

# Robust LLM-based Audio-Visual Speech Recognition with Sparse Modality Alignment and Visual Unit-Guided Refinement

Fei Su<sup>1,2</sup>, Cancan Li<sup>1,2</sup>, Juan Liu<sup>2,1,\*\*</sup>, Wei Ju<sup>4</sup>, Hongbin Suo<sup>4</sup>, Ming Li<sup>3,2,\*\*</sup>

<sup>1</sup> School of Computer Science, Wuhan University, China

<sup>2</sup> School of Artificial Intelligence, Wuhan University, China

<sup>3</sup> School of Artificial Intelligence, The Chinese University of Hong Kong, Shenzhen, China

<sup>4</sup> AI Center, OPPO, China

fei.su@whu.edu.cn, liujuan@whu.edu.cn, mingli369@cuhk.edu.cn

## Abstract

Audio-Visual Speech Recognition (AVSR) integrates acoustic and visual information to enhance robustness in adverse acoustic conditions. Recent advances in Large Language Models (LLMs) have yielded competitive automatic speech recognition performance and shown effectiveness for AVSR. However, prior approaches project audio and visual features independently or apply shallow fusion, limiting cross-modal alignment and complementary exchange while increasing the LLM’s computational load. To address this, we propose AVUR-LLM, an LLM-based Audio-Visual Speech Recognition via Sparse Modality Alignment and Visual Unit-Guided Refinement. Experiments on LRS3 demonstrate state-of-the-art results for AVSR. Under additive-noise conditions at 0 dB SNR, it achieves 37% relative improvement over the baseline system. We release the AVUR-LLM code and related resources.<sup>1</sup>

**Index Terms:** audio-visual speech recognition, visual discrete units, large language models

## 1. Introduction

Audio-Visual Speech Recognition (AVSR) integrates acoustic signals with visual cues to enhance robustness and accuracy under adverse acoustic conditions. Early supervised pipelines, spanning DNN/CTC [1, 2], sequence-to-sequence [3], and RNN-T formulations [4], laid the groundwork but remain vulnerable to noise. Self-supervised learning improves data efficiency and transferability by pretraining strong speech and audio-visual encoders. Self-supervised acoustic models such as wav2vec [5] and WavLM [6] enhance representation quality. Audio-visual pretraining with AV-HuBERT [7] and AV-data2vec [8] further boosts the performance of AVSR, and automatic-label training provides additional gains [9]. In recent years, Large Language Models (LLMs) have been shown effective for speech recognition, with studies integrating them as parametric language models or as multi-modal reasoners alongside speech encoders, achieving competitive recognition performance [10]. In AVSR, recent LLM-based frameworks fuse Whisper [11] with visual features [12] and couple Whisper with LLaMA variants to exploit multimodal context [13, 14].

Prior studies on LLM-based AVSR [15, 16, 17] typically use self-supervised encoders as feature extractors. Acoustic and visual features are temporally compressed and linearly projected to the LLM embedding space, and the LLM is then applied for prediction or shallow fusion. Another approach in LLM-based

speech recognition adopts a rescoring strategy [18, 19, 20, 21], providing acoustic features to the LLM or applying text-only rescoring to reduce recognition error. However, despite these advances, LLM-based speech recognition still faces limitations: Tight coupling to acoustic/visual features increases memory consumption and sensitivity to input noise. Furthermore, multimodal fusion lacks fine-grained control. To reduce the computational cost and improve the robustness, recent works [22, 23, 24] discretize audio or visual embeddings via vector quantization, yielding token sequences that suppress high-frequency variability and shorten LLM prompts. Building on these works, we introduce a lightweight and stable alignment mechanism that enables controlled cross-modal exchange without disturbing the pretrained audio pathway.

Motivated by the above, our work further develops controlled multimodal interaction and adaptive fusion, while additionally employing speech-informative visual discrete tokens to guide LLM-based rescoring. In this paper, we propose AVUR-LLM. The main contributions are: (1) a stability-oriented sparse alignment that inserts lightweight cross-modal blocks into the upper audio encoder layers, (2) a confidence-aware fusion that modulates token-level visual injection, with visual discrete units-based prompt for LLM rescoring, (3) experiments on LRS3 dataset show that our model outperforms AVSR baseline systems and reduces WER under both clean and noisy conditions, achieving a 37% relative reduction at 0 dB SNR.

