

RESIDUAL RECURRENT NEURAL NETWORK FOR SPEECH ENHANCEMENT

Jalal Abdulbaqi, Yue Gu, Shuhong Chen and Ivan Marsic

Rutgers, the State University of New Jersey, USA

{jalal.nazar, yue.guapp, shuhong.chen and marsic}@rutgers.edu

ABSTRACT

Most current speech enhancement models use spectrogram features that require an expensive transformation and result in phase information loss. Previous work has overcome these issues by using convolutional networks to learn the temporal correlations across high-resolution waveforms. These models, however, are limited by memory-intensive dilated convolution and aliasing artifacts from upsampling. We introduce an end-to-end fully recurrent neural network for single-channel speech enhancement. The network structured as an hourglass-shape that can efficiently capture long-range temporal dependencies by reducing the features resolution without information loss. Also, we use residual connections to prevent gradient decay over layers and improve the model generalization. Experimental results show that our model outperforms state-of-the-art approaches in six quantitative evaluation metrics.

Index Terms— Speech enhancement, speech denoising, recurrent neural network, waveform, residual connection

1. INTRODUCTION

Speech enhancement has important applications in voice communication, hearing aids, and automatic speech recognition. Speech enhancement removes background noise from noisy speech signals, increasing speech quality and intelligibility [1], [2]. Early research used non-trainable statistical approaches on spectrograms, such as spectral subtraction [3], Wiener filter [4], statistical model-based methods [5], the subspace method [6], minimum mean-square error estimator, and optimally-modified log-spectral amplitude [7], [8]. These methods showed limited performance on speech with non-stationary noise, which is common in real-life environments. Non-negative matrix factorization has later been widely used for speech separation and enhancement [9], [10].

Recently, deep neural networks have been employed to overcome the non-stationary condition and have improved speech quality and intelligibility. Early models used mapping-based methods, where the enhanced signal is directly predicted from the noisy one. Several such deep learning models have been developed, including denoising autoencoders [11] (using fully-connected layers), recurrent neural networks (RNN) [12] and convolutional neural networks (CNN). Later, a masking-based method was introduced to enhance the signal by applying the noisy signal to the predicted mask [13]–[16].

Most of these methods use time-frequency (T-F) spectrogram features instead of time-domain waveform, since T-F has a reduced resolution. Spectrogram features, however, have certain limitations. First, the pre- and post-processing operations such as discrete

Fourier transform and its inverse are computationally expensive, and cause artifacts in the output signal [1], [2]. Second, these approaches usually only estimate the magnitude, and use the noisy phase to produce the enhanced speech. Research has shown that the phase can enhance the speech quality [17]. Recent research has considered predicting the phase and the magnitude at the cost of model complexity, such as adding special model for phase component [18].

Recently, several studies proposed overcoming previous limitations by working directly on the waveform. Fu et al. [19] compared fully-convolutional networks with fully-connected networks. Pascual et al. [20] implemented a generative adversarial network for speech enhancement (SEGAN), using strided convolutions, residual connections, and an encoder-decoder architecture. Later, a text-to-speech model called Wavenet [21] directly synthesized raw waveforms. Qian et al. [22] and Rethage et al. [23] presented a modified version of WaveNet for speech denoising. The former integrated a Bayesian framework WaveNet, while the latter used a non-causal dilated convolution with residual connections. Germain et al. [24] presented dilated convolutions combined with a feature loss network. Stroller et al. [25] adapted the U-Net [26] model for source separation using dilated convolutions and linear interpolation instead of transposed convolution for upsampling. All these methods used convolutional neural networks due to their ability to capture the samples' dependencies better than fully-connected networks. Because waveform is a sequential datatype, it requires a temporal context as well. Recurrent neural networks are known to capture the long-range temporal sequence information [27] and are used in many sequential applications such as speech recognition, neural machine translation, and spectrogram-based speech enhancement. To our knowledge, only [28] and [29] have applied RNN to process waveform signals. The first one used RNN to denoise a non-speech waveform, while the latter used RNN for speech bandwidth extension, but no one has used it for waveform-based speech enhancement. The reason is that the high resolution of waveforms requires more expensive, deeper, and wider networks. It is difficult to build a deep RNN because of saturating activation function, which causes gradient decay over layers. Also, we found empirically that RNNs sufficiently wide to process the high-resolution waveforms exceeded the available memory capacity. Therefore, we introduce our residual hourglass recurrent neural network for waveform-based single-channel speech enhancement. Our model overcomes the RNN limitations by introducing two techniques. First, the network architecture has an hourglass shape; the layers in the lower pyramid reduce the number of time-steps and increases the number of units (width), while the upper pyramid does the reverse. This architecture allows the RNN to handle high-resolution waveform features without memory overflow. Second, using residual connections between the same-

