

# Personalized Lip Reading: Adapting to Your Unique Lip Movements with Vision and Language

Jeong Hun Yeo<sup>1</sup>, Chae Won Kim<sup>1</sup>, Hyunjun Kim<sup>1</sup>,  
Hyeongseop Rha<sup>1</sup>, Seunghee Han<sup>1</sup>, Wen-Huang Cheng<sup>2</sup>, Yong Man Ro<sup>1\*</sup>

<sup>1</sup>Integrated Vision and Language Lab, KAIST, South Korea

<sup>2</sup>National Taiwan University, Taipei, Taiwan

sedne246@kaist.ac.kr, chaewonkim@kaist.ac.kr, kimhj709@kaist.ac.kr,  
ryool\_1832@kaist.ac.kr, gkstmtdgml211@kaist.ac.kr, wenhuang@csie.ntu.edu.tw, ymro@kaist.ac.kr

## Abstract

Lip reading aims to predict spoken language by analyzing lip movements. Despite advancements in lip reading technologies, performance degrades when models are applied to unseen speakers due to their sensitivity to variations in visual information such as lip appearances. To address this challenge, speaker adaptive lip reading technologies have advanced by focusing on effectively adapting a lip reading model to target speakers in the visual modality. However, the effectiveness of adapting language information, such as vocabulary choice, of the target speaker has not been explored in previous works. Additionally, existing datasets for speaker adaptation have limited vocabulary sizes and pose variations, which restrict the validation of previous speaker-adaptive methods in real-world scenarios. To address these issues, we propose a novel speaker-adaptive lip reading method that adapts a pre-trained model to target speakers at both vision and language levels. Specifically, we integrate prompt tuning and the LoRA approach, applying them to a pre-trained lip reading model to effectively adapt the model to target speakers. Furthermore, to validate its effectiveness in real-world scenarios, we introduce a new dataset, VoxLRS-SA, derived from VoxCeleb2 and LRS3. It contains a vocabulary of approximately 100K words, offers diverse pose variations, and enables the validation of adaptation methods in the wild, sentence-level lip reading for the first time in English. Through various experiments, we demonstrate that the existing speaker-adaptive method also improves performance in the wild at the sentence level. Moreover, we show that the proposed method achieves larger improvements compared to the previous works.

**Code** — <https://bit.ly/Personalized-Lip-Reading>

**Datasets** — <https://bit.ly/VoxLRS-SA-Dataset>

**Extended version** — <https://bit.ly/2409-00986>

## Introduction

Lip reading, also known as Visual Speech Recognition (VSR), aims to predict spoken language by analyzing visual cues from lip movements. It plays a crucial role in enhancing human communication in bustling cafes, crowded streets, or

loud factories where audio signals are compromised. Furthermore, it can support individuals with hearing impairments, enabling them to comprehend speech through visual means. With these practical benefits and the development of deep learning, we can accurately infer what a speaker says without a speaker’s voice by utilizing recent lip reading systems (Shi et al. 2022; Ma et al. 2023; Yeo et al. 2024b).

Despite significant advancements in lip reading technologies, their performance often declines when applied to unseen speakers not included in training data (Assael et al. 2016; Kim, Kim, and Ro 2022). This limitation arises because lip reading models are inherently sensitive to variations in lip appearances, movements, and speaking speeds among different individuals. To address this issue, recent studies (Kim, Kim, and Ro 2023; He et al. 2024; Luo et al. 2023; Wu et al. 2024) have incorporated speaker adaptation techniques into lip reading. These techniques aim to adapt pre-trained lip reading models to unseen speakers by utilizing minimal additional training parameters. Specifically, these works have applied prompt tuning (Brown et al. 2020; Zhou et al. 2022) and Low-Rank Adaptation (LoRA) (Hu et al. 2021) strategies to improve the adaptability and performance of lip reading models on target speakers.

Although recent studies have advanced our understanding of speaker adaptation in lip reading, existing approaches suffer from several limitations. Firstly, previous works have focused solely on visual cues for speaker adaptation, overlooking the role of individual linguistic patterns. While language-level adaptation for target speakers has demonstrated effectiveness in the NLP domain (Kolár, yang, and Shriberg 2007), these effects remain unexplored in speaker-adaptive lip reading. Secondly, previous studies such as (Kim, Kim, and Ro 2022) and (He et al. 2024) primarily utilize the GRID and LRW-ID datasets, which do not adequately capture the complexity of real-world scenarios. As shown in Table 1, these datasets especially lack pose diversity and have a limited vocabulary size.

To address these limitations, we propose a novel method to adapt the model to unseen target speakers in vision and language levels. This methodology comprises two primary components: 1) For vision level adaptation, the proposed method adjusts the pre-trained lip reading model to adapt to the lip appearance, movements, and speaking speed of the

\*Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Dataset	Segment	Split	# speakers	Duration (hrs)	# words	Pose & Background Diversity	Transcription labels
GRID (Cooke et al. 2006)	Baseline & Adaptation	-	34	27	51	✗	Sentence
	Baseline	Train	17,560	155	500		
LRW-ID (Kim, Kim, and Ro 2022)	Baseline	Train(Valid)	20	9.6	500	✓	Word
	Adaptation	Test	20	9.6	500		
	Baseline	Train	9,621	1,716	103,418		
VoxLRS-SA	Baseline	Train	20	18.4	10,912	✓	Sentence
	Adaptation	Valid	20	3.4	4,654		
	Adaptation	Test	20	2.0	3,460		

Table 1: Comparison of our dataset with publicly available speaker adaptive lip reading datasets. This table outlines the differences in dataset configurations across various factors such as the number of words, pose, and background diversity.

target unseen speaker while minimizing the number of training parameters. 2) For language level adaptation, we aim to adapt the pre-trained model to an individual’s unique linguistic patterns such as frequent vocabulary choice at the language level, when predicting the target speaker’s spoken language.

