

Multi-modal Speech Transformer Decoders: When Do Multiple Modalities Improve Accuracy?

Yiwen Guan, Viet Anh Trinh, Vivek Voleti, Jacob Whitehill
Worcester Polytechnic Institute
{yguan2, vtrinh, vsvoleti, jrwhitehill}@wpi.edu

Abstract—Decoder-only discrete-token language models have recently achieved significant success in automatic speech recognition. However, systematic analyses of how different modalities impact performance in specific scenarios remain limited. In this paper, we investigate the effects of multiple modalities on speech recognition accuracy on both synthetic and real-world datasets. Our experiments suggest that: (1) Integrating more modalities can increase accuracy but the benefit depends on the amount of auditory noise. We also show for the first time the benefit of combining audio, image context, and lip information in one speech recognition model. (2) Images as a supplementary modality for speech recognition provide their greatest benefit at moderate audio noise levels; moreover, they exhibit a different trend compared to inherently synchronized modalities like lip movements. (3) Performance improves on both synthetic and real-world datasets when the most relevant visual information is filtered as a preprocessing step.

Index Terms—automatic speech recognition, large language models, multi-modal processing

I. INTRODUCTION

When developing multi-modal automatic speech recognition (ASR) systems, when do the different input modalities (speech, visual context, lip movements of the speaker, etc.) help, and under what conditions is the benefit of each modality the strongest? The past few years have seen rapidly growing interest in ASR based on decoder-only discrete-token language models (e.g., [1–3]). Such models are attractive in part due to their ability to accept multi-modal inputs (e.g., audio, text, images) and generate multi-modal outputs (e.g., text tokens for ASR, audio tokens for speech-to-speech translation, etc.). The ability to process multiple input streams, such as audio that was spoken, lip movements of the speaker, an image of what was spoken about, etc., can enhance robustness of multi-modal ASR systems under challenging conditions, such as noisy environments. Advantages of such models include the ability to pretrain on large-scale audio or audio-visual datasets [4–6], and to leverage advanced language understanding capabilities from LLMs [7] to perform multi-modal tasks. Despite recent development of multi-modal large language models (MLLMs) for speech processing (e.g., Google Gemini), studies that systematically investigate how different modalities impact ASR performance are scarce.

In this work, we systematically analyze the impact of different input modalities (speech audio, visual context, lip movements, text from optical character recognition) in multi-modal ASR models across a range of noise conditions. Our goal is to identify noise regimes under which the modalities

are complementary to each other. As a motivating example, if a speaker is presenting a Powerpoint slide but the speech signal is weak (due to unclear pronunciation or background noise), then the visual information of the slide could be used to help transcribe what the speaker said – *as long as* the speech was at least clear enough that a correspondence between the visual and auditory information can be established. Moreover, if the speaker’s lips are visible, then lip-reading could potentially boost accuracy as well. On the other hand, the use of additional modalities might actually hurt performance due to a longer input length. Moreover, the fact that some modalities (e.g., speech & lips) are *synchronized* with each other whereas others (e.g., speech & image context) are not, may make it more difficult for the model to establish a correspondence between information sources and require different processing architectures.

Our paper investigates the following research questions: (1) Do additional modalities always help ASR accuracy? (2) Does each modality provide a uniform accuracy boost across all noise levels? (3) How does irrelevant visual information affect ASR performance? To this end, we introduce a synthesized, multi-modal dataset (3-Equations) that is highly controllable, on which we can conveniently simulate various tailored situations. This dataset focuses on mathematical scenarios, as specific educational applications are part of the motivation for this work. We further experiment on the real-world SlideAVSR dataset [18]. In addition to these experiments, we also extend the MLLM in [3] to support more input modalities and demonstrate for the first time how audio, image context, and lip information can be combined in one model to increase ASR performance. Beyond multi-modal ASR, our findings have potential applications in areas like video conferencing, and educational AI partners [38], where multi-modal inputs improve noisy environment performance and aid understanding in technical contexts.

II. RELATED WORK

Audio-visual speech recognition: Despite the tremendous success of ASR systems over the past decade, quality and reliability issues still exist in acoustically noisy environments [8]. One strategy to mitigate such noise is to harness complementary visual information, e.g., of the speaker’s lips [5, 9–12]. A state-of-the-art approach is AV-HuBERT [13]: it learns audio-visual speech representations by feature clustering and masked prediction. By extending noise augmentation to AV-

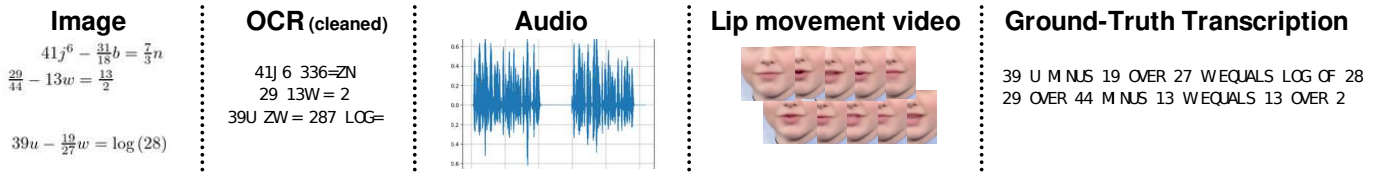


Fig. 1. An example of the 3-Equations dataset. From left to right: the image sample shows 3 mathematical equations; the OCR texts are extracted with EasyOCR and cleaned to retain only numbers, letters, and operators; the audio contains randomly reading 2 out of the 3 equations aloud; the lip movement video displays a lip region reading the corresponding equation sentences; the ground-truth transcription is the plain-text translation (label) of the speech. In this example, the speech reads the third and the second equations in order.

