



# Cocktail-Party Audio-Visual Speech Recognition

Thai-Binh Nguyen<sup>1</sup>, Ngoc Quan Pham<sup>2</sup>, Alexander Waibel<sup>1,2</sup>

<sup>1</sup>Karlsruhe Institute of Technology, Germany

<sup>2</sup>Carnegie Mellon University, USA

thai-binh.nguyen@kit.edu

## Abstract

Audio-Visual Speech Recognition (AVSR) offers a robust solution for speech recognition in challenging environments, such as cocktail-party scenarios, where relying solely on audio proves insufficient. However, current AVSR models are often optimized for idealized scenarios with consistently active speakers, overlooking the complexities of real-world settings that include both speaking and silent facial segments. This study addresses this gap by introducing a novel audio-visual cocktail-party dataset designed to benchmark current AVSR systems and highlight the limitations of prior approaches in realistic noisy conditions. Additionally, we contribute a 1526-hour AVSR dataset comprising both talking-face and silent-face segments, enabling significant performance gains in cocktail-party environments. Our approach reduces WER by 67% relative to the state-of-the-art, reducing WER from 119% to 39.2% in extreme noise, without relying on explicit segmentation cues.

**Index Terms:** audio-visual speech recognition, cocktail-party

## 1. Introduction

The visual information obtained from observing a person speak can alter the way auditory signals are perceived, a phenomenon known as the McGurk effect [1]. In cocktail-party environments, even strong ASR models [2, 3] which mark a significant advance over early efforts [4] in conversational speech, still experience significant performance degradation. In such challenging conditions, combining visual cues, like facial movements, with auditory input significantly improves speech comprehension [5, 6]. Inspired by this interplay, AVSR systems have been developed to leverage visual cues for enhancing speech recognition, particularly in noisy environments. This concept has been explored and validated over the past several decades since its introduction in 1976 [7, 8, 9, 10, 11].

Recent advancements in AVSR have been largely propelled by deep learning models, including the adoption of end-to-end architectures like Transformer [12] and Conformer [13]. These models have been enhanced by improved data utilization, such as pre-training with self-supervised methods like AV-HuBERT [12], or by leveraging pre-trained ASR models to generate transcriptions for unlabeled AV datasets, as demonstrated in [13, 14]. While the combination of audio and visual modalities is expected to make these models robust to noise, our experiments reveal a significant performance decline in cocktail-party environments. For instance, the state-of-the-art (SOTA) AVSR model, Auto-AVSR, achieves an impressive 1.5% Word Error Rate (WER) on the LRS2 dataset [13], but when background speech noise is added, its WER rises drastically to 69%. This performance drop is similarly observed in other SOTA models, as discussed in the experiment section, highlighting critical con-

cerns about their practicality in real-world noisy scenarios.

Applying AVSR to the cocktail-party problem is an active area of research. Studies such as [15, 16, 17] utilize visual information as a query to isolate the target speaker in a mixed audio signal, leveraging the visual modality to focus on and transcribe the speech of a specific individual. Rather than directly outputting the target speaker's transcription, other approaches, like [18, 19, 20], use visual features to extract the target speech signal. A common characteristic of these studies is their reliance on datasets such as LRS2 [21], LRS3 [22], and VoxCeleb2 [23], which are augmented by mixing utterances to create training and evaluation data with perfect alignment. In these datasets, the speech consistently originates from the speaker shown in the video, or the visual input always depicts a talking face. This alignment allows models to reliably identify the target speaker and generate outputs based on the visible speaker. Recent work, such as [24], has attempted to introduce out-of-sync audio-visual pairs to simulate more challenging scenarios. However, even in these cases, evaluations are typically conducted on datasets like LRS3, which do not reflect the complexity of real-world cocktail-party environments.

