 2.96 2.83	SNR 5 COVL 1.95 2.19 2.56 3.19 3.09 3.16	dB SSNR [dB] 1.23 1.82 2.05 4.03 4.73 2.32	STOI [%] 72.2 73.1 74.5 75.9 76.7 75.7	1-WER [%] 35.9 37.6 47.5 62.8 70.9 64.8
Model Wiener Bayesian SEGAN Wavenet CD-Wavenet FSEGAN AeGAN	PESQ 2.20 2.38 2.50 2.73 2.77 2.68 2.87	CSIG 1.85 2.32 2.83 3.74 3.49 3.68 3.91	(b) CBAK 2.38 2.23 2.52 2.91 2.96 2.83 3.03) SNR 5 COVL 1.95 2.19 2.56 3.19 3.09 3.16 3.37	dB SSNR [dB] 1.23 1.82 2.05 4.03 4.73 2.32 3.64	STOI [%] 72.2 73.1 74.5 75.9 76.7 75.7 79.2	1-WER [%] 35.9 37.6 47.5 62.8 70.9 64.8 68.9

IV. RESULTS AND DISCUSSION

Table I and Fig. 3 present the quantitative evaluations of the proposed cross-domain SE frameworks in comparison to approaches utilizing single-domain losses. Additionally, comparisons were carried out against other model-based approaches and DL-based time-domain and TF-domain approaches. The recent Bayesian MMSE model-based approach achieves comparable performance to the SEGAN framework with respect to both the speech quality and intelligibility metrics. The AeGAN framework achieves a significant enhancement of the resultant speech tracks compared to the FSEGAN framework as reflected across all metric scores. This is attributed to the incorporation of the cascaded generator network architecture as well as the feature-based loss function. For instance, at SNR 0 dB improvements from 46.9% to 54.2% and from 2.33 to 2.48 were observed for the WER and PESQ, respectively. With regards to time-domain SE approaches, the Wavenet framework introduces the capability of enhancing speech tracks of variable-lengths. Despite of this, it outperforms the prior SEGAN framework which deals with fixed-length input tracks.

To investigate the impact of cross-domain loss functions in SE, the highest performing networks from both the time and TF-based family of DL techniques were singled out. Namely, the AeGAN and Wavenet frameworks. These approaches were then expanded upon by cross-domain loss functions and compared to their unaltered counterparts. In CD-AeGAN, the output TF-magnitude representations from AeGAN are transformed back to the time-domain before being



Fig. 3. Boxplots of cross-domain models in comparison to their original counterparts over both SNRs (0, 5 dB). The SSNR (left) represents the speech quality. The WER (right) represents the speech intelligibility. The quantity (1-WER) can be negative as in these cases, the accumulated errors exceed the number of words in the reference sentence.

additionally penalized by an L1 loss function. This translates into an improvement of the SSNR score compared to AeGAN by approximately 1.44 and 1.23 dB for SNR 0 and 5 dB, respectively. Also, an enhancement of the PESQ score is observed with a slight improvement in the CSIG, CBAK and COVL. Furthermore, a marginal or no improvement in the speech intelligibility metrics, represented by WER and STOI, is perceived. The second investigated cross-domain approach expands upon the conventional Wavenet framework with an additional pixel-wise loss in the TF-magnitude. In contrast to CD-AeGAN, penalizing the TF-magnitude representations noticeably enhances the speech intelligibility of the resultant tracks. Specifically, the WER was improved by approximately 8% for both SNR values. Conversely, marginal improvement can be observed to the speech quality metrics. On the whole, the CD-AeGAN outperforms the CD-Wavenet in most of the metrics except for the WER where the CD-Wavenet achieves the best score. This agrees with recent literature hypothesizing that penalizing the time domain after the enhancement of TF representations is beneficial to the overall quality metrics [37].

The above results indicate that time-domain and TF-based loss components have distinct strengths. Specifically, penalizing the speech data in the TF-domain positively influences the resultant speech intelligibility, whereas time-domain loss functions directly impact the speech quality. From another perspective, time-domain approaches benefit from implicitly enhancing the speech phase component which positively impacts the quality metrics, whereas TF-based approaches exploit the perceptual features in the magnitude representations, thus, enhancing the ineligibility of the speech tracks. Consequently, cross-domain SE approaches would assist in bridging the gap between the distinct advantages of single-domain frameworks.

V. CONCLUSION

In this work, we present an investigative study regarding the impact of simultaneously enhancing speech signals in both time and TF domains in comparison to single-domain approaches. To achieve this, we expand the conventional Wavenet and AeGAN frameworks with corresponding cross-domain loss functions. Quantitative evaluations have illustrated that penalizing time-domain losses directly influence the resultant speech quality. In contrast, TF-based loss functions impact the intelligibility of the output speech. In the future, this may lead to the development of tailored architectures to resolve the drawbacks of single-domain frameworks. For instance, this could combine the benefit of time-domain approaches with respect to implicitly enhancing the phase component while concurrently exploiting the perceptual information in the TFmagnitude representations.

