

Plug-and-Steer: Decoupling Separation and Selection in Audio-Visual Target Speaker Extraction

Doyeop Kwak^{1,*}, Suyeon Lee^{1,*}, Joon Son Chung¹

¹ Korea Advanced Institute of Science and Technology, South Korea

dobbyk@kaist.ac.kr, syl14356@kaist.ac.kr, joonson@kaist.ac.kr

Abstract

The goal of this paper is to provide a new perspective on audio-visual target speaker extraction (AV-TSE) by decoupling separation and target selection. Conventional AV-TSE systems typically integrate audio and visual features deeply to re-learn the entire separation process, which can act as a fidelity ceiling due to the noisy nature of in-the-wild audio-visual datasets. To address this, we propose Plug-and-Steer, which assigns high-fidelity separation to a frozen audio-only backbone and limits the role of the visual modality strictly to target selection. We introduce the Latent Steering Matrix (LSM), a minimalist linear transformation that re-routes latent features within the backbone to anchor the target speaker to a designated channel. Experiments across four representative architectures show that our method effectively preserves the acoustic priors of diverse backbones, achieving perceptual quality comparable to that of the original backbones. Audio samples are on the demo page.¹

Index Terms: audio-visual target speaker extraction

1. Introduction

Speech separation is the task of isolating individual voices from a recording where multiple people are speaking simultaneously—a challenge often referred to as the “Cocktail Party Problem”. In recent years, the performance of Audio-only Speech Separation (AOSS) models has reached an impressive turning point with advanced architectures [1–4]. These systems can now separate overlapping speech with remarkable clarity, often rivaling or even exceeding theoretical limits of acoustic quality [5]. However, despite these milestones, audio-only models face a fundamental hurdle in real-world deployment: permutation ambiguity. They remain blind to the target’s identity, unable to automatically determine which output channel contains the voice the user intends to hear.

To resolve this ambiguity, conventional Audio-Visual Target Speaker Extraction (AV-TSE) architectures typically integrate audio and visual features deeply within the model via cross-attention [6–9], concatenation [10–12], or tailored architectures for deep fusion [13–16]. This design paradigm assumes that effective extraction requires re-learning the entire separation process to leverage visual cues for both (1) identifying the target and (2) refining the separation quality through lip-sync information. While this joint optimization can be beneficial, we question whether using visual cues to further improve acoustic separation—a task AOSS models already perform with significant proficiency—always yields a net gain. Large-scale audio-visual datasets collected in the wild, such as LRS2 [17]

and VoxCeleb2 [18], often contain intrinsic noise and reverberation, providing sub-optimal supervision for high-fidelity extraction [19]. In such cases, full-parameter training on noisy audio-visual data can act as a fidelity ceiling, potentially limiting the final output quality compared to what is achievable by pure AOSS models trained on studio-quality corpora.

In this work, we propose a shift in perspective: **Plug-and-Steer**. Instead of pursuing entangled multi-modal fusion to achieve separation and selection simultaneously, we decouple these roles to leverage the strengths of pre-trained acoustic engines. We assign high-fidelity separation entirely to a frozen AOSS backbone and limit the visual modality’s role strictly to target selection. This approach is motivated by our observation that frozen AOSS models appear to possess a latent structure where output channels can be reordered at the feature level within a single separator block. We demonstrate that this requires nothing more than a minimalist $C \times C$ linear transformation, termed a Latent Steering Matrix (LSM).

By treating the LSM as a “steering wheel,” we utilize a lightweight visual steering module to anchor the target speaker to a designated output channel. Unlike post-hoc selection—a straightforward strategy that identifies the target from fully decoded outputs—our mechanism is embedded directly within the feature flow. By reusing fine-grained latent features, Plug-and-Steer eliminates the redundant re-encoding or synchronization stages required by external selection pipelines. Furthermore, the LSM bridges the internal feature flow, allowing the routing logic to be optimized directly through signal-level reconstruction losses. This enables a direct gradient flow that ensures more stable target selection than detached classification.