A key limitation of commonly used audio-visual datasets like LRS2 and LRS3 is that they fail to capture the complexities of cocktail-party scenarios. Firstly, these datasets predominantly feature single-speaker samples, which do not reflect the multi-speaker interactions typical of cocktail-party environments. Secondly, in noisy environments, a critical challenge is determining whether a visible speaker is actively talking or not. This necessitates the inclusion of both talking face and silent face within an utterance, a factor largely overlooked in current datasets. For the first issue, simulating multi-speaker noise, has been partially addressed through methods such as randomly mixing utterances or adding artificial noise to samples [15, 16, 17, 24]. However, the second issue, involving scenarios with silent face and ambiguous speech activity, remains underexplored. A notable exception is the Chinese multi-channel audio-visual conversation dataset, MISP [25], which attempts to fill this gap. Nevertheless, studies utilizing the MISP dataset [26, 27] predominantly focus on tasks like audio-visual speaker diarization or audio-visual target speaker extraction (AVTSE) prior to conducting AVSR. While valuable, these approaches address only specific challenges and fall short of tackling the broader set of obstacles faced by AVSR models in realistic cocktail-party environments.

In this study, we focus on developing end-to-end AVSR models tailored to handle cocktail-party scenarios. Our contributions are as follows: (1) we define a novel audio-visual cocktail-party dataset that differs significantly from MISP in three key aspects—it is an English dataset, includes multiple overlapping conversations, and features single-channel audio

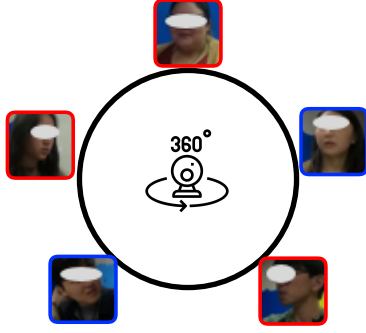


Figure 1: Schematic overview of AVCocktail’s recording scene.

data. The inclusion of overlapping conversations and single-channel audio makes our dataset more challenging and closer to real-world scenarios compared to MISP, which primarily consists of single conversations per recording and uses multi-channel audio [25]. (2) We introduce a 1526-hour AVSR dataset that addresses the limitations of previous datasets by incorporating silent-face utterances, which are crucial for distinguishing active speech. (3) We propose a robust data pipeline for augmenting AVSR datasets, improving their suitability for training models capable of handling cocktail-party environments. (4) Finally, we train a strong baseline AVSR model that demonstrates effective performance on our cocktail-party dataset.

## 2. Cocktail-Party AVSR

### 2.1. Task definition

Given an input sequence of audio  $A = \{a_1, a_2, \dots, a_T\}$  and video  $V = \{v_1, v_2, \dots, v_T\}$ , where  $T$  is the total number of time steps,  $a_t$  represents the audio feature and  $v_t$  represents the visual feature at time step  $t$ . The task is to predict the transcription  $Y_{\text{target}} = \text{AVSR}(A, V) = \{y_1, y_2, \dots, y_N\}$ , where  $N$  is the length of the target speech transcription, AVSR denotes the model that takes both audio and visual features as input to generate the transcription. While the video captures the target speaker, the audio may include overlapping speech, background noise, and uncertainty in speech activity when the target speaker is visible but not speaking.

### 2.2. Baseline

We adopt two off-the-shelf architectures to evaluate the effectiveness of the proposed method through fine-tuning. The first model, AV-HuBERT CTC/Attention (AV1), uses AV-HuBERT [12] as the encoder and the decoder integrates a projection layer and a Transformer decoder with joint CTC/Attention training [29]. The second model is the Conformer CTC/Attention architecture (AV2) proposed by [30], where the encoder consists of two Conformer blocks: one for audio and one for visual feature extraction. The decoder is identical to AV1, employing joint CTC/Attention training.

In addition to fine-tuning the above architectures, we directly evaluate recent AVSR models that have achieved SOTA performance on the LRS2 and LRS3 datasets as strong baselines. These include Auto-AVSR (denoted as AV3) [13], which uses a Conformer with CTC/Attention and is trained on 3448 hours of AVSR data. Two additional variants of AV3 are included for benchmarking: a visual-only model (V1) and an audio-only model (A2). Another baseline is the Muavic-EN (AV4) model [2], which employs AV-HuBERT as the encoder and a Transformer decoder, trained exclusively on the LRS3 dataset with various types of additive noise. Whisper-Flamingo

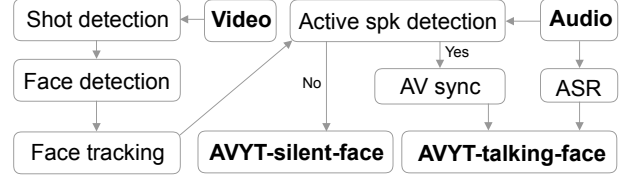


Figure 2: Pipeline to generate AVYT dataset

(AV5) [14] combines AV-HuBERT with Whisper-large and is trained on LRS3, Vox2, and augmented noisy data. Lastly, Whisper-large (A1), trained on 5M hours of diverse audio data, serves as a strong baseline for benchmarking audio-only performance.

