

Speech Denoising with Deep Feature Losses

François G. Germain, Qifeng Chen, and Vladlen Koltun

Abstract—We present an end-to-end deep learning approach to denoising speech signals by processing the raw waveform directly. Given input audio containing speech corrupted by an additive background signal, the system aims to produce a processed signal that contains only the speech content. Recent approaches have shown promising results using various deep network architectures. In this paper, we propose to train a fully-convolutional context aggregation network using a deep feature loss. That loss is based on comparing the internal feature activations in a different network, trained for acoustic environment detection and domestic audio tagging. Our approach outperforms the state-of-the-art in objective speech quality metrics and in large-scale perceptual experiments with human listeners. It also outperforms an identical network trained using traditional regression losses. The advantage of the new approach is particularly pronounced for the hardest data with the most intrusive background noise, for which denoising is most needed and most challenging.

Index Terms—Speech denoising, speech enhancement, deep learning, context aggregation network, deep feature loss

I. INTRODUCTION

SPEECH denoising (or enhancement) refers to the removal of background content from speech signals [1]. Due to the ubiquity of this audio degradation, denoising has a key role in improving human-to-human (e.g., hearing aids) and human-to-machine (e.g., automatic speech recognition) communications. A particularly challenging but common form of the problem is the under-determined case of single-channel speech denoising, due to the complexity of speech processes and the unknown nature of the non-speech material. The complexity is further compounded by the nature of the data, since audio material contains a high density of data samples (e.g., 16,000 samples per second). Challenges also arise in mediated human-to-human communication, as perception mechanisms can make small errors still noticeable by the average user [2].

In this work, we present an end-to-end deep learning approach to speech denoising. Our approach trains a fully-convolutional denoising network using a deep feature loss. To compute the loss between two waveforms, we apply a pretrained audio classification network to each waveform and compare the internal activation patterns induced in the network by the two signals. This compares a multitude of features at different scales in the two waveforms. We perform extensive experiments that compare the presented approach to recent state-of-the-art end-to-end deep learning techniques for denoising. Our approach outperforms them in both objective speech quality metrics and large-scale perceptual experiments with human listeners, which indicate that our approach is more effective than the baselines. The advantages of the presented

approach are particularly pronounced for the hardest, noisiest inputs, for which denoising is most challenging.

A. Related Work

Before the popularization of deep networks, denoising systems relied on spectrogram-domain statistical signal processing methods [1], followed more recently by spectrogram factorization-based methods [3]. Current denoising pipelines instead rely on deep networks for state-of-the-art performance. However, most pipelines still operate in the spectrogram domain [4]–[11]. As such, signal artifacts then arise due to time aliasing when using the inverse short-time Fourier transform to produce the time-domain enhanced signal. This particular issue can be somewhat alleviated, but with increased computational cost and system complexity [12]–[18].

Recently, there has been growing interest in the design of performant denoising pipelines that are optimized end-to-end and directly operate on the raw waveform. Such approaches aim at fully leveraging the expressive power of deep networks while avoiding expensive time-frequency transformations or loss of phase information [19]–[22]. Some of these approaches typically use simple regression loss functions for training the network [19], [20] (e.g., L^1 loss on the raw waveform), while ones with more advanced loss functions have shown limited gains in mismatched conditions [21], [22].

For our loss function, we are inspired by computer vision research, where activations in pretrained classification networks were found to yield effective loss functions for image stylization and synthesis [23], [24]. To compute the loss between two images, these approaches apply a pretrained image classification network to both. Each image induces a pattern of internal activations in the network to be compared, and the loss is defined in terms of their dissimilarity. Such complex training losses have been shown to yield state-of-the-art algorithms without the need for prior expert knowledge or added complexity for the processing network itself. Furthermore, increased performance can be achieved even without task-specific loss networks [25]. Our work develops this idea in the context of speech processing.

