



LipGER: Visually-Conditioned Generative Error Correction for Robust Automatic Speech Recognition

Sreyan Ghosh, Sonal Kumar, Ashish Seth, Purva Chiniya, Utkarsh Tyagi, Ramani Duraiswami, Dinesh Manocha

University of Maryland, College Park, USA

{sreyang, sonalkum, aseth125, pchiniya, utkarsh, ramanid, dmanocha}@umd.edu

Abstract

Visual cues, like lip motion, have been shown to improve the performance of Automatic Speech Recognition (ASR) systems in noisy environments. We propose LipGER (Lip Motion aided Generative Error Correction), a novel framework for leveraging visual cues for noise-robust ASR. Instead of learning the cross-modal correlation between the audio and visual modalities, we make an LLM learn the task of visually-conditioned (generative) ASR error correction. Specifically, we instruct an LLM to predict the transcription from the N -best hypotheses generated using ASR beam-search. This is further conditioned on lip motions. This approach addresses key challenges in traditional AVSR learning, such as the lack of large-scale paired datasets and difficulties in adapting to new domains. We experiment on 4 datasets in various settings and show that LipGER improves the Word Error Rate in the range of 1.1%-49.2%. We also release LipHyp, a dataset with hypothesis-transcription pairs and lip motion cues.¹

Index Terms: audio-visual speech recognition, large language models, noise-robust speech recognition

1. Introduction

Automatic speech recognition (ASR) enables efficient and accurate transcription of spoken languages. However, in real-world scenarios, speech is generally accompanied by noise from several sources, including background sounds, interfering speakers, and reverberation [1]. The performance of ASR systems (trained on clean speech) degrades significantly in such scenarios [2]. Visual lip motion emerges as a natural source for supervision to improve speech understanding in noisy environments due to its tight coupling with audio [3]. This has motivated much work on Audio-Visual Speech Recognition (AVSR), where researchers have proposed novel ways to fuse visual information with speech for noise-robust ASR [4, 3, 5]. However, training AVSR systems that are effective and resilient to acoustic variations like accent, dialect, etc., suffers from a lack of diverse large-scale audio-visual datasets. Acquiring such datasets is prohibitively expensive, as they require to be recorded in controlled environments [3]. On the other hand, existing ASR systems benefit from the relative abundance of speech-only data [6, 7]. Thus, an effective solution would be to employ visual information to refine transcriptions from an existing powerful ASR model. To the best of our knowledge, no existing system has effectively been able to use this hypothesis.

Humans exhibit exceptional resilience to variations in acoustic environments, as our recognition system goes beyond merely processing the received acoustic signals. Typically, we

¹Code and Data: <https://github.com/Sreyan88/LipGER>

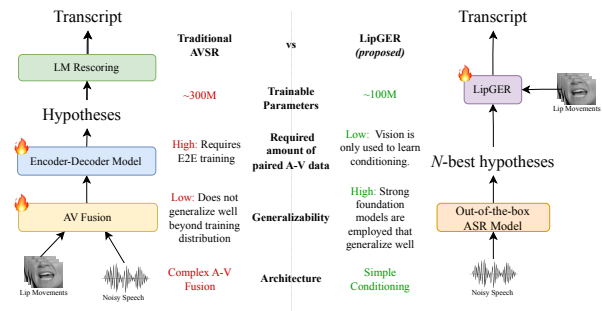


Figure 1: Comparison of traditional AVSR methods and LipGER. LipGER benefits over these methods on several aspects by overcoming the requirement of E2E training.

infer unclear or distorted spoken words by leveraging both visual cues and our intrinsic understanding of language. Very recently, Chen *et al.* [8] proposed HyParadise, a dataset for the task of ASR generative error correction with Large Language Models (LLMs). In contrast to the conventional method of using an LM to first re-score all transcriptions in the N -best hypothesis list with an LM and then select the best candidate, they propose to predict the transcription directly by prompting an LLM with the N -best hypothesis list. The authors argue that their primary motivation stems from the belief that the N -best list contains valuable information, as each hypothesis presents a unique textual interpretation of the input speech [9, 10, 11]. Hence, sentences that are disregarded could still contain accurate tokens essential for predicting the correct transcription. Inspired by prior work in uncertainty-estimation-based ASR [12, 13] and auto-regressive decoding [14], we posit that this hypothesis is more useful in noisy conditions where a higher uncertainty of beam search decoding leads to a more diverse N -best hypotheses list that contain accurate tokens [15].

