MONAURAL SPEECH ENHANCEMENT USING DEEP NEURAL NETWORKS BY MAXIMIZING A SHORT-TIME OBJECTIVE INTELLIGIBILITY MEASURE

Morten Kolbæk, Zheng-Hua Tan, Jesper Jensen

Department of Electronic Systems, Aalborg University, Aalborg, Denmark {mok,zt,jje}@es.aau.dk

ABSTRACT

In this paper we propose a Deep Neural Network (DNN) based Speech Enhancement (SE) system that is designed to maximize an approximation of the Short-Time Objective Intelligibility (STOI) measure. We formalize an approximate-STOI cost function and derive analytical expressions for the gradients required for DNN training and show that these gradients have desirable properties when used together with gradient based optimization techniques.

We show through simulation experiments that the proposed SE system achieves large improvements in estimated speech intelligibility, when tested on matched and unmatched natural noise types, at multiple signal-to-noise ratios. Furthermore, we show that the SE system, when trained using an approximate-STOI cost function performs on par with a system trained with a mean square error cost applied to short-time temporal envelopes. Finally, we show that the proposed SE system performs on par with a traditional DNN based Short-Time Spectral Amplitude (STSA) SE system in terms of estimated speech intelligibility. These results are important because they suggest that traditional DNN based STSA SE systems might be optimal in terms of estimated speech intelligibility.

Index Terms— Speech Enhancement, Deep Neural Networks, Speech Intelligibility, Speech Denoising, Deep Learning.

1. INTRODUCTION

Design and development of Speech Enhancement (SE) algorithms capable of improving speech quality and intelligibility has been a long-lasting goal in both academia and industry [1, 2]. Such algorithms are useful for a wide range of applications e.g. for mobile communications devices and hearing assistive devices [1].

Despite a large research effort for more than 30 years [1–3] modern single-microphone SE algorithms still perform unsatisfactorily in the complex acoustic environments, which users of e.g. hearing assistive devices are exposed to on a daily basis, e.g. traffic noise, cafeteria noise, or competing speakers.

Traditionally, SE algorithms have been divided into at least two groups; statistical-model based techniques and data-driven techniques. The first group encompasses techniques such as spectral subtraction, the Wiener filter and the short-time spectral amplitude minimum mean square error estimator [1–3]. These techniques make statistical assumptions about the probability distributions of the speech and noise signals, that enable them to suppress the noise dominated time-frequency regions of the noisy speech signal. In particularly, for stationary noise types this type of algorithms may perform well in terms of speech quality, but in general these techniques do not improve speech intelligibility [4–6]. The second group encompasses data-driven or machine learning techniques e.g. based on non-negative matrix factorization [7], support vector machines [8], and Deep Neural Networks (DNNs) [9, 10]. These techniques make no statistical assumptions. Instead, they learn to suppress noise by observing a large number of representative pairs of noisy and noise-free speech signals in a supervised learning process. SE algorithms based on DNNs can, to some extent, improve speech intelligibility for hearing impaired and normal hearing people, in noisy conditions, if sufficient *a priori* knowledge is available e.g. the identity of the speaker or the noise type. [11–13].

Although the techniques mentioned above are fundamentally different, they typically share at least two common properties. First, they often aim to minimize a Mean Square Error (MSE) cost function, and secondly, they operate on short frames ($\approx 20 - 30$ ms) in the Short-Time discrete Fourier Transform (STFT) domain [1,2]. However, it is well known [2, 14] that the human auditory system has a non-linear frequency sensitivity, which is often approximated using e.g. a Gammatone or a one-third octave filter bank [2]. Furthermore, it is known that preservation of modulation frequencies below 7 Hz is critical for speech intelligibility [14, 15]. This suggests that SE algorithms aimed at the human auditory system could benefit by incorporating such information. Numerous works exist, e.g. [10, 16-23] and [1, Sec. 2.2.3] and the references therein, where SE algorithms have been designed with perceptual aspects in mind. However, although these algorithms do take some perceptual aspects into account, they do not directly optimize for speech intelligibility.