## 3. Data Preparation

In this study, we utilize four datasets. For training, we use LRS2 (train and pretrain sets), Vox2 (train set), and AVYT. For testing, we evaluate on the LRS2 test set (including a modified version) and AVCocktail. Details of LRS2 test set, AVYT, and AVCocktail are provided in the following subsections. About Vox2, we simply employ Whisper-large [2] to transcribe the audio and retain only the English segments.

### 3.1. Lip Reading Sentences 2 (LRS2)

LRS2 [21] is a widely used AVSR dataset commonly utilized for benchmarking<sup>1</sup>. To evaluate the performance of SOTA AVSR models in cocktail-party settings, we augment LRS2 by randomly adding interfering speakers in the background with varying signal-to-noise ratios (SNR). Specifically, we introduce up to two interfering speakers, with SNR levels controlled at  $-5$ ,  $0$ ,  $5$ , and  $10$  dB. The original LRS2 dataset, without any interfering speakers, corresponds to an SNR of  $\infty$  and zero interferer.

### 3.2. AVCocktail

In our AVCocktail dataset, we focus on a scenario where people gather around a table, forming small groups of 2 to 5 individuals, each group engaged in a discussion topic. Figure 1 illustrates a recording scene with two such groups. A single 360-degree camera is placed at the center of the table to capture all participants’ faces. From the 360-degree footage, we extract  $224 \times 224$  cropped video clips and a single audio channel for each speaker. Each recording session lasts 5 to 7 minutes, resulting in a total evaluation set of approximately 6.1 hours of video from 45 speakers. All individual video then been segmented and transcribe by human.

### 3.3. Automatic AVSR dataset from Youtube (AVYT)

Due to the limitations of existing AVSR datasets, as described in Section 1, we identified the need for an additional dataset to better train AVSR models for cocktail-party scenarios. To address this, we introduce AVYT, derived from the 6,533 hours of YouTube content introduced in [31]. AVYT consists of two subsets: a silent-face set with 77 hours spanning 79k clips, and a talking-face set containing 1449 hours across 666k clips.

Figure 2 illustrates the processing pipeline for constructing the AVYT dataset. The first step, inspired by the LRS2 data processing pipeline [21], involves shot detection, face detection, and face tracking. This process yields cropped video clips containing a single speaker. In the second step, an active

<sup>1</sup>LRS3 was unavailable at the time of this study

Table 1: WER (%) of models on the LRS2 dataset. ★ denotes our fine-tuned model. "AV" (in the modality column) indicates models that utilize both audio and visual features.

Model ID	Model	Modality	Train dataset	Interferer	SNR (dB)					Avg
					-5	0	5	10	$\infty$	
AV1	AV-HuBERT CTC/Attention★	AV	lrs2,vox2,avyt	0						
				1	<b>6.4</b>	<b>3.5</b>	<b>3.4</b>	<b>2.8</b>		
				2	<b>9.0</b>	<b>4.4</b>	<b>3.2</b>	<b>2.8</b>		<b>4.1</b>
AV2	Conformer CTC/Attention★	AV	lrs2,vox2,avyt	0					10.9	
				1	19.6	17.1	18.0	15.7		
				2	20.1	17.6	18.2	16.1		17.0
AV3	Auto-AVSR [13]	AV	lrs2,vox2	0					1.7	
			lrs3	1	56.6	16.6	10.3	4.2		
			avspeech	2	69.6	21.8	11.7	3.5		21.7
AV4	Muavic-EN [2]	AV	lrs3	0					7.2	
				1	18.9	10.8	9.7	8.5		
				2	25.4	12.1	9.8	8.8		12.4
AV5	Whisper-Flamingo [14]	AV	lrs3,vox2	0					6.1	
				1	96.9	37.4	26.2	12.1		
				2	99.6	38.6	30.6	13.4		40.1
A1	Whisper large-v3 [28]	Audio	5M hours	0					3.7	
				1	97.7	30.9	13.2	6.5		
				2	99.9	31.1	14.8	6.5		33.8
A2	Auto-AVSR [13]	Audio	same as AV3	0					<b>1.5</b>	
				1	93.9	30.5	22.7	5.3		
				2	95.8	33.0	23.7	6.2		34.7
V1	Auto-AVSR [13]	Visual	same as AV3	0			15.7			15.7