Main Contributions. We propose LipGER, a simple yet effective framework for noise-robust AVSR. As an alternative to the traditional AVSR learning setup, which leverages parallel audio-visual data to jointly learn the correlation between the two modalities, LipGER leverages the visual cues for conditioning an LLM to learn generative error correction for ASR. Similar to Chen *et al.*, given an N -best hypotheses list from noisy source speech, we construct a prompt from it. We then feed this prompt to an LLM to predict the transcription. The LLM is further conditioned on the corresponding lip motion of the speech through multi-modal adapters. We build LipGER on the hypothesis that visual cues play an essential role in language learning in children [16, 17]. Our proposed paradigm of language-space denoising for transcribing noisy speech simplifies noise-robust ASR that usually employs a speech enhance-

ment front-ends. LipGER relies on language-comprehension, that is much easier to learn, benefits from abundant data and can generalize better. Additionally, LipGER employs a much simpler architecture than most existing AVSR systems, which require more complex architectures for learning the audio-visual correlation. Finally, LipGER can be used with out-of-the-box powerful and general-purpose ASR models, thus overcoming the requirement of large-scale paired audio-visual data for different languages, dialects or accents. To summarize, our contributions are as follows:

1. We present LipGER, a simple yet effective technique for noise-robust AVSR. Different from traditional AVSR systems, LipGER performs visually-conditioned generative error correction for accurately transcribing noisy speech.
2. We experiment with LipGER on 4 datasets under various synthetic and real-world noisy conditions and show that LipGER improves WER on these datasets by 1.1%-49.2%.
3. Additionally, we also open-source LipHyp, a large-scale dataset with hypothesis-transcription pairs equipped with visual lip motion cues to promote further research in this space.

2. Related Work

Noise Robust and Audio-Visual Speech Recognition. Recent noise-robust ASR methods primarily employ speech enhancement front-end to denoise noisy speech before ASR [18]. Other methods include domain adversarial training to learn noise-invariant speech features [19] or various pre-processing steps for denoising [6]. On the other hand, it is widely believed that integrating visual cues, like lip motion activity, can help improve the noise-robustness of ASR models [20]. However, they have not been widely adopted due to the lack of large-scale AV datasets recorded in diverse settings with varying languages, accents, etc. Various AVSR systems have been proposed in the literature that aims to improve the interaction between the audio and visual cues to push ASR performance [21]. However, they have not been extensively tested in noisy conditions.

LM Rescoring and Generative Error Correction (GEC). LM rescoring has been widely employed for improving ASR systems. Typically, an external LM is trained separately and utilized to re-score the N -best list of hypotheses generated by ASR decoding with beam search. Prior work has explored several ways, including shallow fusion, deliberation, component fusion, and cold fusion. LMs are also widely used for the error correction task in different languages [22], leveraging only the 1-best hypothesis generated by the ASR model [23]. Furthermore, more recent works utilize a candidates list after decoding for error correction [24, 23]. To the best of our knowledge, we are the first to explore multi-modal GEC.

3. Methodology

3.1. LipHyp

In this section, we describe the creation of LipHyp, a dataset with 601k+ unique pairs of hypothesis list \mathcal{H}_i and ground truth transcription \mathcal{T}_i for an audio \mathcal{A}_i , where the video \mathcal{V}_i is available for each pair. Specifically, for every $\mathcal{H}_i = \{\mathcal{Y}_0, \dots, \mathcal{Y}_N\}$, we extract the $N+1$ best hypothesis from the beam search decoder of an ASR system, where \mathcal{Y}_0 can be regarded as the ASR prediction for the audio, and the other N hypothesis can be considered as *other hypotheses*. To construct LipHyp, we employ open-source datasets and models described next.

3.1.1. ASR Models

We employ 3 state-of-the-art ASR models for N -best hypothesis generation, namely WavLM [25], Whisper [6] and Auto-AVSR [26]. All 3 models are unimodal speech-only ASR models, as we did not want additional visual cues to influence our hypotheses list. This is also aligned with our final goal of employing speech-only models for N -best hypotheses inference and using only visual cues for post-processing.

WavLM. We employ the ESPnet toolkit [27] for our WavLM-based ASR system. The architecture consists of two blocks: a self-supervised pre-trained front-end model and the ASR model (433 million parameters in total). The front-end, which is WavLM_{Large} consists of 24 transformer-based encoder layers and is pre-trained using a combination of LibriLight [28] (60k hours of data), Gigaspeech [7] (10k hours of data), and VoxPopuli [29] (24k hours of data). The high-level features from the front-end are fed into the ASR back-end for fine-tuning, which consists of 12 layers of conformer-based [30] encoder layers and 6 transformer-based decoder layers. The ASR system is fine-tuned on 960-hours LibriSpeech data.