In this paper we propose an SE system that maximizes an objective speech intelligibility estimator. Specifically, we design a DNN based SE system that maximizes an approximation of the Short-Time Objective Intelligibility (STOI) [24] measure. The STOI measure has been found to be highly correlated with intelligibility as measured in human listening tests [2, 24]. We derive analytical expressions for the required gradients used for the DNN weight updates during training and use these closed-form expressions to identify desirable properties of the approximate-STOI cost function. Finally, we study the potential performance gain between the proposed approximate-STOI cost function with a classical MSE cost function. We note that our goal is not to achieve state-of-the-art STOI improvements per se, but rather to study and compare the proposed approximate-STOI based SE system to existing DNN based enhancement schemes. Further improvement may straightforwardly be achieved with larger datasets and complex models like long shortterm memory recurrent, or convolutional, neural networks [25].

2. SPEECH ENHANCEMENT SYSTEM

In the following we introduce the approximate-STOI measure and we present the DNN framework used to maximize it. Finally, we discuss techniques used to reconstruct the enhanced and approximate-STOI optimal speech signal in the time-domain.

2.1. Approximating Short-Time Objective Intelligibility

Let x[n] be the n^{th} sample of the clean time-domain speech signal and let a noisy observation y[n] be defined as

$$y[n] = x[n] + z[n], \tag{1}$$

where z[n] is an additive noise sample. Furthermore, let x(k,m)and y(k,m), $k = 1, \ldots, \frac{K}{2} + 1$, $m = 1, \ldots M$, be the single-sided magnitude spectra of the K-point Short-Time discrete Fourier Transforms (STFT) of x[n] and y[n], respectively, where M is the number of STFT frames. Also, let $\hat{x}(k,m)$ be an estimate of x(k,m) obtained as $\hat{x}(k,m) = \hat{g}(k,m)y(k,m)$ where $\hat{g}(k,m)$ is an estimated gain value. In this study we use a 10 kHz sample frequency and a 256 point STFT, i.e. K = 256, with a Hann-window size of 256 samples (25.6 ms) and a 128 sample frame shift (12.8 ms). Similarly to STOI [24], we define a short-time temporal envelope vector of the j^{th} one-third octave band for the clean speech signal as

$$\mathbf{x}_{j,m} = [X_j(m-N+1), X_j(m-N+2), \dots, X_j(m)]^T, (2)$$

where

$$X_j(m) = \sqrt{\sum_{k=k_1(j)}^{k_2(j)-1} x(k,m)^2},$$
(3)

and $k_1(j)$ and $k_2(j)$ denote the first and last STFT bin index of the j^{th} one-third octave band, respectively. Similarly, we define $\mathbf{y}_{j,m}$ and $Y_j(m)$ for the noisy observation. Also, let $\hat{\mathbf{x}}_{j,m} = diag(\hat{\mathbf{g}}_{j,m})\mathbf{y}_{j,m}$ be the short-time temporal one-third octave band envelope vector of the enhanced speech signal, where $\hat{\mathbf{g}}_{j,m}$ is a gain vector defined in the j^{th} one-third octave band and $diag(\hat{\mathbf{g}}_{j,m})$ is a diagonal matrix with the elements of $\hat{\mathbf{g}}_{j,m}$ on the main diagonal. We use N = 30 such that the short-time temporal one-third octave band envelope vectors will span a duration of 384 ms, which ensures that important modulation frequencies are captured [24]. In total, J = 15 one-third octave bands are used with the first band having a center frequency of 150 Hz and the last one of approximately 3.8 kHz. These frequencies are chosen such that they span the frequency range in which human speech normally lie [24]. For mathematical tractability, we discard the clipping step¹, otherwise performed by STOI [24], and define the approximated STOI measure as

$$\mathcal{L}(\mathbf{x}_{j,m}, \hat{\mathbf{x}}_{j,m}) = \frac{\left(\mathbf{x}_{j,m} - \mu_{\mathbf{x}_{j,m}}\right)^T \left(\hat{\mathbf{x}}_{j,m} - \mu_{\hat{\mathbf{x}}_{j,m}}\right)}{\|\mathbf{x}_{j,m} - \mu_{\mathbf{x}_{j,m}}\| \|\hat{\mathbf{x}}_{j,m} - \mu_{\hat{\mathbf{x}}_{j,m}}\|}, \quad (4)$$

where $\|\cdot\|$ is the euclidean ℓ^2 -norm and $\mu_{\mathbf{x}_{j,m}}$ and $\mu_{\hat{\mathbf{x}}_{j,m}}$ are the sample means of $\mathbf{x}_{j,m}$ and $\hat{\mathbf{x}}_{j,m}$, respectively. Obviously, $\mathcal{L}(\mathbf{x}_{j,m}, \hat{\mathbf{x}}_{j,m})$ is simply the Envelope Linear Correlation (ELC) between the vectors $\mathbf{x}_{j,m}$ and $\hat{\mathbf{x}}_{j,m}$.