II. METHOD

A. Denoising Network

Let \mathbf{x} be an audio signal corresponding to speech β that is corrupted by an additive background signal \mathbf{n} so that $\mathbf{x} = \beta + \mathbf{n}$. Our goal is to find a denoising operator g such that $g(\mathbf{x}) \approx \beta$. We use a fully-convolutional network architecture based on context aggregation networks [26]. The output signal is synthesized sample by sample as we slide the network along the input. Context aggregation networks have been previously used in the WaveNet architecture for speech

F. Germain is with the Center for CCRMA, Stanford University, Stanford, CA 94305. E-mail: francois@ccrma.stanford.edu. This work was performed while he was interning at Intel Labs. Q. Chen and V. Koltun are with the Intelligent Systems Lab, Intel Labs, Santa Clara, CA 95054.

synthesis [27]. Our architecture is simpler than WaveNet – no skip connections across layers, no conditioning, no gated activations – while our loss function is more advanced, as described in Section II-B.

a) Context aggregation: Our network consists of 16 convolutional layers. The first and last layers (the degraded input signal and the enhanced output signal, respectively) are 1-dimensional tensors of dimensionality $N \times 1$. The number of samples N in the input signal varies and is not given in advance. The signal sampling frequency f_s is assumed to be 16 kHz. Each intermediate layer is a 2-dimensional tensor of dimensionality $N \times W$, where W is the number of feature maps in each layer. (We set $W = 64$.) The content of each intermediate layer is computed from the previous layer via a dilated convolution with 3×1 convolutional kernels [26] followed by an adaptive normalization (see below) and a pointwise nonlinear leaky rectified linear unit (LReLU) [28] $\max(0.2x, x)$. Because of the normalization, no bias term is used for the intermediate layers. We zero-pad all layers so that their “effective” length is constant at N . Our network is then trained to handle the beginning and end of audio files even when speech content is near the sequence edges.

The dilation operator aggregates long-range contextual information without changing sampling frequency across layers [26], [27]. Here, we increase the dilation factor exponentially with depth from 2^0 for the 1st intermediate layer to 2^{12} for the 13th one. We do not use dilation for the 14th and last one. For the output layer, we use a linear transformation (1×1 convolution plus bias with no normalization and no nonlinearity) to synthesize the sample of the output signal. The receptive field of the pipeline is $2^{14} + 1$ samples, i.e., about 1 s of audio for $f_s = 16$ kHz. We thus expect the system to capture context on the time scales of spoken words. A similar network architecture was shown to be advantageous in terms of compactness and runtime for image processing [29].

b) Adaptive normalization: The adaptive normalization operator used in our network matches the one proposed in [29] and improves performance and training speed. It adaptively combines batch normalization and identity mapping of the input x as the weighted sum $\alpha_k x + \beta_k BN(x)$ (where $\alpha_k, \beta_k \in \mathbb{R}$ are scalar weights for the k -th layer and BN is the batch normalization operator [30]). The weights α, β are learned by backpropagation as network parameters.

B. Feature loss

In our experiments, simple training losses (e.g., L^1) led to noticeably degraded output quality at lower signal-to-noise ratios (SNRs). The network seemed to improperly process low-energy speech information of perceptual importance. Instead, we train the denoising network using a deep feature loss that penalizes differences in the internal activations of a pretrained deep network that is applied to the signals being compared. By the nature of layered networks, feature activations at different depths in the loss network correspond to different time scales in the signal. Penalizing differences in these activations thus compares many features at different audio scales.

In computer vision, there are standard classification networks such as VGG-19 [31], pretrained on standard clas-

sification datasets such as ImageNet [32]. Such standard classification networks do not exist in the audio processing field yet, so we design and train our own feature loss network.

a) Feature loss network: We design a simple audio classification network inspired by the VGG architecture in computer vision [31], since it is known as a particularly effective feature loss architecture [25]. The network consists of 15 convolutional layers with 3×1 kernels, batch normalization, LReLU units, and zero padding. Each layer is decimated by 2, halving the length of the subsequent layer compared to the preceding one. The number of channels is doubled every 5 layers, with 32 channels in the first intermediate layer. Each channel in the last feature layer is average-pooled to yield the output feature vector. The receptive field is $2^{15} - 1$ samples. We train the network using backpropagation by feeding its output vector as features to one or more logistic classifiers with a cross-entropy loss for one or more classification tasks.

b) Denoising loss function: Let Φ^m be the m -th feature layer of the feature loss network, with layers at different depths corresponding to features with various time resolutions. The feature loss function is defined as a weighted L^1 loss on the difference between the feature activations induced in different layers of the network by the clean reference signal β and the output $g(\mathbf{x})$ of the denoising network being trained:

$$\mathcal{L}_{\beta, \mathbf{x}}(\theta) = \sum_{m=1}^M \lambda_m \|\Phi^m(\beta) - \Phi^m(g(\mathbf{x}; \theta))\|_1, \quad (1)$$

where θ are the parameters of the denoising network. The weights λ_m are set to balance the contribution of each layer to the loss. They are set to the inverse of the relative values of $\|\Phi^m(\beta) - \Phi^m(g(\mathbf{x}; \theta))\|_1$ after 10 training epochs. (For these first 10 epochs, the weights are set to 1.)