Whisper. We employ the Whisper_{Base} model to generate hypotheses. The encoder-decoder transformer-based architecture has 74 million parameters and is trained on 680,000 hours of multilingual-weakly labeled speech data from the web.

Auto-AVSR. Similarly, the ASR encoder consists of a 1D ResNet-18 [31] followed by a Conformer decoder. The model is trained using joint CTC/attention training [32] and trained using LRS2 [33], LRS3 [33], VoxCeleb2 [34] and AVSpeech [35].

3.1.2. Datasets

We employ 3 datasets for training, namely LRS2, LRS3, and Facestar [36]. For testing, in addition to the 3 datasets we also add EasyCom [37], a real-world noisy dataset recorded in an egocentric view. To generate the N -best hypothesis from these datasets, we simulate noisy speech for all splits. To achieve this, we first add room reverberation by convolving clean speech with impulse responses from the MIT Impulse Response Survey [38]. It consists of 270 impulse responses collected from different locations that volunteers encountered in their daily lives and reflect a variety of standard acoustic settings. Subsequently, we add random audio samples from the VoxCeleb2 dataset to simulate interfering speakers and random audio samples from Audioset [39] to simulate background noise. For the background noise, we adjust the signal-to-noise ratio randomly between 0 and 40 dB with respect to the clean signal.

3.2. LipGER

Fig. 2 illustrates the working of LipGER. First, we input a noisy audio to the ASR system to generate $N + 1$ -best hypotheses via beam search decoding. Next, we construct an instruction from the hypothesis list and input it to the LLM to generate an output. The LLM generation is conditioned on the corresponding lip motion, which is fed as visual information to the model. In the next sections, we describe our architecture in detail.

3.2.1. Lip Motion Encoder

Similar to state-of-the-art lip-reading systems [40], our lip motion encoder L takes M mouth regions of interest (ROIs) as input. These ROIs are obtained by performing mouth cropping on the input video using facial landmark detection and alignment in a method similar to [41]. L consists of a 3D convolu-

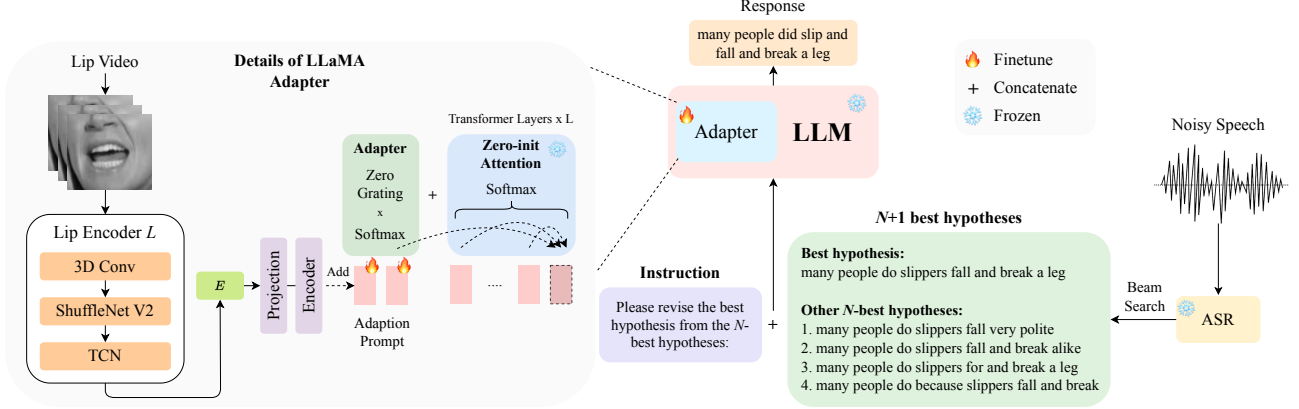


Figure 2: Illustration of **LipGer**. **LipGer** performs generative error correction on the N -best hypotheses generated by beam search decoding from an ASR model. **LipGer** is built on an LLM, which performs multi-modal reasoning by conditioning on lip motions. Specifically, we build a prompt from the $(N+1)$ -best hypotheses list, which instructs the LLM to rewrite the best hypothesis using the other hypotheses. The lip motion is conditioned on the LLM using multi-modal adapters with encodings obtained from a lip encoder L , and only the adapters and L are trained during the fine-tuning stage.

tional layer followed by a ShuffleNet v2 [42] network to extract a time-indexed sequence of feature vectors. They are then processed by a temporal convolutional network (TCN) to extract the final lip motion feature map $E \in \mathbb{R}^{V \times M}$, where V is the fixed sequence length of the feature. E is then fed to the LLM.

