



# Incorporating Linguistic Constraints from External Knowledge Source for Audio-Visual Target Speech Extraction

Wenxuan Wu<sup>1</sup>, Shuai Wang<sup>2,3</sup>, Xixin Wu<sup>1</sup>, Helen Meng<sup>1</sup>, Haizhou Li<sup>2,4</sup>

<sup>1</sup>Department of SEEM, The Chinese University of Hong Kong, Hong Kong SAR, China

<sup>2</sup>SRIBD, School of Data Science, The Chinese University of Hong Kong, Shenzhen, China

<sup>3</sup>School of Intelligence Science and Technology, Nanjing University, Suzhou, China

<sup>4</sup>Department of ECE, National University of Singapore, Singapore

{wwu, wuxx, hmmeng}@se.cuhk.edu.hk, shuaiwang@nju.edu.cn, haizhouli@cuhk.edu.cn

## Abstract

Audio-visual target speaker extraction (AV-TSE) models primarily rely on target visual cues to isolate the target speaker's voice from others. We know that humans leverage linguistic knowledge, such as syntax and semantics, to support speech perception. Inspired by this, we explore the potential of pre-trained speech-language models (PSLMs) and pre-trained language models (PLMs) as auxiliary knowledge sources for AV-TSE. In this study, we propose incorporating the linguistic constraints from PSLMs or PLMs for the AV-TSE model as additional supervision signals. Without introducing any extra computational cost during inference, the proposed approach consistently improves speech quality and intelligibility. Furthermore, we evaluate our method in multi-language settings and visual cue-impaired scenarios and show robust performance gains.

**Index Terms:** Pre-trained language model, multi-modal, cocktail party, target speaker extraction

## 1. Introduction

Audio-visual target speaker extraction (AV-TSE) simulates the human ability to extract the target speaker's utterance in a cocktail party scenario using visual cues. Most existing studies focus on improving the audio-visual fusion mechanisms [1, 2, 3, 4, 5] or addressing visual cue-impaired scenarios [6, 7]. Despite the progress, the role of external knowledge has often been overlooked. A previous study suggests that when humans listen to words with similar semantics, the same neural units in the human auditory area will be activated [8]. In other words, linguistic knowledge may also serve as a useful cue for human's auditory perception. The linguistic knowledge functions as a prior and could be pre-trained or pre-loaded from an external knowledge source.

Such biological insight motivates us to incorporate external linguistic knowledge into AV-TSE systems. Pretrained language models (PLMs) and pretrained speech-language models (PSLMs) encode rich knowledge from massive unlabeled text and speech data through self-supervised learning, making them promising knowledge sources [9, 10]. While previous studies attempted PSLMs to improve TSE performance, most methods directly inject the PSLM features into TSE systems [11, 12]. However, this approach poses several risks: First, PSLMs capture both task-relevant and task-irrelevant knowledge [13], feeding irrelevant knowledge into the extractor may degrade performance or introduce unwanted interference [14]. Second, features derived from PSLMs might not perfectly align with intermediate features in TSE models, requiring additional fusion or alignment mechanisms, which may introduce addi-

tional parameter overhead [11, 3, 4]. Third, the expensive computation cost of PSLM/PLMs makes it infeasible for edge devices when such pretrained models are involved in the inference process.

To address these challenges, we propose to extract linguistic knowledge from PSLMs or PLMs as additional supervision signals during the optimization process. Similar approaches have been explored in some speech-related tasks. In [15], linguistic knowledge from PLMs is distilled into ASR models to enhance speech recognition performance. Additionally, recent studies on neural speech codecs [16, 17, 18] have leveraged PSLMs to improve the modeling of linguistic information, demonstrating the effectiveness of integrating pretrained knowledge into speech compression frameworks. Despite its success in other speech applications, for AV-TSE it remains underexplored. In this study, we investigate different PSLMs and PLMs as external knowledge sources and evaluate their effectiveness in enhancing AV-TSE systems. Key contributions are as follows,

- We systematically investigate the effectiveness of different pretrained language models (PLMs) and pretrained speech-language models (PSLMs) as knowledge sources for AV-TSE and design corresponding adapter modules to effectively integrate their linguistic knowledge.
- We comprehensively evaluate our method's robustness across multiple languages and challenging conditions, such as visual cue-impaired scenarios.
- Our method consistently enhances speech extraction quality in terms of signal similarity, speech intelligibility, and semantic coherence, demonstrating the effectiveness of using linguistic constraints for AV-TSE.