Ultimately, this work serves as a proof-of-concept demonstrating that well-established AOSS backbones can be transformed into a target extraction system at minimal cost. By offering a scalable blueprint for this adaptation, our framework allows extraction performance to scale naturally alongside the rapid evolution of underlying separation engines. Our primary contributions are summarized as follows:

- We analyze the latent structural properties of diverse AOSS architectures, showing that speaker identity is permutable via a simple linear transformation at the feature level.
- We propose Plug-and-Steer, a framework decoupling separation from selection to preserve the high-fidelity acoustic priors of frozen AOSS models.
- We demonstrate that our internal steering mechanism achieves comparable selection accuracy and better computational efficiency than post-hoc selection strategies.
- Experiments across four representative architectures validate that our method transforms diverse AOSS backbones into AV-TSE systems while preserving their perceptual quality.

*These authors contributed equally.

¹<https://plugandsteer.github.io>




Figure 1: Training process of the latent steering matrix W for the i -th separator block. \hat{s}_1 and \hat{s}_2 are the outputs from the first and second channels of the pre-trained audio-only backbone.

2. Method

Our methodology decouples separation from selection by treating the frozen AOSS backbone as a high-fidelity engine steered by visual cues. We first investigate whether a single linear transformation can re-route latent features within the backbone. We then leverage this mechanism to train a visual steering module that anchors the target speaker to a fixed output channel.

2.1. Latent Steering Matrix (LSM): A steering wheel

We introduce a simple feature-level re-routing strategy to control output channel permutations of the AOSS backbone within the latent space. Given an intermediate audio feature from the i -th separator block $f_i \in \mathbb{R}^{C \times T_a}$, we assume that speaker identity reordering can be approximated by a linear transformation $W \in \mathbb{R}^{C \times C}$. We term W the Latent Steering Matrix (LSM), and apply it as a residual transformation as follows:

$$f'_i = (I + g \cdot W)f_i, \quad (1)$$

where I is the identity matrix and $g \in \{0, 1\}$ is a binary gate. When the gate is inactive ($g=0$), the features remain unchanged. Conversely, an active gate ($g=1$) triggers W to induce a latent speaker swap, effectively permuting the output channels.

2.2. Training LSM

To learn the latent transformation required for speaker re-routing, we first train the LSM under a forced-swap condition by fixing $g = 1$ within a frozen AOSS model. Specifically, for a 2-speaker mixture where the pre-trained model outputs (\hat{s}_1, \hat{s}_2) , we train W to manipulate the audio feature f_i to produce the steered output: (\hat{s}_2, \hat{s}_1) . As shown in Fig. 1, the channel-swapped predictions of the pre-trained backbone serve as the training targets. The training objective is the sum of the negative Scale-Invariant Signal-to-Noise Ratio (SI-SNR) [20] across both channels, computed between the steered outputs and the permuted reference signals.

2.3. Visual steering module: Learning to steer

We adapt the AOSS model for AV-TSE by consistently routing the target speaker to a designated output channel, regardless of the backbone’s initial output order. To this end, we learn a gate value to control the LSM based on visual cues. As shown in Fig. 2, we design a lightweight visual steering module that predicts a frame-wise gate value $g_t \in [0, 1]$ based on the latent audio features f_i and the visual embedding $v \in \mathbb{R}^{T_v \times C_v}$ extracted from the target’s lip motion.

The temporal dimension of v is linearly interpolated from T_v to T_a to match the resolution of the audio feature f_i . We then concatenate the two features along the channel dimension. The joint feature is processed through a lightweight modified Temporal Convolutional Network (TCN) adapted from [21], which consists of two blocks, each containing three convolutional layers, followed by a sigmoid-activated gate head. The




Figure 2: Training process of the AV-TSE module by learning the frame-wise gate value g_t . The visual steering module predicts the value g_t to control the degree of steering.

output value g_t dynamically modulates the steering operation based on Eq. 1. For backbones utilizing 2D latent spaces such as frequency-time grids, we first project the channel dimension to a reduced space C_r , and flatten the non-temporal dimensions into a 1D feature representation per time step.