HuBERT pre-training, downstream models attain better robustness to acoustic noise [14]. Other works [15, 16] explore the performance gap between audio-only and audio-visual models in noisy conditions and show that the gap becomes larger as the noise levels increase. Instead of harnessing visual representations of the lips, some works such as SlideSpeech [17] and SlideAVSR [18] have explored using optical character recognition (OCR) to extract information from the subject matter (lecture slides). To leverage the extracted text, SlideSpeech uses cross-attention to combine speech embeddings and contextual phrase embeddings for contextual ASR; SlideAVSR proposes DocWhisper which provides the texts to Whisper as prompts. Both works demonstrate performance improvement by integrating OCR texts.

Multi-modal large language models: Recognizing the powerful language generation, zero-shot transfer, and contextual learning capabilities of large language models (LLMs), significant efforts have been made to harness the knowledge from LLM pre-training to empower multi-modal tasks [19]. Research works such as Qwen-Audio [20] and LLaVA [21] have focused on multi-modal comprehension. Meanwhile, the outstanding generative capability has also inspired many works to extend unimodal LLMs to perform multi-modal generation, such as MiniGPT-5 [22] and NExT-GPT [23], leading to the emergence of multi-modal large language models (MLLMs). Even so, only a few works incorporate visual speech modeling with LLMs, leaving this a relatively unexplored area. VSP-LLM [24] integrates LLMs into visual speech processing, enhancing context modeling for tasks like visual speech recognition (VSR) and translation (VST). Llama-AVSR [25] enables an LLM to perform ASR, VSR, and AVSR tasks by leveraging pre-trained audio-visual encoders. However, both works primarily concentrate on tasks like visual speech processing or AVSR while not exploring combinations with other modalities.

III. DATASETS

A. 3-Equations

To systematically explore the interaction between modalities, we synthesized a multi-modal dataset consisting of images, audio clips, and lip movements. Each example contains multiple randomly generated mathematical equations – reminiscent of the kinds of visual content that appear in lecture slides. The motivation is to design a dataset that forces the model to exploit information across modalities while allowing us to introduce complexity in a controlled manner. In this way, we can study how the model relies on information from

REF:	33 -31 equal s 49 p i -23 t equal s j
HYP:	33 -31 equal s 49 p * * * * * * * * * *
REF:	x +43 f equal s 6 i 32 +35 equal s 10 * * * * * * * * * * * * * * * *
HYP:	x +43 * equal s 6 i 32 +35 equal s 10 49 k +28 a equal s 27 e over 47

Fig. 2. Examples of Gemini’s output on the 3-Equations 2-noise. For a fair comparison, we “help” Gemini spell out some words it recognizes as symbols. REF stands for reference, HYP stands for model’s hypothesis.

each modality separately, especially when some modalities are insufficient or corrupted. An example from this dataset is illustrated in Fig. 1. To simulate an insufficient auditory modality scenario, each audio sample only reads out two of the three equations randomly from the image. This setting encourages the model to rely on both visual and auditory modalities: without auditory information, the model cannot find the correct spoken equations; without visual information, the model will fail in a noisy environment.

The dataset consists of 10,000 examples, each containing one image sample with its OCR text, one audio sample, and one lip movement video sample. Specifically, the images depict three mathematical equations, involving operations such as addition, subtraction, logarithms, fractions and exponentiation, with each image sized 450×200 . The audio part contains 25.2 hours of synthesized speech, averaging 9 seconds per sample, with 20.2 hours for training and 2.5 hours each for development and testing. The lip movement videos are generated from the synthesized audio at 25 FPS using a static portrait image. To create the dataset, we employ pyttxs3 [34] to generate speech utterances, latex command to produce equation images, EasyOCR [35] to obtain OCR texts, and Wav2Lip [36] to generate lip-synced videos.

One advantage of creating this dataset is that, we can freely simulate anticipated situations, such as when multi-modal fusion is important. Hence, we added noise from the MUSAN dataset [37] to the second half of each equation utterance at varying signal-to-noise ratios (SNR). This process produces a dataset we refer to as “2-noise”. The first half of each utterance remains clean, allowing the model to leverage this “incomplete” clean auditory information to locate the correct equation in the image, and subsequently complete the speech transcription with clean visual information.