3.2.2. Lip Motion Conditioned LLM

Large Language Model (LLM). We employ an LLM for the task of GER due to its language generation and reasoning capabilities [43]. We only employ the pre-trained version of an LLM (similar to the general practice of employing a pre-trained LM to rescore the hypotheses) and fine-tune it using our LipHyp dataset to complete the GER task. We build instruction-response pairs from instances in the LipHyp dataset for training and evaluation. Once the model is trained, the evaluation phase is similar to prompting, where we prompt the LLM for the transcription. We next describe how we build training instances from LipHyp for fine-tuning.

Instruction Construction. We resort to the following template for instruction construction:

Below is the best-hypotheses transcribed from a speech recognition system. Please try to revise it using the words that are only included in the other-hypothesis, and write the response for the true transcription.

Best-hypothesis: you a very kind day

Other-hypotheses: you are very kind day, you have very kind day, \dots N hypotheses

Response: you are very kind

Multi-modal Adapters. As described earlier, **LipGer** completes the task of GER on noisy speech with multi-modal reasoning on lip motions. We design our architecture such that the fine-tuning can be performed in a parameter-efficient way, which benefits over existing AVSR systems that require a model to be trained from scratch for taking visual cues into consideration. Several parameter-efficient fine-tuning techniques have been proposed in literature [44], and we resort to adapter-based fine-tuning inspired by Zhang *et al.* [45]. Precisely, for every layer l of the L transformer layers, we first construct two learnable prompts, $P_l^v \in \mathbb{R}^{K \times C}$ and $P_l^a \in \mathbb{R}^{K \times C}$, where K is the

prompt length and C is the feature dimension. Then we project the lip motion encoding E into the same dimension as P_l^v and concat them together along the sequence dimension as follows:

$$I_l = \text{Concat}(\text{Projection}(E), P_l^v) \quad (1)$$

where $I_l \in \mathbb{R}^{(K+V) \times C}$. Next, we pass I_l through a transformer encoder with T layers and slice I_l back to a sequence length of K as follows:

$$G_l = \text{Encoder}(I_l)[:K, :] \quad (2)$$

Finally, we add G_l and P_l^a to obtain the final adaptation prompt $A_l \in \mathbb{R}^{K \times C}$. A_l is then concatenated as a prefix to the l -length word tokens T_l as follows:

$$[A_l; T_l] \in \mathbb{R}^{(K+I) \times C} \quad (3)$$

We fine-tune the LLM using standard cross-entropy loss and keep only the adapters, the lip motion encoder, and the visual encoder trainable ($\approx 200M$ parameters).

4. Experiments and Results

4.1. Experimental Setup

Hyper-parameters. We experiment with the pre-trained only versions of TinyLlama [46] (1.1B parameters) and Phi-1.5 [47] (1.3B parameters) for LLMs. We do not experiment with larger LLMs due to resource constraints. We fine-tune for 2 epochs with a constant learning rate of $5e-3$, a weight decay of 0.02, and a batch size of 32. For adapters, we employ $K=15$, and our vision encoder has 4 transformer encoder layers.

Baselines. We first compare **LipGer** with other rescoring methods, namely (i) GER [8] – **LipGer** without visual cues (ii) LM_{rank} – We use the same LLM as GER for rescoring the N -best hypotheses and finally take the hypothesis with the best score averaged across the LLM and ASR model scores. In an alternative setup for the noisy speech, we also employ a speech enhancement front-end, a HiFi-GAN [48], to denoise the noisy speech before passing it to the ASR model. For comparison with other AVSR methods, we use Auto-AVSR, AV-HuBERT [3], TM-CTC [49], MOCO+Wav2Vec2 [50] and Xu *et al.* [51]. **We employ recent baselines with open-source code available as we re-trained all models for evaluation.**

Table 1: *Result comparison of LipGER with other baselines. Left: Comparison of LipGER with text-only GER, LM-based restoring (+LM_{score}) and speech enhancement front-end (+Enhance). “LR” denotes LipGER trained on the low-resource split of the same dataset as the test set. “All” denotes LipGER trained on the entire LipHyp dataset. LipGER improves over baselines in the range of 1.1%-49.2%. Right: Comparison of LipGER with other AVSR systems. LipGER improves over baselines in the range of 1.9%-18.4%, is more robust to domain shifts and significantly outperforms them in low-resource settings.*