## 2. Methodology

Our proposed framework is illustrated in Fig. 1. The methodology involves sparse modality alignment, adaptive modulated fusion and visual unit-guided refinement.

### 2.1. Sparse Modality Alignment

We propose Sparse Modality Alignment (SMA), which sparsely inserts lightweight alignment blocks into the upper audio encoder. In each block, visual features act as queries and attend to audio context to calibrate the visual representations.

We employ the Whisper encoder [11] for audio and the AV-HuBERT encoder [7] for visual input, extracting features  $\mathbf{X}_a \in \mathbb{R}^{T_a \times D_a}$  and  $\mathbf{X}_v \in \mathbb{R}^{T_v \times D_v}$ , respectively. We then leverage the audio pathway’s higher temporal fidelity and strong pretraining to guide visual alignment and enhancement. Aligning the visual features to the audio temporal resolution mitigates frame-rate mismatch. Restricting alignment to the upper layers of the audio encoder and applying stop-gradient to the audio keys/values preserve acoustic representations, stabilize optimization, and

<sup>\*\*</sup>indicates the corresponding author.

<sup>1</sup><https://github.com/yakumo72/AVUR-LLM>

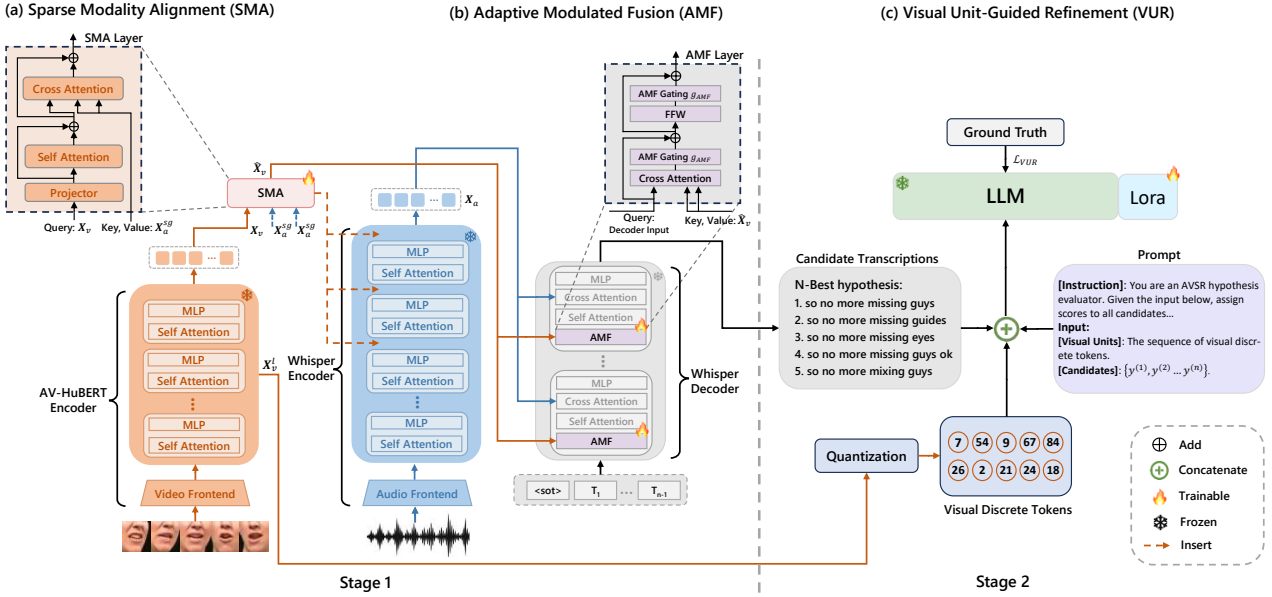


Figure 1: Illustration of AVUR-LLM. (a) The Sparse Modality Alignment (SMA) module, audio-conditioned cross-attention sparsely inserted into the upper audio encoder layers; (b) The Adaptive Modulated Fusion (AMF) module, confidence-aware gating that adaptively modulates visual injection in the decoder; (c) The Visual Unit-Guided Refinement (VUR) module, discretizes mid-layer visual features into tokens for LLM rescoring. Training has two stages: Stage 1 performs SMA+AMF and outputs  $N$ -best hypotheses; Stage 2 extracts and compresses  $X_v^\ell$  into visual discrete tokens and uses a LoRA-adapted LLM to rescore the hypotheses. Here  $\mathbf{X}_a$  and  $\mathbf{X}_v$  denote audio and visual encoder features.  $\hat{\mathbf{X}}_v$  denotes the SMA-refined visual features,  $\mathbf{X}_v^\ell$  denotes visual features extracted from encoder layer  $\ell$ .  $\mathbf{X}_a^{\text{sg}}$  denotes the audio features serving as keys/values with stop-gradient.  $g_{\text{AMF}}$  denotes the AMF gate coefficients.

incur modest computational overhead.