Specifically, we effectively design the vision-level adaptation by employing both padding prompts (Kim, Kim, and Ro 2022) and LoRA (He et al. 2024) in the visual encoder, different from previous speaker adaptive lip reading methods solely relying on prompt or LoRA. Moreover, we focus on two components within the visual encoder. Initially, the padding prompt with LoRA is applied to the spatial encoding stage, which is crucial for adapting to the lip appearances of the target speaker. Subsequently, it is applied to the temporal encoding stage, enabling the model to better account for variations in lip movement speeds and styles specific to each speaker, thus enhancing the model’s accuracy in capturing dynamic visual cues. For language-level adaptation, we apply input prompt tuning to a decoder based on the fact that input-level modification of pre-trained models are effective for adapting them to different tasks (Brown et al. 2020; Zhou et al. 2022). With input prompt tuning, the decoder learns the probability of language modeling specific to the speaker when predicting the spoken language of the target speaker.

To validate our proposed speaker adaptation method in real-world scenarios, we introduce a new dataset named VoxLRS-SA, derived from VoxCeleb2 (Chung, Nagrani, and Zisserman 2018) and LRS3 (Afouras, Chung, and Zisserman 2018) datasets. Together, these sources consist of about 1700 hours of YouTube videos, and as a result, VoxLRS-SA naturally captures diverse speaker poses and contains a rich vocabulary. Moreover, the VoxLRS-SA dataset combines the strengths of its parent datasets: LRS3, which includes text transcription labels but lacks specific speaker information, and VoxCeleb2, which provides detailed speaker information but does not include text transcriptions. To address these gaps, we utilize a pre-trained Automatic Speech Recognition (ASR) model (Radford et al. 2023) to generate automatic text labels (Ma et al. 2023; Yeo et al. 2024c) for the VoxCeleb2 dataset. Additionally, we employ a face recognition technology (Serengil and Ozpinar 2024) to generate speaker information for the LRS3 dataset.

The key contributions of this paper are as follows: 1) To the best of our knowledge, this is the first speaker adaptive lip reading method leveraging multimodal information at vision and language levels and exploring its effectiveness in the wild sentence-level lip reading. Moreover, the proposed method effectively adapts to a target speaker by employing both prompts and LoRA approaches. 2) We introduce the VoxLRS-SA dataset to validate speaker adaptive lip reading methods in the wild, sentence-level lip reading in English. It contains diverse poses and a vocabulary size of approximately 100K. 3) Through comprehensive experimental validation on the VoxLRS-SA dataset, we show that previous speaker-adaptive approaches also improve sentence-level performance in real-world settings. Furthermore, our newly proposed adaptation method surpasses previous methods in improving outcomes when tailored to specific target speakers.

## Related Work

### Lip Reading

Lip reading is the task of predicting spoken language by analyzing visual cues from lip movements. This technology has significantly advanced with the development of deep learning and the availability of large-scale audio-visual lip reading databases.

In word-level lip reading, the introduction of publicly available datasets such as LRW (Chung and Zisserman 2017) and LRW-1000 (Yang et al. 2019) has been crucial. Based on these databases, (Stafylakis and Tzimiropoulos 2017) designed 3D convolutional layers and 2D ResNet architectures as the front-end, with LSTM networks as the back-end. Additionally, two-stream networks, which utilize both raw video and optical flow, have been proposed by (Xiao et al. 2020; Weng and Kitani 2019). Some works (Kim, Yeo, and Ro 2022; Yeo, Kim, and Ro 2023) have used audio signals to enhance the visual representations by learning visual-to-audio mapping in the memory network.

In sentence-level lip reading, the release of the large-scale LRS2 (Afouras et al. 2018) and LRS3 (Afouras, Chung, and Zisserman 2018) databases has been important in English. With the development of these datasets, an end-to-end model trained using Connectionist Temporal Classification (CTC) loss (Graves et al. 2006) was proposed by (Assael et al. 2016). Subsequently, (Son Chung et al. 2017) devel-

oped a model utilizing the Seq2Seq architecture (Sutskever, Vinyals, and Le 2014), and further architectural enhancements were achieved by (Afouras et al. 2018) with the implementation of the Transformer model (Vaswani et al. 2017). Recently, state-of-the-art models (Ma, Petridis, and Pantic 2021) have employed a joint CTC/Attention (Watanabe et al. 2017) loss with the Conformer (Gulati et al. 2020) architecture. (Shi et al. 2022; Haliassos et al. 2022) utilized audio-visual data to train backbone models in a self-supervised manner and proved its effectiveness. In the context of other languages, LIP-RTVE dataset (Gimeno-Gómez and Martínez-Hinarejos 2023) provides a resource for automatic lip reading in continuous Spanish in the wild, featuring both speaker-independent and speaker-dependent partitions. Furthermore, with the development of self-supervised learning using audio-visual data, (Kim et al. 2024, 2023) have begun to explore lip reading technologies for multiple languages.

Although there has been great development in lip reading, these systems suffer from degraded performance when applied to unseen speakers not included in the training dataset, due to individual variations in lip appearances and movements. Therefore, it is necessary to explore speaker adaptation techniques to narrow the performance gap between seen and unseen speakers in lip reading. However, existing datasets (Kim, Kim, and Ro 2022; Cooke et al. 2006) for speaker-adaptive lip reading do not cover real-world scenarios for English, due to their limited vocabulary size and lack of pose diversity. To address this issue, we propose a novel dataset VoxLRS-SA, which contains about 100K vocabulary size and diverse pose videos.

## Speaker Adaptation

Speaker adaptation (Miao, Zhang, and Metze 2015) was primarily developed for Automatic Speech Recognition (ASR). An early work (Liao, McDermott, and Senior 2013) fine-tuned various parts of the network to adapt the model to a target speaker. Other studies (Li and Sim 2010; Seide et al. 2011) suggested adding speaker-dependent layers to a pre-trained model. In (Swietojanski and Renals 2014), speaker-dependent vectors were added to all pre-trained hidden layers to adjust them for the test speaker. To provide additional inputs that vary by speaker, (Abdel-Hamid and Jiang 2013b,a; Xue et al. 2014) proposed speaker codes. Recently, works (Klejšch et al. 2019; Huang et al. 2020) have explored the utilization of meta-learning and speech synthesis for speaker adaptation, respectively.

In contrast to the efforts for speaker adaptation in ASR, only a few works handled speaker adaptation for lip reading. Motivated by the traditional speaker adaptation methods (Anastasakos, McDonough, and Makhoul 1997; Gopinath 1998) of ASR, (Almajai et al. 2016) introduced utilize these works into lip reading. Recently, (Kim, Kim, and Ro 2022) proposed adapting the CNN layer of a visual encoder to the speaker’s unique features, such as the shape of the lips, by using user-dependent padding. Moreover, (Kim, Kim, and Ro 2023; He et al. 2024) developed speaker adaptive lip reading by adapting the visual encoder to target the speaker at the input, spatial, and temporal levels. (Li et al. 2024; Wu

et al. 2024) proposed using lip landmarks to mitigate the sensitivity of lip reading systems to individual variation information such as lip appearances.