speaker detection model [32] identifies talking-face segments and silent-face segments, assigning them to separate sets. The third step further refines the talking-face clips: AV sync [33] ensures proper audio-video synchronization, and the synchronized clips are then transcribed using Whisper-large [2] to create the final talking-face set.

### 3.4. Data augmentation pipeline

#### Algorithm 1 Data Augmentation Pipeline

```

1: Function generate_sample(dialog, silent_face, interferer)
2: for video in [Vox2, LRS2, AVYT-talking-face] do
3:   videos = [video]
4:   if dialog then
5:     videos += n_rand([Vox2, LRS2, AVYT-talking-face])
6:   if silent_face then
7:     videos += n_rand(AVYT-silent-face)
8:   end if
9: end if
10: avsr_sample = concat(shuffle(videos))
11: if interferer then
12:   avsr_sample = augment_speech_noise(avsr_sample)
13: end if
14: avsr_sample = video_transform(avsr_sample)
15: yield avsr_sample
16: end for

```

Algorithm 1 outlines the data augmentation pipeline designed to construct samples for fine-tuning our AVSR model. The pipeline incorporates three key augmentation strategies: dialog augmentation, adding silent-face video clips, and introducing interfering speakers. These strategies can be applied individually or combined, with the option to enable or disable each one as needed. Dialog augmentation merges multiple video clips to simulate conversational scenarios, a technique proven effective in prior studies [34, 35]. This approach is particularly

crucial when combined with the second augmentation strategy, which integrates silent-face video clips. Since silent-face clips are labeled with  $\langle unk \rangle$  transcripts, merging multiple clips into a dialog-like structure prevents the model from relying solely on visual cues (frame sequence differences) to infer transcripts for silent-face segments. The third strategy introduces interferer speakers in the background at varying SNR. Finally, the pipeline applies video transformations, including horizontal flipping, random cropping, and adaptive time masking for both visual and audio streams, to further diversify the training data. The video frames are cropped to the mouth region of interest (ROI) using a  $96 \times 96$  bounding box, while the audio is sampled at a 16 kHz rate, similar to [12, 30].

## 4. Experimental setup

As described in Section 2.2, we fine-tune two model architectures. The AV-HuBERT CTC/Attention (AV1) model uses the AV-HuBERT large [12] as the encoder, which has 24 transformer blocks, each with 16 attention heads. The CTC/Attention decoder is a 6-layer Transformer with the same dimensions and number of attention heads as the encoder. The Conformer CTC/Attention (AV2) [30] consists of 12 encoder layers for both audio and visual inputs, each with 16 attention heads. The decoder CTC/Attention in this model is similar to the first, consisting of a 6-layer Transformer.

The data pre-processing pipeline used to train our model is detailed in Section 3.4. We then conduct several evaluations. First, we compare the performance of our fine-tuned model with current SOTA AVSR models on the LRS2 test set (including a modified version), as described in Section 3.1. This benchmark evaluates how model performance degrades under data distortions, similar to the challenges posed by cocktail-party scenarios. The second experiment assesses these models on the real cocktail-party dataset, AVCocktail. Since each recording lasts 5 to 7 minutes, a segmentation step is required before inference.

Table 2: WER (%) of models on the AVCocktail dataset

Recognition Model ID	Segmentation		
	Active Speaker Detection [32]	Fixed chunk (10s)	Gold
AV1	<b>22.6</b>	<b>39.2</b>	<b>18.2</b>
AV2	48.4	89.5	41.9
AV3	74.6	133.2	67.8
AV4	35.6	119.0	26.1
AV5	70.8	133.3	58.3
A1	67.4	143.9	54.7
A2	75.8	131.7	70.3
V1	56.0	167.9	49.9

We employ three segmentation strategies. (1) a model-based approach using Active Speaker Detection [32], which leverages both audio and visual cues to determine speaking segments. (2) a fixed 10-second chunk-based approach with a sliding window and no overlap and (3) manual segmentation, where segments are labeled by humans. Finally, we conduct an ablation study on the data augmentation techniques introduced in Section 3.4 to identify the most influential factors in AVSR model performance.