2.2. Maximizing the Approximated STOI Measure using DNNs

The approximated STOI measure given by Eq. (4) is defined in a one-third octave band domain and our goal is to find $\hat{\mathbf{x}}_{j,m} = diag(\hat{\mathbf{g}}_{j,m})\mathbf{y}_{j,m}$ such that Eq. (4) is maximized, i.e. finding an optimal gain vector $\hat{\mathbf{g}}_{j,m}$. In this study we estimate these optimal gains using DNNs. Specifically, we use Eq. (4) as a cost function and train multiple feed-forward DNNs, one for each one-third octave band, to estimate gain vectors $\hat{\mathbf{g}}_{j,m}$, such that the approximated



Fig. 1. ℓ^2 -norm of Eq. (5) as function of cost function value.

STOI measure is maximized. For the remainder of this paragraph we omit the subscripts j and m for convenience.

Most modern deep learning toolkits, e.g. Microsoft Cognitive Toolkit (CNTK) [29], perform automatic differentiation, which allow one to train a DNN with a custom cost function, without the need of computing the gradients of the cost function explicitly [25]. Nevertheless, when working with cost functions that have not yet been exhaustively studied, such as the approximated STOI measure, an analytic expression of the gradient can be valuable for studying important properties, such as gradient ℓ^2 -norm. It can be shown (details omitted due to space limitations) that the gradient of Eq. (4), with respect to the desired signal vector $\hat{\mathbf{x}}$, is given by

$$\nabla \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) = \left[\frac{\partial \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})}{\partial \hat{x}_1}, \frac{\partial \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})}{\partial \hat{x}_2}, \dots, \frac{\partial \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})}{\partial \hat{x}_N}\right]^T \quad (5)$$

where

$$\frac{\partial \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})}{\partial \hat{x}_m} = \frac{\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) \left(x_m - \mu_{\mathbf{x}} \right)}{\left(\hat{\mathbf{x}} - \mu_{\hat{\mathbf{x}}} \right)^T \left(\mathbf{x} - \mu_{\mathbf{x}} \right)} - \frac{\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) \left(\hat{x}_m - \mu_{\hat{\mathbf{x}}} \right)}{\left(\hat{\mathbf{x}} - \mu_{\hat{\mathbf{x}}} \right)^T \left(\hat{\mathbf{x}} - \mu_{\hat{\mathbf{x}}} \right)}, \quad (6)$$

is the partial derivative of $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})$ with respect to entry m of $\hat{\mathbf{x}}$.

Furthermore, it can be shown that the ℓ^2 -norm of the gradient as formulated by Eqs. (5) and (6), is given by

$$\left\|\nabla \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})\right\| = \sqrt{1 - \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})^2} \left\|\hat{\mathbf{x}}\right\|^{-1},\tag{7}$$

which is shown in Fig. 1 as function of $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})$ for the complete range [-1, 1], and for $\|\hat{\mathbf{x}}\| = 1$. We see from Fig. 1 that the ℓ^2 norm of $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})$ is a concave function with a global maximum at $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) = 0$ and is symmetric around zero. We also observe that $\|\nabla \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})\|$ is monotonically decreasing when $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) < 0$ and $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}}) > 0$ with $\|\nabla \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})\| = 0$ when \mathbf{x} and $\hat{\mathbf{x}}$ are either perfectly correlated or perfectly anti-correlated. Since $\|\nabla \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})\|$ is large when \mathbf{x} and $\hat{\mathbf{x}}$ are uncorrelated and zero when perfectly correlated, and $\|\nabla \mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})\| \neq 0$ otherwise, Eq. (4) is well suited as a cost function for gradient-based optimization techniques, such as Stochastic Gradient Descent (SGD) [25], since it guarantees nonzero step lengths for all inputs during optimization except at the optimal solution. In practice, to apply SGD we minimize $-\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})$.