Specifically, the audio resolution is preserved, while visual features are upsampled via a lightweight differentiable resampler and projected to the audio feature dimension, yielding  $\mathbf{X}'_v = \mathcal{R}_v(\mathbf{X}_v) W_v + b_v$  with  $\mathbf{X}'_v \in \mathbb{R}^{T_a \times D_a}$ . Here,  $\mathcal{R}_v(\cdot)$  denotes lightweight differentiable resampling that aligns the visual features to the temporal resolution  $T_a$ .  $W_v \in \mathbb{R}^{D_v \times D_a}$  and  $b_v \in \mathbb{R}^{D_a}$  are learnable parameters of a linear projector converting the feature dimension from  $D_v$  to  $D_a$ .

Subsequently, we insert lightweight blocks into the upper audio encoder layers to align visual features with audio cues while keeping the audio features fixed. Let  $\mathbf{X}_a \in \mathbb{R}^{T \times D}$  and  $\mathbf{X}'_v \in \mathbb{R}^{T \times D}$ , where  $T = T_a$  and  $D = D_a$ . Each block performs:

$$\begin{aligned} \mathbf{X}_v^{\text{sa}} &= \text{MHA}(\mathbf{X}'_v, \mathbf{X}'_v, \mathbf{X}'_v), \\ \hat{\mathbf{X}}_v &= \mathbf{X}'_v + \text{MHA}(\mathbf{X}_v^{\text{sa}}, \mathbf{X}_a^{\text{sg}}, \mathbf{X}_a^{\text{sg}}), \end{aligned} \quad (1)$$

where MHA denotes Multi-Head Attention [25], sa denotes self-attended visual features, sg denotes stop-gradient, ensuring that audio features act only as keys/values and remain unchanged. We insert three alignment blocks into the upper layers of the audio encoder, so that alignment focuses on high-level representations and preserves computational efficiency and audio encoder stability.

## 2.2. Adaptive Modulated Fusion

To adaptively exploit complementary visual cues during decoding, we first estimate a token-level acoustic reliability signal via a forward-only probe on the decoder input and audio keys/values. Then, we use this signal to gate the magnitude of visual injection. Inspired by [12], we additionally apply a tanh-based direction

gate that determines whether the injected visual features operate in a positive or negative direction.

Specifically, we adapt Whisper’s decoder by inserting a forward-only acoustic probe at each layer, with gradients stopped through the queries and audio features. Given decoder queries  $\mathbf{Q}_\ell \in \mathbb{R}^{T_y \times D}$  and audio memory  $\mathbf{X}_a \in \mathbb{R}^{T \times D}$ , the probe attention is:

$$A_\ell = \text{softmax} \left( \frac{(\mathbf{Q}_\ell^{\text{sg}} W_Q) (\mathbf{X}_a^{\text{sg}} W_K)^T}{\sqrt{D}} \right). \quad (2)$$

The token-wise acoustic uncertainty  $S_{\ell,t}$  is computed as the entropy of the probe attention over audio frames  $A_\ell^t$ , normalized to  $[0, 1]$  by  $\log T$ . The values of  $S_{\ell,t}$  indicate the concentration of attention on audio frames, where smaller values correspond to more reliable acoustic contexts and larger values reflect higher acoustic uncertainty. Therefore, the amplitude of visual feature injection is defined as:

$$g_{\text{amp}} = \sigma(a S_\ell + b), \quad (3)$$

where  $a, b \in \mathbb{R}$  are learnable scalars controlling slope and offset, and  $\sigma(\cdot)$  denotes the sigmoid function, ensuring  $g_{\text{amp}} \in (0, 1)$ .