## 5. Results

Table 1 presents the WER (%) of baseline models and our fine-tuned models on the LRS2 test set, including both the original and modified versions. The WERs for models evaluated on the original LRS2 test set are shown in the column where SNR =  $\infty$ . Overall, all models perform well on the original clean LRS2 dataset. The best-performing model in this setting is A2, the Conformer CTC/Attention audio-only model, which is expected since the dataset consists of clean speech, and A2 is well-trained with in-domain data.

However, performance degrades significantly as SNR decreases and the number of interfering speakers increases. Audio-only models are the most affected by noise, with A1’s WER rising from 3.7% to 99.9% and A2’s WER increasing from 1.5% to 95.8%. In contrast, the Conformer CTC/Attention visual-only model (V1) remains unaffected by noise, maintaining a constant WER of 15.7%. Among the baseline audio-visual models (AV3, AV4, AV5), despite leveraging visual features and noise augmentation during training, performance still deteriorates significantly under noisy conditions. Notably, AV3 and AV5 suffer the most, with WERs rising from 1.7% and 6.1% to 69.6% and 99.6%, respectively. AV4 demonstrates the highest noise robustness among the baselines, likely due to its augmentation with both speech noise and additional noise types beyond “natural,” “music,” and “babble,” which were used in AV3. Although AV5 employs the same augmentation strategy as AV4, it appears to rely more on audio than visual information, leading to the worst performance under extreme noise conditions.

AV1 and AV2 are our fine-tuned models, trained with in-domain data as described in Section 3.4. AV2 is initialized from AV3’s parameters. After fine-tuning, AV2 achieves overall better performance than AV3 (17.0% WER vs. 21.7%), but its WER on the original LRS2 test set increases significantly from 1.7% to 10.9%. AV1, which employs AV-HuBERT as the encoder and a CTC/Attention decoder, achieves the best performance among all models, with a WER of 2.1% on the original LRS2 test set and an average WER of 4.1% overall.

Unlike LRS2, AVCocktail consists of long video recordings that contain both talking-face and silent-face segments. The choice of segmentation method influences the proportion of

Table 3: Ablation study on the impact of data factors on AV-HuBERT CTC/Attention performance in AVCocktail dataset.

Dataset	Interferer	Dialog	Silent-face	WER
lrs2,vox2				58.8
lrs2,vox2	✓			32.5
+AVYT-talking	✓			30.2
+AVYT-talking	✓	✓		29.4
+AVYT	✓		✓	28.5
+AVYT	✓	✓	✓	22.6

silent-face segments included in the inference data. In the AVCocktail dataset, the total speaking duration accounts for 46.3% of the recordings. Active Speaker Detection (ASD) achieves a precision of 84.4%, recall of 97.8%, and F1-score of 90.6% in detecting speaking segments. This means that with ASD-based segmentation, nearly all talking-face segments are retained, though some silent-face segments are incorrectly detected as speech. In contrast, fixed-length segmentation inherently includes both talking-face (46.3%) and silent-face (53.7%) segments, as it processes the video in uniform chunks without considering whether the target speaker is speaking or silent.

Table 2 presents the performance of different models on the AVCocktail dataset. In general, a higher proportion of silent-face segments leads to worse WER. For baseline models (AV[3-5], A[1-2], and V1), the WER in the extreme case of fixed-length segmentation exceeds 100%, while our AV1 model still achieves a WER of 39.2% in this scenario. When using a segmentation model like ASD, the WER improves significantly. As expected, the best performance is achieved with gold segmentation. Our AV1 model demonstrates strong robustness, achieving the best performance among all baseline models across all types of segmentation.