2.3. Reconstructing Approximate-STOI Optimal Speech

When a gain vector $\hat{\mathbf{g}}_{j,m}$ has been estimated by a DNN, the enhanced speech envelope in the one-third octave band domain can be computed as $\hat{\mathbf{x}}_{j,m} = diag(\hat{\mathbf{g}}_{j,m})\mathbf{y}_{j,m}$. However, what we are really interested in is $\hat{x}(k,m)$, i.e. the estimated speech signal in the STFT domain, since $\hat{x}(k,m)$ can straightforwardly be transformed into the time-domain using the overlap-and-add technique [2]. We therefore seek a mapping from the gain vector $\hat{\mathbf{g}}_{j,m}$ estimated in the one-third octave band domain, to the gain $\hat{g}(k,m)$, for a single STFT coefficient. To do so, let $\hat{g}_j(m)$ denote the gain value estimated by a DNN to be applied to the noisy one-third octave band amplitude in frame m. We can then derive the relationship between the gain value

¹It has been observed empirically, that omitting the clipping step most often does not affect the performance of STOI, e.g. [20, 26–28].

 $\hat{g}_j(m) \ge 0$ in the one-third octave band, and the corresponding gain values $\hat{g}(k,m) \ge 0$ in the STFT domain as

$$\hat{X}_{j}(m) = \hat{g}_{j}(m)Y_{j}(m) = \sqrt{\sum_{k=k_{1}(j)}^{k_{2}(j)-1} (\hat{g}(k,m)y(k,m))^{2}}.$$
 (8)

One solution to Eq (8) is

$$\hat{g}_j(m) = \hat{g}(k,m), \ k = k_1(j), \dots k_2(j) - 1.$$
 (9)

Generally, the solution in Eq. (9) is not unique; many choices of $\hat{g}(k,m)$ exist that give rise to the same estimated one-third octave band $\hat{X}_j(m)$ (and hence the same value of $\mathcal{L}(\mathbf{x}, \hat{\mathbf{x}})$). We choose, for convenience, a uniform gain across the STFT coefficients within a one-third octave band. Since envelope estimates $\hat{X}_j(m)$ are computed for successive values of m, N estimates exist for each $\hat{X}_j(m)$, which are averaged during enhancement. When reconstructing the enhanced speech signal in the time domain, we use the overlap-andadd technique using the phase of the noisy STFT coefficients [2].

3. EXPERIMENTAL DESIGN

To evaluate the performance of the approximate-STOI optimal DNN based SE system we have conducted series of experiments involving multiple matched and unmatched noise types at various SNRs.

3.1. Noisy Speech Mixtures

The clean speech signals used for training all models are from the Wall Street Journal corpus [30]. The utterances used for training and validation are generated by randomly selecting utterances from 44 male and 47 female speakers from the WSJ0 training set entitled si_tr_s. The training and validation sets consist of 20000 and 2000 utterances, respectively, which is equivalent to approximately 37 hours of training data and 4 hours of validation data. The test set is similarly generated using utterances from 16 speakers from the WSJ0 validation set si_dt_05 and evaluation set si_et_05, and consists of 1000 mixtures or approximately 2 hours of data, see [31] for further details. Notice, the speakers in the test set are different from the speakers in the validation and training sets.

We use six different noise types: two synthetic signals and four noise signals recorded in real-life. The synthetic noise signals encompass a stationary Speech Shaped Noise (SSN) signal and a highly non-stationary 6-speaker Babble (BBL) noise. For real-life noise signals we use the street (STR), cafeteria (CAF), bus (BUS), and pedestrian (PED) noise signals from the CHiME3 dataset [32]. The SSN noise signal is Gaussian white noise, shaped according to the long-term spectrum of the TIMIT corpus [33]. Similarly, the BBL noise signal is constructed by mixing utterances from TIMIT. Further details on the design of the SSN and BBL noise signals can be found in [13]. All noise signals are split into non-overlapping sequences with a 40 min. training sequence, a 5 min. validation sequence and a 5 min. test sequence, i.e. there is no overlap between the noise sequences used for training, validation and test.

The noisy speech signals used for training and testing are constructed using Eq. (1), where a clean speech signal x[n] is added to a noise sequence z[n] of equal length. To achieve a certain SNR, the noise signal is scaled based on the active speech level of the clean speech signal as per ITU P.56 [34]. The SNRs used for the training and validation sets are chosen uniformly from [-5, 10] dB. The SNR range is chosen to ensure that SNRs are included where intelligibility ranges from degraded to perfectly intelligible.

Table 1. Training conditions for different SE systems.

