This is a repository copy of *Exploration of audio quality assessment and anomaly localisation using attention models*.

Version: Accepted Version

Proceedings Paper:

Exploration of Audio Quality Assessment and Anomaly Localisation Using Attention Models

Qiang Huang   and   Thomas Hain

Depart of Computer Science, University of Sheffield
Sheffield, UK
{q.huang, t.hain}@sheffield.ac.uk

Abstract

Many applications of speech technology require more and more audio data. Automatic assessment of the quality of the collected recordings is important to ensure they meet the requirements of the related applications. However, effective and high performing assessment remains a challenging task without a clean reference. In this paper, a novel model for audio quality assessment is proposed by jointly using bidirectional long short-term memory and an attention mechanism. The former is to mimic a human auditory perception ability to learn information from a recording, and the latter is to further discriminate interferences from desired signals by highlighting target related features. To evaluate our proposed approach, the TIMIT dataset is used and augmented by mixing with various natural sounds. In our experiments, two tasks are explored. The first task is to predict an utterance quality score, and the second is to identify where an anomalous distortion takes place in a recording. The obtained results show that the use of our proposed approach outperforms a strong baseline method and gains about 5% improvements after being measured by three metrics, Linear Correlation Coefficient and Spearman's Rank Correlation Coefficient, and F1.

Index Terms: quality assessment, attention model, anomaly localisation

1. Introduction

Speech quality assessment aims to find and quantize the differences between original speech signals and the ones with variations. There are two ways to assess speech quality: subjective and objective evaluation. The subjective evaluation is made by a listener's opinion in terms of some pre-defined criterion, e.g., the mean opinion score (MOS). The MOS is generally conducted by computing the arithmetic mean of all individual values on a predefined scale that a subject assigns to one's opinion of the performance of a system quality [1]. As subjective evaluation may be time-consuming and expensive due to the need of human assessors, objective quality evaluation has been widely used to predict the rating scores. For objective evaluation, perceptual evaluation of speech quality (PESQ) [2] can analyze an audio recording sample-by-sample after a temporal alignment of corresponding excerpts of reference and test signal. This means objective evaluation still requires a "golden" reference for each utterance to be evaluated, which considerably restricts the applicability of such assessment tools in real-world scenarios [3]. Accordingly, it is highly desirable to develop a reliable assessment model.

In recent years, due to the rapid development of deep neural networks, some related technologies have been used for speech/voice quality assessment. Spille et al., [4] used a deep neural network to predict speech intelligibility. Soni et al. [5] applied a sub-band autoencoder to first learn features to be used by the following neural-network-based prediction model. Fu et al. [3] developed a non-intrusive speech quality evaluation model to predict PESQ scores using a BLSTM model on audio recordings. Avila et al., [6] investigated the applicability of three neural network-based approaches for non-intrusive audio quality assessment based on mean opinion score(MOS) estimation.

In this paper, our task focuses on two aspects. The first is to assess the quality of audio recordings at utterance level, and the second is to locate when an anomalous distortion takes place in the recording. For this purpose, a frame-level based quality assessment architecture using BLSTM [7] is employed. The structure is designed to compute the frame-level score, and then infer an utterance-level score. Moreover, by calculating frame-level scores and detecting possible anomalous variations, anomaly regions can be thus located in a recording. To further increase the ability of quality assessment against interferences, an attention mechanism is employed. The use of attention aims to allow focuses to certain frames which are related to the target set by users. This means the target related information will probably be given a large weight, while a small weight will probably be allocated to irrelevant features. This is useful to discriminate interferences from desired signals, and thus help to assess the quality of a recording and anomaly localisation. The use of attention mechanisms has led to some state-of-the-art performances in different research fields, e.g., natural language processing [8, 9, 10], speech recognition [11, 12, 13], speaker recognition [14, 15, 16, 17], speech enhancement [18, 19]. However, to our know knowledge, more research in the use of attention on speech quality assessment is needed. The related details of the proposed architecture and how the attention mechanism is used in our work will be presented in following sections.

The rest of paper is organised as follows: Section 2 depicts the details of our proposed approach. In Section 3, the data used for model training and evaluation and experiment setup are introduced. The related experimental results and analysis are given in Section 4, and finally the conclusion is drawn in Section 5.

