

Incorporating Broad Phonetic Information for Speech Enhancement

Yen-Ju Lu¹, Chien-Feng Liao¹, Xugang Lu², Jieh-weih Hung³, Yu Tsao¹

¹Research Center for Information Technology Innovation, Academic Sinica

²National Institute of Information and Communications Technology, Japan

³National Chi Nan University

{r03942063, r06946002}@ntu.edu.tw, xugang.lu@nict.go.jp, jwhung@ncnu.edu.tw,
yu.tsao@citi.sinica.edu.tw

Abstract

In noisy conditions, knowing speech contents facilitates listeners to more effectively suppress background noise components and to retrieve pure speech signals. Previous studies have also confirmed the benefits of incorporating phonetic information in a speech enhancement (SE) system to achieve better denoising performance. To obtain the phonetic information, we usually prepare a phoneme-based acoustic model, which is trained using speech waveforms and phoneme labels. Despite performing well in normal noisy conditions, when operating in very noisy conditions, however, the recognized phonemes may be erroneous and thus misguide the SE process. To overcome the limitation, this study proposes to incorporate the broad phonetic class (BPC) information into the SE process. We have investigated three criteria to build the BPC, including two knowledge-based criteria: place and manner of articulatory and one data-driven criterion. Moreover, the recognition accuracies of BPCs are much higher than that of phonemes, thus providing more accurate phonetic information to guide the SE process under very noisy conditions. Experimental results demonstrate that the proposed SE with the BPC information framework can achieve notable performance improvements over the baseline system and an SE system using monophonic information in terms of both speech quality intelligibility on the TIMIT dataset.

Index Terms: speech enhancement, broad phonetic classes, articulatory attribute

1. Introduction

Speech enhancement (SE) systems aim to transform the distorted speech signal to an enhanced one with improved speech intelligibility and quality. In many real-world applications, such as speech coding [1, 2], assistive listening devices [3, 4] and automatic speech recognition (ASR) [5–7], SE has been widely used as a front-end processor, which effectively reduces noise components and improves the overall performance. Recently, along with the burgeoning of deep learning (DL), applying DL-based models for the SE task has been extensively investigated [8–14]. In particular, the DL-based models provide a strong capability of modeling non-linear and complex transformations from noisy and clean speech, and therefore the DL-based SE approaches have yielded notable improvements over traditional SE methods, especially under very low signal-to-noise-ratio (SNR) and non-stationary noise conditions. Another vital feature of DL-based models is their high flexibility to incorporate heterogeneous data, which may not be easily performed for traditional signal processing-based approaches. Previous studies have confirmed that the face/lip images [15] and symbolic sequences for acoustic signals [16] can be incorporated into an SE system.

The phonetic information generated from an acoustic model (AM) has also been adopted to guide the SE process and achieve improved denoising performance. In [17–19], an AM and an SE system are trained jointly, where one system’s input depends on the other’s output. In [20], the SE result was treated as an adaptive feature to generate a phoneme posteriorgram (PPG) more precisely. Although achieving improved performance under ordinary conditions, the AM may generate inaccurate PPGs under very noisy or acoustically mismatched conditions, which may misguide the SE process to generate even poorer results. To overcome the issue, we proposed to incorporate the broad phonetic class (BPC) posteriorgram (BPPG) in the SE system. We argue that the speech signals within the same BPC share the same noisy-to-clean transformation, and the BPPG can guide the SE process well even under very noisy conditions. The main concept of BPC is clustering phonemes with similar characteristics into a broad class. The criteria for clustering can be either knowledge-based or data-driven. The most common knowledge-based criteria exploit the manner and place of articulations for pronouncing the phonemes [21]. For data-driving criteria, the confusions between phonemes are measured based on some predefined evaluation metrics, and the phonemes that generate the closest metric scores (easily confused) are clustered into the same class. As compared to phoneme-based AM, the labels needed to train the BPC-AM are much more easily accessed, and the recognition accuracy is generally higher, especially in noisy conditions. In the past, the BPC has been widely used in various speech-related applications, including speaking rate estimation [22], multilingual systems [23] and phone recognition [24–27]. Experimental results have confirmed that BPCs can effectively boost recognition rates as a pre-processor or in unsupervised learning tasks.