Prior work has concentrated primarily on visual cues for speaker adaptation, overlooking the importance of individual linguistic patterns. (Kolár, yang, and Shriberg 2007; Guz et al. 2010) have utilized a speaker’s lexical pattern for dialogue act and sentence segmentation tasks. However, leveraging linguistic cues has not yet been integrated into speaker-adaptive lip reading. By incorporating multimodal information, we improve speaker-specific adaptations, providing a more holistic approach to speaker adaptation.

## Method

In this paper, the primary goal of our proposed method is to leverage individual characteristics at both the vision and language levels to improve the understanding of speaker-specific information, thereby offering a more comprehensive approach to speaker adaptation.

Recent work (Yeo et al. 2024a) proves that the LLM’s context modeling capability is particularly effective in lip reading systems. Motivated by this, we employ this framework as our baseline lip reading model. As illustrated in Figure.1, the architecture of the lip reading model consists of two components, a visual encoder for modeling spatio-temporal information from lip movements, and an LLM as a decoder for predicting the sentence by taking visual features as input. To adapt the pre-trained lip reading model to target speakers, we propose a two-fold approach that addresses both the vision and language levels.

## Vision Level Adaptation

The visual encoder of the pre-trained lip reading model consists of two parts: the front-end  $\mathcal{F}$ , which has a CNN-based architecture that encodes spatial information, and the back-end  $\mathcal{B}$ , which has a Conformer-based architecture that models temporal dependencies between lip sequences. Based on these two modules of the visual encoder, we specifically aim to adapt the visual encoder to lip appearances, movements, and the speaking speed of target speakers via vision-level adaptation.

Given the speaker-specific training dataset for speaker adaptation  $\tau = \{(x_i, y_i)\}_{i=1}^N$ , where  $x_i \in \mathbb{R}^{T \times H \times W \times C}$  is the lip movements video of the target unseen speaker,  $T$  is the frame length of the video,  $y_i$  is the text transcriptions corresponding to its lip movements, and  $N$  is the number of training samples. To adapt the appearances of the target speaker, a padding prompt (Kim, Kim, and Ro 2023) and LoRA (Hu et al. 2021) are employed in the CNN layer of the front-end. These learnable parameters enable the front-end to adapt to the target speaker without any modifications to the pre-trained weights. Specifically, the padding prompt  $p$  replaces the zero padding of the CNN layer. At the same time, the CNN layer is reconstructed with the low-rank decomposition matrices. With these two additional parameters, the spatial information encoding process can be formulated as follows:  $f_s = \mathcal{F}_{\theta,p}(x_i)$ , where  $f_s \in \mathbb{R}^{T \times D}$  spatial feature of lip movements,  $D$  is the embedding dimension of the

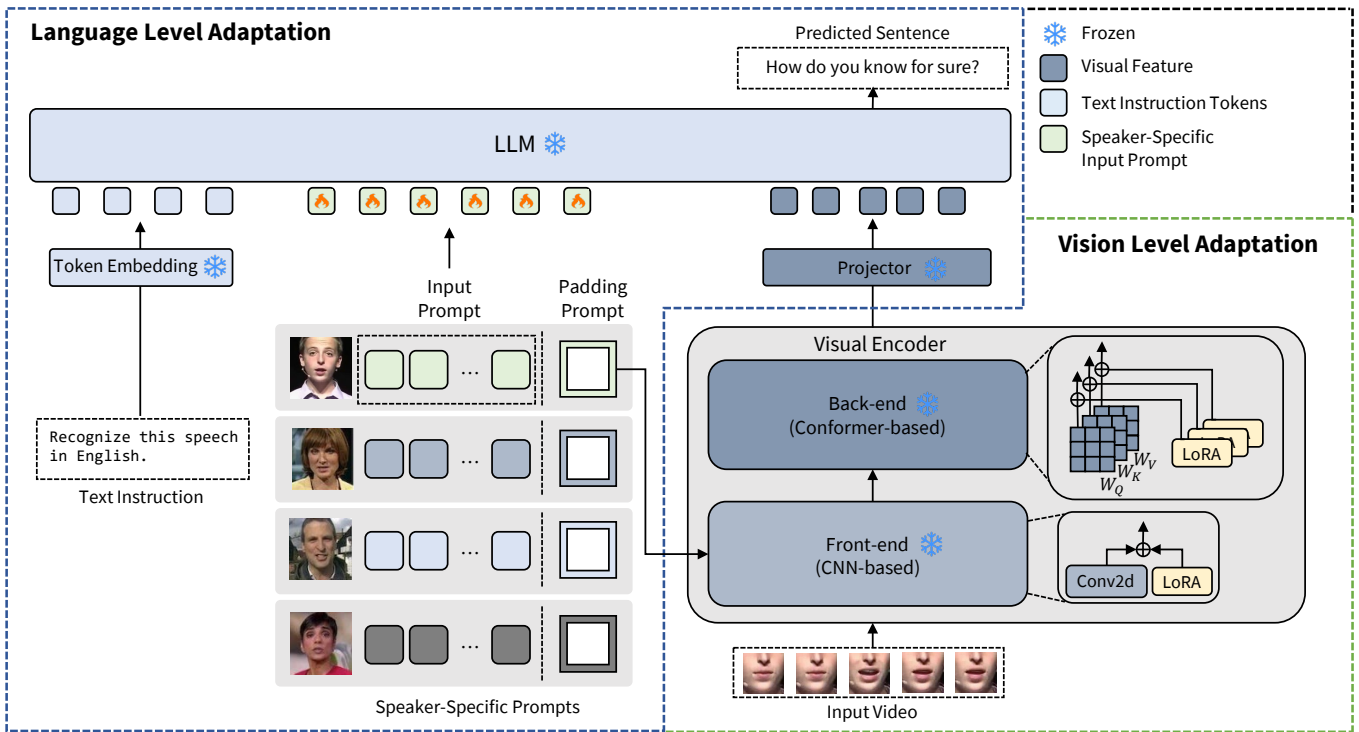


Figure 1: Overview of the proposed lip reading model architecture illustrating dual adaptation strategies. Visual level adaptation employs padding prompts and LoRA at the front-end for unique lip features and at the back-end for lip movement speed and style. Language level adaptation uses input prompt tuning to learn the language modeling probability of each speaker.

front-end,  $\theta$  indicates the weights of the front-end with the trainable parameters constructing the low-rank decomposition matrices, and  $p$  indicates the replacement zero-padding with the padding prompts in the front-end.