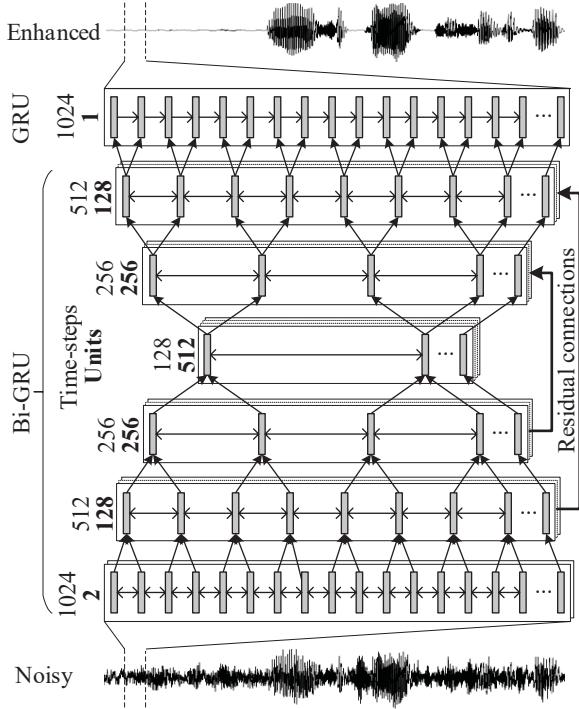


Fig. 1. Our proposed RNN architecture. Seven stacked RNN layers with the numbers on the left representing the number of time steps and the number of units in each layer. Wider layers have fewer units and vice versa. The two bold arrows on the right represent the residual connections.

shaped layers from the lower pyramid to the upper one prevents gradient decay over layers and improves the model generalization. Advantages of our model:

- Uses a raw waveform, without any transformation or handcrafted features.
- Does not need any linear interpolation method for upsampling, which can lose useful information.
- Is a simple end-to-end design that outperforms several more complex neural network approaches.
- We think that this deep RNN architecture can be applied for regression problems other than speech enhancement, which has long-term dependency and high-resolution data.

We evaluated our model using six objective metrics, demonstrating its ability to significantly enhance speech quality and intelligibility. The next section reviews the model architecture. Section 3 describes the dataset we used and the preprocessing operations. Section 4 presents the experimental setup and discusses the results. Section 5 concludes and suggests future work.

2. MODEL ARCHITECTURE

Our model includes seven GRU layers with two residual connections. The first six layers are bidirectional and the last one is a single GRU (Figure 1). The goal of our speech enhancement network is to learn non-linear relationships, so that noisy speech $x(t)$ can be translated into clean speech $y(t)$:

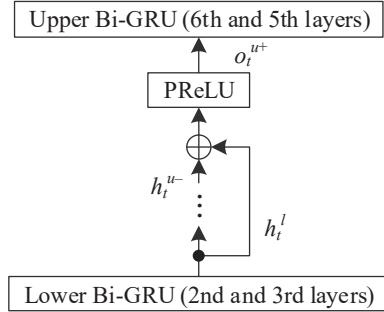


Fig. 2. A high-level view highlighting the residual connections in our proposed model from Figure 1.

$$y(t) = f(x(t)) \quad (1)$$

The input vector $X = (x_1, \dots, x_T)$ represents a T -seconds wide segment from a noisy audio waveform signal.