Table 3 presents the ablation study on the impact of different data factors on AV-HuBERT CTC/Attention performance in the AVCocktail dataset, using ASD for segmentation. A total of six experiments were conducted. Training with only conventional AVSR datasets (LRS2 and Vox2) results in a WER of 58.8%. Adding speech noise improves performance by 44.7% relative (WER reduced to 32.5%). Incorporating the AVYT-talking-face dataset further reduced the WER by 7.1%, reaching 30.2%. Dialog augmentation alone had a minimal impact, slightly decreasing the WER to 29.4%. Using full AVYT (both talking-face and silent-face sets) without dialog augmentation led to a WER of 28.5%. The most substantial improvement was achieved by combining dialog augmentation with the full AVYT dataset, achieving a significantly lower WER of 22.6%.

## 6. Conclusion

In this study, we benchmarked SOTA AVSR models, which perform impressively on conventional datasets like LRS2/LRS3 but struggle with cocktail-party scenarios. We highlighted the gap between conventional datasets and real-world cocktail-party scenarios, where target speakers are not always active. The presence of silent-face segments significantly impacts AVSR model performance, as the model tends to hallucinate output. To address this gap, we introduced the AVYT dataset and a data augmentation pipeline to improve model robustness. Additionally, we created AVCocktail, the first English audio-visual cocktail-party benchmark, to evaluate AVSR performance in realistic multi-speaker settings. All datasets and models are publicly available for further research at: <https://github.com/nguyenvulebinh/AVSRCocktail>

## 7. Acknowledgment

The authors gratefully acknowledge support from Carl Zeiss Stiftung under the project Jung bleiben mit Robotern (P2019-01-002). This work was also partially supported by the European Union's Horizon research and innovation programme (grant No. 101135798, project Meetween), the Volkswagen Foundation project "How is AI Changing Science? Research in the Era of Learning Algorithms" (HiAICS), and KIT Campus Transfer GmbH (KCT) staff in accordance to the collaboration with Carnegie-AI, as well as the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg and BMBF. Part of this research was supported by a grant from Zoom Video Communications.