Given the decoder input  $\mathbf{x}_t$ , the visual context  $\mathbf{c}_{v,t}$  is obtained by MHA with query  $\text{LN}(\mathbf{x}_t)$  and keys/values  $\hat{\mathbf{X}}_v$ . The adaptive modulated fusion (AMF) gating uses two scalars per layer:  $g_{\text{AMF,att}}$  for the cross-attention branch and  $g_{\text{AMF,ff}}$  for the feed-forward branch. Each gate combines a direction term, obtained by applying  $\tanh(\cdot)$  to a learnable scalar, with a token-wise amplitude term  $\sigma(\cdot)$  derived from the acoustic uncertainty  $S_\ell$ . The AMF gates modulate the cross-attention and feed-forward residuals, respectively:

$$\begin{aligned} \mathbf{x}'_t &= \mathbf{x}_t + g_{\text{AMF,att}} \mathbf{c}_{v,t}, \\ \mathbf{y}_t &= \mathbf{x}'_t + g_{\text{AMF,ff}} \text{FFW}(\text{LN}(\mathbf{x}'_t)). \end{aligned} \quad (4)$$

where LN is Layernorm, and FFW is a Multi-Layer Perceptron (MLP). We train using token-level cross-entropy between the model’s predicted transcript and the ground-truth tokens. During optimization, the trainable components are the sparse modality alignment module and the adaptive modulated fusion module. At inference, beam search yields an  $N$ -best list  $Y = \{y_i\}_{i=1}^N$  with sequence scores  $s_{\text{infer}}(y_i)$ .

### 2.3. Visual Unit-Guided Refinement

Since the AV-HuBERT encoder’s visual representations carry speech-relevant structure, we discretize them into visual tokens and use the resulting sequence as conditioning input to an LLM for  $N$ -best rescoring. Unlike prior LLM-based approaches [15, 16, 20] that feed the LLM with continuous audio/visual feature projections or fused representations, VUR provides compact visual token sequences derived from AV-HuBERT as an explicit visual prompt.

Concretely, we use visual embeddings  $\mathbf{X}_v^\ell \in \mathbb{R}^{T_v \times D_v}$  from the visual encoder, with frame vectors  $\mathbf{x}_{v,t}^\ell \in \mathbb{R}^{D_v}$ . An offline  $K$ -means codebook with centroids  $\{\mathbf{k}_1, \dots, \mathbf{k}_K\}$  is learned [7], and each frame is quantized by assigning it to the nearest centroid  $\mathbf{k}_j$ . To exploit temporal stationarity and shorten the token stream, we apply run-length compression. Maximal contiguous spans with identical labels are collapsed into a single token by averaging the frames in each span and retaining the common label [13]. Let  $R_1, \dots, R_M$  denote the maximal contiguous index sets with identical labels. For each  $R_m$  we compute the averaged feature  $\bar{\mathbf{x}}_{v,m}^\ell = |R_m|^{-1} \sum_{t \in R_m} \mathbf{x}_{v,t}^\ell$  and assign the common label  $k_m$ . The resulting sequence of visual discrete units is cached once per utterance and reused for  $N$ -best rescoring.

Given the first-stage inference results, a LLM [26] is adapted with LoRA [27] to assign a scalar score  $r_i$  to each candidate, where LoRA parameters are trainable. Let  $i_{\text{gt}}$  denote the candidate that is closest to the ground-truth transcript. Training minimizes a list-wise softmax:

$$\mathcal{L}_{\text{VUR}} = -\log \frac{\exp(r_{i_{\text{gt}}})}{\sum_{j=1}^N \exp(r_j)}, \quad (5)$$

where  $r_j \in \mathbb{R}$  is the LLM-assigned score for the  $j$ -th candidate,  $N$  is the size of the  $N$ -best list.

## 3. Experimental Results

### 3.1. Experimental Settings

All experiments are conducted on the LRS3 dataset [28], which is a widely used publicly available English audio-visual speech recognition dataset. The dataset comprises approximately 433 hours of transcribed TED talk videos. In addition to the full training set, we also evaluate performance on the 30-hour “train-val” subset. Additionally, following [12], we extend the training data by combining the LRS3 training set with 1326 hours of English-speaking videos from VoxCeleb2 [29]. To assess noise robustness, we add babble noise from the MUSAN dataset [30] at SNRs of 10, 5, 0,  $-5$ , and  $-10$  dB as in [12].