For a better modeling of the speaker-specific temporal information such as speaking speed, the back-end needs to be adjusted. In order to effectively adjust without pre-trained weights of the back-end, LoRA is also utilized in the back-end. Specifically, LoRA is applied to the attention layer to learn the temporal dependencies between any two positions of speaker-specific lip movements. It can be formulated as follows:  $f_v = \mathcal{B}_\theta(f_s)$ , where  $f_v \in \mathbb{R}^{T \times D}$  is a visual feature containing spatio-temporal information of lip movements, and  $\theta$  indicates the weights of the back-end with the trainable parameters constructing the low-rank decomposition matrices. The visual features encoded from the visual encoder are fed into the projector, to align it with the embedding space of the LLM. These aligned visual features are denoted as  $f'_v \in \mathbb{R}^{T \times D_L}$ , where  $D_L$  is the embedding dimension of LLM. With the frozen LLM taking aligned visual features as input, the predicted sentence  $\hat{y}_i$  is obtained. From this prediction, the cross entropy loss  $\mathcal{L} = CE(\hat{y}_i, y_i)$  is calculated, and the visual encoder is optimized to adapt to the target speaker at the vision-level.

### Language Level Adaptation

Speaker-specific characteristics encompass not only visual cues such as the appearance of lips but also linguistic pat-

terns, including preferred vocabulary choices. Since individuals exhibit unique linguistic patterns, the likelihood of specific sequences of words occurring varies from person to person. To learn these speaker-specific linguistic patterns, our approach utilizes an input prompt tuning motivated by (Brown et al. 2020; Zhou et al. 2022; Lester, Al-Rfou, and Constant 2021).

Specifically, given that the visual feature  $f'_v$  is already aligned with the embedding space of the LLM, the input prompts, as trainable parameters, are concatenated with these features along the temporal dimension and then fed into the LLM as input. Consequently, the input prompt is designed to have the same dimension as that of the LLM embedding space dimension and can be denoted as  $p_i \in \mathbb{R}^{N_p \times D_L}$ , where  $N_p$  is the length of the prompts. With the self-attention layers in the LLM and model temporal dependencies between input prompts and visual features, these prompts can influence speaker-specific sentence predictions. It can be formulated as follows:  $\hat{y}_i = \mathcal{D}(p_i \oplus f'_v)$ , where the  $\oplus$  is the concatenate function and  $\mathcal{D}$  represents the LLM functioning as the decoder for sentence prediction. In addition, the input prompts are trained to predict more accurate language modeling probabilities for specific speakers using target speaker data. This objective is achieved by minimizing the cross-entropy loss, which can be formulated as  $\mathcal{L} = CE(\hat{y}_i, y_i)$ , which optimizes the input prompt for language-level adaptation.

Split	Data Information	Speaker Number											
		S1   S2	S3   S4	S5   S6	S7   S8	S9   S10	S11   S12	S13   S14	S15   S16	S17   S18	S19   S20		
Train	# videos	280   261	281   291	268   280	297   283	294   253	308   316	286   304	279   315	302   322	305   317		
	Duration(min)	45.05   45.06	45.09   45.13	45.08   45.12	45.10   45.05	45.12   45.26	45.05   45.00	45.07   45.01	45.10   44.99	45.00   45.05	45.09   45.06		
	# words	1585   1489	1796   1645	1433   1409	1535   1671	1466   1614	1535   1400	1637   1394	1344   1560	1548   1409	1485   1423		
Valid	# videos	94   85	78   83	74   74	68   66	74   63	71   61	68   62	52   60	47   42	38   44		
	Duration(min)	14.07   13.87	13.01   12.01	11.83   11.46	11.51   11.30	11.16   10.77	10.31   9.66	9.46   9.22	8.5   8.27	7.66   6.00	5.92   5.85		
	# words	774   720	800   731	657   641	674   726	655   692	591   569	619   548	483   540	507   444	441   392		
	overlap ratio	0.33   0.27	0.26   0.25	0.27   0.27	0.28   0.25	0.28   0.27	0.24   0.28	0.24   0.25	0.23   0.21	0.22   0.21	0.19   0.20		
Test	# videos	41   35	35   41	35   43	33   50	34   41	37   41	40   34	44   43	35   42	39   46		
	Duration(min)	6.12   6.10	6.15   6.10	6.04   6.09	6.08   6.12	6.09   6.08	6.12   6.10	6.16   6.10	6.16   6.15	6.15   6.15	6.11   6.15		
	# words	461   425	507   459	427   438	425   419	419   491	396   468	499   396	368   425	402   438	473   388		
	overlap ratio	0.18   0.18	0.18   0.17	0.21   0.20	0.17   0.18	0.20   0.19	0.18   0.20	0.16   0.19	0.19   0.18	0.19   0.21	0.20   0.19		

Table 2: Detailed data information of 20 test speakers in VoxLRS-SA

## VoxLRS-SA

To validate the effectiveness of the proposed method in real-world scenarios for English, we develop the VoxLRS-SA dataset based on LRS3 and VoxCeleb2 datasets. While the widely used LRS3 dataset for sentence-level lip reading in the wild includes human-annotated transcription labels, it does not contain Speaker ID information for each video. In contrast, the VoxCeleb2 dataset includes Speaker ID information but does not provide text transcriptions. To address this gap, we generate the speaker ID information for the LRS3 dataset by using face verification (Kim, Kim, and Ro 2022) and make pseudo text labels for the VoxCeleb2 dataset by employing a pre-trained ASR model (Ma et al. 2023; Yeo et al. 2024c).