## 2. Methods

### 2.1. Problem Formulation

Given a mixture speech signal  $x$  and the target speaker's visual cue  $v$ , the goal of an AV-TSE system  $f_\theta$  is to extract a clean speech signal  $\hat{y}$  that approximates the ground truth target speech  $y$  via  $\hat{y} = f_\theta(x, v)$ . In most conventional AV-TSE systems, the model parameter  $\theta$  can be obtained by optimizing the SI-SDR loss  $\mathcal{L}_{\text{SI-SDR}}$  as

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\text{SI-SDR}}(f_\theta(x, v), y) \quad (1)$$

Given an external knowledge source (such as a pretrained language model (PLM) or a pretrained speech-language model (PSLM)), we aim to introduce an additional supervision term to incorporate linguistic constraint into the optimization process. To achieve this, we modify the objective function as follows:

Corresponding author: Shuai Wang

$$\theta^{**} = \arg \min_{\theta} \underbrace{\alpha \mathcal{L}_{\text{SI-SDR}}(f_{\theta}(x, v), y)}_{\text{Signal Reconstruction}} + \underbrace{\beta \mathcal{L}_{\text{LC}}(\hat{z}, z)}_{\text{Linguistic Constraint}} \quad (2)$$

where  $\hat{z}$  is the linguistic representation obtained from the extracted speech  $\hat{y}$ , and  $z$  is the target representation obtained from a pretrained knowledge model.

Note that when a PSLM is used as an external knowledge source,  $z$  represents a sequence of latent vectors derived from the ground truth speech  $y$ , and  $\hat{z}$  is a sequence of the same shape but derived from the predicted speech  $\hat{y}$ . Conversely, when a PLM serves as an external knowledge source,  $z$  is a compressed vector that captures the global linguistic information of the transcript of  $y$ , while  $\hat{z}$  is a vector of the same dimensionality as  $z$ , but obtained by averaging the latent sequence extracted from the PSLM. In this case, as illustrated in Figure.1, additional adapters are required, assuring that  $z$  and  $\hat{z}$  are in the same latent space.

## 2.2. Incorporation of Linguistic Constraints

The linguistic knowledge is encoded either in a PSLM or a PLM:

- A PSLM accepts audio as input and provides phoneme-level semantic and acoustic contextual features.
- A PLM accepts text as input and provides contextualized linguistic information at the word and sentence level.

A high-level illustration of the AV-TSE system incorporated with linguistic constraints is shown in Figure 1, where different forms of integration are also depicted.

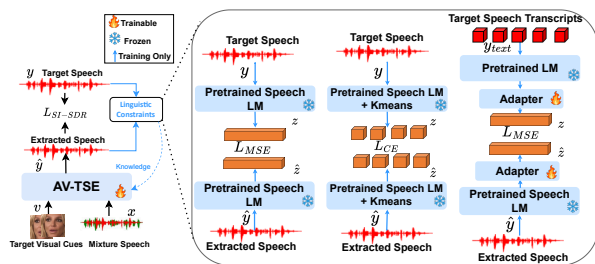


Figure 1: AV-TSE model with linguistic constraints. From left to right: continuous features and discrete tokens of PSLM, continuous features of PLM.

We explore two ways to incorporate PSLMs: continuous features and discrete tokens after quantization. For continuous features, we utilize feature representations extracted from the 24th layer of HuBERT-Large (HuBERT-L-L24) [9], the 24th layer of WavLM-Large (WavLM-L-L24) [19], and the 12th layer of WavLM-Base (WavLM-B-L12). We employ the Mean Squared Error (MSE) loss to minimize discrepancies between the extracted features and the corresponding knowledge representations. For discrete tokens, we utilize tokenized representations extracted from HuBERT-L-L24. The Cross-Entropy (CE) loss is adopted to minimize the discrepancy between the predicted and reference token distributions.

To incorporate linguistic constraints from PLMs, we employ the 12th layer of RoBERTa-Base (RoBERTa-B-L12) [20] and utilize the target speech transcript as input to obtain the corresponding linguistic feature representation as illustrated in Figure. 1. For extracted speech  $\hat{y}$ , we could not use the same

PLMs due to modality mismatch, instead, we adopt a PSLM to encode  $\hat{y}$  and introduce two adapters to bridge the latent space shift, then we calculate the MSE loss between two projected embeddings to enhance linguistic constraints.