We use Whisper Medium [11] as the audio encoder and decoder. Input features are 80-dimensional log-Mel filterbank coefficients. For the visual modality, frame-level representations are obtained from the AV-HuBERT Large [7]. Their weights are frozen during training. As for the LLM, we employ LLaMA2-7B [26]. A LoRA module is inserted into each layer of LLaMA with the rank of 16 for fine-tuning. A linear projection layer is introduced to map the visual discrete tokens to the LLaMA embedding dimension.

Table 1: WER (%) on LRS3 for AVUR-LLM and previous methods in ASR and AVSR under 30 h, 433 h, and LRS3+VoxCeleb2 (1759 h) settings.

Method	Audio Encoder	Visual Encoder	Training Data (h)	WER(%) ↓
<i>Audio-Only</i>				
CM-seq2seq [31]	Conformer	–	433	2.3
Fast Conformer [32]	Conformer	–	435	1.6
AV-HuBERT [7]	Transformer	–	433	1.6
Whisper-finetuned [12]	Transformer	–	433	1.1
Auto-AVSR [9]	Conformer	–	1902/3448	1.0/1.0
Llama-AVSR [15]	Whisper	–	30/433/1759	1.5/1.1/0.79
<b>AVUR-LLM</b>	Whisper	–	30	<b>1.3</b>
<b>AVUR-LLM</b>	Whisper	–	433	<b>1.0</b>
<b>AVUR-LLM</b>	Whisper	–	1759	<b>0.70</b>
<i>Audio-Visual</i>				
CM-seq2seq [31]	Conformer	–	433	2.3
AV-HuBERT [7]	Transformer	–	433	1.4
CMA [33]	Transformer	–	433	1.5
AV-data2vec [8]	Transformer	–	433	2.5
Fast Conformer [32]	Conformer	–	1687	0.9
Auto-AVSR [9]	Conformer	–	1902/3448	1.0/0.9
ViT3D-CM [34]	Conformer	–	90K	1.6
LP Conformer [35]	Conformer	–	100K	0.9
Whisper-Flamingo [12]	Whisper	AV-HuBERT	433/1759	1.0/0.86
Llama-AVSR [15]	Whisper	AV-HuBERT	433/1759	0.95/0.77
Llama-MTSK [17]	Whisper	AV-HuBERT	433	0.9
MMS-LLaMA [16]	Whisper	AV-HuBERT	433/1759	0.9/0.72
<b>AVUR-LLM</b>	Whisper	AV-HuBERT	433	<b>0.75</b>
<b>AVUR-LLM</b>	Whisper	AV-HuBERT	1759	<b>0.68</b>

### 3.2. Training and Inference Details

**SMA and AMF Training with N-best Generation.** The video data is sampled at 25 fps and converted to grayscale. Lip-centered regions of size  $96 \times 96$  are extracted using Dlib and aligned to a mean face template. During training, random  $88 \times 88$  crops are applied with horizontal flipping probability 0.5, while center crops are used during inference. For visual discrete tokens, we quantize frame-level features from the 12th AV-HuBERT layer into 2000 K-means clusters, followed by compression [23]. AdamW optimizer is employed for training with a cosine annealing learning rate scheduler and weight decay of 0.1. The learning rate is set to  $1 \times 10^{-4}$  for training.

**VUR with LLM.** The N-best hypotheses are paired with the visual discrete tokens to form the input for LLM. The prompt is defined as: [Instruction]: You are an AVSR hypothesis evaluator. Based on the visual units and language correctness, assign a score to each candidate transcription. Input: [Visual Units] {tokens} [Candidates]  $\{y_1, y_2, \dots, y_N\}$ , where **tokens** denotes the sequence of visual discrete tokens. LLaMA2-7B is adapted using LoRA fine-tuning with rank 16, a dropout rate of 0.05, and modifications applied to the query, key, value, and output projection layers, while keeping the backbone parameters frozen. Training is performed with a learning rate of  $5 \times 10^{-4}$ . Model performance is evaluated using Word Error Rate (WER) for both ASR and AVSR tasks.

### 3.3. Experimental Results and Analysis

#### 3.3.1. Main Results

Table 1 shows the performance comparisons of the proposed method with prior ASR and AVSR methods. In the audio-only setting, AVUR-LLM attains WERs of 1.3% (30 h setting) and 1.0% (433 h setting), outperforming Llama-

Table 2: WER (%) on LRS3 (433 h setting) under clean and noisy conditions. We report AVSR results at SNRs {10, 5, 0, -5, -10} dB. “-” indicates not reported.