REFERENCES

- L. A. Gatys, A. S. Ecker, and M. Bethge, "Image style transfer using convolutional neural networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 2414–2423.
- [2] C. Godard et al., "Digging into self-supervised monocular depth estimation," in *IEEE/CVF International Conference on Computer Vision* (*ICCV*), 2019, pp. 3827–3837.
- C. P. Schwegmann et al., "Very deep learning for ship discrimination in synthetic aperture radar imagery," in *IEEE International Geoscience* and Remote Sensing Symposium, 2016, pp. 104–107.
- [4] K. Armanious et al., "An adversarial super-resolution remedy for radar design trade-offs, in 27th European Signal Processing Conference (EUŠIPCO), 2019.
- S. Abdulatif et al., "Towards adversarial denoising of radar micro-Doppler signatures," in *International Radar Conference*, 2019. [5]
- [6] M. S. Seyfioğlu, A. M. Özbayoğlu, and S. Z. Gürbüz, "Deep convolutional autoencoder for radar-based classification of similar aided and unaided human activities," IEEE Transactions on Aerospace and Electronic Systems, vol. 54, no. 4, pp. 1709-1723, 2018.
- [7] M. Ulrich et al., "Person recognition based on micro-Doppler and thermal infrared camera fusion for firefighting," in 21st Internat Conference on Information Fusion (FUSION), 2018, pp. 919–926. in 21st International
- D. Ravì et al., "Deep learning for health informatics," IEEE Journal of Biomedical and Health Informatics, vol. 21, no. 1, pp. 4-21, 2017.
- [9] K. Armanious et al., "ipA-MedGAN: Inpainting of arbitrary regions in medical imaging," in *IEEE International Conference on Image Processing (ICIP)*, 2020, pp. 3005–3009.
- [10] D. Wang and J. Chen, "Supervised speech separation based on deep learning: An overview," *IEEE/ACM Transactions on Audio, Speech, and* 2010 June 201 Language Processing, vol. 26, no. 10, pp. 1702-1726, 2018.
- E. Vincent et al., Audio source separation and speech enhancement, John Wiley & Sons, 2018.
- [12] S. Pascual et al., "SEGAN: Speech enhancement generative adversarial network," in *Interspeech*, 2017, pp. 3642–3646.
- [13] S. Pascual et al., "Towards generalized speech enhancement with generative adversarial networks," in *Interspeech*, 2019, pp. 1791–1795.
- [14] H. P. Phan et al., "Improving GANs for speech enhancement," arXiv, vol. abs/2001.05532, 2020.
- [15] D. Rethage et al., "A wavenet for speech denoising," in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), April 2018, pp. 5069-5073.
- [16] F. Weninger et al., "Single-channel speech separation with memory-enhanced recurrent neural networks," in *IEEE International Conference* on Acoustics, Speech and Signal Processing (ICASSP), May 2014, pp. 2010.0011 3709-3713.
- [17] P. S. Huang et al., "Joint optimization of masks and deep recurrent neural networks for monaural source separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, pp. 2136–2147, 2015.
- F. Weninger et al., "Speech enhancement with LSTM recurrent neural networks and its application to noise-robust ASR," in *Latent Variable* [18] Analysis and Signal Separation, Aug. 2015, pp. 91-99
- [19] T. Gao et al., "Densely connected progressive learning for LSTM-based speech enhancement," in *IEEE International Conference on Acoustics*, *Speech and Signal Processing (ICASSP)*, April 2018, pp. 5054–5058.
- [20] A. Nicolson and K. K. Paliwal, "Deep Xi as a front-end for robust automatic speech recognition," *arXiv*, vol. abs/1906.07319, 2019.
 [21] S. Chakrabarty and E. A. P. Habets, "Time-frequency masking based
- online multi-channel speech enhancement with convolutional recurrent neural networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 787-799, 2019.
- [22] S. Chakrabarty et al., "Time-frequency masking based online speech enhancement with multi-channel data using convolutional neural net-works," in 16th International Workshop on Acoustic Signal Enhancement (IWAENC), 2018, pp. 476-480.
- [23] A. Pandey and D. Wang, "TCNN: Temporal convolutional neural network for real-time speech enhancement in the time domain," in *IEEE* International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2019, pp. 6875-6879.