Test Set	Model	Baseline	+LM _{score}	+Enhance	+GER	+ LipGER _{LR}	LipGER _(ours)	LipGER _{All}	
LRS2	<i>test-real</i>	Whisper	33.4	32.8 _{-1.8%}	32.9 _{-1.5%}	29.7 _{-11.8%}	26.3 _{-21.26%}	25.9 _{-22.5%}	25.1 _{-24.8%}
		WavLM	5.9	5.6 _{-5.1%}	6.2 _{+5.1%}	7.0 _{+18.6%}	5.4 _{-8.5%}	5.4 _{-8.5%}	5.3 _{-10.3%}
		Auto-AVSR	1.5	1.5 _{-0.0%}	1.9 _{+26.7%}	1.5 _{-0.0%}	1.5 _{-0.0%}	1.4 _{-6.7%}	1.4 _{-6.7%}
	<i>test-simu</i>	Whisper	74.2	72.6 _{-2.1%}	67.1 _{-9.6%}	48.2 _{-35.0%}	43.7 _{-41.1%}	42.8 _{-42.3%}	40.1 _{-46.0%}
		WavLM	19.1	17.9 _{-6.3%}	18.4 _{-3.7%}	21.7 _{13.6%}	17.9 _{-6.3%}	17.0 _{-11.0%}	16.8 _{-12.04%}
		Auto-AVSR	26.4	25.8 _{-2.3%}	26.2 _{-0.8%}	24.1 _{-8.71%}	21.8 _{-17.42%}	21.1 _{-20.0%}	20.2 _{-23.5%}
LRS3	<i>test-real</i>	Whisper	19.0	18.2 _{-4.2%}	18.8 _{-1.1%}	15.5 _{-18.4%}	14.8 _{-22.1%}	14.6 _{-23.2%}	14.4 _{-24.2%}
		WavLM	4.4	4.4 _{-0.0%}	4.2 _{-4.5%}	4.1 _{-6.8%}	4.0 _{-9.0%}	4.0 _{-9.0%}	3.9 _{-11.4%}
		Auto-AVSR	1.2	1.1 _{-8.3%}	1.3 _{+8.3%}	1.1 _{-8.3%}	1.1 _{-8.3%}	1.1 _{-8.3%}	1.1 _{-8.3%}
	<i>test-simu</i>	Whisper	30.9	28.3 _{-8.4%}	30.2 _{-2.3%}	20.2 _{-34.6%}	17.3 _{-44.0%}	16.1 _{-47.0%}	15.7 _{-49.2%}
		WavLM	17.7	17.5 _{-1.1%}	17.0 _{-4.0%}	18.5 _{+4.5%}	14.7 _{-16.9%}	14.6 _{-17.5%}	14.7 _{-16.9%}
		Auto-AVSR	16.9	16.0 _{-5.3%}	15.5 _{-8.3%}	16.0 _{-5.33%}	12.9 _{-23.7%}	12.2 _{-27.8%}	11.7 _{-30.8%}
Facestar	<i>test-real</i>	Whisper	17.2	16.8 _{-2.3%}	16.1 _{-6.4%}	12.9 _{-25.0%}	12.4 _{-27.9%}	12.5 _{-27.3%}	12.5 _{-27.3%}
		WavLM	25.7	24.0 _{-6.6%}	23.9 _{-7.0%}	18.2 _{-29.2%}	16.8 _{-34.6%}	16.5 _{-35.8%}	15.5 _{-38.75%}
		Auto-AVSR	20.5	19.6 _{-4.4%}	18.2 _{-11.2%}	18.1 _{-11.7%}	14.0 _{-31.7%}	14.2 _{-30.7%}	14.2 _{-30.7%}
	<i>test-simu</i>	Whisper	57.5	55.5 _{-3.5%}	53.6 _{-6.8%}	55.2 _{-4.0%}	50.5 _{-12.3%}	50.5 _{-12.3%}	46.0 _{-20.0%}
		WavLM	53.7	51.1 _{-4.8%}	51.0 _{-5.0%}	50.9 _{-5.2%}	44.0 _{-18.1%}	43.2 _{-19.6%}	40.8 _{-24.0%}
		Auto-AVSR	50.1	49.0 _{-2.2%}	48.7 _{-2.8%}	48.4 _{-3.4%}	43.1 _{-14.0%}	41.5 _{-17.2%}	39.0 _{-22.2%}
EasyCom	<i>test-real</i>	Whisper	102.3	99.4 _{-2.8%}	100.3 _{-2.0%}	85.8 _{-16.1%}	76.7 _{-25.0%}	76.6 _{-25.1%}	76.6 _{-25.1%}
		WavLM	102.1	99.9 _{-2.2%}	100.9 _{-1.2%}	94.7 _{-7.2%}	94.5 _{-7.4%}	94.3 _{-7.6%}	94.1 _{-7.8%}
		Auto-AVSR	79.6	76.0 _{-4.5%}	80.4 _{+1.0%}	68.0 _{-14.6%}	63.9 _{-19.7%}	63.9 _{-19.7%}	62.9 _{-21.0%}