RNNs can efficiently realize temporal features in sequential data, so they have been used widely to process speech data either for speech recognition or enhancement. We chose gated recurrent units (GRU) instead of long short-term memory units (LSTM) or vanilla RNN. Both GRU and LSTM outperform vanilla RNN [28], but GRUs have a simpler structure and train faster than LSTMs. In addition, we chose bidirectional RNNs since in speech enhancement each predicted sample can depend on future as well as past noisy samples. The stacked GRU increases the capacity of the network by sharing the hidden states not only from the same layer but also from the lower layers as well. The stacked bidirectional RNNs share their hidden states, so that the hidden state (h_t^l) of a bi-GRU unit in layer l at time t is obtained by concatenating its forward (\vec{h}_t^l) and backward (\overleftarrow{h}_t^l) hidden states, which depend on the lower layer $l-1$ at time t and this layer at time $t-1$:

$$\vec{h}_t^l = \overrightarrow{GRU}(\vec{h}_{t-1}^{l-1}, \overleftarrow{h}_{t-1}^{l-1}) \quad (2)$$

$$\overleftarrow{h}_t^l = \overleftarrow{GRU}(\overleftarrow{h}_{t-1}^{l-1}, \vec{h}_{t-1}^{l-1}) \quad (3)$$

$$h_t^l = \text{conc}(\vec{h}_t^l, \overleftarrow{h}_t^l) \quad (4)$$

The two pyramids of our hourglass architecture keep the number of trainable parameters within the memory constraints. the bottom pyramid decreases the number time steps while increasing the number of GRU units per layer, and the top pyramid does the reverse. this approach allows for deeper networks. we did not use upsampling techniques, such as linear interpolation, because the information can be lost. instead, we reshape the RNN output to the desired fewer time steps. reshaping the layer output to decrease and increase the time steps prevents losing data, and allows the RNN to have a sufficient size of units. however, while stacking RNNs can increase the capacity of the network, deeper RNNs usually have gradient decay issues due to their saturating activation functions. to address this issue, we used residual connections between the lower and upper layers (figures 1 and 2). the residual connections facilitate training the deep RNN, and provide better generalization by combining the low-level features with the high-level ones in the upper layers. in figure 2, the hidden states of the lower layer (h_t^l) and those of the upper layer before the residual connection (h_t^{u-}) are combined to produce the residual output:

Table 1. Evaluation results of our proposed model compared with other state-of-the-art research work using six objective metrics on the same dataset [30]. Higher scores are better, and the highest scores are boldfaced.

Model	Features type	SSNR	PESQ	STOI	CSIG	CBAK	COVL
No Enhancement (Noisy)	-	1.68	1.97	0.820	3.35	2.44	2.63
SEGAN, 2017 [20]	waveform	7.73	2.16	0.93	3.48	2.94	2.80
CNN-GAN, 2018 [14]	spectrogram	-	2.34	0.93	3.55	2.95	2.92
Wavenet, 2018 [23]	waveform	-	-	-	3.62	3.23	2.98
MMSE-GAN, 2018 [16]	spectrogram	-	2.53	-	3.80	3.12	3.14
DFL, 2019 [24]	waveform	-	-	-	3.86	3.33	3.22
(Our model)	waveform	14.71	3.20	0.98	4.37	4.02	3.82

$$o_t^{u+} = PReLU(h_t^l + h_t^{u-}) \quad (5)$$

where $PReLU$ is the parametric rectified linear unit activation function. Finally, we use a single forward GRU to output the enhanced speech with the same size of the input vector:

$$\vec{h}_t^l = \overrightarrow{GRU} \left(\vec{h}_t^{l-1}, \vec{h}_{t-1}^l \right) \quad (6)$$

Therefore, the output will be created by combining the hidden states for each input segment:

$$Y = (\vec{h}_1^7, \dots, \vec{h}_T^7) \quad (7)$$

where Y denotes the enhanced signal output and \vec{h}_1^7 denotes the hidden state of the last (seventh) layer.