III. TRAINING

A. Feature Loss

a) Tasks: To generate a general-purpose feature loss network, we train it jointly on multiple audio classification tasks (only the logistic classifier parameters are trained as task-dependent). We use two tasks from the DCASE 2016 challenge [33]: the acoustic scene classification task and the domestic audio tagging task. In the first task, we are provided with audio files featuring various scenes (e.g., beach); the goal is to determine the scene type for each file. In the second task, we are given audio files featuring events of interest (e.g., child speaking); the goal is to determine which events took place in each file (with possibly multiple events in one file).

b) Data: For the scene classification task, the training set [34] consists of 30-second-long audio files sampled at 44.1 kHz, split among 15 different scenes (i.e., classes). As we need to develop a feature loss for the reduced sampling frequency of 16 kHz, we resample the data. The audio files are stereo, so we split them into two mono files. The training set contains 2,340 files. For the tagging task, the training set *CHiME-Home-refine* [35] consists of 4-second-long mono audio files sampled at 16 kHz, with 7 different tags (i.e., labels). The training set contains 1,946 files.

c) Training: Network weights are initialized with Xavier initialization [36]. We use the Adam optimizer [37] with a learning rate of 10^{-4} . The model is trained for 2,500 epochs. In each epoch, we iterate over the training data for each task, alternating between files from each task. The order of the files is randomized independently for each epoch. The dataset for the first task is larger than the one for the second task, so we present some of the files in the second dataset (chosen at random) a second time to preserve strict alternation between tasks. 1 epoch consists of 4,680 iterations (1 file per iteration). As a data augmentation procedure, we do not present entire clips, but present a continuous section of minimal duration 2^{15} samples that is culled at random for each iteration.

B. Speech Denoising

a) Data: We use the noisy dataset made available in [38]. To our knowledge, this is the largest available dataset for denoising that provides pre-mixed data with a clearly documented mixing procedure. It also has the benefit of being the dataset used in two recent works that we use as baselines. All details concerning the data can be found in [38]. The training set is generated from the speech data of 28 speakers (14 male/14 female) and the background data of 10 unique background types. Each noise segment is used to generate four files with 0, 5, 10, and 15dB SNR. The published files are sampled at 48kHz and normalized so that the clean speech files have a maximum absolute amplitude of 0.5. We resample them to 16kHz. The complete dataset comprises 11,572 files.

b) Training: Network weights and biases are initialized using the Xavier initialization and to zero, respectively. The adaptive normalization parameters are initialized at $\alpha = 1$ and $\beta = 0$. The feature loss is computed using the first $M = 6$ layers. We use the Adam optimizer with a learning rate of 10^{-4} . We train for 320 epochs (80 h) on a Titan X GPU. In each epoch, we present the entire dataset in randomized order (1 file per iteration) and files are presented in their entirety.

IV. EXPERIMENTAL SETUP

A. Baselines

As baselines, we use a Wiener filtering pipeline with a priori noise SNR estimation (as implemented in [39]), and two recent state-of-the-art methods that use deep networks to perform end-to-end denoising directly on the raw waveform: the Speech Enhancement Generative Adversarial Network (SEGAN) [21] and a WaveNet-based network [20]. This last one is designed around minor modifications to the architecture in [27]. It uses stacked context aggregation modules with gated activation units, skip connections, and a conditioning mechanism. The modifications include training with a regression loss (L^1 on the raw waveform) rather than a classification loss. The number of layers is larger than in our network (30), while the receptive field is smaller ($3 \cdot 2^{11}$ samples), capturing contextual information on more limited time scales. The network architecture is also distinctly more complex than ours. For both deep learning baselines, we use the code and models published by their respective authors. These models are optimized by their authors on the exact same training dataset, allowing fair comparison.