We derive pseudo-labels \bar{g} to supervise the gate by comparing the frozen AOSS model’s output permutations against the desired target order. Let the output of the AO backbone be $\hat{\mathbf{s}} = [\hat{s}_1, \hat{s}_2]^\top$, and the reference signals be $\mathbf{s}^{\text{ref}} = [s^{\text{tgt}}, s^{\text{tiff}}]^\top$ for the target speech s^{tgt} and interference s^{tiff} . Defining $\mathbf{P}_g \in \mathbb{R}^{2 \times 2}$ as the permutation matrix indexed by $g \in \{0, 1\}$ (identity for $g = 0$, swap for $g = 1$), the pseudo-label \bar{g} is defined as:

$$\bar{g} = \arg \max_{g \in \{0, 1\}} \sum_{i=1}^2 \text{SI-SNR}([\mathbf{P}_g \hat{\mathbf{s}}]_i, s_i^{\text{ref}}), \quad (2)$$

and used for the frame-wise binary cross-entropy loss \mathcal{L}_{BCE} . The steered output $\hat{\mathbf{s}}^{\text{LSM}}$ is obtained by applying LSM with predicted g_t . The SI-SNR loss $\mathcal{L}_{\text{SI-SNR}}$ is defined as the negative total SI-SNR between $\hat{\mathbf{s}}^{\text{LSM}}$ and the reference \mathbf{s} . The overall training objective for the visual steering module is given by:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{BCE}}(g_t, \bar{g}) + \lambda \mathcal{L}_{\text{SI-SNR}}(\hat{\mathbf{s}}^{\text{LSM}}, \mathbf{s}). \quad (3)$$

During this phase, only the visual steering module is updated, while the pre-trained backbone and the LSM remain frozen.

3. Experimental Setup

3.1. Datasets

We conduct our AV-TSE experiments on LRS2-2mix [11], a two-speaker speech separation benchmark partitioned into 20k (~ 23 h), 5k, and 3k samples for training, validation, and testing, respectively. Mixtures are synthesized by mixing pairs of utterances from different speakers in LRS2 [17] with random SNRs in $[-5, 5]$ dB. To analyze the effect of acoustic priors, we compare two pre-training configurations: a clean setup using the studio-quality Libri2Mix train-100 subset [22] (~ 58 h) and an in-the-wild, noisier setup using the LRS2-2mix train set. Audio signals are sampled at 16 kHz in mono format and randomly truncated to 3 s during training with variance normalization. Visual inputs are 25 FPS grayscale sequences, obtained by center-cropping the original 224×224 LRS2 frames to 112×112 .

3.2. Training and evaluation details

Baseline models. We conduct our analysis across four representative speech separation models: Conv-TasNet [21], DPRNN [23], TF-GridNet [1], and MossFormer2 [2]. To compare our approach against the established AV-TSE baselines, we use AV-ConvTasNet [24], AV-DPRNN [10], AV-TFGridNet [25], and AV-MossFormer2 [11] with open-sourced




Figure 3: Layer-wise performance preservation rate (%) across different AOSS backbones.

weights pre-trained on LRS2-2mix² for evaluation. Our visual encoder follows the same architecture as the AV-TSE baselines and is initialized with lip-reading pre-trained weights.

Metrics & evaluation. We use the Scale-Invariant Signal-to-Distortion Ratio improvement (SI-SDRi) [20] as a standard metric in speech separation, and DNSMOS [26] and NISQA [27] as non-intrusive metrics to assess the perceptual quality of audio. For the audio-only backbones, we follow the Permutation Invariant Training (PIT) protocol [28], selecting the output-target assignment that maximizes the SI-SDR.