Method	Modality	Clean	SNR Level (dB), WER (%) ↓				
			10	5	0	-5	-10
AV-HuBERT [7]		1.4	2.0	2.6	5.8	16.6	34.9
CMA [33]		1.5	1.8	2.4	4.4	11.9	25.8
Whisper-Flamingo [12]	AV	1.0	1.4	2.0	6.3	28.0	42.6
AVGER [20]		1.1	-	1.9	3.7	12.3	29.1
Llama-AVSR [15]		0.95	-	2.2	4.2	16.9	-
MMS-LLaMA [16]		0.9	-	1.3	2.7	7.4	-
<b>AVUR-LLM</b>	AV	<b>0.75</b>	<b>1.0</b>	1.4	<b>1.7</b>	<b>6.3</b>	<b>12.9</b>

AVSR [15] (1.5%/1.1%). Under the extended supervision setting (LRS3+VoxCeleb2, 1759 h), AVUR-LLM further reduces the audio-only WER to 0.70%, outperforming Llama-AVSR (0.79%) at the same scale. In the audio-visual setting, AVUR-LLM attains a WER of 0.75% with LRS3 (433 h setting), surpassing recent LLM-based systems [15, 16, 17] and establishing a new state of the art among LLM-based AVSR models, outperforming Whisper-Flamingo (1.0%) and the best prior LLM-based result (0.9%) by 25% and 16.7% relative WER reductions, respectively. Under the 1759 h setting, AVUR-LLM achieves 0.68% WER, consistently outperforming prior LLM-based systems. Notably, at the 433 h setting, our result remains competitive with Conformer/Transformer systems [32] that are trained on larger supervised corpora.

### 3.3.2. Noise Robustness Evaluation

In Table 2, we evaluate the robustness of AVUR-LLM across acoustic noise levels. AVUR-LLM attains the lowest WER in most of the SNR conditions. In particular, AVUR-LLM achieves a 37% relative WER reduction over MMS-LLaMA at 0 dB SNR (1.7% vs. 2.7%). Across noisy conditions (5/0/-5 dB), AVUR-LLM achieves an overall relative WER reduction of 14.7%, averaged over the best prior results. Under relatively clean conditions, decoder acoustic attention remains highly concentrated, producing low uncertainty and limiting visual injection. However, as the SNR decreases, acoustic attention becomes diffuse. The AMF increases the visual contribution, and the error rate is reduced substantially. At -5 and -10 dB, AVUR-LLM reduces WER to 6.3% and 12.9%, compared with the best prior results of 7.4% and 25.8%, indicating that the proposed method is effective under severe noise. It should be noted that babble noise was generated from different datasets across studies.

### 3.3.3. Ablation Results

Table 3 shows the results of our ablation study. When all components are enabled, WER is reduced from 1.30% to 0.75% (clean) and from 6.70% to 1.70% at 0 dB SNR, corresponding to relative reductions of 42.3% and 74.6%, respectively. Using AMF+VUR without SMA already yields a large gain (2.10%), while the SMA-only setting (w/o AMF and VUR) improves little under noise (6.70%), indicating that confidence-aware fusion and refinement module are the primary drivers of noise robustness. Adding SMA on top of AMF+VUR further reduces the 0 dB WER from 2.10% to 1.70%, suggesting that sparse alignment provides complementary benefits by aligning multimodal features and mitigating cross-modal mismatch.

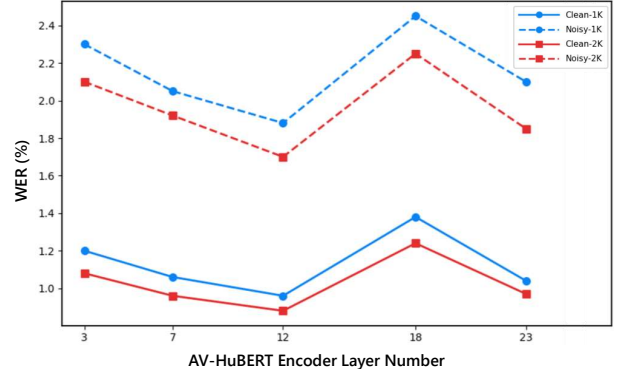


Figure 2: Effect of visual discrete tokens extraction depth and codebook size on WER (%). Solid lines denote clean, dashed lines denote 0 dB SNR. Blue curves denote K-means codebook size  $K=1000$ , red curves denote  $K=2000$ .

