

Deep MOS Predictor for Synthetic Speech Using Cluster-Based Modeling

Yeunju Choi, Youngmoon Jung, Hoirin Kim

School of Electrical Engineering, KAIST, Daejeon, Republic of Korea

{wkadldppy, dudans, hoirkim}@kaist.ac.kr

Abstract

While deep learning has made impressive progress in speech synthesis and voice conversion, the assessment of the synthesized speech is still carried out by human participants. Several recent papers have proposed deep-learning-based assessment models and shown the potential to automate the speech quality assessment. To improve the previously proposed assessment model, MOSNet, we propose three models using cluster-based modeling methods: using a global quality token (GQT) layer, using an Encoding Layer, and using both of them. We perform experiments using the evaluation results of the Voice Conversion Challenge 2018 to predict the mean opinion score of synthesized speech and similarity score between synthesized speech and reference speech. The results show that the GQT layer helps to predict human assessment better by automatically learning the useful quality tokens for the task and that the Encoding Layer helps to utilize frame-level scores more precisely.

Index Terms: speech synthesis, speech quality assessment, cluster-based modeling, Encoding Layer, global quality token

1. Introduction

Recent advances in deep learning have led to significant growth in various speech processing fields [1–4]. However, in contrast to speech recognition task, there is no “right answer” in speech generation tasks such as text-to-speech (TTS) or voice conversion (VC). For this reason, subjective measures such as the mean opinion score (MOS) and similarity score have been used to evaluate naturalness and similarity, respectively [5]. That is, the quality measurement of the synthesized speech is still carried out by many human subjects, which is expensive and time-consuming [6]. Moreover, the results may change depending on several factors, such as human subjects and audio hardware.

There are many objective measures of speech quality to reflect human perception [7–10]. The most widely used measures are Mel-cepstral distance (MCD) [7] and the perceptual evaluation of speech quality (PESQ) [8]. However, these are full-reference measures in that they need ground-truth speech as reference. There are also no-reference measures such as ANIQUE [9] or ITU-T recommendation P. 563 [10]. However, most of these measures are targeted at detecting artifacts caused by lossy compression and transmission in telephony, not at evaluating the quality of synthetic speech.

With the advances in deep learning, researchers have recently proposed deep-learning-based models that can evaluate the quality of synthetic speech without reference speech [11–14]. Patton *et al.* [12] proposed AutoMOS, based on long short-term memory (LSTM), to predict MOS values. Fu *et al.* [13] proposed Quality-Net based on bidirectional LSTM (BLSTM) to predict the frame-level PESQs. Recently, Lo *et al.* [14] proposed MOSNet that generates frame-level MOSs from the features of convolutional neural network-BLSTM (CNN-BLSTM) and predicts the utterance-level MOS using the frame-

level scores. Moreover, they modified MOSNet to evaluate the similarity score for VC and extended the deep-learning-based quality assessment area to similarity score prediction.

These researches have shown the potential to automate the assessment of the synthesized speech using deep neural networks. However, it is difficult to understand the assessment criteria for speech quality evaluation performed by humans. Everyone has a different perspective on speech quality, and even the same person has a different perspective each time. Therefore, we propose a model using a global quality token (GQT) layer that can automatically learn the criteria as soft clusters.

Furthermore, even though Quality-Net [13] and MOSNet [14] showed performance improvement by assuming that the quality score of an utterance is an average of frame-level scores, the exact relation between the utterance- and frame-level scores remains poorly understood. With the inspiration that humans will determine an utterance-level score in a more sophisticated way, we propose a model using an Encoding Layer to aggregate frame-level scores by considering not only the simple average of the frame-level scores, but also their distribution information.

2. Relation to prior work

Lo *et al.* [14] proposed and compared three different model architectures for MOS prediction. All of them consist of a feature extractor, two fully-connected (FC) layers, and a global average pooling (GAP) layer. The main difference between the three models is the architecture of the feature extractor; CNN, BLSTM, and CNN-BLSTM. In this paper, we use the CNN-BLSTM-based MOSNet as our baseline, which showed the best performance among them. The architecture of MOSNet is shown in Table 1. First, the feature extractor generates frame-level feature vectors from an input magnitude spectrogram. Then, the following FC layers map each frame-level feature vector into a frame-level score. Finally, the GAP layer outputs the utterance-level score by averaging the frame-level scores.