**Automatic Labels in VoxCeleb2** Recent works (Ma et al. 2023; Yeo et al. 2024c) have demonstrated that a lip reading model trained using automatically generated transcriptions through a pre-trained ASR model performs similarly to one trained with human-annotated labels. Motivated by these works, we utilize the Whisper (Radford et al. 2023) ASR model to generate automatic labels for the videos classified as the English portion of the VoxCeleb2 dataset. Then, we choose 20 speakers who have more than 50 minutes of video content to construct the test and adaptation (or validation) sets. Additionally, we manually clean the text transcriptions of the test set to accurately evaluate the effectiveness of the speaker adaptation method.

**Speaker ID information annotation** Since the pseudo-labels are generated using the pre-trained ASR model, their transcription quality may be worse than human-annotated labels. To mitigate this issue, we combine VoxCeleb2 with the LRS3 dataset, which contains manually annotated transcriptions corresponding to 433 hours of videos.

Since the LRS3 dataset does not include speaker ID information, we first generate the speaker ID information for both datasets using a method similar to the one described in (Kim, Kim, and Ro 2022), employing a face recognition model. Since there may be overlapping speakers between the two datasets, we take the following approach: For each video in LRS3 and VoxCeleb2, we divide the video into three equal segments and extract one face image from each segment.

Then, we employ the Facenet (Serengil and Ozpinar 2024) face recognition model to encode the facial images into features of speaker information. To identify overlapping speakers between the two datasets, we compare the feature vectors from LRS3 and VoxCeleb2. This comparison is done using a similarity measure, such as cosine similarity, to identify the overlapping speakers.

## Experimental Setup

### Dataset

**VoxLRS-SA** is divided into two parts corresponding to each purpose. Firstly, the Baseline part of VoxLRS-SA is designed to train the baseline lip reading model. This training dataset is composed of 1.7K hours of video data from about 9K speakers. Secondly, the Adaptation part is proposed to adapt the pre-trained lip reading model to the target speaker and evaluate the effectiveness of the adaptation method. For this purpose, we construct the Adaptation part with 20 speakers, which do not overlap with the 9K speakers of the Baseline part. The *train* split contains video data varying from 45 minutes per speaker, and the *test* and *val* sets contain about 5 minutes of data for each of the 20 speakers. The detailed data information is shown in Table 2.

**VoxCeleb2** (Chung, Nagrani, and Zisserman 2018) is a speaker verification dataset. It contains only speaker information for over 6,000 speakers and 2,442 hours of multilingual videos. By referring to the (Shi et al. 2022), we utilize an English portion, which amounts to 1,326 hours.

**LRS3** (Afouras, Chung, and Zisserman 2018) is a widely used sentence-level lip reading dataset comprising 433 hours of English talking face videos and human-annotated transcriptions. These talking face videos are collected from TED and TEDx talks, thus contain diverse poses.

### Implementation details

**Preprocessing** By following a preprocessing process (Ma et al. 2023), we crop the mouth region from talking face videos into a size of 96 pixels in both width and height. We employ RetinaFace (Deng et al. 2020) to detect facial landmarks and crop the mouth region based on these landmarks.

For data augmentation (Shi et al. 2022), horizontally flipping and randomly cropping are used.

**Architecture** The design of our visual encoder incorporates a conformer architecture, configured with an embedding dimension of 768, 12 attention heads, a feed-forward dimension of 3072, and 12 blocks. The configuration of the transformer decoder, which is used for pre-training the visual encoder, includes an embedding dimension of 768, 8 attention heads, a feed-forward dimension of 3072, and 9 blocks. We employ the LLaMA3-8B (Touvron et al. 2023) model as the LLM for our proposed method.

**Baseline Lip Reading Model** We initialize the visual encoder with weights from a pre-trained lip reading model trained on the LRS3 dataset (Ma et al. 2023). The visual encoder is fine-tuned alongside a transformer decoder from scratch on the Baseline part of the VoxLRS-SA dataset. During the first stage of training, we employ a tri-stage learning rate with a peak learning rate of  $1e-3$ , warmup steps of 10K, decay steps of 20K, and utilize 8 NVIDIA RTX 3090 GPUs with a maximum frame count of 1800 and gradient accumulation set to 4. In a further step of our experiment, we replace the transformer decoder with the frozen LLaMA3-8B model, continuing the training on the same datasets. In this second stage of training, we adopt a cosine learning rate strategy with a learning rate of  $5e-5$ , total training steps of 30K, and a warmup period of 0.5K steps, using a batch size of 1 and increasing the gradient accumulation to 8.

**Speaker-Adaptive Lip Reading Model** For vision-level adaptation, a number of training steps is set to 300, and use a cosine learning rate scheduler without warmup. The learning rate begins at  $1e-4$  and gradually decreases to a minimum of  $1e-5$  over 5000 update periods. The LoRA configuration includes a rank of 8 and a scaling factor of 16. The target modules for LoRA adaptation are the convolutional layer, a query, a key, and a value in the self-attention layer. For language-level adaptation, the visual encoder is kept frozen to focus adaptation efforts on the language model. The LoRA configuration is applied to the attention layers of the LLM, specifically targeting the query, key, and value components. The number of training updates is set to 70.

## Experimental Results

### Baselines for comparisons

Since this is the first work to explore the effectiveness of speaker-adaptive lip reading at the wild sentence level in English, we have established comparison methods including previous word-level speaker-adaptive methods and full fine-tuning methods to validate the effectiveness of the proposed method.

**LoRA** (Hu et al. 2021) is a currently widely used technique for computationally efficient fine-tuning. We set three types of adaptation methods based on LoRA. The V LoRA and L LoRA are adapting the visual encoder at vision level adaptation and LLM at language-level adaptation, respectively.

**Padding Prompt** (Kim, Kim, and Ro 2022) utilize additional speaker-specific input in the convolution layer of the visual encoder. It enables the speaker-specific input to memorize the visual characteristics of the target speaker.

Method	Trainable Params		WER(%)
	Encoder	Decoder	
Baseline	181.3M*	8B*	47.3
<b>Vision-Level Adaptation</b>			
V LoRA	+0.6M	-	42.5
Padding Prompt	+0.1M	-	42.9
Fintune-F	11.2M	-	42.8
Fintune-B	170.1M	-	40.2
Finetune-F&B	181.3M	-	40.0
Ours	+0.7M	-	41.5
<b>Language-Level Adaptation</b>			
L LoRA	-	+4.7M	44.7
Ours	-	+0.04M	44.1
<b>Vision-and-Language-Level Adaptation</b>			
Finetune-F&B	181.3M	+0.04M	39.6
Ours	+0.7M	+0.04M	40.9

Table 3: This table compares the WER performance ( $\downarrow$ ) for various speaker-adaptive training strategies, detailing changes in trainable parameters for the encoder and decoder. \* indicates the number of total parameters.