Inference. We apply a threshold $\tau = 0.5$ to the averaged gate value g during inference. For AV-TSE methods, we sequentially extract each speaker’s speech with the corresponding video.

Training configurations. AOSS models are pre-trained using Adam [29] with an initial learning rate of 5×10^{-4} . The learning rate is halved every 5 plateau epochs, with training terminating below 10^{-6} . The total batch size is 4, and the training objective is $\mathcal{L}_{\text{SI-SNR}}$. After pre-training audio backbones, we optimize our Plug-and-Steer framework with a cosine annealing learning rate scheduler. With the AOSS backbone frozen, training is conducted in two stages: (1) LSM is trained for 10k steps without a warmup period; and (2) the visual steering module is trained for 100k steps, including a 1k-step warmup. C_r for DPRNN and TF-GridNet is set to 16, and λ in Eq. 3 is set to 0.1. All experiments, including inference, were conducted on NVIDIA RTX 4090 GPUs with an AMD EPYC 7543 CPU.

3.3. Alternative adaptation strategies for AV-TSE

To evaluate the proposed steering mechanism, we compare it against two alternative strategies that adapt a pre-trained AOSS backbone for target extraction via residual connections. These variants replace the gate head with a visual adapter that transforms v_{tcn} into a residual feature: $f'_i = f_i + \text{Adapter}(v_{\text{tcn}})$. Unlike our approach, which simply re-routes existing representations, this method refines the acoustic feature itself. We consider both full fine-tuning, where the entire backbone is unfrozen, and partial adaptation, where the backbone remains frozen. For full fine-tuning specifically, the backbone learning rate is set to 5×10^{-5} to mitigate catastrophic forgetting, while the adapter follows the standard training protocol.

4. Results and Analysis

4.1. Layer-wise performance preservation

To identify the optimal location for latent manipulation, we evaluate the performance preservation rate across all separator blocks $i \in [1, \dots, L]$. This rate is defined as the ratio of the SI-SDRi achieved after applying the LSM relative to the

²https://github.com/modelscope/ClearerVoice-Studio/tree/main/train/target_speaker_extraction

Method	AO ✳	# params.	SI-SDRi (dB)	DNSMOS	NISQA
<i>Libri2Mix GT</i>	-	-	-	3.16	3.93
<i>LRS2-2mix GT</i>	-	-	-	2.38	3.19
<i>AO-oriented</i>					
Conv-TasNet [21]		5.1M	7.12	2.35	2.31
+ Residual	✗	6.6M	11.72	2.31	2.57
+ Residual	✓	1.5M	9.37	2.13	2.27
+ LSM (Ours)	✓	1.5M	6.82	2.28	2.22
DPRNN [23]		2.6M	7.77	2.35	2.37
+ Residual	✗	5.2M	11.66	2.28	2.43
+ Residual	✓	2.5M	7.90	2.04	1.71
+ LSM (Ours)	✓	2.0M	7.74	2.35	2.36
TF-GridNet [1]		14.4M	14.81	2.80	4.32
+ Residual	✗	17.0M	13.38	2.36	3.17
+ Residual	✓	2.6M	13.09	2.56	3.20
+ LSM (Ours)	✓	2.0M	14.79	2.79	4.29
MossFormer2 [2]		55.7M	12.68	2.79	3.47
+ Residual	✗	57.3M	15.54	2.51	3.02
+ Residual	✓	1.6M	14.02	2.44	2.95
+ LSM (Ours)	✓	1.6M	12.65	2.79	3.47
<i>AV-TSE Baselines</i>					
AV-ConvTasNet [24]		10.3M	11.51	2.44	2.49
AV-DPRNN [10]		4.1M	11.97	2.40	2.58
AV-TFGridNet [25]		9.6M	15.10	2.51	3.53
AV-MossFormer2 [11]		57.3M	15.52	2.64	3.06