They formulate a loss function using both utterance-level mean squared error (MSE) and frame-level MSE as follows:

$$L = \frac{1}{S} \sum_{s=1}^S [(\hat{Q}_s - Q_s)^2 + \frac{\alpha}{T_s} \sum_{t=1}^{T_s} (\hat{Q}_s - q_{s,t})^2], \quad (1)$$

where S is the number of training utterances, and T_s is the number of the frames of the s -th utterance. \hat{Q}_s and Q_s are the ground-truth and predicted value for the utterance-level score of the s -th utterance, respectively. $q_{s,t}$ is the predicted frame-level score at time t for the s -th utterance, and α is a weighting factor for the frame-level MSE.

Furthermore, they extended CNN-based MOSNet to predict the similarity score between a pair of utterances. Two utterances of an input pair have the same length by zero-padding and share the convolution layers among the modified CNN-based MOSNet. Two CNN feature maps from the input pair are concatenated and fed into the following FC layers and a GAP layer

Table 1: Configuration of the model architectures. + GQT and + EL denote using the GQT layer and the Encoding Layer, respectively. The convolutional layer parameters are denoted as $\text{conv}\{\text{receptive field size}\}\text{-}\{\text{number of channels}\}\{\text{stride}\}$. N is the number of frames. SC is skip connection. GAP is global average pooling. K is the number of the codewords for the Encoding Layer. For + EL and + GQT + EL, an additional FC layer of a pooling layer is omitted for simplicity.

model	MOSNet	+ GQT	+ EL	+ GQT + EL
input	$N \times 257$ magnitude spectrogram			
conv. layers	$\left\{ \begin{array}{l} \text{conv}3 - (\text{channels})/1 \\ \text{conv}3 - (\text{channels})/1 \\ \text{conv}3 - (\text{channels})/3 \end{array} \right\} \times 4$ $\text{channels} = [16, 16, 32, 32]$			
GQT layer with SC	-	# GQTs = 10	-	# GQTs = 10
recurrent layer	BLSTM-128			
FC layers	FC-128, ReLU, dropout <hr/> FC-1 (<i>frame-level scores</i>)			
pooling layer	GAP layer (<i>utterance-level score</i>)		Encoding Layer ($K=10$) & GAP layer (<i>utterance-level score</i>)	

to generate a similarity score. In this work, we propose SIMNet by modifying CNN-BLSTM-based MOSNet in a similar way, and used this as a baseline model for similarity score prediction. SIMNet concatenates two different feature maps of the shared CNN from the utterance pair and uses the result as an input of the following BLSTM layer.

Two main contributions of our work are the following. First, motivated by a global style token (GST) [15], we propose a model using a global quality token (GQT) layer which learns the tokens that reflect the criteria for speech quality evaluation. The GQTs operate in the same way as the GSTs, but we call them global quality tokens because they are learned in terms of speech quality. Second, we propose a model using an Encoding Layer which aggregates the frame-level scores by considering the distribution of them. Although Quality-Net and MOSNet showed that using frame-level scores improves MOS prediction performance, they only considered average information of the frame-level scores. We consider not only average information but also distribution information of frame-level scores.

3. Proposed models

To improve MOSNet, we propose two models based on the GQT layer and the Encoding Layer, respectively. Throughout this paper, ‘+ GQT’ and ‘+ EL’ stand for using the GQT layer and using the Encoding Layer, respectively. Our third model is a combination of the first and second models. For MOSNet, we use 16, 16, 32, and 32 filters for each convolutional block containing three convolutional layers. The architectures of the four models are described in Table 1.