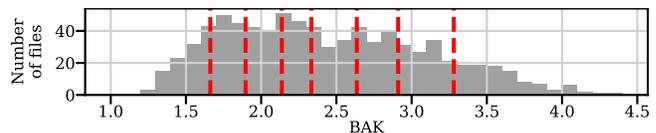


Fig. 1. Distribution of the test set in terms of composite background score. The test set was partitioned into 8 tranches, demarcated by red dashed lines.

B. Data

All our testing is done in mismatched conditions. The data source is the same as in Section III-B. The speech is obtained from 2 speakers (1 male/1 female). The background data is obtained from 5 distinct background types. Neither the speakers nor the backgrounds used at test time were seen during training. Each background segment is used to generate four files with 2.5, 7.5, 12.5, and 17.5 dB SNR. The complete test set comprises 824 files. Our denoising pipeline needs about 12 ms to process every 1 s of audio in our configuration. The denoised files for our pipeline and the baselines are available as supplementary material at <http://ieeexplore.ieee.org>.

C. Quantitative measures

a) Objective quality metrics: To evaluate each system, we compare its output to the ground-truth speech signal (i.e., the clean speech alone). The common metrics to measure speech quality given ground-truth are compared in [1]. We use here the composite scores from [39] that were found to be best correlated with human listener ratings. These consist of the overall (OVL), the signal (SIG), and the background (BAK) scores, each on a scale from 1.0 to 5.0, and corresponding respectively to the measure of overall signal quality, the measure of quality when considering speech signal degradation alone, and the measure of quality when considering background signal intrusiveness alone [40]. We also report the SNR [41], as a raw measure of the relative energies of the residual background and the speech in a given signal, quantified in decibel (dB). We use the implementations in [1]. For all metrics, higher scores denote better performance.

The test dataset is divided into 4 mixing SNR subgroups (see Section IV-B). We argue that the dataset should be rather considered as a continuous distribution of degradation, since SNR correlates poorly with human perception of the degradation level [1]. The continuum of degradation levels is better represented in the distribution of the background intrusiveness BAK score. (The SIG score is less informative since the undistorted speech signal is added.) To evaluate performance as a function of input degradation magnitude, we partition the test set into 8 tranches of equal size, corresponding to the 8 octiles of the BAK score distribution as shown in Figure 1, with tranches representing a different denoising difficulty.

TABLE I
PERFORMANCE FOR DIFFERENT APPROACHES ACCORDING TO OBJECTIVE QUALITY MEASURES. (HIGHER IS BETTER.)

	SNR	SIG	BAK	OVL
Noisy	8.45	3.34	2.44	2.63
Wiener	12.28	3.23	2.68	2.67
SEGAN	14.82	3.21	2.76	2.56
WaveNet	18.18	2.87	3.08	2.43
Ours	19.00	3.86	3.33	3.22

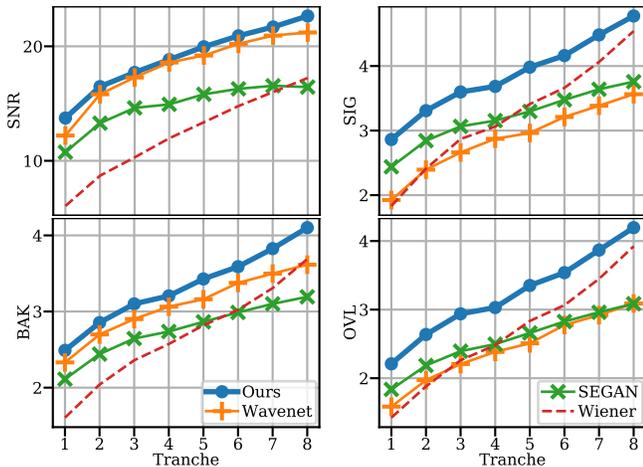


Fig. 2. Performance of different denoising approaches according to 4 objective quality measures (SNR, SIG, BAK, and OVL), plotted for each tranche in the test set. For all measures, higher is better.

b) Results: Table I reports these metrics for our approach and the baselines, evaluated over the test set. Our method outperforms all the baselines according to all measures by a comfortable margin. The plots in Figure 2 further show that our network yields the best quality for all levels of background intrusiveness separated in tranches, with a particularly significant margin according to perceptually-motivated composite measures. Table II shows the benefit of using a feature loss compared to training the same denoising network, by the same procedure on the same data, using an L^1 or an L^2 loss. Training with a feature loss outperforms networks trained with other losses. In particular, while an L^1 loss achieves a similar SNR score as our feature loss, the feature loss shows definite improvement for the BAK and OVL metrics. It also scores well for the SIG metric, especially in the noisier tranches, demonstrating the ability to capture meaningful features when important cues are hidden in the noise.

