



Modality-Specific Speech Enhancement and Noise-Adaptive Fusion for Acoustic and Body-Conduction Microphone Framework

Yunsik Kim¹, Yoonyoung Chung^{1,2,3*}

¹Department of Electrical Engineering, ²Department of Semiconductor Engineering, ³Center for Semiconductor Technology Convergence, Pohang University of Science and Technology (POSTECH), Pohang, Korea

ys.kim@postech.ac.kr, ychung@postech.ac.kr

Abstract

Body-conduction microphone signals (BMS) bypass airborne sound, providing strong noise resistance. However, a complementary modality is required to compensate for the inherent loss of high-frequency information. In this study, we propose a novel multi-modal framework that combines BMS and acoustic microphone signals (AMS) to achieve both noise suppression and high-frequency reconstruction. Unlike conventional multi-modal approaches that simply merge features, our method employs two specialized networks: a mapping-based model to enhance BMS and a masking-based model to denoise AMS. These networks are integrated through a dynamic fusion mechanism that adapts to local noise conditions, ensuring the optimal use of each modality's strengths. We performed evaluations on the TAPS dataset, augmented with DNS-2023 noise clips, using objective speech quality metrics. The results clearly demonstrate that our approach outperforms single-modal solutions in a wide range of noisy environments.

Index Terms: multi-modal speech enhancement, body-conduction microphone

1. Introduction

Recent advances in deep learning-based speech enhancement have demonstrated remarkable progress through two dominant methodologies: masking-based and mapping-based approaches [1]. Masking-based methods [2, 3, 4, 5, 6, 7] primarily operate in the time-frequency (TF) domain, where noisy speech spectrograms—computed via short-time Fourier transform (STFT)—are multiplied by estimated ideal ratio masks to suppress noise components. These methods excel at preserving speech phase information and enabling interpretable TF bin-wise processing. In contrast, mapping-based methods [8, 9, 10, 11, 12] directly estimate clean speech features from noisy inputs. Among them, time-domain end-to-end approaches operate directly on raw waveforms, jointly optimizing both magnitude and phase to reduce spectral decomposition artifacts [8, 12]. TF-domain mapping methods, on the other hand, estimate clean spectral magnitudes while reusing the noisy phase, allowing a simpler structure [13].

Despite their strengths, both paradigms face limitations in extreme noise scenarios. Masking-based models suffer from irreversible spectral distortions when noise dominates speech energy [6], while time-domain networks struggle to separate overlapping speech and noise temporal features [12]. These limitations stem from their reliance on acoustic microphone signals, which inherently capture both speech and noise. Body-conduction microphones, which detect speech-induced vibrations through body tissues, provide inherent noise resistance by avoiding airborne sound transmission [14, 15, 16, 17].

However, body-conducted speech experiences severe high-frequency attenuation above 2 kHz due to tissue filtering, resulting in muffled and less intelligible speech [18, 19]. Although mapping-based enhancement techniques [8, 9, 12] can partially restore high-frequency content, the irreversible spectral loss reduces their effectiveness. Consequently, existing solutions struggle to perform consistently across the full range of noise conditions, from moderate to extreme noise environments, underscoring the need for a multi-modal approach that dynamically adapts to varying noise environments encountered in real-world scenarios [20].

Recent multi-modal frameworks have demonstrated promising performance by leveraging the complementary strengths of body-conduction microphone signals (BMS) and acoustic microphone signals (AMS). Yu et al.[21] proposed a convolutional neural network (CNN) that uses BMS as an auxiliary modality to enhance AMS. Wang et al.[22] introduced an attention-based fusion mechanism within a convolutional recurrent network, enabling efficient speech enhancement in the complex spectral domain. Wang et al.[23] reported an attention-based fusion approach with a semi-supervised CycleGAN framework, effectively utilizing limited paired AMS and BMS data. Unlike prior approaches that rely on implicit fusion of BMS and AMS through concatenation or attention-based fusion, we emphasize the need for modality-specific processing to achieve effective spectral restoration and noise suppression.