In this paper, we proposed to combine the BPPG into an SE system, termed BPSE. Three criteria that we used to define the BPC are tested. They are based on manner of articulation, place of articulations, and data-driven, thus termed BPC(M), BPC(P) and BPC(D). We evaluated the proposed BPSE on the TIMIT dataset. Two standard evaluation metrics, speech quality (PESQ) [28] and short-time objective intelligibility (STOI) [29], are used to evaluate the SE performance. For comparison, we implemented an SE system with monophonic PPG and two baseline SE systems. Experimental results first validate the effectiveness of incorporating the BPC information for the SE system. These results also confirm that the BPSE system outperforms an SE with monophonic PPG, especially under very noisy conditions, where PPG may be inaccurate.

The rest of this paper is organized as follows. Section 2 introduces the criteria used to define BPC. Section 3 details the proposed BPSE system. Section 4 presents the experimental setup and results. Finally, Section 5 concludes this paper.

2. Broad Phonetic Classes

In this section, we introduce two knowledge-based and one data-driven criteria used for defining BPC.

2.1. Knowledge-based criterion

In the phonetic field, consonants can be readily classified based on the manner and place where the vocal tract obstructs the airstream. Accordingly, the entire consonants can be divided into different place/manner-wise articulation classes. Distinct characteristics are found between the different manners/places of articulation, and different manner can even be easily revealed by the shape of waveform [21]. It is also found in [27] that phonemes belonging to the same manner/place of articulation contain very similar spectral characteristics, and may generate high confusion results when performing speech recognition.

In this study, based on the manner of articulation, we cluster all the 60 phonemes into 5 clusters: vowels, stops, fricatives, nasals and silence as suggested in [27], where the diphthongs and semi-vowels are merged into the vowel class. On the other hand, according to the place of articulation, we divide phonemes into 9 clusters: bilabial, labiodental, dental, alveolar, postalveolar, velar, glottal, vowels, and silence, as suggested in [21]. Please note that all the vowels are clustered into one distinct class in both manner and place articulations. The underlying reason is that these two classification criteria are connected with the manner/place where the vocal tract obstructs the airstream, while the vowels do not have such properties.

2.2. Data-driven criterion

In contrast to the knowledge-based criteria, the data-driven criterion conducts phoneme clustering through the phoneme similarity measured by an ASR. Based on [30], the confusion matrix, \mathbf{M} , contains information about the similarities between each pair of phonemes, where the entry \mathbf{M}_{ij} denotes the number of the event for phoneme i being mistakenly recognized as phoneme j . A symmetric similarity matrix, \mathbf{S} , can be computed from the confusion matrix, \mathbf{M} , where \mathbf{S}_{ij} , the similarity between phonemes i and j , is computed by (1):

$$\mathbf{S}_{ij} = \mathbf{S}_{ji} = \sum_{k \neq i, j} \min(\mathbf{M}_{ik}, \mathbf{M}_{jk}) \quad (1)$$

where k is the phone index and $k \neq i, j$. From Eq. (1), we note that each entry in the similarity matrix presents the distance between a pair of phonemes based on how often they are confused by the ASR results. When applying the similarity metric to cluster phonemes, we initially define each phoneme as one distinct cluster. Then, we gradually reduce the cluster number by merging the nearest clusters. This process repeats until the cluster number meets our expectations. In Table 1, we list the clustering results (9 clusters, which are recommended in [30]) obtained by the data-driven criterion on the TIMIT dataset.

3. Proposed BPSE system

Figure 1 shows the overall architecture of the proposed BPSE system, which consists of a Transformer-based SE model [31] and a BPPG extractor.

3.1. PPG and BPPG extractors

To extract PPG and BPPG, we trained acoustic models (AMs) with noisy speech data as input, and phoneme labels and BPC

Table 1: Clustering results by data-driven on TIMIT.