3.1. Global quality tokens for MOS and similarity score

A GST model is proposed in [15] for style-expressive end-to-end TTS. It consists of a reference encoder, style token layer, and sequence-to-sequence model. The reference encoder ex-

tracts reference embedding of the reference utterance to refer the style. The GSTs of the style token layer, which are shared by all training sequences, become the soft clusters of reference embedding through training. The goal of the style token layer is to calculate the style embedding, which represents the style of the reference utterance, as a weighted sum of the GSTs. The weights assigned to the GSTs are learned by a multi-head attention module [16]. The style token layer consisting of GSTs and the attention module is randomly initialized and then trained jointly with the whole GST model. Therefore, the GSTs can become useful soft clusters for style modeling, and the GST model can synthesize speech with a specific style of the reference utterance by extracting the style embedding from it.

Although Wang *et al.* [15] suggested the GST layer for TTS, they also showed that GSTs could be used as speaker classification features. There is also a report [17] that the GST layer helps to improve speech recognition performance. In this paper, we report for the first time that the GST layer also helps to improve the performance of the quality assessment task. Here, we refer to the GST as the GQT, as we mentioned in the previous section. Besides, the reference encoder, GQTs, and multi-head attention module are collectively called the GQT layer.

As the GSTs are targeted for a separate reference utterance, they need a separate reference encoder. Unlike GSTs, the GQTs are targeted for the same input utterance of MOS prediction. Therefore, we design our reference encoder to share the convolutional layers with MOSNet. A gated recurrent unit (GRU) [18] layer follows the convolutional layers and the last hidden state of the GRU layer serves as the reference embedding. Then, the quality embedding is calculated in the same way to calculate the style embedding in [15]. As a kind of skip connection, the quality embedding is added to all the frame-level feature vectors of the CNN. We use the resulting representation as an input of the BLSTM layer.

When we use a GQT layer for a similarity score prediction task, we apply the shared GQT layer to generate two quality embeddings from an utterance pair. Then we add each quality embedding to the corresponding CNN feature map before the concatenation of two CNN feature maps.

3.2. An Encoding Layer for MOS and similarity score

The Encoding Layer [19] was proposed for texture recognition by learning the inherent visual codewords directly from the loss function. The codewords are learned from the distribution of the CNN features. The Encoding Layer also acts as a pooling layer, which converts feature vectors of any size into a fixed-length representation. Given N feature vectors, $X = \{x_1, \dots, x_N\}$, and K codewords, $C = \{c_1, \dots, c_K\}$, the output representation of the Encoding Layer is $e = \{e_1, \dots, e_K\}$, called residual encoding vector. The residual vector r_{ik} is calculated by $r_{ik} = x_i - c_k$ and the assigning weight for r_{ik} is given by

$$w_{ik} = \frac{\exp(-s_k \|r_{ik}\|^2)}{\sum_{j=1}^K \exp(-s_j \|r_{ij}\|^2)}, \quad (2)$$

where s_k is a learnable smoothing factor for c_k . Then the residual encoding for the k -th codeword c_k is calculated as follows:

$$e_k = \sum_{i=1}^N e_{ik} = \sum_{i=1}^N w_{ik} r_{ik}. \quad (3)$$

In the speaker recognition task, Cai *et al.* [20] and Jung *et al.* [21] used the Encoding Layer for aggregating the frame-level

speaker features to generate speaker embedding and improved the performance. Motivated by this, we apply the Encoding Layer in the speech quality assessment model. However, we use the Encoding Layer on frame-level scores, not frame-level feature vectors. Moreover, we utilize both GAP layer and Encoding Layer together to combine information from both layers. In ground terrain recognition task, Xue *et al.* [22] showed it is better to use both the Encoding Layer and the GAP layer.

Specifically, we put an Encoding Layer parallel to the GAP layer. The outputs from the Encoding Layer and the GAP layer are the residual encoding vector and the average score, respectively. They are concatenated and used as an input of the following FC layer to predict the utterance-level score.

When we use an Encoding Layer for similarity score prediction, the Encoding Layer aggregates the frame-level similarity scores as in MOS prediction.

4. Experiments

4.1. Dataset

As in [14], we use the MOS and similarity evaluation results from the Voice Conversion Challenge (VCC) 2018 [23]. The challenge comprised two tasks: Hub task (parallel VC) and Spoke Task (non-parallel VC). The VCC 2018 dataset is based on the device and production speech (DAPS) dataset [24], which includes recordings of professional US English speakers. There are a total of eight source speakers and four target speakers. A total of 23 teams submitted systems to the Hub task, with 11 of them additionally participating in the Spoke task. There are a total of 38 evaluated systems, including the source speaker, the target speaker, and two baseline systems for the two tasks.