3. DATASET AND PREPROCESSING

The dataset used for training and evaluating our model has been set up in [30]. We chose this dataset because it is large, has different types of non-stationary noise, and is public so that we can compare our results with other published work. This dataset is an excerpt of the Voice Bank corpus [31] with 28 speakers (14 male and 14 female) of the same accent region (England) and another 56 speakers (28 male and 28 female) of other accent regions (Scotland and United States).

The noisy data used for training are two artificially generated (speech shaped noise and babble) and eight real noise recordings from the Demand database [32]. The noises are from different environments such as kitchens, offices, public spaces, transportation stations, and streets. The training set includes 11,572 utterances with four signal-to-noise (SNR) values: 15 dB, 10 dB, 5 dB, and 0 dB. The noisy data used for testing include two other speakers of the same corpus from England (a male and a female), and five other noises from the Demand database. The chosen noises include a living room, an office, a bus, and street noise. The testing set includes 824 utterances with four SNR: 17.5 dB, 12.5 dB, 7.5 dB and 2.5 dB. We downsampled the audio signals to 16kHz, getting a reasonable dataset size for recognizing speech. Our preprocessing included slicing both noisy and clean speech signals into 1024 samples (~64 ms) with 25% overlap during training and without

overlap during the evaluation. We did not use any other preprocessing, such as pre-emphasis.

4. EXPERIMENT SETUP AND RESULTS

Our architecture uses seven GRU layers. The first six are bi-directional, while the last one is single-directional to produce the enhanced signal (Figure 1). The number of units per layer are: 2, 128, 256, 512, 256, 128, and 1; the size of the time steps per layer are: 1024, 512, 256, 128, 256, 512, and 1024. Two residual connections link the second and third layers with the sixth and fifth, respectively. The PReLU activation function is used with residual connections, as it does not saturate the negative values compared to Leaky-ReLU and has been shown to improve model fitting [33]. The model has 2 million trainable parameters, which is small with respect to Wavenet which has 6.3 million. We use the Xavier normal initializer [34] for the kernel weights, with zero-initialized biases. Xavier initialization keeps the values of the weights in a reasonable range, preventing the inputs from shrinking or growing more than needed through the layers. It determines the initialization values with respect to the number of input and output neurons. The initializer for the recurrent states is a random orthogonal matrix [35], which helps the RNN stabilize by avoiding vanishing or exploding gradients. The stability occurs because the orthogonal matrix has an absolute eigenvalue of one, which avoids the gradients from exploding or vanishing due to repeated matrix multiplication.

We use the log-cosh loss function, a regression loss function that takes on the behavior of squared-loss when the loss is small, and absolute loss when the loss is large; this reduces the influence of wrong predictions. The optimization algorithm used is RMSprop [36], which helps the training of large neural networks on large redundant datasets. In addition, Keras [37] documentation recommends using this algorithm with RNN. We trained the model until the validation loss converged with a batch size of 512, using two NVIDIA GTX-1080 GPUs. We used different learning rates during training, starting at 10^{-4} and gradually decreasing to 10^{-8} . The library used to implement this work was Keras with TensorFlow [38] as a backend. The training process took about 20 hours for 50 epochs. To evaluate our model, we computed six objective metrics using an open-source implementation^{1,2}:

- Segmental signal-to-noise ratio (SSNR) [1]: computed by dividing the clean and enhanced signals into segments and