## 8. References

- [1] H. McGurk and J. MacDonald, "Hearing lips and seeing voices," *Nature*, vol. 264, no. 5588, pp. 746–748, 1976.
- [2] M. Anwar, B. Shi, V. Goswami, W.-N. Hsu, J. Pino, and C. Wang, "Muavic: A multilingual audio-visual corpus for robust speech recognition and robust speech-to-text translation," in *Interspeech 2023*, 2023, pp. 4064–4068.
- [3] T.-S. Nguyen, S. Stüker, and A. Waibel, "Super-human performance in online low-latency recognition of conversational speech," *arXiv preprint arXiv:2010.03449*, 2020.
- [4] T. Zeppenfeld, M. Finke, K. Ries, M. Westphal, and A. Waibel, "Recognition of conversational telephone speech using the janus speech engine," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*. IEEE, 1997.
- [5] P. Duchowski, U. Meier, and A. Waibel, "See me, hear me: integrating automatic speech recognition and lip-reading," in *3rd International Conference on Spoken Language Processing (ICSLP 1994)*, 1994, pp. 547–550.
- [6] A. Waibel, M. T. Vo, P. Duchowski, and S. Manke, "Multimodal interfaces," *Artificial Intelligence Review*, vol. 10, 1996.
- [7] R. Stiefelwagen, J. Yang, and A. Waibel, "Modeling focus of attention for meeting indexing," in *Proceedings of the seventh ACM international conference on Multimedia (Part 1)*, 1999, pp. 3–10.
- [8] U. Bub, M. Hunke, and A. Waibel, "Knowing who to listen to in speech recognition: Visually guided beamforming," in *1995 International Conference on Acoustics, Speech, and Signal Processing*, vol. 1. IEEE, 1995, pp. 848–851.
- [9] J. Yang, R. Stiefelwagen, U. Meier, and A. Waibel, "Visual tracking for multimodal human computer interaction," in *Proceedings of the SIGCHI conference on Human factors in computing systems*, 1998.
- [10] R. Stiefelwagen, J. Yang, and A. Waibel, "Estimating focus of attention based on gaze and sound," in *Proceedings of the 2001 workshop on Perceptive user interfaces*, 2001, pp. 1–9.
- [11] A. Waibel, H. Steusloff, R. Stiefelwagen *et al.*, "Chil: Computers in the human interaction loop," 2005.
- [12] B. Shi, W.-N. Hsu, and A. Mohamed, "Robust self-supervised audio-visual speech recognition," in *Interspeech 2022*, 2022.
- [13] P. Ma, A. Haliassos, A. Fernandez-Lopez, H. Chen, S. Petridis, and M. Pantic, "Auto-avsr: Audio-visual speech recognition with automatic labels," in *ICASSP*, 2023, pp. 1–5.
- [14] A. Rouditchenko, Y. Gong, S. Thomas, L. Karlinsky, H. Kuehne, R. Feris, and J. Glass, "Whisper-flamingo: Integrating visual features into whisper for audio-visual speech recognition and translation," in *Interspeech 2024*, 2024, pp. 2420–2424.
- [15] G.-L. Chao, W. Chan, and I. Lane, "Speaker-targeted audio-visual models for speech recognition in cocktail-party environments," in *Interspeech 2016*, 2016, pp. 2120–2124.
- [16] Y. Wu and *et al.*, "Audio-visual multi-talker speech recognition in a cocktail party," in *Interspeech*, 2021.
- [17] J. Li, C. Li, Y. Wu, and Y. Qian, "Robust audio-visual asr with unified cross-modal attention," in *ICASSP*, 2023, pp. 1–5.
- [18] A. Ephrat, I. Mosseri, O. Lang, T. Dekel, K. Wilson, A. Hassidim, W. T. Freeman, and M. Rubinstein, "Looking to listen at the cocktail party: a speaker-independent audio-visual model for speech separation," *ACM Trans. Graph.*, vol. 37, no. 4, Jul. 2018. [Online]. Available: <https://doi.org/10.1145/3197517.3201357>
- [19] R. Gao and K. Grauman, "Visualvoice: Audio-visual speech separation with cross-modal consistency," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2021, pp. 15 490–15 500.
- [20] 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 *ICASSP*, 2024, pp. 12 632–12 636.
- [21] J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Lip reading sentences in the wild," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 3444–3453.
- [22] T. Afouras, J. S. Chung, and A. Zisserman, "Lrs3-ted: a large-scale dataset for visual speech recognition," 2018. [Online]. Available: <https://arxiv.org/abs/1809.00496>
- [23] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," in *Interspeech 2018*, 2018, pp. 1086–1090.
- [24] J. Li, C. Li, Y. Wu, and Y. Qian, "Unified cross-modal attention: Robust audio-visual speech recognition and beyond," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, pp. 1941–1953, 2024.
- [25] H. Chen and *et al.*, "The first multimodal information based speech processing (misp) challenge: Data, tasks, baselines and results," in *ICASSP*, 2022, pp. 9266–9270.
- [26] H. Chen, S. Wu, Y. Dai, Z. Wang, J. Du, C.-H. Lee, J. Chen, S. Watanabe, S. M. Siniscalchi, O. Scharenborg, D.-Y. Liu, B.-C. Yin, J. Pan, J.-Q. Gao, and C. Liu, "Summary on the multimodal information based speech processing (misp) 2022 challenge," in *ICASSP*, 2023, pp. 1–2.
- [27] H. Chen and *et al.*, "Summary on the multimodal information-based speech processing (misp) 2023 challenge," in *ICASSPW*, 2024, pp. 123–124.
- [28] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," 2022. [Online]. Available: <https://arxiv.org/abs/2212.04356>
- [29] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [30] P. Ma, S. Petridis, and M. Pantic, "End-to-end audio-visual speech recognition with conformers," in *ICASSP*, 2021, pp. 7613–7617.
- [31] F. Retkowski and A. Waibel, "From text segmentation to smart chaptering: A novel benchmark for structuring video transcriptions," in *EACL*, Mar. 2024, pp. 406–419. [Online]. Available: <https://aclanthology.org/2024.eacl-long.25>
- [32] J. Liao, H. Duan, K. Feng, W. Zhao, Y. Yang, and L. Chen, "A light weight model for active speaker detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 22 932–22 941.
- [33] J. S. Chung and A. Zisserman, "Out of time: Automated lip sync in the wild," in *ACCV Workshops*, 2016. [Online]. Available: <https://api.semanticscholar.org/CorpusID:26294509>
- [34] T.-B. Nguyen and A. Waibel, "Synthetic conversations improve multi-talker asr," in *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2024, pp. 10 461–10 465.
- [35] T.-B. Nguyen and Waibel, "Msa-asr: Efficient multilingual speaker attribution with frozen asr models," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.