To achieve this, we introduce the Body-Acoustic Fusion Network (BAF-Net), a multi-modal framework designed to leverage the unique strengths of BMS and AMS through task-specific architectures and a noise-adaptive fusion. Unlike previous approaches, the BAF-Net uses two specialized networks: a mapping-based SE-conformer [12] to restore high-frequency components in BMS and a masking-based DCCRN [2] to suppress noise in AMS. Critically, the BAF-Net introduces a mask-guided fusion mechanism that dynamically adjusts modality contributions based on noise conditions inferred from the DC-CRN's output, eliminating the need for explicit signal-to-noise ratio (SNR) estimation. This design allows optimal exploitation of BMS's inherent noise robustness in high-noise regions and AMS's spectral richness in low-noise conditions.

The framework's principal innovations include: (1) task-specialized enhancement, where separate networks explicitly handle BMS's high-frequency reconstruction and AMS's noise suppression, and (2) adaptive signal-level fusion, which utilizes fusion coefficients estimated by a lightweight CNN from the noise mask, enabling seamless modality switching. To evaluate our model, we used a simulated dataset built from the TAPS corpus [24] and DNS-2023 noise data [25]. Our results demonstrate that BAF-Net consistently outperforms single-modal baselines in terms of objective speech quality metrics.

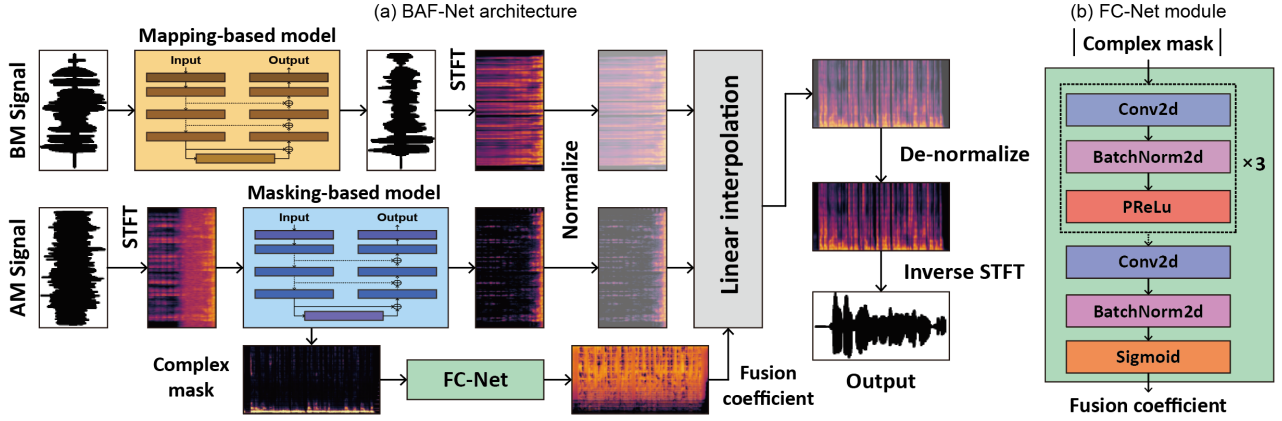


Figure 1: (a) Network structure of the proposed BAF-Net. (b) FC-Net module for the fusion coefficient estimation.

2. BAF-Net

2.1. Problem Settings

The BAF-Net adaptively combines BMS and AMS using a noise-dependent fusion strategy. The fusion strategy is designed to select the enhanced BMS in regions dominated by severe noise, while leveraging denoised AMS in relatively cleaner regions. This allows the framework to exploit the inherent robustness of BMS against ambient noise and the rich spectral content provided by AMS in noise-free or low-noise conditions. The framework comprises two specialized networks:

1. Mapping-based model (f_{map}): Enhances raw BMS x_{bm} by reconstructing high-frequency components lost due to tissue filtering, yielding the spectrogram $\hat{X}_{bm} \in \mathbb{C}^{T \times F}$.

$$\hat{X}_{bm} = \text{STFT}(\hat{x}_{bm}), \quad \hat{x}_{bm} = f_{map}(x_{bm}) \quad (1)$$