Table 1: Performance of AV-TSE methods on LRS2-2mix grouped by audio-only (AO) backbone, where the AO backbones are pre-trained on Libri2Mix. AO ✳ indicates whether the pre-trained AO backbone is frozen or not. # params. denotes the number of trainable parameters. Bold and underlined values represent global and backbone-specific best scores.

original AO performance. As illustrated in Fig. 3, the capacity for near-lossless swapping increases significantly with model depth. Applying the LSM at the final separator block yields the highest preservation rates across all architectures: 96.22% for Conv-TasNet, 99.67% for DPRNN, 99.91% for TF-GridNet, and 99.43% for MossFormer2. Notably, modern architectures such as TF-GridNet and MossFormer2 demonstrate a remarkable ability to maintain disentangled speaker features throughout the network, retaining 86.13% and 95.44% of their original performance even at their first separator blocks. These observations suggest that in advanced separation backbones, speaker identity is already distinctly encoded and manipulable from the early stages of processing, well before reaching the final output layers. Consequently, we utilize the final layer of each backbone for all subsequent AV-TSE experiments to ensure maximum fidelity.

4.2. AV-TSE performance

4.2.1. Analysis of acoustic prior preservation

We first conduct a controlled analysis to investigate how Plug-and-Steer preserves acoustic priors compared to conventional adaptation strategies, specifically full and partial residual fine-tuning. As shown in Tab. 1, while residual-based methods may improve intrusive metrics such as SI-SDRi by aligning features with the LRS2-2mix distribution, they often suffer from degraded perceptual quality (DNSMOS and NISQA). This suggests that such methods compromise high-fidelity priors to accommodate the noisier ground truth of audio-visual datasets. In contrast, our LSM approach achieves target selectivity while largely retaining the original AO characteristics. In the case of TF-GridNet, residual adaptation results in suboptimal performance, even falling below the original AO backbone, which likely stems from architectural mismatches between the 2D feature space and the TCN-based steering module. These results indicate that conventional adaptation may require sophisticated,

Method	AO ✳	SI-SDRi (dB)	DNSMOS	NISQA
MossFormer2 [2]		15.42	2.88	3.88
+ Residual	✗	17.11	2.53	3.28
+ Residual	✓	16.24	2.51	3.21
+ LSM (Ours)	✓	15.40	2.88	3.87
AV-MossFormer2 [11]		15.52	2.64	3.06

Table 2: AV-TSE performance on LRS2-2mix using a powerful backbone. The baseline (AV-MossFormer2) is trained solely on LRS2-2mix.

Method	AO ✳	SI-SDRi (dB)	DNSMOS	NISQA
Conv-TasNet [21]		11.50	2.30	2.32
+ Residual	✗	11.33	2.29	2.33
+ Residual	✓	11.28	2.28	2.31
+ LSM (Ours)	✓	11.21	2.26	2.25
DPRNN [23]		12.60	2.38	2.55
+ Residual	✗	12.59	2.34	2.51
+ Residual	✓	11.64	2.24	1.97
+ LSM (Ours)	✓	12.54	2.36	2.53
TF-GridNet [1]		15.90	2.52	3.61
+ Residual	✗	14.01	2.45	3.20
+ Residual	✓	14.93	2.47	3.33
+ LSM (Ours)	✓	15.87	2.52	3.60
MossFormer2 [2]		15.07	2.52	3.15
+ Residual	✗	15.54	2.51	3.02
+ Residual	✓	15.03	2.52	3.13
+ LSM (Ours)	✓	15.02	2.52	3.14

Table 3: In-domain AV-TSE performance on LRS2-2mix. Both AO and AV models are trained on the same data.

model-specific designs to bridge such structural gaps. Conversely, since the objective of our approach is to simply steer existing latent features toward the designated output channel, it eliminates the need for complex feature manipulation and provides a robust, architecture-agnostic solution.