**Finetune** By following (Kim, Kim, and Ro 2023), we fine-tune three parts. Finetune-F and -B are fine-tuning the front- and back-ends of the visual encoder, respectively. Finetune-F&B is fine-tuning the whole part of the visual encoder.

### Comparison with the previous method at Vision and Language Levels Adaptation

Table 3. provides details of experiments conducted to evaluate different training strategies in speaker adaptive lip reading. The baseline method consists of a considerable number of parameters, with 181.3 million in the encoder and 8 billion in the decoder, resulting in a Word Error Rate (WER) of 47.3%.

In the vision-level adaptation, we compare the proposed method with V LoRA, Padding Prompt, and three types of full fine-tuning to validate the effectiveness. V LoRA and Padding Prompt added a small amount of trainable parameters (0.6M and 0.1M, respectively) but achieved improvements in WER, reducing it to 42.5% and 42.9%, respectively. The full-fine-tuning methods Finetune F, B, and F&B, obtained substantial reductions, with results of 42.8%, 40.2%, and 40.0%, respectively. While these methods adapting the back-end modules show superior performances, they require training over 170M parameters. In contrast, our proposed method achieves a WER of 41.5% by training only 0.7M parameters. It is worth noting that the proposed method outperforms the Finetune-F method, which utilizes 11.2M parameters.

In the language-level adaptation, we have kept the parameters of the visual encoder frozen to focus on verifying the effectiveness of adapting speaker-specific language information. We compare the proposed method, which utilizes in-

Adapt Min	Speaker Number																		Mean(↓)
	S1   S2	S3   S4	S5   S6	S7   S8	S9   S10	S11   S12	S13   S14	S15   S16	S17   S18	S19   S20									
<b>Baseline</b>	42.2   43.1	59.5   48.2	32.7   83.8	41.9   80.4	32.3   64.9	49.3   37.0	66.9   31.8	33.5   48.2	53.3   31.0	29.6   35.9	47.3								
<b>1 min</b>	43.7   44.9	53.0   44.1	39.3   71.5	44.1   76.4	31.5   58.4	47.5   36.6	57.3   29.4	37.8   48.0	56.0   31.7	36.6   35.7	46.2								
<b>5 min</b>	40.9   42.7	50.5   41.3	30.9   68.5	43.9   70.5	29.8   56.5	44.1   37.7	57.5   28.4	33.4   45.2	52.6   28.8	29.4   34.4	43.4								
<b>15 min</b>	39.6   39.4	50.8   42.0	28.3   69.9	41.8   71.1	29.1   55.2	45.7   36.7	53.9   27.5	32.0   44.4	51.4   28.3	29.0   33.1	42.5								
<b>30 min</b>	37.9   38.7	51.9   41.7	26.8   67.9	40.2   66.7	29.7   55.9	42.1   34.8	55.1   28.3	32.8   43.9	48.1   26.0	28.2   32.2	41.4								
<b>45 min</b>	35.4   38.1	48.6   39.6	27.5   64.7	41.2   67.1	29.9   55.0	44.1   34.0	54.1   27.5	31.6   43.2	50.2   27.1	28.5   30.7	40.9								

Table 4: WER performance (↓) across different adaptation durations: This table shows the WER for 20 speakers, across different adaptation durations from 1 up to 45 minutes. Each duration denotes the total length of the video sample used for adaptation.

put prompt tuning, with the LoRA-based approach. By employing the LoRA approach in the LLM for language level adaptation, we obtain a WER of 44.7%. With the proposed method, a WER of 44.1% is achieved by using only 0.04M learnable parameters. Despite the proposed method utilizing fewer parameters compared to the LoRA approach, it shows better performance.

In the vision-and-language level adaptation, we compare the proposed method with the Finetune F&B, which shows the best performance in the vision-level adaptation. For a fair comparison, we also apply the language-level adaptation on Finetune F&B using our input prompt. The results are shown in the lower part of Table 3. With the Finetune F&B adaptation method, a WER of 39.6% is achieved, which is the best performance in these experiments. The proposed method achieves a marginally higher WER of 40.9% than Finetune F&B by utilizing adaptations on both vision and language levels for the target speaker. Despite showing marginally lower performance compared to Finetune F&B, the proposed method has the advantage of using only 0.74M parameters, which is less than 0.1% of the number of parameters of Finetune F&B.

### Analysis of Adaptation Duration Effects on Lip Reading Performances

One of the challenges in speaker-adaptive lip reading is that it is not easy to collect sufficient data for the target speaker. Therefore, we validate the effectiveness of the proposed method by using a small amount of adaptation data corresponding to 1 and 5 minutes. Moreover, in order to verify the effectiveness according to the amount of adaptation data, we additionally conduct experiments by utilizing 15, 30, and 45 minutes of adaptation data.

Table 4 illustrates the effect on target speaker lip reading performance according to variations of the amount of speaker-specific data, across 20 speakers. The WER performances of a baseline across 20 speakers range widely, with the lowest WER observed at 29.6 and the highest at 83.8, demonstrating variability in pre-trained baseline lip reading model accuracy across different speakers. The overall mean WER for the baseline is calculated at 47.3.

When we apply the proposed method across 20 speakers, the average performance is increased in all settings, regardless of the amount of adaptation data. Specifically, the

speaker-adaptive lip reading model trained using 1-minute speaker-specific data achieves an average of 46.2% WER, which is 1.2% better than the baseline. However, the performance of some speakers is worse than the baseline. This result is different from (Kim, Kim, and Ro 2022), which shows that 1 minute of adaptation data is enough to improve the performance of all speakers in word-level lip reading. In sentence-level speaker adaptive lip reading, it seems that there is a need to be over 1 minute of speaker-specific data, to adapt the pre-trained model to the target speaker.

By analyzing the results when using 5 minutes of speaker-specific data, the performance of the adaptive lip reading model outperforms that of the baseline for 19 out of 20 speakers. In other words, utilizing 5 minutes of speaker-specific data is enough to adapt the pre-trained model to target speakers. The adapted model achieves an average WER of 43.4%, which is a 3.9% improvement to the baseline model. Overall, employing more speaker-specific training data up to 45 minutes improves the average WER without saturation of performances, as seen in the average WERs of 42.5%, 41.4%, and 40.9% for 15, 30, and 45 minutes, respectively.