ID:	SO	S1	S2	S 3	S4	S5	S6	S 7	S 8	S9
Cost:	ELC	ELC	ELC	ELC	ELC	EMSE	EMSE	EMSE	EMSE	EMSE
Noise:	SSN	BBL	CAF	STR	ALL	SSN	BBL	CAF	STR	ALL

3.2. Model Architecture and Training

To evaluate the performance of the proposed SE system a total of ten systems, identified as S0 – S9, have been trained using different cost functions and noise types as presented in Table 1. Five systems (S0–S4) have been trained using the ELC loss from Eq. (4) and five systems (S5–S9) have been trained using a standard MSE loss, denoted as Envelope MSE (EMSE), since it operates on short-time temporal one-third octave band envelope vectors. This is to investigate the potential performance difference between models trained with an approximate-STOI loss and models trained with the commonly used MSE loss. Eight systems (S0–S3 and S5–S8) are trained as noise type specific systems, i.e. they are trained as noise type general systems, i.e. they are trained on all noise types (Noise: "ALL" in Table 1). This is to investigate the performance drop, if any, when a single system is trained to handle multiple noise types.

Each DNN consists of three hidden layers with 512 units with ReLU activation functions and a sigmoid output layer. The DNNs are trained using SGD with the backpropagation technique and batch normalization [25]. The DNNs are trained for a maximum of 200 epochs with a minibatch size of 256 randomly selected short-time temporal one-third octave band envelope vectors and the learning rates were set to 0.01, and $5 \cdot 10^{-5}$ per sample initially, for S0–S4, and S5–S9, respectively. The learning rates were scaled down by 0.7 when the training cost increased on the validation set. The training was terminated when the learning rate was below 10^{-10} . The different learning rates for the systems trained with the ELC cost function and the systems trained with the EMSE cost functions were found from preliminary experiments. All models were implemented using CNTK [29] and the script files needed to reproduce the reported results can be found in [31].

4. EXPERIMENTAL RESULTS

We have evaluated the performance of the ten systems based on their average ELC and STOI scores computed on the test set. The STOI score is computed using the enhanced and reconstructed timedomain speech signal, whereas the ELC score is computed using short-time one-third octave band temporal envelope vectors.

4.1. Matched and Unmatched Noise Type Experiments

In Table 2 we compare the ELC scores for the noise type specific systems trained using the ELC (S0–S4), and EMSE (S5–S8) cost functions, and tested in matched noise-type conditions (SSN, BBL, CAF, and STR) at an input SNR of -5, 0, and 5 dB. Results covering the SNR range from -10 to 20 dB can be found in [31]. All models achieve large improvements in ELC with an average improvement of approximately 0.15-0.20, for all SNRs and noise types, compared to the ELC score of the noisy, unprocessed signals (denoted UP. in Tables 2 to 4). We also see that, as expected, models trained with the ELC cost function (S0–S4) in general achieve similar or slightly higher ELC scores compared to the models trained with EMSE (S5–S8). In Table 3 we report the STOI scores for the systems in Table 2 tested in identical conditions. We see moderate to large improvements in STOI in all conditions with an average improvement from 0.07–0.13. We also observe that the systems trained with the EMSE

Table 2. ELC results for S0 - S9 tested with SSN, BBL, CAF, and STR

-			SSN					BBL					CAF					STR		
SNR [dB]	UP.	S0 (ELC)	S5 (EMSE)	S4 (ELC)	S9 (EMSE)	UP.	S1 (ELC)	S6 (EMSE)	S4 (ELC)	S9 (EMSE)	UP.	S2 (ELC)	S7 (EMSE)	S4 (ELC)	S9 (EMSE)	UP.	S3 (ELC)	S8 (EMSE)	S4 (ELC)	S9 (EMSE)
-5 0 5	0.36 0.52 0.66	0.66 0.77 0.82	0.65 0.76 0.81	0.64 0.75 0.80	0.63 0.74 0.79	0.34 0.50 0.64	0.50 0.69 0.78	0.51 0.69 0.77	0.48 0.67 0.77	0.48 0.67 0.77	0.43 0.57 0.68	0.61 0.73 0.79	0.59 0.71 0.78	0.58 0.72 0.79	0.58 0.70 0.77	0.45 0.58 0.69	0.70 0.78 0.82	0.68 0.76 0.80	0.68 0.77 0.81	0.66 0.75 0.79
Avg.	0.51	0.75	0.74	0.73	0.72	0.49	0.66	0.66	0.64	0.64	0.56	0.71	0.69	0.70	0.68	0.57	0.77	0.75	0.75	0.73