2. Proposed Architecture

Figure 1 shows the proposed architecture using deep neural networks and an attention model. Given the input spectrogram $S \in \Re^{F \times T}$ of an utterance U, the proposed model aims to compute the quality $Q_{f_t} \in \Re^{1 \times T}$ of each frame $f_t \in \Re^{F \times 1}$ of utterance U, and then infer an utterance-level score $Q_U \in \Re^1$. As shown in Figure 1, the proposed structure consists of three parts. In the first part, a one-dimensional CNN (1DCNN) layer
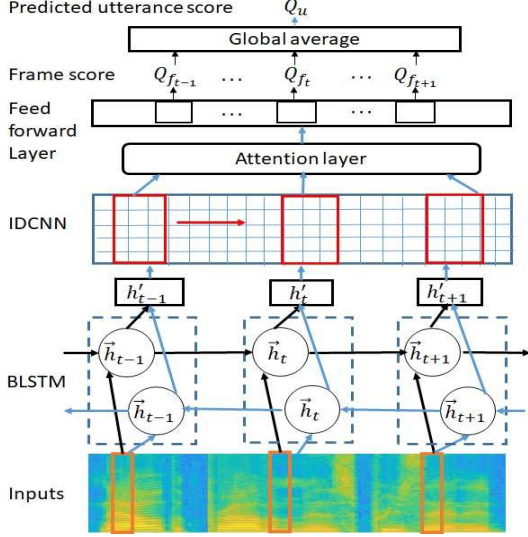
Figure 1: *Architecture of speech quality assessment and anomaly localisation using BLSTM, 1DCNN and attention model.*

[20] is cascaded with a BLSTM layer to learns features using the contextual information in the time and frequency domains. An attention layer is used in the second part. The third part computes the frame-level value Q_{f_t} using a fully connected layer and finally infers the utterance-level score Q_U by averaging over all frames of utterance U.

2.1. Information Acquisition

The spectrogram $S = (f_1 \cdots f_T)$ of an input utterance U is processed by a BLTM layer and then by a one-dimensional CNN (1DCNN) layer.

$$h^{blstm} = \text{BLSTM}(f_{1..T}) \tag{1}$$

$$h^{1dcnn} = \text{1DCNN}(h^{blstm}) \tag{2}$$

The BLSTM is an improvement over LSTM in that it captures both the previous timesteps (past features) and the future time steps (future features) via forward and backward states, respectively. It can be implemented by:

$$
\begin{aligned}
\overrightarrow{h}_t &= \text{LSTM}(f_t, \overrightarrow{h}_{t-1}) \\
\overleftarrow{h}_t &= \text{LSTM}(f_t, \overleftarrow{h}_{t+1}) \\
h_t^{blstm} &= \overrightarrow{h}_t + \overleftarrow{h}_t
\end{aligned}
\tag{3}
$$

where the output of BLSTM layer $h^{blstm} \in \Re^{L \times T}$ is formed by concatenating the forward hidden state vector $\overrightarrow{h}$ and the backward hidden state vector $\overleftarrow{h}$, and L is the vector dimension.

The 1DCNN as described in [21] makes use of:

$$h_n^{cnn} = K_n \odot h^{blstm} \tag{4}$$

where $\odot$ denotes the convolution between the nth kernel $K_n \in \Re^{F \times 3}$ ($n \in [1..N]$) and the output of BLSTM h^{blstm}. $h^{cnn} \in \Re^{N \times T}$ represents the output of 1DCNN layer. The use of 1DCNN in this proposed architecture instead of 2DCNN is mainly because the 1DCNN has two advantages relating to feature extraction and computation efficiency [22]. These advan-

tages make it relatively easy to train and offer the small computational complexity while achieving state-of-the-art performance levels [23].

The use of the BLSTM is to mimic the human auditory perception system, as a decision made by a human generally needs to consider the possible effects caused by contextual information, especially our final aim is to compute an utterance-level score. Although the use of context information might be helpful to the quality estimation of the current frame, it may bring some negative impacts caused by the future or past frames if they are are corrupted by noise. To more accurately predict frame quality and locate an anomaly, the use of an attention model might be an effective way.