2. Masking-based model (f_{mask}): Denoises AMS by estimating a complex ideal ratio mask (cIRM) $M \in \mathbb{C}^{T \times F}$. The denoised AMS spectrogram \hat{X}_{am} is obtained by element-wise multiplication with the noisy AMS spectrogram X_{am} :

$$\hat{X}_{am} = X_{am} \odot M, \quad M = f_{mask}(X_{am}). \quad (2)$$

The enhanced spectrograms from both modalities, \hat{X}_{am} and \hat{X}_{bm} , are then combined using a linear interpolation scheme:

$$\hat{X} = \alpha \cdot \hat{X}_{am} + (1 - \alpha) \cdot \hat{X}_{bm}, \quad (3)$$

where the fusion coefficient $\alpha \in (0, 1)^{T \times F}$ adaptively balances the enhanced BMS and the denoised AMS. The coefficient is estimated by the fusion coefficient network (FC-Net) f_F using the cIRM magnitude $|M|$, which reflects the local SNR:

$$\alpha = f_F(|M|). \quad (4)$$

In strong noise regions, the FC-Net predicts $\alpha \rightarrow 0$, thereby favoring the enhanced BMS. Conversely, in lower noise regions, the FC-Net predicts $\alpha \rightarrow 1$, favoring the denoised AMS. This adaptive fusion approach ensures that the optimal modality is dynamically selected based on the instantaneous noise conditions at each TF bin.

To train the model, multi-resolution STFT [26] and L1 losses are jointly minimized:

$$\mathcal{L} = \mathbb{E}[\|\hat{x} - x\|_1] + \sum_{k=1}^K \mathbb{E}[\| |\text{STFT}_k(\hat{x})| - |\text{STFT}_k(x)| \|_1], \quad (5)$$

where STFT_k computes spectral magnitudes with varying window/hop sizes. The L1 term minimizes waveform differences, while the multi-STFT loss matches spectral content across resolutions.

2.2. BAF-Net architecture

The architecture operates in the TF domain and consists of three core components: modality-specific processing networks for enhancing BMS and AMS, an FC-Net for noise-adaptive signal fusion, and a synthesis stage that reconstructs the final waveform. Fig. 1a illustrates the BAF-Net architecture.

The BAF-Net initially processes BMS and AMS independently. An SE-conformer [12] is used for BMS enhancement, while a DCCRN [2] is used for AMS denoising. The SE-conformer restores high-frequency information in BMS, and the DCCRN separates speech from noisy AMS, producing \hat{X}_{bm} and \hat{X}_{am} , respectively. FC-Net estimates the fusion coefficient α using $|M|$. Before fusion, \hat{X}_{bm} and \hat{X}_{am} are normalized using their respective mean energies $E(\hat{X}_{bm})$ and $E(\hat{X}_{am})$:

$$E(X) = \frac{1}{T \cdot F} \sum_{t=1}^T \sum_{f=1}^F \left(\text{Re}(X(t, f))^2 + \text{Im}(X(t, f))^2 \right), \quad (6)$$

$$\tilde{X}_{bm} = \frac{\hat{X}_{bm}}{\sqrt{E(\hat{X}_{bm})}}, \quad \tilde{X}_{am} = \frac{\hat{X}_{am}}{\sqrt{E(\hat{X}_{am})}}. \quad (7)$$

This normalization ensures that the fusion coefficient α balances the contributions of BMS and AMS based on their relative energy levels. The fusion coefficient α is then used to interpolate the normalized spectrograms linearly:

$$\tilde{X} = \alpha \cdot \tilde{X}_{am} + (1 - \alpha) \cdot \tilde{X}_{bm}. \quad (8)$$

Table 1: *Speech enhancement performance on the TAPS dataset at various SNR levels, evaluated via wideband PESQ, STOI, and CER.*

Input	Model	PESQ					STOI					CER				
		-20 dB	-10 dB	0 dB	10 dB	15 dB	-20 dB	-10 dB	0 dB	10 dB	15 dB	-20 dB	-10 dB	0 dB	10 dB	15 dB
BMS	SE-conformer [12]	—	—	1.971	—	—	—	0.892	—	—	—	—	0.244	—	—	—
Noisy AMS	SE-conformer [12]	1.116	1.305	1.706	2.109	2.240	0.523	0.715	0.839	0.893	0.900	0.886	0.736	0.465	0.308	0.275
	AMSDCCRN [2]	1.114	1.259	1.610	2.062	2.289	0.498	0.656	0.768	0.811	0.822	0.861	0.699	0.440	0.261	0.243
Noisy AMS and BMS	FCN-LF [21]	1.053	1.063	1.159	1.338	1.421	0.326	0.462	0.606	0.685	0.696	0.919	0.842	0.610	0.401	0.351
	FCN-EF [21]	1.228	1.345	1.413	1.464	1.482	0.775	0.809	0.828	0.842	0.845	0.578	0.483	0.430	0.390	0.381
	Attention-DC-CRN-LF [22]	1.448	1.590	1.701	1.758	1.775	0.822	0.840	0.856	0.866	0.869	0.361	0.323	0.282	0.254	0.243
	Attention-DC-CRN-EF [22]	1.546	1.789	2.000	2.127	2.156	0.840	0.871	0.891	0.902	0.905	0.386	0.290	0.225	0.197	0.185
	BAF-Net	2.082	2.260	2.491	2.751	2.875	0.904	0.914	0.929	0.944	0.949	0.222	0.206	0.187	0.170	0.167