For MOS evaluation, 267 people rated the naturalness of 20,580 submitted utterances with a score ranging from 1 (“Completely unnatural”) to 5 (“Completely natural”). The corresponding number of evaluation results is 82,304, and the ground-truth MOS of each utterance was obtained by averaging all the MOS ratings of the utterance. Among the 20,580 <audio, ground-truth MOS> pairs, we use 15,580, 3,000, and 2,000 pairs for training, validation, and testing, respectively. The MOS of each system is obtained by averaging all the MOS values of the utterances from the system.

The same 267 people also rated the similarity between two utterances with a score among 1 (“Same, absolutely sure”), 2 (“Same, not sure”), 3 (“Different, not sure”), and 4 (“Different, absolutely sure”). An utterance pair consists of an anchor utterance, which can be either converted speech or human speech, and a reference utterance, which is an utterance from either the source or the target speaker of the anchor speech with the same linguistic context. There are a total of 30,864 evaluation results, and the ground-truth similarity of each utterance pair is obtained from the average of the scores received. Among a total of 21,608 <audio pair, ground-truth similarity> pairs, we use 17,286 for training, 2,161 for validation, and 2,161 for testing. We consider a pair <anchor system, reference system> as one system pair, then the corresponding number of system pairs for similarity score prediction is 76.

We also use the MOS evaluation results of the VCC 2016 [25] to test the generalization ability of the models trained on the VCC 2018 training set. The VCC 2016 comprised only a parallel voice conversion task, and there are a total of 26,028 utterances from 20 systems. Each system has 1,600 utterance-level evaluation results without any description of the utterances. Therefore, we can report only system-level performance

Table 2: Results of different models. + GQT, + EL, + GQT + EL stand for using the GQT layer, the Encoding layer, and both of them, respectively. The best results are highlighted in bold.

Model	VCC 2018						VCC 2016		
	utterance-level			system-level			system-level		
	MSE	LCC	SRCC	MSE	LCC	SRCC	MSE	LCC	SRCC
MOSNet	0.448	0.651	0.619	0.039	0.966	0.924	0.316	0.896	0.858
+ GQT	0.447	0.654	0.621	0.041	0.968	0.931	0.242	0.921	0.853
+ EL	0.444	0.656	0.617	0.031	0.974	0.938	0.242	0.908	0.855
+ GQT + EL	0.447	0.656	0.616	0.032	0.967	0.940	0.246	0.885	0.839

on the evaluation results of the VCC 2016.

4.2. Implementation details

We implement all the models using PyTorch and train them on a single NVIDIA GTX 1080 Ti GPU with four different random seeds. We set α , the weighting factor for the frame-level MSE, to 0.8. When we use the GQT layer, we use 10 GQTs and 8 heads for a multi-head attention module. When we use the Encoding Layer, we use 10 codewords. We use a batch size of 16 for MOSNet + GQT + EL and 32 for the rest of the models. We use the Adam optimizer with a learning rate of 0.0001 and set a dropout rate to 0.3. We use the validation set to select the model with the lowest MSE during 200 epochs. We report the average MSE, linear correlation coefficient (LCC) [26], and Spearman’s rank correlation coefficient (SRCC) [27] of the models trained with four random seeds.

4.3. Experiments on MOS prediction

First, we discuss the results on MOS prediction with the VCC 2018 test set, shown in Table 2. MOSNet + GQT made improvements with all the metrics except the system-level MSE. To directly interpret the role of each GQT, we should adjust the weights of GQTs and observe the change in the predicted MOSs for various input utterances. However, this process is practically impossible since it requires listening to a lot of utterances and analyzing the factors that affect the MOS, which are expensive and subjective. Instead, from the MOS prediction results, we can infer that the GQTs become useful soft clusters for MOS evaluation.

Using the Encoding Layer improved all the metrics except the utterance-level SRCC as the model aggregated the frame-level scores using their distribution and learned the embeddings that are useful for the aggregation. It achieved the lowest MSE and highest LCC at both the utterance and system level. From the fact that MOSNet + EL shows better performance than MOSNet + GQT, we can say that considering the distribution of frame-level MOSs is more important than learning the quality embeddings for MOS evaluation.