## 3. Experimental Setups

### 3.1. Dataset

In this study, several experimental settings are considered:

- **Training Set:** A two-speaker mixture training set is simulated following previous work [1, 6, 2, 3]. The training set contains 800 speakers with 20,000 utterances. Speakers appearing in the test sets are excluded from the training data and SNR of interfering utterances are set randomly sampled from -10db to 10db. Among training utterances, around 76% English (EN), 6.4% French (FR), 0.5% Italian (IT), 1.4% Spanish (ES), 0.14% Portuguese (PT). Also includes others like 1.2% Dutch (NL) 7.7% German (DE), etc. All speech segments are clipped to 4 seconds.
- **Test Set:** Multiple test sets are designed to evaluate the proposed methods under diverse conditions:
  - **Core:** A two-speaker mixture set is created following [1, 6, 2, 3], with 118 randomly selected speakers and 3,000 simulated mixtures. Each mixture contains speakers of the same or different languages, with EN comprising the largest proportion. During inference, all segments are set to 6 seconds, shorter utterances are zero-padded.
  - **Monolingual:** Five two-speaker mixture test sets are generated for IT, PT, EN, ES, and FR. Each mixture contains speakers from the same language. Due to insufficient test data for non-majority languages in VoxCeleb2 [22], 118 speakers are selected from the development set to simulate 2,000 monolingual mixtures per language, all segments are set to 6 seconds, shorter utterances are zero-padded.
  - **Visual Cue Impaired:** Three scenarios are simulated: partial occlusion, low resolution, and full visual missing, the first two following [23], where we randomly select an obstacle for partial occlusion and apply a Gaussian blur or Gaussian noise to represent low resolution. Full visual missing frames are set to zero value. For each scenario, the impairment ratio ranges from 0% to 100%, uniformly distributed.
  - **Cross-Domain:** Generalization is assessed by evaluating the model on the Lip Reading Sentences 3 (LRS3) test set [24], following the test simulation protocol in [1].

### 3.2. Evaluation Metric

Speech quality is evaluated using the Scale-Invariant Signal-to-Distortion Ratio (SI-SDR) [25] and its improvement (SI-SDR<sub>i</sub>). Perceptual quality and intelligibility are assessed with the Perceptual Evaluation of Speech Quality (PESQ) [26] and the Short-Term Objective Intelligibility (STOI) [27]. To measure semantic similarity between extracted speech and ground truth, SpeechBERTScore with HuBERT Base [28] is employed. Higher values in all metrics indicate better performance.

### 3.3. Implementation Details

AV-Sepformer [2] is selected as the AV-TSE backbone. For adapter modules utilized for PLMs, we employed two fully connected layers. We follow [29] to discrete HuBERT continuous features with 500 cluster centers. A two-stage training frame-

Table 1: The study of incorporating linguistic constraints from different knowledge sources, that is evaluated on the core test set. “EP” denotes whether extra parameters are required during inference.

Knowledge Source	EP	Core Test Set					
		SI-SDR (↑)	SI-SDR <sub>i</sub> (↑)	SDR (↑)	PESQ (↑)	STOI (↑)	SpeechBERTScore (↑)
- (baseline)	✗	12.03	12.07	13.35	2.78	0.88	0.8247
MAR* [3]	✓	13.18	13.23	13.46	3.05	0.89	0.8563
AVSepChain [21]	✓	-	13.60	-	2.72	-	-
CTC**	✗	13.41	13.46	13.73	3.00	0.90	0.8521
HuBERT-L-L24 (token)	✗	13.01	13.06	13.33	2.94	0.89	0.8442
HuBERT-L-L24	✗	13.05	13.10	13.38	2.94	0.89	0.8438
WavLM-B-L12	✗	13.02	13.07	13.35	2.94	0.89	0.8439
WavLM-L-L24	✗	13.34	13.38	13.67	3.00	0.90	0.8510
RoBERTa-B-L12	✗	<b>13.60</b>	<b>13.65</b>	<b>13.93</b>	<b>3.05</b>	<b>0.90</b>	<b>0.8567</b>

\* MAR and \*\* CTC don't rely on external pretrained models

work is employed. In the first stage, AV-Sepformer is pretrained for 100 epochs. Training is halted if no validation improvement is observed for 6 epochs, with the learning rate halved after 3 consecutive epochs without improvement. Only  $L_{SI-SDR}$  is used as the training objective. In the second stage, training continues with the same settings, incorporating both  $L_{SI-SDR}$  and  $L_{LC}$ . The weighting factors  $\alpha$  and  $\beta$  are set to 1 and 10, respectively.