¹https://www.crcpress.com/downloads/K14513/K14513_CD_Files.zip

²<http://ceestaal.nl/stoi.zip>

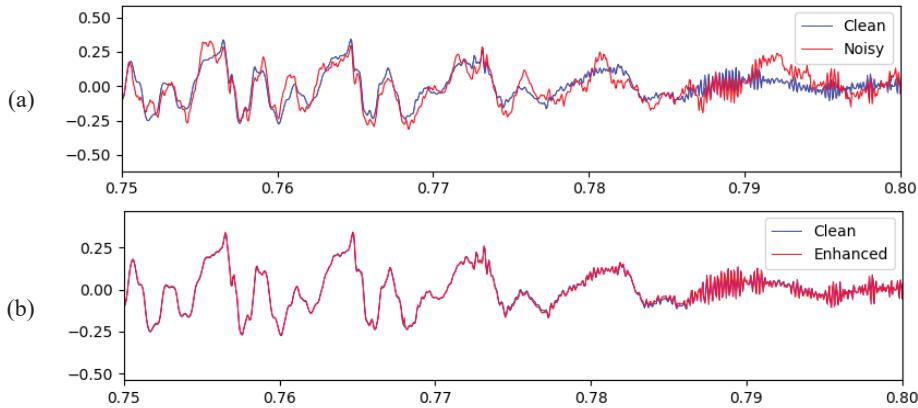


Fig. 3. An illustration of speech enhancement using our model using speech samples with SNR = 2.5 dB and duration of 50 msec from the test dataset. (a) The sample number 232_052. The blue lines represent the clean speech and the red lines represent the noisy speech. (b) The corresponding enhanced speech (red line) compared with the clean input speech (blue line).

computing the segment energies and SNRs, and then returning the mean segmental SNR (dB). The values range from -10 to 35.

- Perceptual evaluation of speech quality (PESQ) [1]: a more complex metric to capture a wider range of distortions. PESQ is the most common metric to evaluate the speech quality, calculated by comparing the enhanced and clean speech. The values range -0.5 to 4.5.
- Short-time objective intelligibility (STOI) [39]: reflects the improvement in speech intelligibility with a score range from 0 to 1.
- Three objective version of mean opinion scores (MOSSs): CSIG for signal distortion evaluation, CBAK for noise distortion evaluation, and COVL for overall quality evaluation. We used their mathematical representations, and their scores range from 1 to 5 [1].

For all these metrics, higher values mean better performance. Two speech test samples (small segment 50 ms) are illustrated in Figure 3. Both samples include non-stationary noise with people talking in background (“cocktail party”) and music playing. For each segment, the foreground speaker talks (high frequency) in the first half, while foreground speaker stops talking (low frequency) in the second half. The enhanced speech signal tracks the clean in both cases, which shows the model ability to capture the clean speech in all speaker events.

Table 1 shows the metrics scores for our model with respect to the other architectures. The comparison results with several current architectures such as SEGAN [20], Mask-based GAN model (CNN-GAN) [14], wavenet model for denoising [23], another masking-based GAN model [16] and finally the speech denoising with deep feature losses (DFL) [25]. All these architectures use the same dataset and the metrics that we used to train and evaluate our model. Therefore, their results are taken directly from their above work. In this results, our model decrease the speech degradation (CSIG) by 13.2% and decrease the background noise intrusiveness (CBAK) by 20.7%, and increase the overall signal quality (COVL) by 18.6% with respect to the best previous architecture DFL [24]. Also, the speech quality is increased by 26.5% with respect to the masking-based GAN model [16].

5. CONCLUSION

We introduced a novel end-to-end fully-recurrent neural network for single-channel speech enhancement. Our recurrent layers are designed in an hourglass shape to reduce the speech signal dimension and assist recognition of the long-term dependencies. The results show that our simple and efficient model outperforms most of the current approaches with more complex architectures. We will evaluate this model with other datasets and apply to other sequential applications.

6. ACKNOWLEDGEMENTS

This research has been supported by the National Library of Medicine of the National Institutes of Health under grant number 2R01LM011834-05 and by the National Science Foundation under grant number IIS-1763509.

7. REFERENCES

- [1] P. C. Loizou, *Speech Enhancement: Theory and Practice*, 2nd ed. Boca Raton, FL, USA: CRC Press, Inc., 2013.
- [2] J. Benesty, S. Makino, and J. Chen, “Speech enhancement,” Springer, 2005, p. 406.
- [3] M. Berouti, R. Schwartz, and J. Makhoul, “Enhancement of speech corrupted by acoustic noise,” in *ICASSP '79. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1979, vol. 4, pp. 208–211, doi: 10.1109/ICASSP.1979.1170788.
- [4] J. S. Lim and A. V. Oppenheim, “All-Pole Modeling of Degraded Speech,” *IEEE Trans. Acoust.*, vol. 26, no. 3, pp. 197–210, Jun. 1978, doi: 10.1109/GES00795.1.
- [5] Y. Ephraim, “Statistical-model-based speech enhancement systems,” *Proc. IEEE*, vol. 80, no. 10, pp. 1526–1555, 1992, doi: 10.1109/5.168664.
- [6] M. Dendrinos, S. Bakamidis, and G. Carayannis, “Speech enhancement from noise: A regenerative approach,” *Speech Commun.*, vol. 10, no. 1, pp. 45–57, Feb. 1991, doi: 10.1016/0167-6393(91)90027-Q.
- [7] Y. Ephraim and D. Malah, “Speech Enhancement Using a Minimum Mean-Square Error Short-Time Spectral Amplitude