2.2. Attention Model

The hidden state h^{blstm} of the BLSTM, computed by equation 3, is used as the input of an attention layer. An attention matrix A is formed by computing the similarity between the hidden state h_t and $h_{t'}$ corresponding to frame f_t and $f_{t'}$ at timesteps t and t', respectively. The attention mechanism is implemented as follows:

$$
\begin{aligned}
h_{t,t'} &= \tanh(h_t^T W_t + h_{t'}^T W_{t'} + b_t) \\
e_{t,t'} &= \sigma(W_a h_{t,t'} + b_a) \\
\mathbf{a}_t &= \text{softmax}(\mathbf{e}_t) \\
l_t &= \sum_{t'} a_{t,t'} \cdot h_{t'}
\end{aligned}
\tag{5}
$$

where σ is the element-wise sigmoid function, W_t and $W_{t'}$ are the weight matrices corresponding to the hidden states h_t and $h_{t'}$; W_a is the weight matrix corresponding to their non-linear combination; b_t and b_a are the bias vectors. The attention-focused hidden state representation l_t of a frame at timestep t is given by the summation of the product of $h_{t'}$ of all other frames at timesteps t' and their similarity $a_{t,t'}$ to the hidden state representation h_t of the current frame.

2.3. Loss function

The score loss L is defined by the summation of an utterance-level loss ($L_\mathbf{u}$) and the loss ($L_\mathbf{f}$) averaged over all frames [3]:

$$
\begin{aligned}
L &= L_\mathbf{u} + L_\mathbf{f} \\
L_\mathbf{u} &= MSE(Q_u, Q_u') \\
L_\mathbf{f} &= \frac{1}{T} \sum_t (Q_u - Q_{f_t}')^2
\end{aligned}
\tag{6}
$$

where Q_u is the target score of utterance U and Q_u' is its predicted value. Q_{f_t}' represents the predicted quality score of the tth frame.

3. Experiment Setup

3.1. Data

In our experiments, the TIMIT dataset [24] was used as a comparison with the methods developed in [3]. About 700 utterances in its training set were used to train the proposed model, and 143 utterances randomly selected from its test set were used for evaluation.

Noise corrupted recordings are generated by mixing the clean recordings with various natural sounds at five signal-noise ratio (SNR) levels (-10dB, -5dB, 5dB, 10dB, 20dB). The noise

signals used were from the the general noise portion of the MU-SAN dataset [25], which contains six hours of various natural sounds, ranging from fax machine, car idling, thunder, wind, footsteps, paper rustling, rain, and birdsong, etc.

In all experiments, spectrograms are used as input features. All of the audio streams are segmented using a 32-ms sliding window with a 16-ms shift. A 512-point FFT was then used to convert each segment into a 257-dimension vector.

3.2. Pseudo score

SNR (dB)	Pseudo Score
-10	1
-5	2
5	4
10	5
20	7
original clean	8

Table 1: *Pseudo scores defined in terms of SNR.*

In Table 1, a set of scores are defined by linking to SNR values. This is to mimic the definition of scores used for MOS, but does not require human assessors' to mark each recording as the SNR value of a recording can be set when precisely mixing the original recording with noise signals. In addition, using a set of scores as assessment target might be able to mitigate the impact caused by the use of noise corrupted target values e.g., using the estimated PESQ values as targets in [3].

The pseudo scores ($\{1, 2, 4, 5, 7\}$) are allocated to the noise-corrupted utterances, whose SNR range from -10dB to 20dB with a 5dB shift. In this experiment, the score of original clean speech is set to "8".

3.3. Structure Configuration

Table 2 shows the configuration of the proposed approach, consisting of seven layers. In the first three layers, the dimension of input frame vector is 257, the output size of BLSTM is 200, and 250 kernels (size=3) used in 1DCNN. The $Frame\_score$ layer computes the frame-based prediction scores using a time-distributed Dense layer, and the $Utterance\_score$ layer outputs the utterance-level prediction using a GlobalAverage layer [20].

3.4. Implementation

Relying on the designed structure, experiments were conducted using two proposed approaches and one baselines. The first proposed approach LC_ATT uses the structure as presented in table 2, and in the second proposed approach, L_ATT, the same structure as LC_ATT is employed without the 1DCNN layer. The method developed in [3] is used as a baseline, which did not use 1DCNN and the attention mechanism in comparison with our proposed approaches. In experiments, RMSprop [26] was used as an optimiser and the initial learning rate was set to 0.001 with 0.95 decay every epoch.

As both utterance-level and frame-level qualities are estimated from different layers, as shown in Figure 1, regression instead of classification was used in our implementations. This is also to compare with the baseline method (Baseline1) [3], which used the same way to compute the utterance-level score. In addition, the use of regression also enables the proposed model to evaluate a recording, whose SNR is not listed in table 1, such as 15dB. To compute F1, a threshold is set (threshold=7.1) in terms of the results obtained on the training data.