1000000000000000000000000000000000000

			SSN					BBL					CAF					STR		
SNR	LIP	S0	S 5	S4	S9		S1	S6	S 4	S9		S2	S 7	S 4	S9	LIP	S 3	S 8	S4	S9
[dB]	01.	(ELC)	(EMSE)	(ELC)	(EMSE)		(ELC)	(EMSE)	(ELC)	(EMSE)		(ELC)	(EMSE)	(ELC)	(EMSE)	01.	(ELC)	(EMSE)	(ELC)	(EMSE)
-5	0.61	0.78	0.78	0.76	0.76	0.59	0.66	0.67	0.65	0.65	0.67	0.76	0.76	0.75	0.75	0.68	0.81	0.82	0.80	0.80
0	0.74	0.88	0.88	0.87	0.87	0.72	0.82	0.82	0.81	0.81	0.78	0.86	0.86	0.85	0.86	0.78	0.88	0.89	0.88	0.88
5	0.85	0.93	0.93	0.92	0.92	0.83	0.90	0.90	0.89	0.90	0.87	0.91	0.92	0.91	0.92	0.87	0.92	0.93	0.92	0.92
Avg.	0.73	0.86	0.86	0.85	0.85	0.71	0.79	0.80	0.78	0.79	0.77	0.84	0.85	0.84	0.84	0.78	0.87	0.88	0.87	0.87

 Table 4. ELC and STOI for S4 and S9 tested with BUS and PED.

			ELC				STOI		
		BUS		PED		BUS		PED	
SNR	UP.	S 4	S9 UP.	S 4	S9 UP.	S4	S9 UP.	S4	S 9
-5	0.56	0.71	0.68 0.35	0.55	0.53 0.77	0.84	0.84 0.60	0.71	0.71
0	0.66	0.79	0.76 0.50	0.70	0.68 0.85	0.90	0.90 0.72	0.83	0.83
5	0.74	0.83	0.81 0.64	0.78	0.76 0.91	0.94	0.94 0.83	0.90	0.90
Avg.	0.65	0.78	0.75 0.50	0.68	0.66 0.84	0.89	0.89 0.72	0.81	0.81

cost function achieve similar improvement in STOI as the systems trained with the ELC cost function. In Table 4, the ELC and STOI scores for the noise type general systems (S4 and S9) tested with the unmatched BUS and PED noise types are summarized. We see average improvement in the order of 0.1-0.18 in terms of ELC score and 0.05 - 0.09 in terms of STOI. We also see the performance gap between the S4 system (trained with ELC cost function) is small compared to the S9 system (trained with EMSE cost function) and that noise specific systems perform slightly better than the noise general one. The results in Tables 2 to 4 are interesting since they show roughly identical global behavior as measured by ELC and STOI for systems trained with the ELC cost functions.

4.2. Gain Similarities Between ELC and EMSE Based Systems

We now study to which extent ELC and EMSE based systems behave similarly on a more detailed level. Specifically, we compute correlation coefficients between the gain vectors produced by each of the two types of systems, for SSN, BBL, and STR noise types, and summarize them in Table 5. In Table 5 we observe that high sample correlations (> 0.90) are achieved for all noise types and both SNRs, which indicates that the gains produced by a system trained with the ELC cost function are quite similar to the gains produced by a system trained with the EMSE cost function, which supports the findings in Sec. 4.1. Similar conclusions can be drawn for the remaining noise types (results omitted due to space limitations, see [31]).

4.3. Approximate-STOI Optimal DNN vs. Classical SE DNN

As a final study we compare the performance of an approximate-STOI optimal DNN based SE system with classical Short-Time Spectral Amplitude (STSA) DNN based enhancement systems that estimate $\hat{g}(k,m)$ directly for each STFT frame (see e.g. [35, 36]). Similarly to S0–S9 these systems are three-layered feed-forward

Table 5 . Sample linear corre-lation between gain vectors.					f able 6 . ical DN	STOI	score and score	for clas 1 BBL.
SNR [dB]	SSN	BBL	STR		SNR [dB]	UP.	# u 512	nits 4096
-5 5	0.93 0.94	0.91 0.96	0.92 0.92		-5 5	0.59 0.83	0.64 0.91	0.66 0.92

DNNs and use 30 STFT frames as input, but differently from S0– S9, they minimize the MSE between STFT magnitude spectra, i.e. across frequency. The DNNs estimate five STFT frames per timestep and overlapping frames are averaged to construct the final gain. We have trained two of these classical systems, with 512 units and 4096 units, respectively, in each hidden layer, using the BBL noise corrupted training set. The results are presented in Table 6.