Clusters	TIMIT phoneme label
Cluster1	bcl, dcl, epi, gcl, kcl, pau, pcl, q, tcl
Cluster2	b, d, dh, f, g, k, p, t, th, v
Cluster3	y
Cluster4	hh, hv
Cluster5	dx, em, en, m, n, ng, nx
Cluster6	aa, ae, ah, ao, aw, ax, ax-h, axr, ay, eh, el, er, ey, ih, ix, iy, l, ow, oy, r, uh, uw, ux, w
Cluster7	ch, jh, s, sh, z, zh
Cluster8	eng
Cluster9	h#

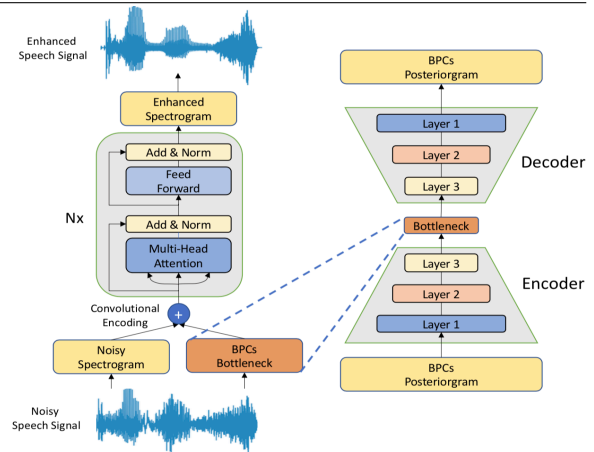
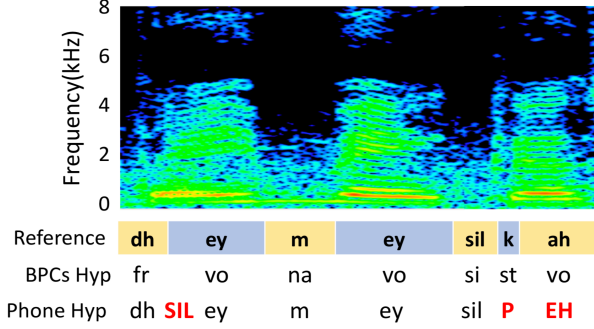


Figure 1: The architecture of the proposed BPSE system

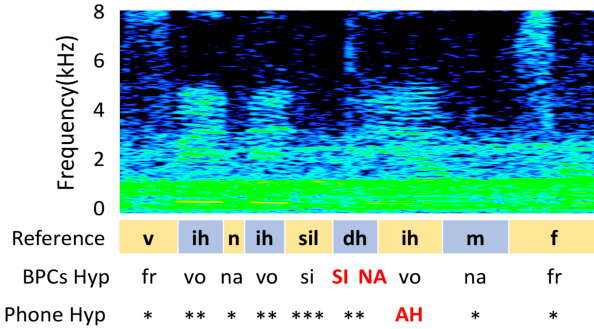
labels as output, respectively. All AMs were trained via an off-the-shelf Kaldi recipe [32] with the same DNN-HMM architecture; the DNN has 7 hidden layers, each layer containing 1024 neurons. With the trained AM and given a speech utterance as the respective input, we collected the sequence of triphone state posteriors to form the PPG and BPPG. As for the case of phone unit, since the total number of triphone states is huge, we used the mono-phone state PPG instead. In order to keep the vector dimensions of PPG and BPPG to be moderate, an auto-encoder (AE) model, as shown in the right part of Figure 1, was used to reduce the size of the BPPG vector. In particular, the input and output of the AE model are both BPPG vectors, and the latent representations with a dimension of 96 were used as the final BPPG features. Finally, the BPPG features are appended to the noisy spectrogram vectors and then served as the input to the SE system.

3.2. Transformer-based SE Network

Transformer is an attention-based deep neural network, originally proposed for machine translation [31] and later explored in many other natural language processing tasks, exhibiting stark improvements over the well-known recurrent neural networks (RNNs). Recently, it was further exploited in the SE task [33], in which the Transformer model was also shown to outperform convolutional-based and recurrent-based models. The self-attention mechanism of Transformer has been well studied in the SE literature [34–36], and one of its particularities is to allow the model to learn long-range dependencies within the sequence efficiently. Also, since Transformer can process the



(a) Spectrogram and recognition result at 10dB SNR



(b) Spectrogram and recognition result at 0dB SNR level

Figure 2: Monophonic and BPC recognition results of a sample utterance in two SNR levels. The length of slots in Reference denotes the start/end time. Phone Hyp and BPC Hyp mean the recognition results by phone-based and BPC-based AMs, respectively. The BPC labels include vowels(vo), stops(st), fricatives(fr), nasals(na) and silence(si). The uppercase phone/BPC labels represent recognition errors, including both insertions and substitutions. The star labels represent deletions.

whole sequence in parallel, the respective computational time is reduced relative to RNNs.