D. Perceptual Experiments

a) Experimental design: Objective metrics are known to only partially correlate with human audio quality ratings [1]. Hence, we also conduct carefully designed perceptual experiments with human listeners. The procedure is based on A/B tests deployed at scale on the Amazon Mechanical Turk platform. The A/B tests are grouped into Human Intelligence Tasks (HITs). Each HIT consists of 100 “ours vs baseline” pairwise comparisons. Each comparison presents two audio clips that can be played in any order by the worker, any number of times. One of the clips is the output of our approach and one is the output of one of the baselines, for the same input from the test set. The files are presented in random order (both within each pair and among pairs), so the worker is given no information as to the provenance of the clips. The worker is asked to select, within each pair, the clip with the cleaner speech. Each HIT includes 10 additional ‘sentry’ comparisons in which the right answer is obvious to guard against negligent or inattentive workers. These sentry pairs are mixed into the HIT in random order. If a worker gives an incorrect answer to two or more sentry pairs, the entire HIT is discarded. Each HIT then contains a total of 110 pairwise comparisons. A worker

TABLE II
TRAINING THE SAME NETWORK WITH DIFFERENT LOSS FUNCTIONS. FOR ALL METRICS, HIGHER IS BETTER.

	SNR	SIG	BAK	OVL
Noisy	8.45	3.34	2.44	2.63
L2	18.46	3.70	3.21	3.07
L1	18.98	3.75	3.27	3.11
Feature loss	19.00	3.86	3.33	3.22

is given 1 hour to complete a HIT. Each HIT is completed by 10 distinct workers.

b) Results: The results are summarized in Table III. This table presents the fraction of blind pairwise A/B comparisons in which the listener rated a clip denoised by our network as cleaner than the clip denoised by a baseline. The preference rates are presented versus each baseline across 4 tranches. The most notable results are for the hardest tranche, where the output of our approach was rated cleaner than the output of recent state-of-the-art deep networks in more than 83% of the comparisons. All results are statistically significant with $p < 10^{-3}$. This demonstrates that our algorithm is more robust in this regime, in which degradation from the background signal is much more noticeable, and for which denoising is particularly useful. For easier tranches, with lower levels of degradation in the input, both our method and the baselines generally perform satisfactorily and listeners can experience more difficulty distinguishing between the different processed files, but the preference rate for our approach remains well above chance (50%), at statistically significant levels, for all baselines across all tranches.

V. CONCLUSION

We presented an end-to-end speech denoising pipeline that uses a fully-convolutional network, using a deep feature loss network pretrained on several relevant audio classification tasks for training. This approach allows the denoising system to capture speech structure at various scales and achieve better denoising performance without added complexity in the system itself or expert knowledge in the loss design. Experiments demonstrate that our approach significantly outperforms recent state-of-the-art baselines according to objective speech quality measures as well as large-scale perceptual experiments with human listeners. In particular, the presented approach is shown to perform much better in the noisiest conditions where speech denoising is most challenging. Our paper validates the combined use of convolutional context aggregation networks and feature losses to achieve state-of-the-art performance.

TABLE III
RESULTS OF PERCEPTUAL EXPERIMENTS. EACH CELL LISTS THE FRACTION OF BLIND RANDOMIZED PAIRWISE COMPARISONS IN WHICH THE LISTENER RATED THE OUTPUT OF OUR APPROACH AS CLEANER THAN THE OUTPUT OF A BASELINE. EACH ROW LISTS RESULTS FOR A SPECIFIC BASELINE. EACH COLUMN LISTS RESULTS FOR A TRANCHE OF THE TESTING SET. (CHANCE IS AT 50%, HIGHER IS BETTER.)