When we combine MOSNet with both the GQT layer and Encoding Layer, the performance of MOSNet + GQT + EL is better than MOSNet + GQT but worse than MOSNet + EL. In other words, the Encoding Layer helps MOSNet + GQT, but the GQT layer does not help MOSNet + EL. As will be described in Section 4.5, we conjecture that GQTs prevent the embeddings of the utterances with similar scores from getting too far from each other, which results in that MOSNet + EL cannot separate the embeddings according to the frame-level scores as before.

With the test results using the VCC 2016 data, we can conclude that either the GQT layer or the Encoding Layer also improves the generalization ability of MOSNet. The Encoding Layer helps generalization through the sophisticated aggregation of frame-level scores. Moreover, the GQT layer directly

Table 3: Results of similarity prediction. ACC denotes accuracy. The best results are highlighted in bold.

Model	Level	MSE	LCC	SRCC	ACC
SIMNet	utterance	0.774	0.552	0.549	0.687
	system	0.052	0.925	0.905	-
+ GQT	utterance	0.763	0.558	0.554	0.687
	system	0.047	0.931	0.913	-
+ EL	utterance	0.770	0.555	0.554	0.684
	system	0.049	0.929	0.911	-
+ GQT + EL	utterance	0.761	0.560	0.558	0.689
	system	0.045	0.934	0.916	-

improves the generalization ability of the model by learning the universal criteria for speech quality evaluation.

4.4. Experiments on similarity score prediction

The performance of the similarity score prediction models is evaluated with the MSE, LCC, SRCC, and accuracy, as shown in Table 3. As we use a scalar for both the ground-truth and predicted similarity score, we regard the scores lower than 2.5 as the answer ‘‘Same’’ and the scores higher or equal to 2.5 as the answer ‘‘Different.’’ Then we calculate the accuracy of a model as the ratio of cases when both answers from the model and humans are the same. All the proposed models show improvements in all the metrics except the accuracy.

Note that, unlike in MOS prediction, using the GQT layer shows better results than using the Encoding Layer. We can infer that finding the criteria for voice similarity evaluation is more important than using the distribution of the frame-level similarity scores. Furthermore, SIMNet + GQT + EL achieved the best performance among the four models, which means that SIMNet + EL takes advantage of the GQT layer. Considering this result and the fact that we judge the voice similarity throughout the whole utterance rather than specific frames, we can infer that SIMNet + EL does not separate the embeddings as far as in MOSNet + EL, according to the frame-level scores.

Finally, we test how well our model, SIMNet + GQT + EL, approximates human assessment in terms of an evaluation method used in VCC 2018 [23]. As mentioned earlier, the utterance-level similarity score can be classified into ‘‘Same’’ or ‘‘Different.’’ According to the method, the similarity score of a system is the ratio of the utterances that received ‘‘Same’’ compared to the target speech. Then we obtain the MSE, LCC, and SRCC by comparing the scores using the answers of our model and human, which are 0.037, 0.714, and 0.696, respectively.

4.5. Embeddings learned from the proposed models

To discuss the effect of the GQT layer and Encoding Layer to the embeddings, we visualized the embeddings of four MOS prediction models using t-distributed stochastic neighbor embedding (t-SNE) [28] in Figure 1. We displayed a frame-level feature of CNN-BLSTM as a single dot. We consider two systems submitted by one team as the same system because they are based on similar algorithms, which results in a total of 26 systems. We use 390 random utterances from the VCC 2018 test set so that an average of 15 utterances per system exists. We evenly select half of 26 systems after sorting the systems according to the MOS. The color in Figure 1 indicates the system that generates the corresponding utterance. Note that each system has its own system-level MOS in that different systems have different MOSs and the utterance-level MOSs are similar