Dataset	LRS2	LRS3	Facestar	LRS3 → EasyCom
AV-HuBERT	28.1	20.6	48.7	94.3
AV-HuBERT _{LR}	30.2	21.7	51.6	97.6
Auto-AVSR	23.0	10.3	47.8	82.3
Auto-AVSR _{LR}	24.7	16.1	56.1	101.5
TM-CTC	29.9	19.4	58.4	93.1
TM-CTC _{LR}	35.8	24.0	62.3	98.5
MOCO+Wav2Vec2	27.3	17.5	54.0	85.8
MOCO+Wav2Vec2 _{LR}	34.0	26.4	58.7	89.6
Xu <i>et al.</i>	30.7	25.9	56.1	101.8
Xu <i>et al.</i> _{LR}	38.4	33.4	65.4	103.8
LipGer (ours)	21.1	12.2	41.5	63.9
LipGer _{LR} (ours)	21.8	12.9	43.1	66.7

Dataset Settings. For training our models, we use either the entire dataset or a low-resource setting as follows: LRS3 (407 hours/30 hours), LRS2 (195 hours/29 hours), Facestar (10 hours/2 hours). For evaluation, we test on both the original (*test-real*) and the simulated noisy test sets (*test-simu*) for all the above datasets together with EasyCom. All methods have been evaluated on the standard Word Error Rate (WER) metric.

4.2. Results

Table 1 (left) compares LipGER with 3 baselines, namely, text-only GER, LM-based rescoring (+LM_{score}) and speech enhancement front-end (+Enhance). We report scores for 3 versions of LipGER – LipGER_{LR} is trained on low-resource train splits of the same dataset the Test Set is taken from, LipGER is trained on the entire train split of the of the same dataset the Test Set is taken from and LipGER_{All} is trained on the entire LipHyp (described in Section 3.1). For EasyCom, we employ LipGER and LipGER_{LR} trained on LRS3. All (speech-only) ASR models employed are pre-trained (training settings are described in Section 3.1) and we don’t fine-tune them on the downstream dataset. LipGER outperforms all our baselines by 6.3%-44.0%. LipGER_{All} improves over LipGER by 0.4%-1.3%. LipGER_{LR} also outperforms all our baselines with only ≈1.5% drop in WER over LipGER, thereby proving its efficacy in low-resource scenarios with limited lip-motion cues available for learning AVSR. On an average, test-real splits have lower WER performance than test-sim splits with EasyCom having the highest WER due to significant motion blur and distant speakers.

Table 1 (right) compares LipGER with other AVSR approaches. LipGER outperforms all baselines by 1.9%-18.4%, showing the most significant improvements on the EasyCom dataset. This dataset, unlike LRS3, features speakers at a distance recorded from a first-person perspective, leading to a notable difference in the evaluation domain. Despite this, LipGER exhibits superior resilience to such domain shifts compared to other AVSR models. This enhanced robustness is attributed to LipGER’s reliance on adapting to language-space rather than speech-space, which inherently offers greater stability against domain variations.

I - Best: the demand of the mark wednesday	Gold
Other N-best Hypotheses:	because of the demand of the market
1. of the demand of the market wednesday	Prediction
2. because of the demand of the mark wednesday	because of the demand of the market
3. this is the demand of the market	
II - Best: he saw it written down	Gold
Other N-best Hypotheses:	he saw it written down somewhere
1. he saw it written down someway	Prediction
2. he saw it written down some	he saw it written down somewhere
3. he saw it written down some one	
III - Best: why i was given over this	Gold
Other N-best Hypotheses:	why are we skipping over this
1. why i was skipping over this	Prediction
2. while i was given over this	why i was given over this
3. why have we given over this	