4.2.2. Scaling with powerful AO backbones

The true potential of Plug-and-Steer is further evidenced when integrated with a more powerful backbone. We employ a MossFormer2 model pre-trained on a 107-hour high-fidelity corpus consisting of VCTK [30], LibriTTS [31], and internal TTS data³. As shown in Tab. 2, applying our framework to this high-performance engine yields SI-SDRi levels comparable to established AV-TSE baselines while maintaining significantly higher perceptual fidelity. Although conventional fine-tuning can achieve higher SI-SDRi, it causes perceptual quality to decline to levels typical of models trained solely on noisy audio-visual data. By protecting acoustic latents from noisy supervision, Plug-and-Steer shows that the extraction quality of the system scales directly with the strength of the AO backbone, achieving high-fidelity extraction that conventional adaptation struggles to maintain.

4.2.3. In-domain adaptation

Tab. 3 summarizes the in-domain adaptation performance on LRS2-2mix, where AO backbones are pre-trained and adapted within the same data distribution. In this scenario, the performance of the proposed LSM remains closely anchored to the original AO results, implying that the separation quality of the backbone acts as both a performance floor and an upper bound. This characteristic ensures that the inherent separation capabilities are effectively preserved, with the final extraction fidelity naturally defined by the pre-trained AO engine. Unlike conventional residual-based fine-tuning, which can be sensitive to

³https://huggingface.co/alibabasglab/MossFormer2_SS_16K

Routing	Method	Accuracy (%)		FLOPs	RTF
		1 epoch	Best epoch		
Internal	\mathcal{L}_{BCE}	50.00	99.73	256.82 G	0.147
	$\mathcal{L}_{SI-SNR} + \text{LSM}$	97.03	99.60		
	$\mathcal{L}_{BCE} + \mathcal{L}_{SI-SNR} + \text{LSM}$	99.01	99.93		
Post-hoc	LSE-C [32, 33]		99.80	308.56 G	0.209
	LSE-D [32, 33]		<u>99.83</u>		

Table 4: Internal routing vs. post-hoc selection. ‘Internal’ denotes our routing module with various loss combinations; ‘Post-hoc’ refers to SyncNet-score-based selection.

specific architectures and often necessitates specialized tuning to avoid degradation, our approach maintains consistent stability across all backbones with minimal parameter updates. Although LSM may not always yield the absolute peak performance, its structural simplicity and minimal training requirements position it as a robust and practical baseline for transforming any pre-trained AO engine into a stable TSE system.

4.3. Routing strategy and optimization

To validate our internal routing mechanism, we compare it against a post-hoc selection pipeline, which represents one of the most straightforward strategies for selecting a target speaker without retraining the AOSS backbone. In this cascaded setup, separator outputs are decoded and subsequently evaluated by an off-the-shelf SyncNet-based model [32]. The target is determined based on lip-sync consistency between the decoded output and the video, as measured by LSE-C and LSE-D scores [33]. As shown in Tab. 4, while our method achieves slightly higher routing accuracy, its primary advantage lies in system efficiency. Unlike post-hoc approaches that necessitate additional decoding, re-encoding, and potential preprocessing, Plug-and-Steer reuses fine-grained latent features directly from the backbone for routing. This elimination of redundant stages reduces the computational burden from 308.56 to 256.82 GFLOPs and improves the real-time factor from 0.209 to 0.147.

Furthermore, ablation results indicate that joint optimization with classification and reconstruction losses is essential for stable and accurate routing. While \mathcal{L}_{BCE} alone suffers from unstable initial convergence and \mathcal{L}_{SI-SNR} yields a slightly lower final accuracy, combining both objectives ensures both rapid convergence and the highest final accuracy. Crucially, the LSM serves as a functional bridge within the latent feature flow, enabling the use of reconstruction-based objectives for the routing module. Without this structural link, gradients from the final output could not propagate back to the steering logic, reducing target selection to a detached classification task. By facilitating this direct gradient flow, the LSM ensures that visual steering is optimized specifically for signal-level reconstruction, resulting in faster convergence and superior extraction performance.