Layer name	Output shape	#Param.
InputLayer	(None, None, 257)	0
BLSTM	(None, None, 200)	286,400
1DCNN	(None, None, 250)	150,250
ATT	(None, None, 250)	16,065
Dense	(None, None, 50)	12,550
Frame_score	(None, None, 1)	51
Utterance_score	(None,1)	0

Table 2: *Configuration of the proposed network structure.*

3.5. Evaluation Metrics

The three metrics used to assess performance are Linear Correlation Coefficient (LCC) [27], Spearman's Rank Correlation Coefficient (SRCC) [28], and F1 [29]. The first two metrics are used to measure the strength of the linear relationship between two variables. The use of F1 is to measure the accuracy of distinguishing the clean utterances from noise-corrupted ones.

4. Results

Figure 2 shows the predicted utterance-level quality scores obtained using the baseline method (figure 2(a)) and our two approaches (figure 2(b) and 2(c)) in the condition of different distortions. The x-axis in each figure denotes the test utterance index and y-axis represents the predicted utterance-level scores. The three figures (figure 2(a)-2(c)) show that the predicted scores obtained using the proposed approaches is closer to the target scores than Baseline1. Moreover, the corresponding statistics are also displayed in 3(a)-3(c). The error bars shown in the three figures represent the mean values and the range of variance obtained using the proposed approaches and the baseline method in different conditions. It can be found that the more signals are corrupted by noise, the higher variances are generated. This means it is hard to identify the quality of the audio signals in poor conditions. In comparison with the baseline method, the use of our approach can clearly reduce the predicted score variance and the deviation between the target scores and the predicted scores.

Table 3 lists LCC, SRCC, and F1 values obtained on the test data using the baseline method and our proposed approaches. The results show that the use of our approaches can yield better performance than the baseline method.

Method	LCC	SRCC	Precision	Recall	F1
Baseline1 [3]	0.876	0.876	0.728	0.777	0.752
L_1DCNN	0.858	0.863	0.926	0.691	0.792
L_Att	0.892	0.894	0.957	0.709	0.815
LC_Att	**0.919**	**0.914**	0.927	0.781	**0.848**

Table 3: *Metric values of LCC, SRCC, and F1 (larger is better) obtained on the test data corrupted by noise using the baseline method and two proposed approaches.*

Since various natural sounds from MUSAN [25] were used as noise to mix with clean signals, these noise signals might take place at different time and have various effects in frequency domain. To mitigate possible bias on LCC, SRCC, and F1 values, the experiments were repeated for eight times and the average values are used as the final results. To further compare the structure using an attention model and not using, the third approach L_1DCNN was also conducted. It has the same structure as L_ATT, but without the attention layer. The LCC and SRCC
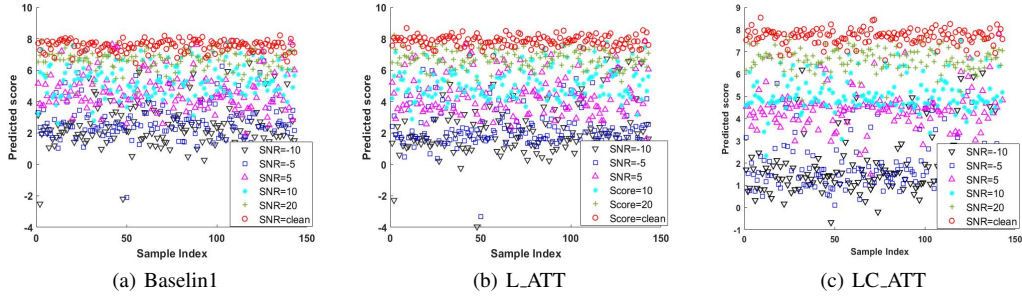
(a) Baselin1 (b) L_ATT (c) LC_ATT

Figure 2: *Predicted scores obtained on the test data using Baseline_1, L_ATT and LC_ATT in noise conditions.*
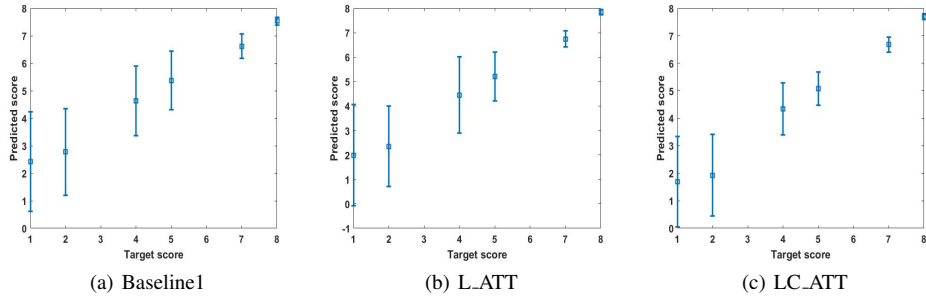

(a) Baseline1 (b) L_ATT (c) LC_ATT

Figure 3: *Mean and variance of predicted scores obtained on the test data using Baseline1, L_ATT and LC_ATT in noise conditions.*

values listed in Table 3 show that LC_ATT can yield consistent advantages over the baseline method, L_1DCNN, and L_ATT.