## 4. Results and Analysis

In this section, linguistic knowledge from different pretrained models is evaluated and analyzed on various test sets. For a more comprehensive comparison, we also include two systems that do not explicitly incorporate pretrained models as external knowledge source but have been validated in previous studies as beneficial for enhancing the model’s linguistic modeling capability. The first is the Mask-And-Recover (MAR) [3] strategy, which incorporates inter- and intra-modality contextual correlations into AV-TSE models by randomly masking mixture and recovering it with MAR block. The original paper implemented the MAR strategy on MuSE [1], but it could also be adapted to other backbones. In this study, we implemented the MAR strategy on AV-Sepformer and retrained it using the same training settings. Another is jointly optimizing AV-Sepformer with Connectionist Temporal Classification (CTC) loss [30], which assists in the temporal alignment between acoustic features and textual logits. In this study, we use Wav2Vec2ForCTC [31] to implement CTC between target speech transcripts and AV-Sepformer outputs.

Additionally, a system explicitly integrates PSLM into AVSepformer to enhance model’s linguistic modeling capability. AVSepChain explicitly integrates HuBERT and AV-HuBERT with AVSepformer, focusing on visual-speech semantic matching [21]. We also use AVSepChain results from the original paper for comparison.

### 4.1. Main experiments on the Core Test Set

Results on the core test set are shown in Table 1. In general, optimizing with language constraints from any of the pretrained knowledge sources consistently improves performance compared to the baseline. Here are some key observations:

#### When using PSLM as external knowledge source,

- The improvements achieved by discrete tokens and continuous features are comparable, as observed from the two versions of HuBERT-L-L24.
- Within the PSLMs, WavLM Base and HuBERT Large exhibit similar performance. However, scaling WavLM Base to

WavLM Large leads to a significant further improvement.

#### When using PLM as external knowledge source,

- RoBERTa-B-L12 outperforms all PSLMs, demonstrating the effectiveness of applying linguistic constraints supervision directly from the textual modality.
- To further validate the contribution of PLM, an additional baseline is introduced, where the original outputs are jointly optimized using CTC loss. Since CTC optimization can also be interpreted as introducing direct textual supervision, its impact is analyzed.
- The results show that CTC optimization significantly enhances performance, even surpassing all PSLMs. However, a performance gap remains compared to PLM, further validating the effectiveness of our method.
- Optimization with linguistic knowledge from PLM is more effective at a low cost. MAR and AVSepChain enhance performance through improved linguistic modeling abilities, but both incur additional inference costs. MAR requires three additional transformer encoder layers and two linear projection layers [3], while AVSepChain needs one cross-attention layer and three Conv1D layers, and integrates AV-HuBERT [32] as the visual frontend [21] to extract visual cues, also integrates HuBERT [9] as backend to resynthesis predicted speech. In contrast, optimization with linguistic knowledge required no extra cost during inference and could perform best with RoBERTa-B-L12.

### 4.2. Evaluation on the Monolingual Test Set

The results are shown in Table. 2. We observe the following: AV-Sepformer demonstrates relatively modest performance on all five monolingual test sets compared to its performance on the core test set, despite the model being exposed to all five languages. In particular, the performance for PT, the SI-SDR is 11.35, and for ES, the SI-SDR is 11.38, likely due to these languages having a small proportion in the training set. Notably, while English accounts for the majority of the training set, its monolingual performance is lower than that of the core set. Mixtures of the same language may pose more challenges for extractions.

With linguistic knowledge from RoBERTa-B-L12, model consistently performs best across all five monolingual test sets. For PT and ES, improvements of approximately 1.5 dB in SI-SDR were achieved without additional fine-tuning. It is promising and can be extended to even lower-resource languages, where collecting multimodal data poses significant challenges.

Table 2: The study of incorporating linguistic constraints from different knowledge sources, that is evaluated on the test sets of different languages. The percentage in parentheses indicates the proportion of data for each language in the training set, as mentioned in Section 3.1.