Tranche:	1 (Hard)	3 (Medium)	5 (Easy)	7 (Very easy)
Ours > Wiener	96.1%	89.4%	81.7%	90.2%
Ours > SEGAN	83.5%	70.5%	64.1%	61.4%
Ours > WaveNet	83.9%	67.0%	61.4%	55.8%

REFERENCES

- [1] P. C. Loizou, *Speech Enhancement: Theory and Practice*, 2nd ed. CRC Press, 2013.
- [2] M. Bosi and R. E. Goldberg, *Introduction to Digital Audio Coding and Standards*. Springer, 2002.
- [3] P. Smaragdis, C. Fevotte, G. J. Mysore, N. Mohammadiha, and M. Hoffman, “Static and dynamic source separation using nonnegative factorizations: A unified view,” *IEEE Signal Processing Magazine*, vol. 31, no. 3, 2014.
- [4] Y. Wang and D. Wang, “Cocktail party processing via structured prediction,” in *Neural Information Processing Systems (NIPS)*, 2012.
- [5] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, “Speech enhancement based on deep denoising autoencoder,” in *Interspeech*, 2013.
- [6] A. Narayanan and D. Wang, “Ideal ratio mask estimation using deep neural networks for robust speech recognition,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2013.
- [7] F. Weninger, J. R. Hershey, J. L. Roux, and B. Schuller, “Discriminatively trained recurrent neural networks for single-channel speech separation,” in *IEEE Global Conference on Signal and Information Processing*, 2014.
- [8] 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, 2015.
- [9] A. Kumar and D. Florencio, “Speech enhancement in multiple-noise conditions using deep neural networks,” *arXiv:1605.02427*, 2016.
- [10] X.-L. Zhang and D. Wang, “A deep ensemble learning method for monaural speech separation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 5, 2016.
- [11] J. Chen and D. Wang, “Long short-term memory for speaker generalization in supervised speech separation,” *Journal of the Acoustical Society of America*, vol. 141, no. 6, 2017.
- [12] J. L. Roux and E. Vincent, “Consistent Wiener filtering for audio source separation,” *IEEE Signal Processing Letters*, vol. 20, no. 3, 2013.
- [13] F. G. Germain, G. J. Mysore, and T. Fujioka, “Equalization matching of speech recordings in real-world environments,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2016.
- [14] T. Gerkmann, M. Krawczyk-Becker, and J. L. Roux, “Phase processing for single-channel speech enhancement: History and recent advances,” *IEEE Signal Processing Magazine*, vol. 32, no. 2, 2015.
- [15] Y. Wang and D. Wang, “A deep neural network for time-domain signal reconstruction,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [16] H. Erdogan, J. R. Hershey, S. Watanabe, and J. L. Roux, “Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015.
- [17] 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, 2017.
- [18] J. A. Moorer, “A note on the implementation of audio processing by short-term Fourier transform,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2017.
- [19] S.-W. Fu, Y. Tsao, X. Lu, and H. Kawai, “Raw waveform-based speech enhancement by fully convolutional networks,” *arXiv:1703.02205*, 2017.
- [20] D. Rethage, J. Pons, and X. Serra, “A WaveNet for speech denoising,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2018.
- [21] S. Pascual, A. Bonafonte, and J. Serrà, “SEGAN: Speech enhancement generative adversarial network,” in *Interspeech*, 2017.
- [22] K. Qian, Y. Zhang, S. Chang, X. Yang, D. Florencio, and M. Hasegawa-Johnson, “Speech enhancement using Bayesian WaveNet,” in *Interspeech*, 2017.
- [23] J. Johnson, A. Alahi, and L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *European Conference on Computer Vision (ECCV)*, 2016.
- [24] Q. Chen and V. Koltun, “Photographic image synthesis with cascaded refinement networks,” in *International Conference on Computer Vision (ICCV)*, 2017.
- [25] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, “The unreasonable effectiveness of deep features as a perceptual metric,” in *Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [26] F. Yu and V. Koltun, “Multi-scale context aggregation by dilated convolutions,” in *International Conference on Learning Representations (ICLR)*, 2016.
- [27] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, “WaveNet: A generative model for raw audio,” *arXiv:1609.03499*, 2016.
- [28] 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.
- [29] Q. Chen, J. Xu, and V. Koltun, “Fast image processing with fully-convolutional networks,” in *International Conference on Computer Vision (ICCV)*, 2017.
- [30] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning (ICML)*, 2015.
- [31] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [32] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li, “ImageNet large scale visual recognition challenge,” *International Journal on Computer Vision (IJCV)*, vol. 115, no. 3, 2015.
- [33] A. Mesaros, T. Heittola, E. Benetos, P. Foster, M. Lagrange, T. Virtanen, and M. D. Plumbley, “Detection and classification of acoustic scenes and events: Outcome of the DCASE 2016 challenge,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 2, 2018.
- [34] A. Mesaros, T. Heittola, and T. Virtanen, “TUT database for acoustic scene classification and sound event detection,” in *European Signal Processing Conference (EUSIPCO)*, 2016.
- [35] P. Foster, S. Sigtia, S. Krstulovic, J. Barker, and M. D. Plumbley, “CHiMe-Home: A dataset for sound source recognition in a domestic environment,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2015.
- [36] X. Glorot and Y. Bengio, “Understanding the difficulty of training deep feedforward neural networks,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2010.
- [37] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2015.
- [38] C. Valentini-Botinhao, X. Wang, S. Takaki, and J. Yamagishi, “Investigating RNN-based speech enhancement methods for noise-robust text-to-speech,” in *ISCA Speech Synthesis Workshop*, 2016.
- [39] Y. Hu and P. C. Loizou, “Subjective comparison of speech enhancement algorithms,” in *IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, 2006.
- [40] ITU-T, “Subjective test methodology for evaluating speech communication systems that include noise suppression algorithm,” ITU-T Recommendation P.835, Tech. Rep., 2003.
- [41] S. R. Quackenbush, T. P. Barnwell, and M. A. Clements, *Objective Measures of Speech Quality*. Prentice Hall, 1988.