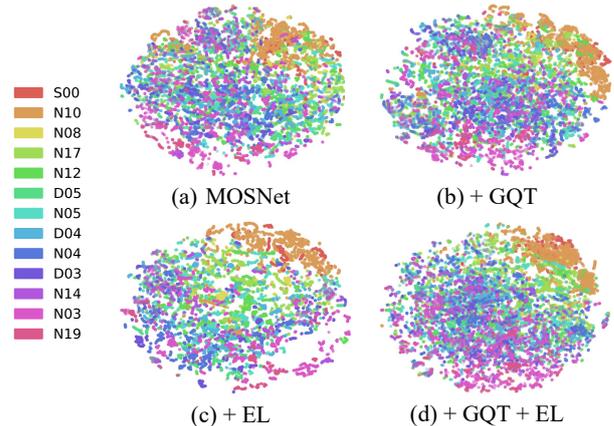


Figure 1: Visualization of embeddings for (a) MOSNet, (b) + GQT, (c) + EL, and (d) + GQT + EL

within each system. The red dots are for the source speech, and the orange dots are for the voice conversion system with the highest MOS.

Compared to (a), (b) shows that the GQT layer prevents the embeddings of the same system from getting apart from each other. Comparing (a) and (c), we see that the Encoding Layer separates the embeddings according to the system. This shows that the Encoding Layer automatically learns the embeddings that are more useful for the aggregation, by utilizing the distribution of the frame-level scores. In (d), we see that the embeddings of the same systems get close to each other while embeddings of different systems get apart from each other. Specifically, + GQT + EL learns distinguishable embeddings between different systems that have higher MOSs, including S00, N10, and N17. However, the other systems having lower MOSs are less distinguishable from each other than in MOS + EL. We infer that this is the main reason for the degradation of performance.

5. Conclusion

We proposed three deep-learning-based speech quality assessment models using cluster-based modeling, which improved MOSNet and SIMNet using a GQT layer, an Encoding Layer, and both of them, respectively. With experimental results on MOS and similarity score prediction, we showed that the GQT layer learns the criteria of speech quality evaluation as soft clusters, and the Encoding Layer utilizes the frame-level scores in a more sophisticated way. For future work, we will apply our models to approximate other speech quality assessments, such as PESQ. Furthermore, we will use our models to guide current TTS models to learn human perception, by using a perceptual loss in training. Finally, we will figure out how to create the synergy between using the GQT layer and using the Encoding Layer for MOS prediction.

6. Acknowledgements

This material is based upon work supported by the Ministry of Trade, Industry & Energy (MOTIE, Korea) under Industrial Technology Innovation Program (No. 10080667, Development of conversational speech synthesis technology to express emotion and personality of robots through sound source diversification).