The original Transformer consists of encoder and decoder networks for sequence-to-sequence learning. In our method, the decoder part is omitted since the input and output sequences have the same length during the SE process. Another modification is that we use convolutional layers to replace positional encoding in order to inject relative location information to the frames in the sequence, which is distinct from [33]. The causal implementation on the transformer was applied for the real-time processing scenario. The rest of the architecture is implemented as a standard Transformer shown in the left part of Figure 1, which is composed of N attention blocks. In each attention block, the first sub-layer is the multi-head self-attention (MHSA), followed by a feed-forward network containing two fully-connected layers. Both sub-layers are followed by a residual connection to the input and a layer normalization [37]. Finally, the Transformer output is projected back to the frequency dimension using a fully-connected layer with ReLU activation, and the corresponding mean-absolute-error relative to the clean speech is computed.

Table 2: Monophonic and BPC recognition results at different SNR levels

Phoneme/BPCs Accuracy(in %)				
SNR	Mono	BPC(D)	BPC(M)	BPC(P)
-5	10.3	26.8	33.3	27.3
0	16.7	39.8	44.1	37.7
5	24.8	51.4	53.9	46.6
10	32.3	59.5	61.2	55.2
15	42.9	66.7	65.9	60.7
Avg	25.4	48.8	51.7	45.5

4. Experiment

4.1. Experimental setup

The experiments were conducted using the TIMIT database [38] together with other noise sources [39] for the SE models. The AM models were trained on clean speech from TIMIT. The 3696 utterances from the TIMIT training set (excluding SA files) were used and randomly corrupted with 100 noise types from [39] at 32 different SNR levels, amounting to 3696 noisy-clean paired training utterances. The test set was constructed by adding 5 unseen noise types into all of the core test set of the TIMIT database (including 192 utterances) at 5 SNR levels (15 dB, 10 dB, 5 dB, 0 dB and -5 dB). The speakers in the test set are different from those in the training set.

The speech waveforms were recorded at 16 kHz sampling rate. The short-time Fourier transform (STFT) with a Hamming window size of 32 ms and a hop size of 16 ms was applied to convert the speech waveforms into spectral features, each with 257 dimensions. The \log_{1p} function ($\log_{1p}(x) = \log(1+x)$) was adopted on the magnitude spectrogram to enforce the range of each coefficient to be positive. During training, a segment of 64 frames was used as the input unit to the SE system. During testing, the enhanced spectral features were synthesized back to the waveform signals via the inverse STFT and an overlap-add procedure. The phases of the noisy signals were used for the waveform reconstruction. The Adam optimizer [40] was used to set the learning rate, and an early stopping was performed according to the validation set to prevent overfitting.

The autoencoder consisted of three fully connected layers in the encoder network, with the dimensions of [512, 256, 96], and the decoder was symmetric to it. The 96-dimensional bottleneck feature from the encoder was processed by a sigmoid function to limit the value range and further concatenated with the noisy spectrogram and then fed into the SE module. For the Transformer, four 1-D convolutional layers were used to encode the input feature. The channel sizes were [1024, 512, 256, 128]. The filter size and stride were set to 3 and 1, respectively. We set the number of attention blocks N to 8. The MHSA consisted of 8 heads and 64 neurons for each head. The two layers of the feed-forward network consisted of 512 and 256 neurons, respectively. Leaky ReLU was used as the activation function.

Based on the model architecture shown in Fig. 1, we implemented three systems using three types of BPPG, namely BPPG(M), BPPG(P), and BPPG(D). For comparison, we built a PPG system based on the same architecture in Fig. 1, and this system is termed PPG(Mono). Two baseline SE systems were implemented: one was based on a Long-short term memory (LSTM), which was single-directional and fitted the causal implementation, and the other was the Transformer-based SE without incorporating PPG and BPPG information.