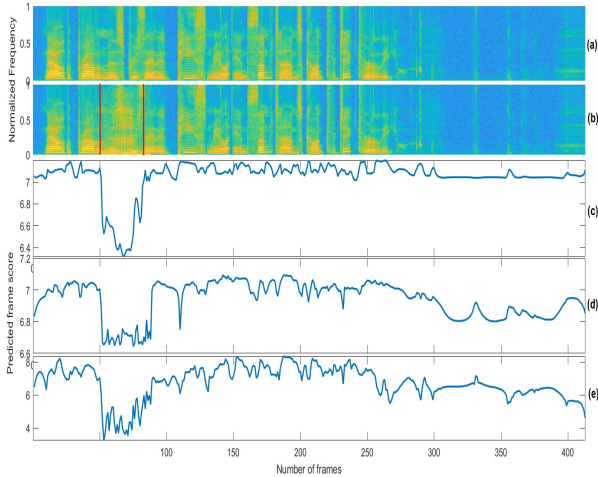Figure 4: *Clean spectrograms (a), partially noise-corrupted spectrum(b), and the predicted frame-level scores obtained using LC_ATT (c), L_ATT (d), and Baseline1 (e).*

In addition to predicting scores at utterance level, identifying when an anomaly occurs in an audio recording is also explored. We demonstrate how an anomalous distortion can be located. In Figure 4(a)-(e), the spectrogram of a recording and its frame-level prediction scores are shown. From top to bottom, Figure 4(a) is the spectrogram of a clean utterance. The noise-corrupted spectrogram is shown in Figure 4(b), where an anomalous distortion (SNR=15dB) takes place from the $37th$ frame to the $87th$ frame, within the range of two solid red lines.

The next three figures, Figure 4(c)-(e), indicate the frame-level prediction sores obtained using LC_ATT, L_ATT, and Baseline1, respectively.

It is clear that all of the three methods can find where the distortion is. However, the use of Baseline1 generates a high score variation within the range where the signal frames are corrupted by noise and outside. Compared to Baseline1, the use of LC_ATT keeps a relative smooth over all audio frames, by which the distortion range is able to precisely located. This case is probably related to the use of 1DCNN and the attention mechanism. The use of 1DCNN might mitigate the possible sudden variations by taking into account the context information. The use of attention mechanism might be able to enlarge the difference between clean signals and anomaly signals by highlighting the target related frame features.

5. Conclusion and future work

A novel structure for audio quality assessment was designed by using the BLSTM, one-dimensional convolutional neural networks and an attention mechanism. It can assess the quality of audio recordings at utterance level and identify the location of a distortion by computing frame-level scores. The obtained results, measured using three metrics, LCC, SRCC, and F1, have shown that the use of attention model can yield better performances than a strong baseline method whether for utterance-level score prediction or for anomaly distortion localisation.

In future, work in three aspects will be taken into account. Firstly, some advanced neural network technologies, such as multi-head attention model, will be used to assess audio quality. Secondly, the assessment technologies will be evaluated on large-sized speech datasets and in various acoustic conditions. Thirdly, the efficiency of assessment technologies will be also evaluated to make it work in some practical applications.

6. References

[1] S. W. Streijl, Robert C. and D. S. Hands., "Mean opinion score (mos) revisited: methods and applications, limitations and alternatives," *Multimedia Systems*, pp. 213–227, 2016.

[2] ITU. (accessed on Feb. 02, 2020) Itu-t recommendation p.862: Perceptual evaluation of speech quality (pesq). [Online]. Available: https://www.itu.int/rec/T-REC-P.862/en

[3] S.-W. Fu, Y. Tsao, H.-T. Hwang, and H.-M. Wang, "Quality-net: An end-to-end non-intrusive speech quality assessment model based on blstm," in *Interspeech*, 2018, pp. 3168–3172,.