The fused output \tilde{X} is then re-scaled using the average root energy of \hat{X}_{bm} and \hat{X}_{am} to restore the overall energy level:

$$\hat{X} = \tilde{X} \cdot \frac{\sqrt{E(\hat{X}_{bm})} + \sqrt{E(\hat{X}_{am})}}{2}. \quad (9)$$

Finally, the enhanced speech \hat{x} is synthesized by applying inverse STFT to \hat{X} :

$$\hat{x} = \text{iSTFT}(\hat{X}). \quad (10)$$

For the mapping-based enhancement, we utilized the SE-conformer, an encoder-decoder architecture that combines convolutional neural networks with a Conformer-based sequence modeling network [27]. We adopted a variant of the SE-conformer [24], optimized for BMS enhancement. For the masking-based denoising, we employed DCCRN, a complex-valued encoder-decoder network. DCCRN operates directly on complex spectrograms, enabling joint estimation of magnitude and phase through complex-valued operations. As illustrated in Fig. 1b, the FC-Net adopts a convolutional architecture composed of three sequential blocks, each containing a Conv2D layer (16 channels, kernel size 7×7), Batch Normalization, and PReLU activation. The final layer replaces PReLU with a Sigmoid activation to constrain the fusion coefficient α to the range $[0, 1]$.

3. Experiments

3.1. Experiment on the TAPS dataset

We evaluated the BAF-Net using the Throat-Acoustic Pairing Speech (TAPS) dataset [24], which includes paired body-conduction microphone and acoustic microphone speech recordings. The dataset was divided into 4,000 training, 1,000 validation, and 1,000 evaluation pairs. To simulate noisy and reverberant conditions, we augmented the data with noise clips and room impulse responses (RIRs) from the DNS-2023 challenge dataset [25]. DNS-2023 provides 65,000 noise clips (150 classes) and 60,000 RIRs for simulating diverse acoustic environments. The noise clips and RIRs used in the test set were not included in the training set, ensuring that the evaluation accurately reflects the model’s performance in unseen acoustic conditions.

For training, noisy AMS was synthesized by mixing clean AMS with a subset of DNS noise clips at SNRs ranging from -15 dB to 20 dB. To simulate reverberation, 75% of training samples were convolved with a randomly selected subset of DNS RIRs, while the remaining 25% retained anechoic conditions. Each training sample consisted of a triplet: BMS, noisy

AMS, and clean AMS. For evaluation, we created a test set using 2,000 noise clips and 1,000 RIRs from DNS-2023 that were entirely excluded from training. Each evaluation sample was generated at five SNR levels: -20 dB, -10 dB, 0 dB, 10 dB, and 15 dB, with reverberation applied to 50% of the samples. This resulted in 10,000 test utterances (2,000 noise clips \times 5 SNRs, split evenly between reverberant and anechoic conditions).

3.2. Training setup and baselines

For comparison, several single-modal baseline models were trained:

- **DCCRN** [2]: A masking-based model trained on noisy AMS.
- **SE-conformer** [12]: A mapping-based model with two variants: one trained on noisy AMS and the other trained on BMS.

We also compared our BAF-Net with the following multi-modal baselines in previous studies:

- **FCN** [21]: A time-domain Fully Convolutional Network (FCN) utilizing two fusion strategies:
 - *Early Fusion (EF)*: Concatenates BMS and AMS at the input stage and processes the concatenated signal with the FCN.
 - *Late Fusion (LF)*: Processes BMS and AMS independently through separate FCNs and subsequently processes the concatenated output with another FCN.
- **Attention-DC-CRN** [22]: A complex-domain Densely Connected Convolutional Recurrent Network (DC-CRN) employing an attention mechanism in two fusion strategies:
 - *Early Fusion (EF)*: Integrates AMS and BMS spectrograms at the input stage using attention mechanisms, followed by the DC-CRN module.
 - *Late Fusion (LF)*: Processes each modality through separate DC-CRN modules and combines the outputs with an attention-based fusion.