Table 3: Ablation of AVUR-LLM on LRS3 (433 h) under clean and noisy (0 dB SNR) conditions.

Method	SMA	AMF	VUR	Clean	Noisy
Whisper-Flamingo [12]	-	-	-	1.10	6.30
Llama-AVSR [15]	-	-	-	0.95	4.20
<b>Staged ablation</b>					
AVUR-LLM (full)	✓	✓	✓	<b>0.75</b>	<b>1.70</b>
- w/o VUR	✓	✓	✗	0.97	5.50
- w/o AMF (SMA only)	✓	✗	✗	1.30	6.70
- w/o SMA (AMF+VUR)	✗	✓	✓	0.85	2.10

### 3.3.4. Further Analysis

Fig. 2 depicts the effect of visual discrete units extraction depth and codebook size on WER under clean and 0 dB conditions. Performance peaks at a mid-level layer (12): with 2K clusters, WER reaches 0.75% (clean) and 1.70% (0 dB), consistently lower than the 1K setting. Shallow (layer 3) and deep (layer 18) features yield higher error, indicating that mid-level representations balance visual-phonetic detail and invariance. Increasing codebook granularity strengthens VUR’s rescoring by providing more discriminative visual tokens, yielding consistent gains across noise conditions.

## 4. Conclusion

In this paper, we presented AVUR-LLM, an LLM-based AVSR framework that combines multimodal fusion with visual unit-guided refinement. The method stabilizes cross-modal interaction in the audio pathway, adaptively regulates visual contributions during decoding, and compresses visual features into discrete tokens for rescoring with a LoRA-adapted LLM. A two-stage training scheme preserves the strengths of the pre-trained audio and visual encoders while limiting the burden on the LLM. Experiments demonstrate improvements under both clean and noisy conditions and provide analyses across SNRs and discretization settings, indicating enhanced robustness and data efficiency.

## 5. Acknowledgments

This research is funded in part by the National Natural Science Foundation of China (62571223) and Yangtze River Delta Science and Technology Innovation Community Joint Research Project (2024CSJGG01100) and OPPO. Many thanks for the computational resource provided by the Advanced Computing East China Sub-Center.

## 6. Generative AI Use Disclosure

Large Language Models (LLMs) were used solely for manuscript polishing (e.g., rephrasing and grammar checks) to improve clarity and readability. The LLMs were not used for ideation, methodology, experimental design, data analysis, or result interpretation. All scientific content was produced and verified by the authors.