[4] C. Spille, S. Ewert, B. Kollmeier, and B. Meyer, "Predicting speech intelligibility with deep neural networks," *Computer Speech and Language*, vol. 48, pp. 51–66, October 2017.

[5] M. H. Soni and H. A. Patil, "Novel deep autoencoder features for non-intrusive speech quality assessment," in *The 24th European Signal Processing Conference (EUSIPCO)*, 2016, pp. 2315–2319.

[6] N. Shah, H. Patil, and M. Soni, "Time-frequency mask-based speech enhancement using convolutional generative adversarial network," in *Conference: Asia-Pacific Signal and Information Processing Association Annual Summit and Conference*, August 2018.

[7] M. Schuster and K. K. Paliwal., "Bidirectional recurrent neural networks," vol. 45, pp. 2673–2681, 1997.

[8] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *ICLR*, vol. abs/1409.0473, 2014.

[9] T. L. Denny Britz, Anna Goldie and Q. Le, "Massive exploration of neural machine translation architectures," in *International Conference on Empirical Methods in Natural Language Processing*, 2017, pp. 1442–1451.

[10] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *International Conference on Neural Information Processing Systems*, 2017, pp. 6000–6010.

[11] J. K. Chorowski, D. Bahdanau, D. Serdyuk, K. Cho, and Y. Bengio, "Attention-based models for speech recognition," in *Advances in Neural Information Processing Systems*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds., 2015, pp. 577–585.

[12] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-endspeech recognition," *IEEE JOURNAL OF SELECTED TOPICS IN SIGNAL PROCESSING*, pp. 1240–1253, December 2017.

[13] Z. Tian, J. Yi, J. Tao, Y. Bai, and Z. Wen, "Self-attention transducers for end-to-end speech recognition," in *International Conference on InterSpeech*, September 2019, pp. 4395–4399.

[14] T. Bian, F. Chen, and L. Xu, "Self-attention based speaker recognition using cluster-range loss," *Neural Computing*, pp. 59–68, November 2019.

[15] Y. Zhu, T. Ko, D. Snyder, B. Mak1, and D. Povey, "Self-attentive speaker embeddings for text-independent speaker verification," in *International Conference on InterSpeech*, September 2019, pp. 3573–3577.

[16] M. India, P. Safari, and J. Hernando, "Self multi-head attention for speaker recognition," in *International Conference on InterSpeech*, September 2019, pp. 4305–4309.

[17] Y. Shi, Q. Huang, and T. Hain, "H-vectors: Utterance-level speaker embedding using a hierarchical attention model," in *International Conference on InterSpeech*, May 2020.

[18] ——, "Robust speaker recognition using speech enhancement and attention model," in *The Speaker and Language Recognition Workshop*, November 2020.

[19] 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*, September 2016, pp. 146–152.

[20] F. Chollet, "keras," https://github.com/fchollet/keras, 2015.

[21] S. Kiranyaz, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman. (2019, May) 1d convolutionalneural networks andapplications–a survey. [Online]. Available: https://arxiv.org/ftp/arxiv/papers/1905/1905.03554.pdf

[22] misslink.ai. (accessed on Feb. 02, 2020) Keras conv1d: Working with 1d convolutional neural networks in keras. [Online]. Available: https://missinglink.ai/guides/keras/keras-conv1d-working-1d-convolutional-neural-networks-keras/

[23] S. Kiranyaz1, O. Avci, O. Abdeljaber, T. Ince, M. Gabbouj, and D. J. Inman. (2019) 1d convolutionalneural networks andapplications–a survey. [Online]. Available: https://arxiv.org/ftp/arxiv/papers/1905/1905.03554.pdf

[24] J. Garofolo, L. Lamel, W. Fisher, J. Fiscus, D. Pallett, N. Dahlgren, and V. Zue, "Timit acoustic-phonetic continuous speech corpus," *Linguistic Data Consortium*, November 1992.

[25] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," October 2015.

[26] G. Hinton, "Neural networks for machinelearning, lecture 6a, overview of mini-batch gradient descent," accessed Feb. 02, 2020. [Online]. Available: https://www.cs.toronto.edu/tijmen/csc321/slides/lecture-slides-lec6.pdf

[27] C. F. Dietrich, *Uncertainty, Calibration and Probability: The Statistics of Scientific and Industrial Measurement*. A. Higher, 1991.

[28] W. W. Daniel, "Spearman rank correlation coefficient," pp. 358–365, 1990.

[29] C. J. Van Rijsbergen, *Information Retrieval*. Butterworth-Heinemann, 1979.