All the models, with the exception of Attention-DC-CRN, were trained using a 25 ms window length, 6.25 ms hop size, and 512-point FFT for spectrogram extraction. Attention-DC-CRN was trained using a 32 ms window length, 16 ms hop size, and 512-point FFT for spectrogram extraction. The Adam optimizer was used with a learning rate of 3×10^{-4} and momentum parameters $\beta_1 = 0.9$ and $\beta_2 = 0.99$.

For the BAF-Net, the mapping and masking models were pre-trained separately for 200 epochs. The pre-trained weights were then transferred into the BAF-Net, and the full model was fine-tuned for another 200 epochs. All models, includ-




Figure 2: Spectrogram analysis of the enhancement process. (a) Body-conduction microphone signal, (b) noisy acoustic microphone signal, (c) enhanced body-conduction microphone signal, (d) denoised acoustic microphone signal, (e) final output, and (f) clean signal.

ing baselines, employed a combined multi-resolution STFT loss ($K = 3$) and L1 loss, with multi-STFT configurations set to $\{512, 50, 240\}$, $\{1024, 120, 600\}$, and $\{2048, 240, 1200\}$ (FFT size, window size, and frame shift). Early stopping based on the validation set ensured a fair comparison across architectures. The model architecture and training code for BAF-Net, as well as baseline models, are available at the public repository¹.

4. Experimental results and discussion

The BAF-Net was evaluated using the simulated TAPS dataset and three metrics. First, we used the wideband Perceptual Evaluation of Speech Quality (PESQ) [28], which ranges from -0.5 to 4.5 . Second, we measured the Short-Time Objective Intelligibility (STOI) [29], which ranges from 0 to 1. Finally, to quantify the recovery of speech information, we assessed the Character Error Rate (CER) using the Whisper-large-v3-turbo automatic speech recognition (ASR) model [30], which had been fine-tuned² on the Korean Zeroth dataset³. Enhanced speech samples were transcribed using the ASR model, and CER was computed by comparing transcriptions with ground-truth labels.

As shown in Table 1, at low SNRs (-20 dB to 0 dB), single-modal BMS-based models outperformed AMS-based models, as noise suppression in AMS inevitably removes speech components. Conversely, AMS-based models achieved superior results at high SNRs (10 and 15 dB), likely due to the spectral limitations of BMS. The BAF-Net outperformed single-modal models across all noise levels, demonstrating the effectiveness of its adaptive fusion in leveraging the complementary strengths of both modalities. In addition, the BAF-Net exhibited clear advantages compared to other multi-modal baselines. At extremely low SNR conditions (-20 and -10 dB), other multi-modal approaches showed limited improvements due to severe noise contamination in AMS. In contrast, BAF-Net leveraged BMS’s inherent noise resistance, resulting in significant performance gains. At higher SNRs (10 and 15 dB), the adaptive fusion mechanism utilized AMS’s spectral richness, leading to notable enhancements in performance.

Spectrogram analysis of the enhanced speech demonstrates BAF-Net’s ability to effectively utilize modality-specific strengths. Fig. 2a displays the spectrogram of the BMS input. For a noisy AMS containing white noise in the low frequency region (Fig. 2b), the mapping-based model restored high-frequency components in BMS (Fig. 2c). In contrast, the masking-based model effectively suppressed noise in AMS (Fig. 2d). The fused output (Fig. 2e) seamlessly integrates enhanced BMS and denoised AMS through adaptive fusion. In the noise-free region (Fig. 2, yellow dashed boxes), the denoised AMS retained nearly noise-free characteristics, causing the fusion mechanism to prioritize AMS. As a result, the fu-




Figure 3: Visualization of the fusion mechanism in the BAF-Net. (a, b) Noisy acoustic microphone signal inputs, (c, d) complex mask magnitudes estimated by the masking-based model, and (e, f) fusion coefficients derived from the complex mask magnitudes at SNR levels of 10 dB and -20 dB, respectively.

sion mechanism produced an output close to the clean reference. Conversely, the fusion mechanism favored the enhanced BMS in the noise-dominant region (Fig. 2, white boxes).