Estimator,” *IEEE Trans. Acoust.*, vol. 32, no. 6, pp. 1109–1121, Dec. 1984, doi: 10.1109/TASSP.1984.1164453.

[8] I. Cohen and B. Berdugo, “Speech enhancement for non-stationary noise environments,” *Signal Processing*, vol. 81, no. 11, pp. 2403–2418, Nov. 2001, doi: 10.1016/S0165-1684(01)00128-1.

[9] K. W. Wilson, B. Raj, P. Smaragdis, and A. Divakaran, “Speech denoising using nonnegative matrix factorization with priors,” in *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2008, pp. 4029–4032, doi: 10.1109/ICASSP.2008.4518538.

[10] J. L. Durrieu, A. Ozerov, C. Févotte, G. Richard, and B. David, “Main instrument separation from stereophonic audio signals using a source/filter model,” in *European Signal Processing Conference*, 2009, pp. 15–19, doi: 10.1099/13500872-142-12-3337.

[11] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, “A regression approach to speech enhancement based on deep neural networks,” *IEEE/ACM Trans. Audio, Speech Lang. Process.*, vol. 23, no. 1, pp. 7–19, 2015.

[12] Z. Zhang, F. Ringeval, J. Han, J. Deng, E. Marchi, and B. Schuller, “Facing Realism in Spontaneous Emotion Recognition from Speech: Feature Enhancement by Autoencoder with LSTM Neural Networks,” *Interspeech 2016*, pp. 3593–3597, 2016.

[13] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, “Deep Recurrent Networks for Separation and Recognition of Single-Channel Speech in Nonstationary Background Audio,” in *New Era for Robust Speech Recognition*, Cham: Springer International Publishing, 2017, pp. 165–186.

[14] N. Shah, H. A. Patil, and M. H. Soni, “Time-Frequency Mask-based Speech Enhancement using Convolutional Generative Adversarial Network,” in *Proceedings, APSIPA Annual Summit and Conference*, 2018, vol. 2018, pp. 12–15.

[15] D. S. Williamson, Y. Wang, and D. Wang, “Complex ratio masking for joint enhancement of magnitude and phase,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 5220–5224, doi: 10.1109/ICASSP.2016.7472673.

[16] M. H. Soni, N. Shah, and H. A. Patil, “Time-Frequency Masking-Based Speech Enhancement Using Generative Adversarial Network,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5039–5043.

[17] K. Paliwal, K. Wójcicki, and B. Shannon, “The importance of phase in speech enhancement,” *Speech Commun.*, vol. 53, no. 4, pp. 465–494, Apr. 2011, doi: 10.1016/J.SPECOM.2010.12.003.

[18] N. Takahashi, P. Agrawal, N. Goswami, and Y. Mitsufuji, “PhaseNet: Discretized Phase Modeling with Deep Neural Networks for Audio Source Separation,” in *Proc. Interspeech 2018*, 2018, pp. 2713–2717.

[19] S.-W. Fu, Y. Tsao, X. Lu, and H. Kawai, “Raw waveform-based speech enhancement by fully convolutional networks,” in *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, 2017, pp. 006–012, doi: 10.1109/APSIPA.2017.8281993.

[20] S. Pascual, A. Bonafonte, and J. Serra, “SEGAN: Speech enhancement generative adversarial network,” in *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2017, vol. 2017-Augus, pp. 3642–3646, doi: 10.21437/Interspeech.2017-1428.

[21] A. van den Oord *et al.*, “WaveNet: A Generative Model for Raw Audio,” in *9th ISCA Speech Synthesis Workshop*, 2016, p. 125.