From Table 6 we see, for example, that such classical STSA-DNN based SE systems trained and tested with BBL noise achieve a maximum STOI score of 0.66 at an input SNR of -5 dB, which is equivalent to the STOI score of 0.66 achieved by S1 in Table 3. We also see that the classical system performs on par with S1 at an input SNR of 5 dB SNR with a STOI score of 0.92 compared to 0.90 achieved by S1. Although surprising, this is an interesting result since it indicates that no improvement in STOI can be gained by a DNN based SE system that is designed to maximize an approximate-STOI measure using short-time temporal one-third octave band envelope vectors. The important implication of this is that traditional STSA-DNN based SE systems may be close to optimal from an estimated speech intelligibility perspective.

5. CONCLUSION

In this paper we proposed a Speech Enhancement (SE) system based on Deep Neural Networks (DNNs) that optimizes an approximation of the Short-Time Objective Intelligibility (STOI) estimator. We proposed an approximate-STOI cost function and derived closed-form expressions for the required gradients. We showed that DNNs designed to maximize approximate-STOI, achieve large improvement in STOI when tested in matched and unmatched noise types at various SNRs. We also showed that approximate-STOI optimal systems do not outperform systems that minimize a mean square error cost. Finally, we showed that approximate-STOI DNN based SE systems perform on par with classical DNN based SE systems. Our findings suggest that a potential speech intelligibility gain of approximate-STOI optimal systems over MSE based systems is modest at best.

6. REFERENCES

- [1] R. C. Hendriks, T. Gerkmann, and J. Jensen, "DFT-Domain Based Single-Microphone Noise Reduction for Speech Enhancement: A Survey of the State of the Art," *Synth. Lect. on Speech and Audio Process.*, vol. 9, no. 1, pp. 1–80, Jan. 2013.
- [2] P. C. Loizou, *Speech Enhancement: Theory and Practice*. CRC Press, 2013.
- [3] Y. Ephraim and D. Malah, "Speech enhancement using a minimum-mean square error short-time spectral amplitude estimator," *IEEE Trans. Acoust., Speech, and Sig. Process.*, vol. 32, no. 6, pp. 1109–1121, 1984.
- [4] Y. Hu and P. C. Loizou, "A comparative intelligibility study of single-microphone noise reduction algorithms," J. Acoust. Soc. Am., vol. 122, no. 3, pp. 1777–1786, Sep. 2007.
- [5] H. Luts *et al.*, "Multicenter evaluation of signal enhancement algorithms for hearing aids," *J. Acoust. Soc. Am.*, vol. 127, no. 3, pp. 1491–1505, 2010.
- [6] J. Jensen and R. Hendriks, "Spectral Magnitude Minimum Mean-Square Error Estimation Using Binary and Continuous Gain Functions," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 20, no. 1, pp. 92–102, Jan. 2012.
- [7] E. M. Grais and H. Erdogan, "Single channel speech music separation using nonnegative matrix factorization and spectral masks," in *Proc. ICDSP*, 2011, pp. 1–6.
- [8] Y. Wang and D. Wang, "Towards Scaling Up Classification-Based Speech Separation," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 21, no. 7, pp. 1381–1390, Jul. 2013.
- [9] Y. Xu et al., "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, Jan. 2015.
- [10] E. W. Healy *et al.*, "An algorithm to increase speech intelligibility for hearing-impaired listeners in novel segments of the same noise type," *J. Acoust. Soc. Am.*, vol. 138, no. 3, pp. 1660–1669, 2015.
- [11] J. Chen *et al.*, "Large-scale training to increase speech intelligibility for hearing-impaired listeners in novel noises," *J. Acoust. Soc. Am.*, vol. 139, no. 5, pp. 2604–2612, 2016.
- [12] E. W. Healy *et al.*, "An algorithm to increase intelligibility for hearing-impaired listeners in the presence of a competing talker," *J. Acoust. Soc. Am.*, vol. 141, no. 6, pp. 4230–4239, 2017.
- [13] M. Kolbæk, Z. H. Tan, and J. Jensen, "Speech Intelligibility Potential of General and Specialized Deep Neural Network Based Speech Enhancement Systems," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 25, no. 1, pp. 153–167, 2017.
- [14] B. Moore, An Introduction to the Psychology of Hearing, 6th ed. Brill, 2013.
- [15] T. M. Elliott and F. E. Theunissen, "The Modulation Transfer Function for Speech Intelligibility," *PLOS Computational Biology*, vol. 5, no. 3, p. e1000302, Mar. 2009.
- [16] Y. Hu and P. C. Loizou, "A perceptually motivated approach for speech enhancement," *IEEE Trans. Speech, Audio, Process.*, vol. 11, no. 5, pp. 457–465, 2003.
- [17] Y. Ephraim and D. Malah, "Speech enhancement using a minimum mean-square error log-spectral amplitude estimator," *IEEE Trans. Acoust., Speech, and Sig. Process.*, vol. 33, no. 2, pp. 443–445, 1985.