Table 3: Averaged STOI scores for the BPSE with BPPG(M), BPPG(P), and BPPG(D) and monophone PPG. The results obtained using the ground-truth BPPG information were also listed, denoted as GT-PPG(Mono) and GT-BPPG(M). The results of original noisy, and enhanced speech generated by LSTM and Transformer were listed for comparison. The boldface numbers indicate the best results among the testing condition.

SNR	Noisy	LSTM	Transformer	PPG(Mono)	Broad Phone Class			Ground Truth	
					BPPG(P)	BPPG(M)	BPPG(D)	GT-PPG(Mono)	GT-BPPG(M)
-5	0.595	0.548	0.620	0.616	0.629	0.627	0.628	0.679	0.708
0	0.701	0.686	0.755	0.759	0.765	0.765	0.763	0.796	0.808
5	0.800	0.815	0.851	0.859	0.860	0.861	0.859	0.876	0.879
10	0.880	0.900	0.912	0.917	0.918	0.918	0.917	0.924	0.925
15	0.935	0.946	0.948	0.950	0.951	0.950	0.951	0.953	0.953
Avg	0.782	0.779	0.817	0.820	0.824	0.824	0.823	0.846	0.855

4.2. Experimental result

4.2.1. Phonetic/BPC recognition accuracies

To validate the assumption that additional PPG/BPPG can improve the SE process, we first analyze the correctness of PPG and BPPG generated by the AMs. Given the testing data, the phoneme/BPC accuracies were shown in Table 2, where monophonic and BPC (based on manner of articulation, place of articulation, and data-driven criteria) recognition results are denoted as "Mono", "BPC(M)", "BPC(P)", and "BPC(D)", respectively.

From Table 2, we note that the "Mono" system (the overall accuracy is 25.4%) is very sensitive to noise as compared to the three BPC systems (the overall accuracies range from 45.5% to 51.7%). In particular, at the SNR of 0 dB, the accuracy rate of "Mono" drops to 16.7% while the accuracy rates of BPCs' were around 40%. These results show that BPC-based systems do not need to distinguish phonemes, some of whose acoustic properties might be highly overlapped, and thus can provide robust recognition performance. To qualitatively analyze the recognition results, Figure 2 shows the spectrograms of sample utterances at two SNR levels (10 dB and 0 dB). The phoneme and BPC(M) recognition results, along with the correct transcription are listed. The results show that at 10 dB-SNR condition, both phoneme and BPC(M) were reliable. However, at the SNR drops to 0 dB, many misrecognized phonemes can be observed, while the BPCs recognition results are still acceptable. Results of Table 2 and Figure 2 suggest that the PPG(Mono) may contain incorrect information that could misguide the SE process.

4.2.2. Speech enhancement results

Table 3 presents the STOI scores of the proposed BPSE system and the comparative methods. From the table, we first confirm PPG(Mono) outperforms Transformer baseline in higher SNR conditions, confirming that the monophonic PPG can improve the SE performance. However, we note that PG(Mono) underperforms Transformer baseline in low SNR levels, which may be owing to the incorrect PPG information somehow deteriorates the original SE performance. Next, we note that all of the three BPPG systems provide consistent improvements over Transformer baseline across different SNR levels and outperform the PPG(Mono) systems. The results confirm the robustness of the BPPG that can provide beneficial guide for SE even under very low SNR conditions. Among the three BPPG systems, BPPG(M) slightly outperforms the other two systems. The reason is yet to be further investigated.

To further identify the effect of PPG and BPPG, we conducted an additional experiment, where the original phoneme labels from the TIMIT corpus was to get the PPG/BPPG infor-

mation. More specifically, the input to the AE in the right side of Fig 1 is an one-hot vector, whose dimension is the number of monophone/BPCs. Here the BPC(M) was used as a representative. This set of results was denoted as Ground Truth (GT-PPG(Mono) and GT-BPPG(M)) in Table 3. From the results of PPG(Mono) and BPPG(M) in the Ground Truth column, BPPG(M) still performed obviously better than PPG(Mono), suggesting that BPPG may be more suitable than PPG to further improve the SE performance.