### Conclusion

This study has introduced a novel speaker-adaptive lip reading method that can be adapted to specific target speakers at both the vision and language levels. Specifically, we have adapted the visual encoder to the visual information of the target speaker, such as lip appearances and speaking speed, by using padding prompts and LoRA. Moreover, we have also adapted the LLM to speaker-specific language information by utilizing input prompt tuning for language-level adaptation. To validate the effectiveness of the proposed method, we have introduced a novel VoxLRS-SA dataset that expanded the vocabulary size and poses variations, compared to previous speaker-adaptive datasets in English. The experimental results have shown the necessity of speaker-adaptive lip reading in real-world scenarios. Moreover, our method not only improves sentence-level lip reading by using only a small amount of data but also outperforms existing speaker-adaptive methods. These findings suggest that incorporating both visual adaptation and language-specific characteristics of speakers can substantially benefit lip reading technologies.

## Acknowledgments

This work was partly supported by two funds: the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. NRF-2022R1A2C2005529) and IITP grant funded by the Korea government (MSIT) (No.2020-0-00004, Development of Previsional Intelligence based on Long-Term Visual Memory Network)

## References

- Abdel-Hamid, O.; and Jiang, H. 2013a. Fast speaker adaptation of hybrid NN/HMM model for speech recognition based on discriminative learning of speaker code. In *Proc. ICASSP*, 7942–7946. IEEE.
- Abdel-Hamid, O.; and Jiang, H. 2013b. Rapid and effective speaker adaptation of convolutional neural network based models for speech recognition. In *INTERSPEECH*, 1248–1252.
- Afouras, T.; Chung, J. S.; Senior, A.; Vinyals, O.; and Zisserman, A. 2018. Deep audio-visual speech recognition. *IEEE transactions on pattern analysis and machine intelligence*, 44(12): 8717–8727.
- Afouras, T.; Chung, J. S.; and Zisserman, A. 2018. LRS3-TED: a large-scale dataset for visual speech recognition. *arXiv preprint arXiv:1809.00496*.
- Almajai, I.; Cox, S.; Harvey, R.; and Lan, Y. 2016. Improved speaker independent lip reading using speaker adaptive training and deep neural networks. In *Proc. ICASSP*, 2722–2726. IEEE.
- Anastasakos, T.; McDonough, J.; and Makhoul, J. 1997. Speaker adaptive training: A maximum likelihood approach to speaker normalization. In *Proc. ICASSP*, volume 2, 1043–1046. IEEE.
- Assael, Y. M.; Shillingford, B.; Whiteson, S.; and De Freitas, N. 2016. Lipnet: End-to-end sentence-level lipreading. *arXiv preprint arXiv:1611.01599*.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Chung, J. S.; Nagrani, A.; and Zisserman, A. 2018. Voxceleb2: Deep speaker recognition. *arXiv preprint arXiv:1806.05622*.
- Chung, J. S.; and Zisserman, A. 2017. Lip reading in the wild. In *Computer Vision—ACCV 2016: 13th Asian Conference on Computer Vision, Taipei, Taiwan, November 20–24, 2016, Revised Selected Papers, Part II 13*, 87–103. Springer.
- Cooke, M.; Barker, J.; Cunningham, S.; and Shao, X. 2006. An audio-visual corpus for speech perception and automatic speech recognition. *The Journal of the Acoustical Society of America*, 120(5): 2421–2424.
- Deng, J.; Guo, J.; Ververas, E.; Kotsia, I.; and Zafeiriou, S. 2020. Retinaface: Single-shot multi-level face localisation in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5203–5212.
- Gimeno-Gómez, D.; and Martínez-Hinarejos, C.-D. 2023. LIP-RTVE: An audiovisual database for continuous Spanish in the wild. *arXiv preprint arXiv:2311.12457*.
- Gopinath, R. A. 1998. Maximum likelihood modeling with Gaussian distributions for classification. In *Proc. ICASSP*, volume 2, 661–664. IEEE.
- Graves, A.; Fernández, S.; Gomez, F.; and Schmidhuber, J. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *Proceedings of the 23rd international conference on Machine learning*, 369–376.
- Gulati, A.; Qin, J.; Chiu, C.-C.; Parmar, N.; Zhang, Y.; Yu, J.; Han, W.; Wang, S.; Zhang, Z.; Wu, Y.; et al. 2020. Conformer: Convolution-augmented transformer for speech recognition. *arXiv preprint arXiv:2005.08100*.
- Guz, U.; Cuendet, S.; Hakkani-Tur, D.; and Tur, G. 2010. Multi-View Semi-Supervised Learning for Dialog Act Segmentation of Speech. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(2): 320–329.
- Haliassos, A.; Ma, P.; Mira, R.; Petridis, S.; and Pantic, M. 2022. Jointly Learning Visual and Auditory Speech Representations from Raw Data. In *The Eleventh International Conference on Learning Representations*.
- He, Y.; Yang, L.; Wang, H.; Zhu, Y.; and Wang, S. 2024. Speaker-Adaptive Lipreading Via Spatio-Temporal Information Learning. In *Proc. ICASSP*, 10411–10415. IEEE.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Huang, Y.; He, L.; Wei, W.; Gale, W.; Li, J.; and Gong, Y. 2020. Using personalized speech synthesis and neural language generator for rapid speaker adaptation. In *Proc. ICASSP*, 7399–7403. IEEE.
- Kim, M.; Kim, H.; and Ro, Y. M. 2022. Speaker-adaptive lip reading with user-dependent padding. In *European Conference on Computer Vision*, 576–593. Springer.
- Kim, M.; Kim, H.-I.; and Ro, Y. M. 2023. Prompt tuning of deep neural networks for speaker-adaptive visual speech recognition. *arXiv preprint arXiv:2302.08102*.
- Kim, M.; Yeo, J.; Park, S. J.; Rha, H.; and Ro, Y. M. 2024. Efficient Training for Multilingual Visual Speech Recognition: Pre-training with Discretized Visual Speech Representation. In *ACM Multimedia*.
- Kim, M.; Yeo, J. H.; Choi, J.; and Ro, Y. M. 2023. Lip reading for low-resource languages by learning and combining general speech knowledge and language-specific knowledge. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15359–15371.
- Kim, M.; Yeo, J. H.; and Ro, Y. M. 2022. Distinguishing homophenes using multi-head visual-audio memory for lip reading. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 1174–1182.
- Klejch, O.; Fainberg, J.; Bell, P.; and Renals, S. 2019. Speaker adaptive training using model agnostic meta-learning. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 881–888. IEEE.