- [24] C. Donahue et al., "Exploring speech enhancement with generative adversarial networks for robust speech recognition," in *IEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, April 2018, pp. 5024-5028.
- [25] D. Michelsanti and Z.-H. Tan, "Conditional generative adversarial net-works for speech enhancement and noise-robust speaker verification," in *Interspeech*, 2017, pp. 2008–2012.
- [26] S. Fu et al., "End-to-End waveform utterance enhancement for direct evaluation metrics optimization by fully convolutional neural networks, IEEE/ACM Transactions on Audio, Speech, and Language Processing, vol. 26, no. 9, pp. 1570–1584, April 2018.
- [27] M. H. Soni, N. Shah, and H. A. Patil, "Time-frequency masking-based speech enhancement using generative adversarial network," in IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2018, pp. 5039-5043.
- S. Abdulatif et al., "AeGAN: Time-frequency speech denoising via generative adversarial networks," in 28th European Signal Processing Conference (EUSIPCO), 2020, pp. 451–455.
- [29] P. C. Loizou, *Speech enhancement: Theory and practice*, CRC Press, Inc., USA, 2nd edition, Feb. 2013.
- 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. [30] D. S. Williamson and D. Wang,
- [31] D. S. Williamson et al., "Complex ratio masking for monaural speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language* Processing, vol. 24, no. 3, pp. 483-492, 2016.
- D. Yin et al., "PHASEN: A phase-and-harmonics-aware speech en-hancement network.," in AAAI, 2020, pp. 9458–9465. [32]
- J. Kim et al., "Multi-domain processing via hybrid denoising networks for speech enhancement," *arXiv*, vol. abs/1812.08914, 2018. [33]
- A. Defossez, G. Synnaeve, and Y. Adi, "Real time speech enhancement in the waveform domain," in *Interspeech*, 2020, pp. 3291–3295. [34]
- [35] U. Isik et al., "PoCoNet: Better speech enhancement with frequencypositional embeddings, semi-supervised conversational data, and biased loss," in *Interspeech*, 2020, pp. 2487–2491.
- Y. Wang and D. Wang, "A deep neural network for time-domain signal reconstruction," in *IEEE International Conference on Acoustics, Speech* [36] and Signal Processing (ICASSP), 2015, pp. 4390-4394.
- [37] C. Tang et al., "Joint time-frequency and time domain learning for speech enhancement," in *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, 2020, pp. 3816–3822.
- decoders," in *Advances in neural information processing systems*, 2016, pp. 4790–4798. A. Van den Oord et al., "Conditional image generation with PixelCNN [38]
- K. K. Paliwal and L. D. Alsteris, "On the usefulness of STFT phase spectrum in human listening tests," *Speech Communication*, vol. 45, no. 2, pp. 153 170" 2005 [39] , pp. 153 – 170", 2005.
- [40] M. Kolbæk et al., "On loss functions for supervised monaural timedomain speech enhancement," *IEEE/ACM Transactions on Audio,* Speech, and Language Processing, vol. 28, pp. 825–838, 2020.
- J. Yamagishi et al., "CSTR VCTK corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92)," 2019. [41]
- [42] J. Thiemann et al., "The diverse environments multi-channel acoustic noise database: A database of multichannel environmental noise recordings," The Journal of the Acoustical Society of America, vol. 133, no. 5, pp. 3591–3591, 2013.
- [43] D. B. Dean et al., "The QUT-NOISE-TIMIT corpus for the evaluation of voice activity detection algorithms," in *Interspeech*, 2010.
- [44] 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 Conference Proceedings*, May 1996, vol. 2, pp. 629-632
- G. Enzner and P. Thüne, "Bayesian MMSE filtering of noisy speech by SNR marginalization with global PSD priors," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 12, pp. 2289–2304, Dec. 2018. [45]
- S. R. Quackenbush et al., Objective measures of speech quality, Prentice [46] Hall. 1988.
- [47] A. W. Rix et al., "Perceptual evaluation of speech quality (PESQ)-a new method for speech quality assessment of telephone networks and codecs," in *IEEE International Conference on Acoustics, Speech and* Signal Processing (ICASSP), May 2001, vol. 2, pp. 749-752
- Y. Hu and P. C. Loizou, "Evaluation of objective quality measures for speech enhancement," *IEEE Transactions on Audio, Speech, and* [48] Language Processing, vol. 16, no. 1, pp. 229-238, Jan. 2008
- A. Hannun et al., "Deep Speech: Scaling up end-to-end speech recogni-tion," *arXiv*, vol. abs/1412.5567, 2014. [49]
- C. H. Taal et al., "A short-time objective intelligibility measure for [50] time-frequency weighted noisy speech," in 2010 IEEE International Conference on Acoustics, Speech, and Signal Processing, March 2010, pp. 4214-4217.