We analyzed the cIRM magnitude and fusion coefficient at 10 dB and -20 dB to verify the adaptive fusion. The noisy AMS spectrograms are shown in Fig. 3a and Fig. 3b. At 10 dB, the mask selectively attenuated noise-dominant regions (Fig. 3c), while the fusion coefficient prioritized AMS in clean areas and BMS in noisy areas (Fig. 3e). At -20 dB, the mask suppressed noise across broader regions (Fig. 3d), and the fusion coefficient depended primarily on BMS except in high-frequency bins (Fig. 3f). This dynamic fusion confirms the BAF-Net’s ability to compensate for information loss in noisy regions using BMS while preserving AMS fidelity in cleaner areas.

5. Conclusion

In conclusion, this study presents the BAF-Net, a multi-modal framework that combines BMS and AMS using task-specific enhancement and adaptive fusion. It restores high-frequency components in BMS using a mapping-based network and reduces noise in AMS using a masking-based model. This approach overcomes the weaknesses of each modality while maximizing their strengths. The BAF-Net employs a simple yet effective fusion method, utilizing enhanced BMS for noisy segments and denoised AMS for clearer parts. Experimental results across various SNR levels show that the BAF-Net surpasses previous single- and multi-modal models, highlighting the importance of modality-specific processing and adaptive fusion.

¹<https://github.com/yskim3271/BAF-Net>

²<https://huggingface.co/ghost613/whisper-large-v3-turbo-korean>

³<https://openslr.org/40>

6. Acknowledgements

This research was supported by the National Research Foundation of Korea grants (RS-2025-00516311, RS-2020-NRF047144) and by the Institute of Information & Communications Technology Planning & Evaluation grant (RS-2019-III191906, Artificial Intelligence Graduate School Program).