APPENDIX

This appendix provides additional details on the denoising and feature loss network architectures presented in Section II.

A. Denoising Network

a) *Layer structure*: We denote the 16 (consecutive) network layers by $\Lambda^0, \dots, \Lambda^{15}$. Λ^0 and Λ^{15} are 1-dimensional tensors of dimensionality $N \times 1$ and correspond to the degraded input signal and the enhanced output signal, respectively. The number of samples N is not given in advance. Each intermediate layer $\Lambda^k \in \{\Lambda^1, \dots, \Lambda^{15}\}$ is a 2-dimensional tensor of dimensionality $N \times W$, where W is the width of (i.e., the number of feature maps in) each layer. For $k = 1, \dots, 14$, the content of each intermediate layer Λ^k is computed from the previous layer Λ^{k-1} via the operation

$$\Lambda_i^k = \Psi \left(\Gamma^k \left(\sum_j \Lambda_j^{k-1} *_r \mathbf{K}_{i,j}^k \right) \right), \quad (2)$$

where Λ_i^k is the i -th feature map of layer Λ^k , Λ_j^{k-1} is the j -th feature map of layer Λ^{k-1} , $\mathbf{K}_{i,j}^k$ is a learned 3×1 convolutional kernel, Γ^k is the adaptive normalization operator and Ψ is a pointwise nonlinearity. Because of the presence of adaptive normalization, no bias term is used for these layers. The operator $*_r$ is a dilated convolution [26], i.e.,

$$(\Lambda_j *_r \mathbf{K}_{i,j})[n] = \sum_{m=-1}^{+1} \mathbf{K}_{i,j}[m] \Lambda_j[n - rm]. \quad (3)$$

The dilation factor for the k -th layer is set at $r_k = 2^{k-1}$ for $k \in \{1, \dots, 13\}$. Between layer Λ^{13} and Λ^{14} , we do not use dilation (i.e., $r_{14} = 1$). For the output layer Λ^{15} , we use a linear transformation (1×1 convolution with no nonlinearity) in order to synthesize the sample of the output signal so that

$$\Lambda^{15} = \sum_j \Lambda_j^{14} \times \mathbf{K}_j^{14} + b, \quad (4)$$

where b is a learned bias term. The receptive field of the network is $2^{14} + 1 = 16385$ samples.