Knowledge Source	EN (76%)		PT(0.14%)		ES (1.4%)		IT (0.5%)		FR (6.4%)	
	SI-SDRi	PESQ	SI-SDRi	PESQ	SI-SDRi	PESQ	SI-SDRi	PESQ	SI-SDRi	PESQ
-(baseline)	11.79	2.72	11.34	2.71	11.38	2.74	11.67	2.76	11.80	2.56
MAR* [3]	12.72	2.90	12.25	2.89	12.20	2.90	12.51	2.93	12.93	2.72
CTC**	13.03	2.92	12.60	2.90	12.59	2.94	12.84	2.94	13.18	2.77
HuBERT-L-L24 (token)	12.63	2.86	12.26	2.85	12.24	2.89	12.47	2.89	12.84	2.71
HuBERT-L-L24	12.65	2.86	12.20	2.84	12.30	2.88	12.49	2.89	12.87	2.71
WavLM-B-L12	12.66	2.86	12.09	2.83	12.22	2.88	12.37	2.87	12.77	2.70
WavLM-L-L24	12.99	2.92	12.56	2.90	12.72	2.90	12.83	2.94	13.18	2.77
RoBERTa-B-L12	<b>13.23</b>	<b>2.97</b>	<b>12.84</b>	<b>2.94</b>	<b>12.88</b>	<b>2.94</b>	<b>12.99</b>	<b>2.94</b>	<b>13.42</b>	<b>2.82</b>

\* MAR and \*\* CTC don't rely on external pretrained models

### 4.3. Evaluation in Visual Cue-Impaired Scenarios

In this section, we investigate whether linguistic constraints also help when primary visual cues are impaired. We directly evaluate the AV-TSE models with visual cue-impaired test set, including three visual cue-impaired scenarios: full visual missing, partial visual occlusion, and low resolution. For a fair comparison, we also employ ImagineNET as a comparison baseline, an AV-TSE system especially tailored for the visual cue missing scenario, by explicitly recovering target visual cue embedding with several visual decoders [6]. We employ two model settings: ImagineNET and ImagineNET-Sys.2, the former is re-trained on our AV-TSE training set, while the latter is from origin checkpoints, which has been trained on the visual full missing dataset.

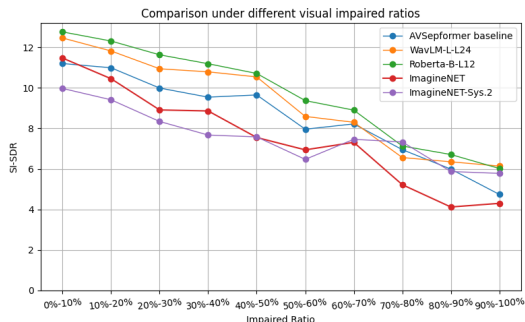


Figure 2: SI-SDR comparison under different visual cue-impaired ratios.

Specifically, we compute the average performance on three visual impaired scenarios with impairment ratios ranging from 0% to 100%, as illustrated in Figure. 2. We observe that RoBERTa-B-L12, and WavLM-L-L24 consistently outperform the vanilla AV-Sepformer across nearly all impairment ratios, and RoBERTa-B-L12 still achieves the best performance. In contrast, ImagineNET demonstrates modest performance at lower impairment ratios, but with high impaired ratios, it shows comparable performance to RoBERTa-B-L12 and WavLM-L-L24. This is likely due to its training bias towards fully missing visual scenarios, as well as model size limitations.

The overall SI-SDR performance of different AV-TSE models on three visual cue-impaired scenarios is reported in Table. 3. Note that AV-Sepformer with RoBERTa-B-L12 or WavLM-L-L24 were not exposed to any visual cue-impaired scenario, but still outperformed ImagineNET Sys.2 across all impair ratios, where the latter have seen visual full missing conditions during training. This finding suggests the proposed linguistic

Table 3: SI-SDR performance under three visual impaired scenarios, “Full occ” denotes full visual missing, “Part occ” denotes partial visual occlusion, “Low reso” denotes low resolution, and Avg denotes average performance on three scenarios.

Knowledge Source	Full occ	Part occ	Low reso	Avg
-(baseline)	5.15	11.26	10.72	9.04
ImagineNET	2.92	11.38	10.40	8.22
ImagineNET-Sys.2 [6]	5.36	9.75	10.40	8.01
WavLM-L-L24	5.51	12.47	12.00	9.98
RoBERTa-B-L12	<b>6.07</b>	<b>12.87</b>	<b>12.32</b>	<b>10.42</b>

constraints’ practical applicability in real-world scenarios.