Figure 3: *Illustration of predictions by LipGER. Instances I. and II. show positive cases where LipGER predicts the correct transcription by revising the best hypothesis with accurate tokens hidden in the other N-best hypotheses. Case III. shows a failure case highlighting the limitation of LipGER, which fails to generate the correct transcription due to a lack of sufficient contextual cues in the other N-best hypotheses.*

5. Conclusion and Limitations

In this paper, we present LipGER, a novel system for learning robust audio-visual speech recognition. In contrast to prior methods, which employ visual cues for learning speech recognition, LipGER leverages visual cues for learning ASR Generative Error Correction. To achieve this, LipGER employs an LLM conditioned on lip motions using multi-modal adapters. LipGER is simple, requires minimal training parameters, and overcomes several limitations in the current audio-visual ASR learning paradigm. We show that LipGER outperforms most baselines from the literature on noise robust speech recognition by significant margins in low and high-resource settings.

Some limitations of LipGER include: (1) the requirement of an LLM for the pipeline and (2) the predictions might encode biases learned by the LLM during pre-training. As part of future work, we would like to find ways to overcome these limitations.

6. References

- [1] J. Li, L. Deng, Y. Gong, and R. Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, 2014.
- [2] G. Krishna, C. Tran, J. Yu, and A. H. Tewfik, "Speech recognition with no speech or with noisy speech," in *IEEE ICASSP 2019*.
- [3] B. Shi, W.-N. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," *arXiv preprint arXiv:2201.02184*, 2022.
- [4] J. Hong *et al.*, "Watch or listen: Robust audio-visual speech recognition with visual corruption modeling and reliability scoring," *arXiv preprint arXiv:2303.08536*, 2023.
- [5] J. Li *et al.*, "Robust audio-visual asr with unified cross-modal attention," in *IEEE ICASSP 2023*.
- [6] A. o. Radford, "Robust speech recognition via large-scale weak supervision," in *ICML 2023*.
- [7] G. Chen *et al.*, "Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio," *arXiv preprint arXiv:2106.06909*, 2021.
- [8] C. CHEN *et al.*, "Hyporadise: An open baseline for generative speech recognition with large language models," in *NerulPS 2023*.
- [9] F. Peng *et al.*, "Search results based n-best hypothesis rescoring with maximum entropy classification," in *IEEE ASRU 2013*.
- [10] E. Variani *et al.*, "Neural oracle search on n-best hypotheses," in *IEEE ICASSP 2020*.
- [11] M. Li *et al.*, "Improving spoken language understanding by exploiting asr n-best hypotheses," *arXiv preprint arXiv:2001.05284*, 2020.
- [12] V. Stouten, P. Wambacq *et al.*, "Model-based feature enhancement with uncertainty decoding for noise robust asr," *Speech communication*, vol. 48, no. 11, pp. 1502–1514, 2006.
- [13] D. T. Tran, E. Vincent, and D. Jouvét, "Fusion of multiple uncertainty estimators and propagators for noise robust asr," in *IEEE ICASSP 2014*.
- [14] A. Malinin *et al.*, "Uncertainty estimation in autoregressive structured prediction," *arXiv preprint arXiv:2002.07650*, 2020.
- [15] Y. Hu *et al.*, "Large language models are efficient learners of noise-robust speech recognition," in *ICLR 2024*.
- [16] A. N. Meltzoff and M. K. Moore, "Imitation of facial and manual gestures by human neonates," *Science*, vol. 198, no. 4312, pp. 75–78, 1977.
- [17] R. Davies, E. Kidd, and K. Lander, "Investigating the psycholinguistic correlates of speechreading in preschool age children," *International Journal of Language & Communication Disorders*, vol. 44, no. 2, pp. 164–174.
- [18] S.-W. Fu, C.-F. Liao, Y. Tsao, and S.-D. Lin, "Metricgan: Generative adversarial networks based black-box metric scores optimization for speech enhancement," in *ICML 2019*, pp. 2031–2041.
- [19] A. Prasad *et al.*, "An investigation of end-to-end models for robust speech recognition," in *IEEE ICASSP 2021*.
- [20] J. Huang and B. Kingsbury, "Audio-visual deep learning for noise robust speech recognition," in *IEEE ICASSP 2013*.