- [18] N. Virag, "Single channel speech enhancement based on masking properties of the human auditory system," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 7, no. 2, pp. 126– 137, 1999.
- [19] P. C. Loizou, "Speech Enhancement Based on Perceptually Motivated Bayesian Estimators of the Magnitude Spectrum," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 13, no. 5, pp. 857–869, 2005.
- [20] L. Lightburn and M. Brookes, "SOBM a binary mask for noisy speech that optimises an objective intelligibility metric," in *Proc. ICASSP*, 2015, pp. 5078–5082.
- [21] W. Han *et al.*, "Perceptual weighting deep neural networks for single-channel speech enhancement," in *Proc. (WCICA*, 2016, pp. 446–450.
- [22] P. G. Shivakumar and P. Georgiou, "Perception Optimized Deep Denoising AutoEncoders for Speech Enhancement - Semantic Scholar," in *INTERSPEECH*, 2016, pp. 3743–3747.
- [23] Y. Koizumi *et al.*, "DNN-based source enhancement selfoptimized by reinforcement learning using sound quality measurements," in *Proc. ICASSP*, 2017, pp. 81–85.
- [24] C. H. Taal *et al.*, "An Algorithm for Intelligibility Prediction of Time-Frequency Weighted Noisy Speech," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [25] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [26] J. Jensen and C. H. Taal, "An Algorithm for Predicting the Intelligibility of Speech Masked by Modulated Noise Maskers," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 11, pp. 2009–2022, 2016.
- [27] A. H. Andersen *et al.*, "Predicting the Intelligibility of Noisy and Nonlinearly Processed Binaural Speech," *IEEE/ACM Trans. Audio, Speech, Lang. Process.*, vol. 24, no. 11, pp. 1908–1920, 2016.
- [28] C. H. Taal, R. C. Hendriks, and R. Heusdens, "Matching pursuit for channel selection in cochlear implants based on an intelligibility metric," in *Proc. EUSIPCO*, 2012, pp. 504–508.
- [29] A. Agarwal *et al.*, "An introduction to computational networks and the computational network toolkit," Microsoft Technical Report {MSR-TR}-2014-112, Tech. Rep., 2014.
- [30] J. Garofolo *et al.*, "CSR-I (WSJ0) Complete LDC93s6a," 1993, philadelphia: Linguistic Data Consortium.
- [31] M. Kolbæk, Z.-H. Tan, and J. Jensen, "Supplemental Material."[Online]. Available: http://kom.aau.dk/~mok/icassp2018
- [32] J. Barker *et al.*, "The third 'CHiME' Speech Separation and Recognition Challenge: Dataset, task and baselines," in *Proc. ASRU*, 2015.
- [33] J. S. Garofolo *et al.*, "DARPA TIMIT Acoustic Phonetic Continuous Speech Corpus CDROM," 1993.
- [34] ITU, "Rec. P.56 : Objective measurement of active speech level," 1993, https://www.itu.int/rec/T-REC-P.56/.
- [35] F. Weninger *et al.*, "Discriminatively trained recurrent neural networks for single-channel speech separation," in *GlobalSIP*, 2014, pp. 577–581.
- [36] M. Kolbæk, Z. H. Tan, and J. Jensen, "Speech Enhancement using Long Short-Term Memory based Recurrent Neural Networks for Noise Robust Speaker Verification," in *Proc. SLT*, 2016, pp. 305–311.