7. References

- [1] K. J. Han, R. Prieto, K. Wu, and T. Ma, “State-of-the-art speech recognition using multi-stream self-attention with dilated 1d convolutions,” *arXiv preprint arXiv:1910.00716*, 2019.
- [2] A. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “WaveNet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [3] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Skerry-Ryan, R. A. Saurous, Y. Agiomyrgiannakis, and Y. Wu, “Natural TTS synthesis by conditioning WaveNet on Mel spectrogram predictions,” in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 4779–4783.
- [4] N. Li, S. Liu, Y. Liu, S. Zhao, and M. Liu, “Neural speech synthesis with Transformer network,” in *Proc. of the AAAI Conference on Artificial Intelligence*, 2019, pp. 6706–6713.
- [5] S. King, “Measuring a decade of progress in text-to-speech,” *Loquens*, vol. 1, no. 1, 2014.
- [6] P. Wagner, J. Beskow, S. Betz, J. Edlund, J. Gustafson, G. E. Henter, S. L. Maguer, Z. Malisz, E. Székely, C. Tännander, and J. Voße, “Speech synthesis evaluation — state-of-the-art assessment and suggestion for a novel research program,” in *Proc. 10th Speech Synthesis Workshop (SSW10)*, 2019.
- [7] R. Kubichek, “Mel-cepstral distance measure for objective speech quality assessment,” in *Proc. IEEE Pacific Rim Conference on Communications Computers and Signal Processing*, 1993, pp. 125–128.
- [8] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, “Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs,” in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2001, pp. 749–752.
- [9] D.-S. Kim, “ANIQUE: an auditory model for single-ended speech quality estimation,” *IEEE Transactions on Speech and Audio Processing*, vol. 13, pp. 821–831, 2005.
- [10] L. Malfait, J. Berger, and M. Kastner, “P. 563: The ITU-T standard for single-ended speech quality assessment,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, pp. 1924–1934, 2006.
- [11] T. Yoshimura, G. E. Henter, O. Watts, M. Wester, J. Yamagishi, and K. Tokuda, “A hierarchical predictor of synthetic speech naturalness using neural networks,” in *Proc. of Interspeech*, 2016, pp. 342–346.
- [12] B. Patton, Y. Agiomyrgiannakis, M. Terry, K. W. Wilson, R. A. Saurous, and D. Sculley, “AutoMOS: Learning a non-intrusive assessor of naturalness-of-speech,” in *Proc. of NIPS End-to-end Learning for Speech and Audio Processing Workshop*, 2016.
- [13] S. Fu, Y. Tsao, H. Hwang, and H. Wang, “Quality-Net: An end-to-end non-intrusive speech quality assessment model based on blstm,” in *Proc. of Interspeech*, 2018, pp. 1873–1877.
- [14] C. Lo, S. Fu, W. Huang, X. Wang, J. Yamagishi, Y. Tsao, and H. Wang, “MOSNet: Deep learning based objective assessment for voice conversion,” in *Proc. of Interspeech*, 2019, pp. 1541–1545.
- [15] Y. Wang, D. Stanton, Y. Zhang, R. Skerry-Ryan, E. Battenberg, J. Shor, Y. Xiao, F. Ren, Y. Jia, and R. A. Saurous, “Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis,” *arXiv preprint arXiv:1803.09017*, 2018.
- [16] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” in *Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 6000–6010.
- [17] D. Liu, C. Yang, S. Wu, and H. Lee, “Improving unsupervised style transfer in end-to-end speech synthesis with end-to-end speech recognition,” in *Proc. of Spoken Language Technology Workshop (SLT)*, 2018, pp. 640–647.
- [18] K. Cho, B. Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, “Learning phrase representations using RNN encoder-decoder for statistical machine translation,” in *Proc. of Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1724–1734.
- [19] H. Zhang, J. Xue, and K. Dana, “Deep ten: Texture encoding network,” in *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 708–717.
- [20] W. Cai, Z. Cai, X. Zhang, X. Wang, and M. Li, “A novel learnable dictionary encoding layer for end-to-end language identification,” in *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5189–5193.
- [21] Y. Jung, Y. Kim, H. Lim, Y. Choi, and H. Kim, “Spatial pyramid encoding with convex length normalization for text-independent speaker verification,” in *Proc. of Interspeech*, 2019, pp. 4030–4034.
- [22] J. Xue, H. Zhang, and K. Dana, “Deep texture manifold for ground terrain recognition,” in *Proc. of Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 558–567.
- [23] J. Lorenzo-Trueba, J. Yamagishi, T. Toda, D. Saito, F. Villavicencio, T. Kinnunen, and Z. Ling, “The voice conversion challenge 2018: Promoting development of parallel and nonparallel methods,” in *Proc. of Odyssey The Speaker and Language Recognition Workshop*, 2018, pp. 195–202.
- [24] G. J. Mysore, “Can we automatically transform speech recorded on common consumer devices in real-world environments into professional production quality speech?—a dataset, insights, and challenges,” *IEEE Signal Processing Letters*, vol. 22, no. 8, pp. 1006–1010, 2015.
- [25] T. Toda, L. Chen, D. Saito, F. Villavicencio, M. Wester, Z. Wu, and J. Yamagishi, “The Voice Conversion Challenge 2016,” in *Proc. of Interspeech*, 2016, pp. 1632–1636.
- [26] K. Pearson, “Notes on the history of correlation,” *Biometrika*, vol. 13, no. 1, pp. 25–45, 1920.
- [27] C. Spearman, “The proof and measurement of association between two things,” *The American Journal of Psychology*, vol. 15, no. 1, pp. 72–101, 1904.
- [28] L. V. D. Maaten and G. Hinton, “Visualizing data using t-SNE,” *Journal of machine learning research*, no. 9, pp. 2579–2605, 2008.