5. Conclusion

In this work, we introduced Plug-and-Steer, a framework that decouples acoustic separation from target selection by redefining the visual modality as a selective router. By using a frozen audio-only backbone and a latent steering matrix, we achieve target selection while preserving the high-fidelity acoustic priors of the original engine. Our results demonstrate that while traditional fine-tuning suffers from noisy audio-visual ground truths, our steering approach effectively shields signal integrity. We expect that this decoupled strategy will provide a robust and scalable blueprint for high-fidelity extraction that adapts seamlessly to evolving speech separation engines.

6. Acknowledgements

This work was supported by the Institute of Information & Communications Technology Planning & Evaluation (IITP) grants funded by the Korean government (RS-2022-II220989, Development of Artificial Intelligence Technology for Multi-speaker Dialog Modeling).

7. Use of Generative AI Disclosure

During manuscript preparation, the authors used OpenAI's ChatGPT and Google's Gemini for language editing and stylistic refinement. These tools were not involved in conceptualizing the methods, designing experimental protocols, or interpreting scientific results. The authors reviewed and verified all suggestions and maintain full responsibility for the integrity and accuracy of the final content.

8. References

- [1] Z.-Q. Wang, S. Cornell, S. Choi, Y. Lee, B.-Y. Kim, and S. Watanabe, "TF-GridNet: Making time-frequency domain models great again for monaural speaker separation," in *Proc. ICASSP*, 2023.
- [2] S. Zhao, Y. Ma, C. Ni, C. Zhang, H. Wang, T. H. Nguyen, K. Zhou, J. Q. Yip, D. Ng, and B. Ma, "MossFormer2: Combining transformer and RNN-free recurrent network for enhanced time-domain monaural speech separation," in *Proc. ICASSP*, 2024.
- [3] K. Li, G. Chen, R. Yang, and X. Hu, "SPMamba: State-space model is all you need in speech separation," *arXiv preprint arXiv:2404.02063*, 2024.
- [4] X. Jiang, C. Han, and N. Mesgarani, "Dual-path Mamba: Short and long-term bidirectional selective structured state space models for speech separation," in *Proc. ICASSP*, 2025.
- [5] S. Lutati, E. Nachmani, and L. Wolf, "SepIt: Approaching a single channel speech separation bound," in *Proc. Interspeech*, 2022.
- [6] S. Lee, C. Jung, Y. Jang, J. Kim, and J. S. Chung, "Seeing through the conversation: Audio-visual speech separation based on diffusion model," in *Proc. ICASSP*, 2024.
- [7] K. Li, R. Yang, F. Sun, and X. Hu, "IIANet: An intra-and inter-modality attention network for audio-visual speech separation," in *Proc. ICML*, 2024.
- [8] J. Lin, X. Cai, H. Dinkel, J. Chen, Z. Yan, Y. Wang, J. Zhang, Z. Wu, Y. Wang, and H. Meng, "AV-SepFormer: Cross-attention SepFormer for audio-visual target speaker extraction," in *Proc. ICASSP*, 2023.
- [9] H. Sato, T. Ochiai, K. Kinoshita, M. Delcroix, T. Nakatani, and S. Araki, "Multimodal attention fusion for target speaker extraction," in *IEEE Spoken Language Technology Workshop*, 2021.
- [10] Z. Pan, M. Ge, and H. Li, "USEV: Universal speaker extraction with visual cue," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 30, pp. 3032–3045, 2022.
- [11] S. Zhao, Z. Pan, and B. Ma, "ClearerVoice-Studio: Bridging advanced speech processing research and practical deployment," in *Proc. Interspeech*, 2025.
- [12] J. Li, M. Ge, R. Cao, L. Wang, J. Dang, S. Zhang *et al.*, "Rethinking the visual cues in audio-visual speaker extraction," *arXiv preprint arXiv:2306.02625*, 2023.