In addition to the STOI scores, Figure 3 presents the average PESQ scores of the SE systems. We note very similar trends to the STOI results as shown in Table 3. All of the three BPPG systems achieve notably better performance in terms of Transformer SE baseline. For the Ground-truth situation, BPPG(M) still yields higher PESQ scores over the PPG(Mono) system.

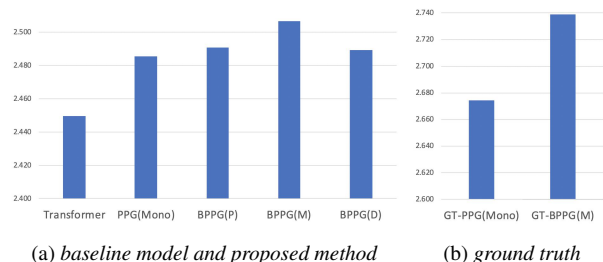


Figure 3: Average PESQ scores baseline model and proposed method

5. Conclusion

In this paper, we proposed to use the BPC information to guide the SE process to achieve better denoising performance. Three clustering criteria were investigated, and the results confirmed that incorporating the BPC information the SE performance can be notably improved over different SNR conditions. The performance even outperforms the SE with monophonic PPG system. The contribution of this work includes: (1) This is the first attempt that incorporates the BPPG to the SE system and obtain promising results. (2) We have verified that both knowledge-based and data-driven criteria can be applied to cluster phonemes into BPCs, both providing beneficial information to the SE system. From the ground-truth results, we note that there are still rooms for further improvements. We will further explore how to more effectively use these BPC information for the SE task in the future study.