- Kolár, J.; yang, I.; and Shriberg, E. 2007. Speaker adaptation of language models for automatic dialog act segmentation of meetings. In *INTERSPEECH*, volume 4, 1621–1624.
- Lester, B.; Al-Rfou, R.; and Constant, N. 2021. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*.
- Li, B.; and Sim, K. C. 2010. Comparison of discriminative input and output transformations for speaker adaptation in the hybrid NN/HMM systems. In *Interspeech*, volume 10, 526–529.
- Li, Y.; Xue, F.; Wu, L.; Xie, Y.; and Li, S. 2024. Generalizing sentence-level lipreading to unseen speakers: a two-stream end-to-end approach. *Multimedia Systems*, 30(1): 1–10.
- Liao, H.; McDermott, E.; and Senior, A. 2013. Large scale deep neural network acoustic modeling with semi-supervised training data for YouTube video transcription. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, 368–373. IEEE.
- Luo, S.; Yang, S.; Shan, S.; and Chen, X. 2023. Learning Separable Hidden Unit Contributions for Speaker-Adaptive Lip-Reading. *arXiv preprint arXiv:2310.05058*.
- Ma, P.; Haliassos, A.; Fernandez-Lopez, A.; Chen, H.; Petridis, S.; and Pantic, M. 2023. Auto-avs: Audio-visual speech recognition with automatic labels. In *Proc. ICASSP*, 1–5. IEEE.
- Ma, P.; Petridis, S.; and Pantic, M. 2021. End-to-end audio-visual speech recognition with conformers. In *Proc. ICASSP*, 7613–7617. IEEE.
- Miao, Y.; Zhang, H.; and Metze, F. 2015. Speaker adaptive training of deep neural network acoustic models using i-vectors. *Proc. TASLP*, 23(11): 1938–1949.
- Radford, A.; Kim, J. W.; Xu, T.; Brockman, G.; McLeavey, C.; and Sutskever, I. 2023. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, 28492–28518. PMLR.
- Seide, F.; Li, G.; Chen, X.; and Yu, D. 2011. Feature engineering in context-dependent deep neural networks for conversational speech transcription. In *2011 IEEE Workshop on Automatic Speech Recognition & Understanding*, 24–29. IEEE.
- Serengil, S.; and Ozpinar, A. 2024. A Benchmark of Facial Recognition Pipelines and Co-Usability Performances of Modules. *Bilim Teknolojileri Dergisi*, 17(2): 95–107.
- Shi, B.; Hsu, W.-N.; Lakhota, K.; and Mohamed, A. 2022. Learning audio-visual speech representation by masked multimodal cluster prediction. *arXiv preprint arXiv:2201.02184*.
- Son Chung, J.; Senior, A.; Vinyals, O.; and Zisserman, A. 2017. Lip reading sentences in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6447–6456.
- Stafylakis, T.; and Tzimiropoulos, G. 2017. Combining residual networks with LSTMs for lipreading. *arXiv preprint arXiv:1703.04105*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Swietojanski, P.; and Renals, S. 2014. Learning hidden unit contributions for unsupervised speaker adaptation of neural network acoustic models. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, 171–176. IEEE.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Watanabe, S.; Hori, T.; Kim, S.; Hershey, J. R.; and Hayashi, T. 2017. Hybrid CTC/attention architecture for end-to-end speech recognition. *IEEE Journal of Selected Topics in Signal Processing*, 11(8): 1240–1253.
- Weng, X.; and Kitani, K. 2019. Learning spatio-temporal features with two-stream deep 3d cnns for lipreading. *arXiv preprint arXiv:1905.02540*.
- Wu, L.; Zhang, X.; Zhang, Y.; Zheng, C.; Liu, T.; Xie, L.; Yan, Y.; and Yin, E. 2024. Landmark-Guided Cross-Speaker Lip Reading with Mutual Information Regularization. *arXiv preprint arXiv:2403.16071*.
- Xiao, J.; Yang, S.; Zhang, Y.; Shan, S.; and Chen, X. 2020. Deformation flow based two-stream network for lip reading. In *2020 15th IEEE international conference on automatic face and gesture recognition (FG 2020)*, 364–370. IEEE.
- Xue, S.; Abdel-Hamid, O.; Jiang, H.; Dai, L.; and Liu, Q. 2014. Fast adaptation of deep neural network based on discriminant codes for speech recognition. *Proc. TASLP*, 22(12): 1713–1725.
- Yang, S.; Zhang, Y.; Feng, D.; Yang, M.; Wang, C.; Xiao, J.; Long, K.; Shan, S.; and Chen, X. 2019. LRW-1000: A naturally-distributed large-scale benchmark for lip reading in the wild. In *2019 14th IEEE international conference on automatic face & gesture recognition (FG 2019)*, 1–8. IEEE.
- Yeo, J. H.; Han, S.; Kim, M.; and Ro, Y. M. 2024a. Where Visual Speech Meets Language: VSP-LLM Framework for Efficient and Context-Aware Visual Speech Processing. *arXiv preprint arXiv:2402.15151*.
- Yeo, J. H.; Kim, M.; Choi, J.; Kim, D. H.; and Ro, Y. M. 2024b. Akvsr: Audio knowledge empowered visual speech recognition by compressing audio knowledge of a pretrained model. *IEEE Transactions on Multimedia*.
- Yeo, J. H.; Kim, M.; and Ro, Y. M. 2023. Multi-temporal lip-audio memory for visual speech recognition. In *Proc. ICASSP*, 1–5. IEEE.
- Yeo, J. H.; Kim, M.; Watanabe, S.; and Ro, Y. M. 2024c. Visual Speech Recognition for Languages with Limited Labeled Data Using Automatic Labels from Whisper. In *Proc. ICASSP*, 10471–10475. IEEE.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Conditional prompt learning for vision-language models. In *Proc. CVPR*, 16816–16825.