[22] K. Qian, Y. Zhang, S. Chang, X. Yang, D. Flornicu, and M. Hasegawa-Johnson, “Speech Enhancement Using Bayesian Wavenet,” in *Proc. Interspeech 2017*, 2017, pp. 2013–2017, doi: 10.21437/Interspeech.2017-1672.

[23] D. Rethage, J. Pons, and X. Serra, “A Wavenet for Speech Denoising,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5069–5073, doi: 10.1109/ICASSP.2018.8462417.

[24] F. G. Germain, Q. Chen, and V. Koltun, “Speech Denoising with Deep Feature Losses,” in *Proc. Interspeech 2019*, 2019, pp. 2723–2727, doi: 10.21437/Interspeech.2019-1924.

[25] D. Stoller, S. Ewert, and S. Dixon, “Wave-U-Net: A Multi-Scale Neural Network for End-to-End Audio Source Separation,” *Int. Symp. Music Inf. Retr.*, pp. 334–340, 2018.

[26] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” *Med. image Comput. Comput. Assist. Interv.*, pp. 234–241, 2015.

[27] A. Graves, “Generating Sequences With Recurrent Neural Networks,” Aug. 2013.

[28] H. Shen, D. George, E. A. Huerta, and Z. Zhao, “Denoising Gravitational Waves with Enhanced Deep Recurrent Denoising Auto-encoders,” in *ICASSP, IEEE International Conference on Acoustics, Speech and Signal Processing - Proceedings*, 2019, vol. 2019-May, pp. 3237–3241, doi: 10.1109/ICASSP.2019.8683061.

[29] Z.-H. Ling, Y. Ai, Y. Gu, and L.-R. Dai, “Waveform Modeling and Generation Using Hierarchical Recurrent Neural Networks for Speech Bandwidth Extension,” *IEEE Trans. Audio, Speech, Lang. Processing*, vol. 26, pp. 883–894, 2018.

[30] C. Valentini-botinhao, X. Wang, S. Takaki, and J. Yamagishi, “Investigating RNN-based speech enhancement methods for noise-robust Text-to-Speech,” *9th ISCA Speech Synth. Work.*, pp. 159–165, 2016, doi: 10.21437/SSW.2016-24.

[31] C. Veaux, J. Yamagishi, and S. King, “The voice bank corpus: Design, collection and data analysis of a large regional accent speech database,” in *2013 International Conference Oriental COCOSDA Held Jointly with 2013 Conference on Asian Spoken Language Research and Evaluation, O-COCOSDA/CASLRE 2013*, 2013, pp. 1–4, doi: 10.1109/ICSDA.2013.6709856.

[32] J. Thiemann, N. Ito, and E. Vincent, “The diverse environments multi-channel acoustic noise database: A database of multichannel environmental noise recordings,” *J. Acoust. Soc. Am.*, vol. 133, no. 5, pp. 3591–3591, May 2013, doi: 10.1121/1.4806631.

[33] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: Surpassing human-level performance on imagenet classification,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1026–1034.

[34] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” *Proc. Int. Conf. Artif. Intell. Stat. (AISTATS'10)*. Soc. Artif. Intell. Stat., 2010.

[35] A. M. Saxe, J. L. McClelland, and S. Ganguli, “Exact solutions to the nonlinear dynamics of learning in deep linear neural networks,” *International Conference on Learning Representations*, 2014.

[36] T. Tieleman and G. Hinton, “Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude,” *COURSERA Neural networks Mach. Learn.*, vol. 4, no. 2, pp. 26–31, 2012.

[37] F. Chollet and others, “Keras.” 2015.

[38] M. Abadi *et al.*, “Tensorflow: a system for large-scale machine learning,” in *OSDI*, 2016, vol. 16, pp. 265–283.

[39] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, “An Algorithm for Intelligibility Prediction of Time-Frequency Weighted Noisy Speech,” *IEEE Trans. Audio, Speech, Lang. Processing*, vol. 19, no. 7, pp. 2125–2136, 2011, doi: 10.1109/TASL.2011.2114881.