7. References

- [1] A. R. Yuliani, M. F. Amri, E. Suryawati, A. Ramdan, and H. F. Pardede, "Speech enhancement using deep learning methods: A review," *Jurnal Elektronik dan Telekomunikasi*, vol. 21, no. 1, pp. 19–26, 2021.
- [2] Y. Hu, Y. Liu, S. Lv, M. Xing, S. Zhang, Y. Fu, J. Wu, B. Zhang, and L. Xie, "Dccrn: Deep complex convolution recurrent network for phase-aware speech enhancement," *arXiv preprint arXiv:2008.00264*, 2020.
- [3] S. Lv, Y. Hu, S. Zhang, and L. Xie, "Dccrn+: Channel-wise sub-band dccrn with snr estimation for speech enhancement," *arXiv preprint arXiv:2106.08672*, 2021.
- [4] X. Hao, X. Su, R. Horaud, and X. Li, "Fullsubnet: A full-band and sub-band fusion model for real-time single-channel speech enhancement," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6633–6637.
- [5] J. Chen, Z. Wang, D. Tuo, Z. Wu, S. Kang, and H. Meng, "Fullsubnet+: Channel attention fullsubnet with complex spectrograms for speech enhancement," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7857–7861.
- [6] S.-W. Fu, C.-F. Liao, Y. Tsao, and S.-D. Lin, "Metricgan: Generative adversarial networks based black-box metric scores optimization for speech enhancement," in *International Conference on Machine Learning (ICML)*, 2019.
- [7] S.-W. Fu, C. Yu, T.-A. Hsieh, P. Plantinga, M. Ravanelli, X. Lu, and Y. Tsao, "Metricgan+: An improved version of metricgan for speech enhancement," *arXiv preprint arXiv:2104.03538*, 2021.
- [8] A. Defossez, G. Synnaeve, and Y. Adi, "Real time speech enhancement in the waveform domain," in *Interspeech*, 2020.
- [9] D. Rethage, J. Pons, and X. Serra, "A wavenet for speech denoising," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5069–5073.
- [10] T.-A. Hsieh, H.-M. Wang, X. Lu, and Y. Tsao, "Wavecrn: An efficient convolutional recurrent neural network for end-to-end speech enhancement," *IEEE Signal Processing Letters*, vol. 27, pp. 2149–2153, 2020.
- [11] K. Wang, B. He, and W.-P. Zhu, "Tstnn: Two-stage transformer based neural network for speech enhancement in the time domain," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7098–7102.
- [12] E. Kim and H. Seo, "Se-conformer: Time-domain speech enhancement using conformer," in *Interspeech*, 2021, pp. 2736–2740.
- [13] K. Tan and D. Wang, "Learning complex spectral mapping with gated convolutional recurrent networks for monaural speech enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 380–390, 2019.
- [14] D. D. Mehta, J. H. Van Stan, M. Zaňartu, M. Ghassemi, J. V. Guttag, V. M. Espinoza, J. P. Cortés, H. A. Cheyne, and R. E. Hillman, "Using ambulatory voice monitoring to investigate common voice disorders: Research update," *Frontiers in bioengineering and biotechnology*, vol. 3, p. 155, 2015.
- [15] Y. Zheng, Z. Liu, Z. Zhang, M. Sinclair, J. Droppo, L. Deng, A. Acero, and X. Huang, "Air-and bone-conductive integrated microphones for robust speech detection and enhancement," in *2003 IEEE Workshop on Automatic Speech Recognition and Understanding (IEEE Cat. No. 03EX721)*. IEEE, 2003, pp. 249–254.
- [16] T. F. Quatieri, K. Brady, D. Messing, J. P. Campbell, W. M. Campbell, M. S. Brandstein, C. J. Weinstein, J. D. Tardelli, and P. D. Gatewood, "Exploiting nonacoustic sensors for speech encoding," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 14, no. 2, pp. 533–544, 2006.
- [17] S. Kang, S. Cho, R. Shanker, H. Lee, J. Park, D.-S. Um, Y. Lee, and H. Ko, "Transparent and conductive nanomembranes with orthogonal silver nanowire arrays for skin-attachable loudspeakers and microphones," *Science advances*, vol. 4, no. 8, p. eaas8772, 2018.
- [18] H. S. Shin, H.-G. Kang, and T. Fingscheidt, "Survey of speech enhancement supported by a bone conduction microphone," in *Speech Communication; 10. ITG Symposium*. VDE, 2012, pp. 1–4.
- [19] P. K. Tran, T. R. Letowski, and M. E. McBride, "The effect of bone conduction microphone placement on intensity and spectrum of transmitted speech items," *The Journal of the Acoustical Society of America*, vol. 133, no. 6, pp. 3900–3908, 2013.
- [20] M. Wang, J. Chen, X. Zhang, Z. Huang, and S. Rahardja, "Multi-modal speech enhancement with bone-conducted speech in time domain," *Applied Acoustics*, vol. 200, p. 109058, 2022.
- [21] C. Yu, K.-H. Hung, S.-S. Wang, Y. Tsao, and J.-w. Hung, "Time-domain multi-modal bone/air conducted speech enhancement," *IEEE Signal Processing Letters*, vol. 27, pp. 1035–1039, 2020.
- [22] H. Wang, X. Zhang, and D. Wang, "Attention-based fusion for bone-conducted and air-conducted speech enhancement in the complex domain," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7757–7761.
- [23] —, "Fusing bone-conduction and air-conduction sensors for complex-domain speech enhancement," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 3134–3143, 2022.
- [24] Y. Kim, Y. Song, and Y. Chung, "Taps: Throat and acoustic paired speech dataset for deep learning-based speech enhancement," *arXiv preprint arXiv:2502.11478*, 2025.
- [25] H. Dubey, A. Aazami, V. Gopal, B. Naderi, S. Braun, R. Cutler, H. Gamper, M. Golestaneh, and R. Aichner, "Icassp 2023 deep noise suppression challenge," in *ICASSP*, 2023.
- [26] R. Yamamoto, E. Song, and J.-M. Kim, "Parallel wavegan: A fast waveform generation model based on generative adversarial networks with multi-resolution spectrogram," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 6199–6203.
- [27] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *arXiv preprint arXiv:2005.08100*, 2020.
- [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 *2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221)*, vol. 2. IEEE, 2001, pp. 749–752.
- [29] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. Jensen, "A short-time objective intelligibility measure for time-frequency weighted noisy speech," in *2010 IEEE international conference on acoustics, speech and signal processing*. IEEE, 2010, pp. 4214–4217.
- [30] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," *arXiv preprint arXiv:2212.04356*, 2022.