b) *Nonlinear units*: For the pointwise nonlinearity Ψ , we use the leaky rectified linear unit (LReLU) [28]:

$$\Psi(x) = \max(\delta x, x) \text{ with } \delta = 0.2. \quad (5)$$

c) *Adaptive normalization*: Γ^k corresponds to the adaptive normalization operation described in Section II-A. For $k \in \{1, \dots, 13\}$, the operator adaptively combines batch normalization and identity mapping as

$$\Gamma^k(x) = \alpha_k x + \beta_k BN(x), \quad (6)$$

where $\alpha_k, \beta_k \in \mathbb{R}$ are learned scalar weights and BN is the batch normalization operator [30].

d) *Zero padding*: Our algorithm uses zero-padding at each layer so that the “effective” length of each layer tensor is constant and identical to N .

e) *Training loss*: The network is trained through backpropagation using our deep feature loss as described in Section II-B (see in particular Equation 1). The feature loss classification network is further detailed in the next section.

B. Feature Loss Network

a) *Feature layer structure*: As mentioned in Section II-B, the network is inspired by the VGG architecture from computer vision. We denote its 15 (consecutive) layers by Φ^0, \dots, Φ^{14} . The first layer Φ^0 is a 1-dimensional tensor of dimensionality $N \times 1$ and corresponds to the input signal. The number of samples N is not given in advance. Each intermediate layer $\Phi^m \in \{\Phi^1, \dots, \Phi^{14}\}$ is a 2-dimensional tensor of dimensionality $\frac{N}{2^m} \times W_m$, where W_m is the width of each layer, set to $W_m = 32 \times 2^{\lfloor \frac{m-1}{5} \rfloor}$ (i.e., the number of features is doubled every 5 layers). The content of each intermediate layer Φ^m is computed from the previous layer Φ^{m-1} through the following operation:

$$\tilde{\Phi}_i^m = \Psi \left(BN \left(\sum_j \Phi_j^{m-1} * \mathbf{L}_{i,j}^m \right) \right), \quad (7)$$

where $\tilde{\Phi}_i^m$ is the i -th feature map of layer Φ^m prior to the decimation operation, Φ_j^{m-1} is the j -th feature map of layer Φ^{m-1} , $\mathbf{L}_{i,j}^m$ is a learned 3×1 convolutional kernel, BN is the batch normalization operator, and Ψ is the same pointwise linearity as in Equation 5. Because of the presence of batch normalization, no bias term is used for these layers. This is followed by the decimation operation

$$\Phi_i^m[n] = \tilde{\Phi}_i^m[2n], \quad (8)$$

following which the length of the subsequent layer is half the length of the preceding one. The receptive field of the network is $2^{15} - 1 = 32767$ samples. The network is zero-padded as necessary for each layer so that $\tilde{\Phi}^m$ and Φ^{m-1} have the same “effective” length.

b) *Classification layer*: To perform the p -th classification task of interest, we first average-pool each channel in the last feature layer $\tilde{\Phi}^{14}$ to yield an output feature vector $\Phi^{15,p}$ of dimensionality $1 \times W_{14}$. This vector is fed to a linear layer to form a logit vector $\Phi^{15,p}$ of dimensionality $1 \times C_p$ (with C_p the number of classes associated with the p -th task) such that

$$\Phi_i^{16,p} = \sum_j \Phi_j^{15,p} \times \mathbf{L}_{i,j}^{16,p} + \tilde{b}_i^p, \quad (9)$$

where $\tilde{\mathbf{L}}_{i,j}^{16,p}$ is a learned scalar weight and \tilde{b}_i^p is a learned bias term. We finally get the output classification vector $\Phi^{17,p}$ of the network through the operation

$$\Phi^{17,p} = \Delta(\Phi^{16,p}), \quad (10)$$

where Δ is the logistic nonlinearity associated with the type of multi-label classification for the p -th task (i.e., vector softmax nonlinearity if the task asks for a unique label for each audio file, pointwise sigmoid if the task allows for any number of labels for each audio file). $\Phi^{17,p}$ is of dimension $1 \times C_p$ and its elements are in the range $[0, 1]$.

c) *Training loss*: Training is done through backpropagation using a cross-entropy loss between the vector $\Phi^{17,p}$ associated with the current file (for task p) and its corresponding ground truth classification vector (i.e., the vector of dimension $1 \times C_p$ in which the c -th element is 1 if the c -th classification label is associated with the file, 0 otherwise).