- [21] A. Das, S. Patikar, and K. Medhi, "A survey on audio-visual speech recognition system," in *2023 IEEE I3CS*, 2023.
- [22] J. Wirth and R. Peinl, "Automatic speech recognition in german: A detailed error analysis," in *2022 IEEE COINS*.
- [23] Y. Leng *et al.*, "Fastcorrect: Fast error correction with edit alignment for automatic speech recognition," *NeurIPS 2021*.
- [24] A. Mani, S. Palaskar, N. V. Meripo, S. Konam, and F. Metze, "Asr error correction and domain adaptation using machine translation," in *IEEE ICASSP 2020*.
- [25] S. Chen *et al.*, "Wavlm: Large-scale self-supervised pre-training for full stack speech processing," *IEEE JSTSP 2022*.
- [26] P. Ma *et al.*, "Auto-avsr: Audio-visual speech recognition with automatic labels," in *IEEE ICASSP 2023*.
- [27] S. Watanabe *et al.*, "Espnet: End-to-end speech processing toolkit," *arXiv preprint arXiv:1804.00015*, 2018.
- [28] J. Kahn *et al.*, "Libri-light: A benchmark for asr with limited or no supervision," in *IEEE ICASSP 2020*. IEEE, 2020. [Online]. Available: <http://dx.doi.org/10.1109/ICASSP40776.2020.9052942>
- [29] C. Wang *et al.*, "Voxpopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation," *arXiv preprint arXiv:2101.00390*, 2021.
- [30] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *arXiv preprint arXiv:2005.08100*, 2020.
- [31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE CVPR 2016*.
- [32] S. Kim *et al.*, "Joint ctc-attention based end-to-end speech recognition using multi-task learning," *IEEE ICASSP 2017*.
- [33] S. *et al.*, "Son chung, joon and others," in *CVPR 2017*.
- [34] J. S. Chung, A. Nagrani, and A. Zisserman, "Voxceleb2: Deep speaker recognition," *arXiv preprint arXiv:1806.05622*, 2018.
- [35] A. Ephrat *et al.*, "Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation," *arXiv preprint arXiv:1804.03619*, 2018.
- [36] K. Yang, D. Markovic, S. Krenn, V. Agrawal, and A. Richard, "Audio-visual speech codecs: Rethinking audio-visual speech enhancement by re-synthesis," in *CVPR 2022*.
- [37] D. *et al.*, "Easycom: An augmented reality dataset to support algorithms for easy communication in noisy environments," 2021.
- [38] J. Traer and J. H. McDermott, "Statistics of natural reverberation enable perceptual separation of sound and space," *Proceedings of the National Academy of Sciences*, vol. 113, no. 48, pp. E7856–E7865, 2016.
- [39] J. F. Gemmeke *et al.*, "Audio set: An ontology and human-labeled dataset for audio events," in *IEEE ICASSP 2017*.
- [40] B. Martinez, P. Ma, S. Petridis, and M. Pantic, "Lipreading using temporal convolutional networks," in *IEEE ICASSP 2020*.
- [41] P. Ma *et al.*, "Visual speech recognition for multiple languages in the wild," *Nature Machine Intelligence*, vol. 4, no. 11, pp. 930–939, 2022.
- [42] N. Ma *et al.*, "Shufflenet v2: Practical guidelines for efficient cnn architecture design," in *ECCV 2018*.
- [43] W. X. Zhao *et al.*, "A survey of large language models," *arXiv preprint arXiv:2303.18223*, 2023.
- [44] L. Xu, H. Xie, S.-Z. J. Qin, X. Tao, and F. L. Wang, "Parameter-efficient fine-tuning methods for pretrained language models: A critical review and assessment," *arXiv preprint arXiv:2312.12148*, 2023.
- [45] R. *et al.*, "LLaMA-adaptor: Efficient fine-tuning of large language models with zero-initialized attention," in *ICLR 2024*.
- [46] P. Zhang, G. Zeng, T. Wang, and W. Lu, "Tinyllama: An open-source small language model," 2024.
- [47] Y. Li *et al.*, "Textbooks are all you need ii: **phi-1.5** technical report," *arXiv preprint arXiv:2309.05463*, 2023.
- [48] J. Su, Z. Jin, and A. Finkelstein, "Hifi-gan: High-fidelity denoising and dereverberation based on speech deep features in adversarial networks," *arXiv preprint arXiv:2006.05694*, 2020.
- [49] T. Afouras, J. S. Chung, A. Senior, O. Vinyals, and A. Zisserman, "Deep audio-visual speech recognition," *IEEE TAPMI 2018*.
- [50] X. Pan *et al.*, "Leveraging unimodal self-supervised learning for multimodal audio-visual speech recognition," in *ACL 2022*, S. Muresan, P. Nakov, and A. Villavicencio, Eds.
- [51] B. Xu *et al.*, "Discriminative multi-modality speech recognition," in *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 2020, pp. 14 433–14 442.