### 4.4. Cross-dataset Evaluation on LRS3

In this section, we directly evaluate the proposed models on the LRS3 test set for cross-domain evaluation. The LRS3 dataset is collected from TED videos, containing less background noise and music compared to VoxCeleb2 [24], and it is widely used in visual speech recognition tasks [33]. We select knowledge sources that perform best on speech and text modality, respectively, i.e. WavLM-L-L24 and RoBERTa-B-L12 for evaluation. We also compare our results with AVSepChain on the LRS3 test set, as reported in the original paper using two-second speech segments [21]. The overall results are shown in Table. 4.

Table 4: Cross domain evaluation on LRS3 test set.

Knowledge Source	SI-SDRi	PESQ
-(baseline)	13.83	2.77
AVSepChain [21]	15.3	<b>3.26</b>
WavLM-L-L24	15.21	2.99
RoBERTa-B-L12	<b>15.42</b>	3.04

We observe that all models achieve better performance compared to the results on Voxceleb2, likely due to fewer noise interferences. Furthermore, RoBERTa-B-L12 achieves the best performance in terms of SI-SDRi, demonstrating the effectiveness of linguistic constraints

## 5. Conclusion

In this study, we propose to derive linguistic constraints from pre-trained language models and speech-language models to optimize AV-TSE. Evaluated on multilingual, visual cue-impaired, and cross-domain test sets, the proposed paradigm consistently outperforms the vanilla AV-TSE backbone, while the linguistic constraints from the language model surpass those of the speech model. In future work, we aim to explore other methods of injecting language information and semantics to further improve AV-TSE performance.

## 6. Acknowledgements

This work was supported by National Natural Science Foundation of China, Grant No. 62401377, Shenzhen Science and Technology Program (Shenzhen Key Laboratory, Grant No. ZDSYS20230626091302006), Shenzhen Science and Technology Research Fund (Fundamental Research Key Project, Grant No. JCYJ20220818103001002), Program for Guangdong Introducing Innovative and Entrepreneurial Teams, Grant No. 2023ZT10X044, Yangtze River Delta Science and Technology Innovation Community Joint Research Project (2024CSJGG01100).