- [13] R. Gao and K. Grauman, "VisualVoice: Audio-visual speech separation with cross-modal consistency," in *Proc. CVPR*, 2021.
- [14] K. Li, F. Xie, H. Chen, K. Yuan, and X. Hu, "An audio-visual speech separation model inspired by cortico-thalamo-cortical circuits," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2024.
- [15] S. Pegg, K. Li, and X. Hu, "RTFS-Net: Recurrent time-frequency modelling for efficient audio-visual speech separation," in *Proc. ICLR*, 2024.
- [16] R. Tao, X. Qian, Y. Jiang, J. Li, J. Wang, and H. Li, "Audio-visual target speaker extraction with reverse selective auditory attention," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, 2025.
- [17] J. S. Chung, A. W. Senior, O. Vinyals, A. Zisserman *et al.*, "Lip reading sentences in the wild," in *Proc. CVPR*, 2017.
- [18] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Interspeech*, 2018.
- [19] J.-C. Chou, C.-M. Chien, and K. Livescu, "AV2Wav: Diffusion-based re-synthesis from continuous self-supervised features for audio-visual speech enhancement," in *Proc. ICASSP*, 2024.
- [20] J. Le Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "SDR-half-baked or well done?" in *Proc. ICASSP*, 2019.
- [21] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time-frequency magnitude masking for speech separation," *IEEE/ACM Trans. on Audio, Speech, and Language Processing*, vol. 27, no. 8, pp. 1256–1266, 2019.
- [22] J. Cosentino, M. Pariente, S. Cornell, A. Deleforge, and E. Vincent, "LibriMix: An open-source dataset for generalizable speech separation," *arXiv preprint arXiv:2005.11262*, 2020.
- [23] Y. Luo, Z. Chen, and T. Yoshioka, "Dual-path RNN: Efficient long sequence modeling for time-domain single-channel speech separation," in *Proc. ICASSP*, 2020.
- [24] J. Wu, Y. Xu, S.-X. Zhang, L.-W. Chen, M. Yu, L. Xie, and D. Yu, "Time domain audio visual speech separation," in *IEEE Automatic Speech Recognition and Understanding workshop*, 2019.
- [25] Z. Pan, G. Wichern, Y. Masuyama, F. G. Germain, S. Khurana, C. Hori, and J. Le Roux, "Scenario-aware audio-visual TF-GridNet for target speech extraction," in *IEEE Automatic Speech Recognition and Understanding workshop*, 2023.
- [26] C. K. Reddy, V. Gopal, and R. Cutler, "DNSMOS: A non-intrusive perceptual objective speech quality metric to evaluate noise suppressors," in *Proc. ICASSP*, 2021.
- [27] G. Mittag, B. Naderi, A. Chehadi, and S. Möller, "NISQA: A deep CNN-self-attention model for multidimensional speech quality prediction with crowdsourced datasets," in *Proc. Interspeech*, 2021.
- [28] D. Yu, M. Kolbæk, Z.-H. Tan, and J. Jensen, "Permutation invariant training of deep models for speaker-independent multi-talker speech separation," in *Proc. ICASSP*, 2017.
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.
- [30] J. Yamagishi, C. Veaux, and K. MacDonald, "CSTR VCTK Corpus: English multi-speaker corpus for CSTR voice cloning toolkit (version 0.92)," *The Rainbow Passage which the speakers read out can be found in the International Dialects of English Archive*, 2019.
- [31] H. Zen, V. Dang, R. Clark, Y. Zhang, R. J. Weiss, Y. Jia, Z. Chen, and Y. Wu, "LibriTTS: A corpus derived from librispeech for text-to-speech," in *Proc. Interspeech*, 2019.
- [32] J. S. Chung and A. Zisserman, "Out of time: automated lip sync in the wild," in *Workshop on Multi-view Lip-reading, ACCV*, 2016.
- [33] K. Prajwal, R. Mukhopadhyay, V. P. Nambodiri, and C. Jawahar, "A lip sync expert is all you need for speech to lip generation in the wild," in *Proc. ACM MM*, 2020.