6. References

- [1] R. Martin and R. V. Cox, "New speech enhancement techniques for low bit rate speech coding," in *Proc. Workshop on Speech Coding Proceedings. Model, Coders, and Error Criteria 1999*.
- [2] A. J. Accardi and R. V. Cox, "A modular approach to speech enhancement with an application to speech coding," in *Proc. ICASSP 1999*.
- [3] D. Wang, "Deep learning reinvents the hearing aid," *IEEE spectrum*, vol. 54, no. 3, pp. 32–37, 2017.
- [4] Y.-H. Lai, F. Chen, S.-S. Wang, X. Lu, Y. Tsao, and C.-H. Lee, "A deep denoising autoencoder approach to improving the intelligibility of vocoded speech in cochlear implant simulation," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 7, pp. 1568–1578, 2016.
- [5] J. Li, L. Deng, Y. Gong, and R. Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, 2014.
- [6] H. Erdogan, J. R. Hershey, S. Watanabe, and J. Le Roux, "Phase-sensitive and recognition-boosted speech separation using deep recurrent neural networks," in *Proc. ICASSP 2015*.
- [7] F. Weninger, H. Erdogan, S. Watanabe, E. Vincent, J. Le Roux, J. R. Hershey, and B. Schuller, "Speech enhancement with lstm recurrent neural networks and its application to noise-robust asr," in *Proc. LVA/ICA*, 2015.
- [8] Y. Wang, A. Narayanan, and D. Wang, "On training targets for supervised speech separation," *IEEE/ACM transactions on audio, speech, and language processing*, vol. 22, no. 12, pp. 1849–1858, 2014.
- [9] B. Xia and C. Bao, "Wiener filtering based speech enhancement with weighted denoising auto-encoder and noise classification," *Speech Communication*, vol. 60, pp. 13–29, 2014.
- [10] X. Lu, Y. Tsao, S. Matsuda, and C. Hori, "Speech enhancement based on deep denoising autoencoder," in *Proc. Interspeech 2013*.
- [11] Y. Xu, J. Du, L.-R. Dai, and C.-H. Lee, "A regression approach to speech enhancement based on deep neural networks," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 1, pp. 7–19, 2014.
- [12] M. Kolbæk, Z.-H. Tan, and J. Jensen, "Speech intelligibility potential of general and specialized deep neural network based speech enhancement systems," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 1, pp. 153–167, 2016.
- [13] K. Tan, X. Zhang, and D. Wang, "Real-time speech enhancement using an efficient convolutional recurrent network for dual-microphone mobile phones in close-talk scenarios," in *Proc. ICASSP 2019*.
- [14] J. Qi, J. Du, S. M. Siniscalchi, and C. Lee, "A theory on deep neural network based vector-to-vector regression with an illustration of its expressive power in speech enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, 2019.
- [15] J. Hou, S. Wang, Y. Lai, Y. Tsao, H. Chang, and H. Wang, "Audio-visual speech enhancement using multimodal deep convolutional neural networks," *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 2, no. 2, pp. 117–128, 2018.
- [16] C.-F. Liao, Y. Tsao, X. Lu, and H. Kawai, "Incorporating Symbolic Sequential Modeling for Speech Enhancement," in *Proc. Interspeech 2019*.
- [17] Z. Chen, S. Watanabe, H. Erdogan, and J. R. Hershey, "Speech enhancement and recognition using multi-task learning of long short-term memory recurrent neural networks," in *Proc. Interspeech 2015*.
- [18] Z. Wang and D. Wang, "A joint training framework for robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 4, pp. 796–806, 2016.
- [19] T. Gao, J. Du, L. Dai, and C. Lee, "Joint training of front-end and back-end deep neural networks for robust speech recognition," in *Proc. ICASSP 2015*.
- [20] Z. Du, M. Lei, J. Han, and S. Zhang, "Pan: Phoneme-aware network for monaural speech enhancement," in *Proc. ICASSP 2020*.
- [21] P. Ladefoged and K. Johnson, *A course in phonetics*. Nelson Education, 2014.
- [22] J. Yuan and M. Liberman, "Robust speaking rate estimation using broad phonetic class recognition," in *Proc. ICASSP 2010*.
- [23] A. Žgank, B. Horvat, and Z. Kačič, "Data-driven generation of phonetic broad classes, based on phoneme confusion matrix similarity," *Speech Communication*, vol. 47, no. 3, pp. 379–393, 2005.
- [24] Y. Tsao, J. Li, and C.-H. Lee, "A study on separation between acoustic models and its applications," in *Proc. Interspeech*, 2005.
- [25] J. Morris and E. Fosler-Lussier, "Conditional random fields for integrating local discriminative classifiers," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 3, pp. 617–628, 2008.
- [26] Y.-T. Lee, X.-B. Chen, H.-S. Lee, J.-S. R. Jang, and H.-M. Wang, "Multi-task learning for acoustic modeling using articulatory attributes," in *Proc. APSIPA ASC*, 2019.
- [27] P. Scanlon, D. P. Ellis, and R. B. Reilly, "Using broad phonetic group experts for improved speech recognition," *IEEE transactions on audio, speech, and language processing*, vol. 15, no. 3, pp. 803–812, 2007.
- [28] 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. ICASSP*, 2001.
- [29] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "An algorithm for intelligibility prediction of time-frequency weighted noisy speech," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2125–2136, 2011.
- [30] C. Lopes and F. Perdigão, "Broad phonetic class definition driven by phone confusions," *EURASIP Journal on Advances in Signal Processing*, vol. 2012, no. 1, p. 158, 2012.
- [31] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, 2017.
- [32] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *Proc. ASRU*, 2011.
- [33] J. Kim, M. El-Khomy, and J. Lee, "T-gsa: Transformer with gaussian-weighted self-attention for speech enhancement," in *Proc. ICASSP 2020*.
- [34] X. Hao, C. Shan, Y. Xu, S. Sun, and L. Xie, "An attention-based neural network approach for single channel speech enhancement," in *Proc. ICASSP 2019*.
- [35] C. Yang, J. Qi, P. Chen, X. Ma, and C. Lee, "Characterizing speech adversarial examples using self-attention u-net enhancement," in *Proc. ICASSP 2020*.
- [36] Y. Koizumi, K. Yaiabe, M. Delcroix, Y. Maxuxama, and D. Takeuchi, "Speech enhancement using self-adaptation and multi-head self-attention," in *Proc. ICASSP 2020*.
- [37] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [38] J. S. Garofolo, "Timit acoustic phonetic continuous speech corpus," *Linguistic Data Consortium, 1993*, 1993.
- [39] G. Hu, "100 nonspeech environmental sounds," *The Ohio State University, Department of Computer Science and Engineering*, 2004.
- [40] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.