## 7. References

- [1] Z. Pan, R. Tao, C. Xu, and H. Li, "Muse: Multi-modal target speaker extraction with visual cues," *ICASSP 2021*.
- [2] 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 *ICASSP 2023*.
- [3] W. Wu, X. Chen, X. Wu, H. Li, and H. Meng, "Target speech extraction with pre-trained av-hubert and mask-and-recover strategy," *CVPR 2024 Sight and Sound Workshop*.
- [4] W. Wu, X. Chen, X. Wu, H. Li, and H. Meng, "Target speech extraction with pre-trained av-hubert and mask-and-recover strategy," in *IJCNN 2024*.
- [5] W. Wu, X. Chen, S. Wang, J. Wang, L. Meng, X. Wu, H. Meng, and H. Li, " $c^2$ av-tse: Context and confidence-aware audio visual target speaker extraction," *IEEE Journal of Selected Topics in Signal Processing 2025*.
- [6] Z. Pan, W. Wang, M. Borsdorf, and H. Li, "Imaginet: Target speaker extraction with intermittent visual cue through embedding inpainting," in *ICASSP 2023*.
- [7] J. Wang, Z. Pan, M. Zhang, R. T. Tan, and H. Li, "Restoring speaking lips from occlusion for audio-visual speech recognition," in *AAAI 2024*.
- [8] M. Jamali, B. Grannan, J. Cai, A. Khanna, W. Muñoz, I. Caprara, A. Paulk, S. Cash, E. Fedorenko, and Z. Williams, "Semantic encoding during language comprehension at single-cell resolution," *Nature*, vol. 631, pp. 1–7, 07 2024.
- [9] W.-N. Hsu, B. Bolte, Y.-H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM TASLP 2021*.
- [10] L. Meng, S. Hu, J. Kang, Z. Li, Y. Wang, W. Wu, X. Wu, X. Liu, and H. Meng, "Large language model can transcribe speech in multi-talker scenarios with versatile instructions," in *ICASSP 2025*.
- [11] J. Peng, M. Delcroix, T. Ochiai, O. Plchot, S. Araki, and J. H. ernocký, "Target speech extraction with pre-trained self-supervised learning models," *ICASSP 2024, 2024*.
- [12] J. Peng, M. Delcroix, T. Ochiai, O. Plchot, T. Ashihara, S. Araki, and J. Černocký, "Probing self-supervised learning models with target speech extraction," in *ICASSPW 2024*.
- [13] S.-w. Yang, H.-J. Chang, Z. Huang, A. T. Liu, C.-I. Lai, H. Wu, J. Shi, X. Chang, H.-S. Tsai, W.-C. Huang *et al.*, "A large-scale evaluation of speech foundation models," *IEEE/ACM TASLP 2024*.
- [14] S. Qiu, B. Han, D. C. Maddix, S. Zhang, Y. Wang, and A. G. Wilson, "Transferring knowledge from large foundation models to small downstream models," in *ICML 2024*.
- [15] M. Han, F. Chen, J. Shi, S. Xu, and B. Xu, "Knowledge transfer from pre-trained language models to cif-based speech recognizers via hierarchical distillation," in *Interspeech 2023*.
- [16] X. Zhang, D. Zhang, S. Li, Y. Zhou, and X. Qiu, "Speechtokenizer: Unified speech tokenizer for speech large language models," *ICLR 2024*.
- [17] Y. Guo, Z. Li, C. Du, H. Wang, X. Chen, and K. Yu, "LSCoDec: Low-bitrate and speaker-decoupled discrete speech codec," *arXiv preprint arXiv:2410.15764, 2024*.
- [18] D. Yang, S. Liu, H. Guo, J. Zhao, Y. Wang, H. Wang, Z. Ju, X. Liu, X. Chen, X. Tan, X. Wu, and H. Meng, "ALMTokenizer: A Low-bitrate and Semantic-rich Audio Codec Tokenizer for Audio Language Modeling," in *ICML 2025*.
- [19] S. Chen, C. Wang, Z. Chen, Y. Wu, S. Liu, Z. Chen, J. Li, N. Kanda, T. Yoshioka, X. Xiao, J. Wu, L. Zhou, S. Ren, Y. Qian, Y. Qian, M. Zeng, and F. Wei, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE JSTSP 2021*.
- [20] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *ArXiv, 2019*.
- [21] Z. Mu and X. Yang, "Separate in the speech chain: cross-modal conditional audio-visual target speech extraction," in *IJCAI 2024*.
- [22] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *INTERSPEECH, 2018*.
- [23] J. Hong, M. Kim, J. Y. Choi, and Y. M. Ro, "Watch or listen: Robust audio-visual speech recognition with visual corruption modeling and reliability scoring," *2023 CVPR*.
- [24] T. Afouras, J. S. Chung, and A. Zisserman, "Lrs3-ted: a large-scale dataset for visual speech recognition," *ArXiv, 2018*.
- [25] J. L. Roux, S. Wisdom, H. Erdogan, and J. R. Hershey, "Sdr – half-baked or well done?" *ICASSP 2019*.
- [26] S. Basterrech, G. Rubino, and M. Varela, "Single-sided Real-time PESQ Score Estimation," *arXiv e-prints*, p. arXiv:1212.6350, Dec. 2012.
- [27] C. H. Taal, R. C. Hendriks, R. Heusdens, and J. R. Jensen, "An algorithm for intelligibility prediction of time–frequency weighted noisy speech," *IEEE TASLP 2011*.
- [28] T. Saeki, S. Maiti, S. Takamichi, S. Watanabe, and H. Saruwatari, "SpeechBERTScore: Reference-aware automatic evaluation of speech generation leveraging nlp evaluation metrics," in *Inter-speech 2024*.
- [29] P. Mousavi, J. Duret, S. Zaiem, L. Della Libera, A. Ploujnikov, C. Subakan, and M. Ravanelli, "How should we extract discrete audio tokens from self-supervised models?" in *Interspeech 2024, 2024*.
- [30] A. Graves, S. Fernández, F. J. Gomez, and J. Schmidhuber, "Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural networks," in *ICML 2006*.
- [31] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *NIPS 2020*.
- [32] B. Shi, W.-N. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," in *ICLR 2024*.
- [33] J. Hong, S. J. Park, and Y. M. Ro, "Intuitive multilingual audio-visual speech recognition with a single-trained model," in *EMNLP 2023*.