## 7. References

- [1] K. Noda, Y. Yamaguchi, K. Nakadai, H. G. Okuno, and T. Ogata, "Audio-visual speech recognition using deep learning," *Applied intelligence*, vol. 42, pp. 722–737, 2015.
- [2] A. Graves, S. Fernández *et al.*, "Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks," in *Proc. ICML*, 2006, p. 369–376.
- [3] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 12, pp. 8717–8727, 2018.
- [4] T. Makino, H. Liao *et al.*, "Recurrent neural network transducer for audio-visual speech recognition," in *Proc. ASRU*, 2019, pp. 905–912.
- [5] A. Baevski, Y. Zhou *et al.*, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [6] S. Chen, C. Wang *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 16, no. 6, pp. 1505–1518, 2022.
- [7] B. Shi, W.-N. Hsu, and A. Mohamed, "Robust Self-Supervised Audio-Visual Speech Recognition," in *Proc. Interspeech*, 2022, pp. 2118–2122.
- [8] J. Lian, A. Baevski *et al.*, "Av-data2vec: Self-supervised learning of audio-visual speech representations with contextualized target representations," in *Proc. ASRU*, 2023, pp. 1–8.
- [9] P. Ma, A. Haliassos *et al.*, "Auto-avsr: Audio-visual speech recognition with automatic labels," in *Proc. ICASSP*, 2023, pp. 1–5.
- [10] C. Chen, Y. Hu *et al.*, "Hyporadise: An open baseline for generative speech recognition with large language models," *Advances in Neural Information Processing Systems*, vol. 36, pp. 31 665–31 688, 2023.
- [11] A. Radford, J. W. Kim *et al.*, "Robust speech recognition via large-scale weak supervision," in *Proc. ICML*, 2023, pp. 28 492–28 518.
- [12] A. Rouditchenko, Y. Gong *et al.*, "Whisper-flamingo: Integrating visual features into whisper for audio-visual speech recognition and translation," in *Proc. Interspeech*, 2024, pp. 2420–2424.
- [13] J. Yeo, S. Han *et al.*, "Where visual speech meets language: Vsp-llm framework for efficient and context-aware visual speech processing," in *Proc. EMNLP*, 2024, pp. 11 391–11 406.
- [14] U. Cappellazzo, M. Kim *et al.*, "Scaling and enhancing llm-based avsr: A sparse mixture of projectors approach," in *Proc. Interspeech*, 2025, pp. 1823–1827.
- [15] U. Cappellazzo, M. Kim, H. Chen, P. Ma *et al.*, "Large language models are strong audio-visual speech recognition learners," in *Proc. ICASSP*, 2025, pp. 1–5.
- [16] J. H. Yeo, H. Rha *et al.*, "Mms-llama: Efficient llm-based audio-visual speech recognition with minimal multimodal speech tokens," in *Proc. ACL*, 2025, pp. 20 724–20 735.
- [17] U. Cappellazzo, M. Kim *et al.*, "Adaptive audio-visual speech recognition via matryoshka-based multimodal llms," *ArXiv*, vol. abs/2503.06362, 2025.
- [18] S. Radhakrishnan, C.-H. H. Yang *et al.*, "Whispering llama: A cross-modal generative error correction framework for speech recognition," in *Proc. ACL*, 2023, pp. 10 007–10 016.
- [19] C. Chen, R. Li *et al.*, "It's never too late: Fusing acoustic information into large language models for automatic speech recognition," in *Proc. ICLR*, 2024, pp. 437–453.
- [20] R. Liu, H. Yuan, G. Gao, and H. Li, "Listening and seeing again: Generative error correction for audio-visual speech recognition," *Information Fusion*, vol. 120, p. 103077, 2025.
- [21] S. Ghosh, S. Kumar *et al.*, "Lipger: Visually-conditioned generative error correction for robust automatic speech recognition," in *Proc. Interspeech*, 2024, pp. 1920–1924.
- [22] B. Shi, W.-N. Hsu, K. Lakhota, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," in *Proc. ICLR*, 2022.
- [23] M. Kim, J. Yeo *et al.*, "Efficient training for multilingual visual speech recognition: Pre-training with discretized visual speech representation," in *Proc. ACM MM*, 2024, pp. 1311–1320.
- [24] P. Mousavi, J. Duret *et al.*, "How should we extract discrete audio tokens from self-supervised models?" in *Proc. Interspeech*, 2024, pp. 2554–2558.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] H. Touvron, L. Martin *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.
- [27] E. J. Hu, Y. Shen *et al.*, "LoRA: Low-rank adaptation of large language models," in *Proc. ICLR*, 2022.
- [28] T. Afouras, J. S. Chung, and A. Zisserman, "Lrs3-ted: a large-scale dataset for visual speech recognition," *arXiv preprint arXiv:1809.00496*, 2018.
- [29] J. S. Chung, A. Nagrani, and A. Zisserman, "VoxCeleb2: Deep speaker recognition," in *Proc. Interspeech*, 2018, pp. 1086–1090.
- [30] D. Snyder, G. Chen, and D. Povey, "Musan: A music, speech, and noise corpus," *arXiv preprint arXiv:1510.08484*, 2015.
- [31] P. Ma, S. Petridis, and M. Pantic, "End-to-end audio-visual speech recognition with conformers," in *Proc. ICASSP*, 2021, pp. 7613–7617.
- [32] M. Burchi, K. C. Puvvada *et al.*, "Multilingual audio-visual speech recognition with hybrid ctc/rnn-t fast conformer," in *Proc. ICASSP*, 2024, pp. 10 211–10 215.
- [33] S. Kim, K. Jang *et al.*, "Learning video temporal dynamics with cross-modal attention for robust audio-visual speech recognition," in *Proc. SLT*, 2024, pp. 447–454.
- [34] D. Serdyuk, O. Braga, and O. Siohan, "Transformer-based video front-ends for audio-visual speech recognition for single and multi-person video," in *Proc. Interspeech*, 2022, pp. 2833–2837.
- [35] O. Chang, H. Liao *et al.*, "Conformer is all you need for visual speech recognition," in *Proc. ICASSP*, 2024, pp. 10 136–10 140.