Another advantage of using this dataset is that, given that our image data consists of mathematical symbols, the visual information could be difficult for standard image codecs to extract. This allows us to study the impact of an imperfect image codec on the ASR accuracy. In particular, if we consider

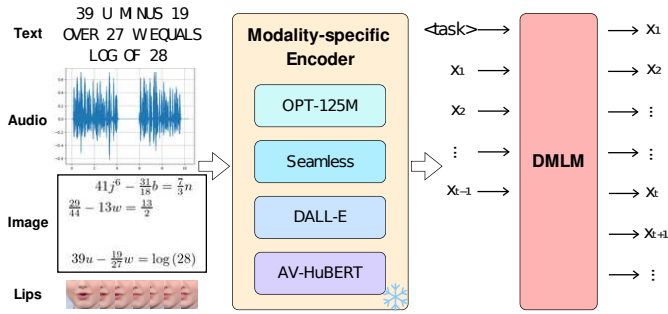


Fig. 3. An overview of Discrete Multimodal Language Model (DMLM). The inputs are encoded by modality-specific encoders, and the encoded multimodal tokens are concatenated to a task prompt, and then passed to DMLM for processing in next-token prediction fashion.

an “oracle OCR” that perfectly transcribes the 3 equation sentences in the image (but does not select which 2 out of the 3 were actually spoken), then the sequence: no vision, image encoding, real OCR, oracle OCR can be considered as an ascending order of image representation quality. Therefore, we can systematically explore how the quality of the visual modality, which is a supplementary modality for the task, affects speech recognition performance.

Interestingly, we identify a particular failure mode of Gemini when feeding 3-Equations 2-noise data into Gemini-1.5 Flash for transcription. Examples of Gemini’s output are shown in Fig. 2. It appears to overly rely on a single modality, whereby prone to generate transcriptions for 1 or 3 equations instead of the correct 2 spoken equations. This finding also illustrates how this synthetic dataset can be useful for analyzing even very powerful models.

B. SlideAVSR

To generalize our findings to the real world, we leverage SlideAVSR [18], an audio-visual dataset of paper explanation videos collected from YouTube. The dataset provides manually transcribed speech, synchronized visual slides, and preprocessed OCR keywords. Since the videos contain many AI technical terms, accurate transcription is difficult without referring to the slides, which makes the dataset useful for multi-modal experiments. The dataset comprises 245 hours of audio data, with 195 hours allocated for training, 20 hours for development, and 30 hours for testing. Since most of the videos in SlideAVSR are recorded from the presenter’s laptop, the audio quality is generally very high. To simulate a broader range of noise conditions, we thus also add noise to the clean audio at different SNRs.

IV. EXPERIMENTAL SETUP

A. Model

Our multi-modal speech recognition experiments are based on a Discrete Multi-modal Language Model (DMLM, depicted in Fig. 3) [3], which is a discrete token-based Transformer decoder model using OPT [26] as the backbone. DMLM tokenizes the input data by frozen modality-specific encoders to form a discrete token sequence. In particular, DMLM

uses the Seamless codec [27] to convert audio waveforms to discrete speech tokens, and the DALL-E encoder [28] to convert images to discrete image tokens. In addition, to extend DMLM to accept lip movement input, we employ AV-HuBERT to obtain lip tokens as well as text tokens representing the lip-reading hypothesis [13, 14]. The input token sequence is concatenated to a task description (e.g. ASR, lip-to-text) and then passed to DMLM for processing. Thanks to the attention mechanism, DMLM is able to learn how to combine information from different modality tokens. For example, in Fig. 3, when the speaker is saying “thirty-nine”, the model learns to attend to the corresponding audio input (i.e. the beginning of the first audio segment), image input (i.e. the bottom left of the image), and lip input (i.e. the beginning of the lip movements) to complete the transcription. Due to the parallelism of the input tokens from all modalities, the multi-modality can have a mild impact on the time complexity.

B. Model Configurations and Training Details

Following prior work [3], we employ the OPT model with 125M parameters as the backbone LLM of DMLM. To enable DMLM to solve multi-modal problems, we fine-tune DMLM on a mixture of multi-modal tasks, including speech recognition, speech translation, image generation, and image captioning, using LibriSpeech-960 [29], CVSS [30], CoVoST2 [31], and COCO [32]. We use this fine-tuned DMLM as a pre-trained model, and further fine-tune it on desired tasks such as multi-modal ASR. We extend the length-normalized trimodal loss function proposed in [3] to accept more modalities. The model is trained using AdamW [33] with $\beta=(0.9, 0.999)$, weight decay of $1e-4$, $lr=1e-6$, on a single NVIDIA A100 GPU with a batch size of 4 and patience of 5.

C. Evaluation Metrics

We report the Word Error Rates (WER) for each dataset, using the Whisper English text normalizer [4] to clean and standardize the text by removing extraneous characters, normalizing spaces, and converting text to lowercase. We also report a relative WER benefit of adding extra modalities to the audio task, which is calculated by $(WER_A - WER_{X+A})/WER_A$, where A denotes audio and X stands for any additional modalities. This metric indicates better performance with a larger value.

V. RESULTS AND DISCUSSION

A. Do additional modalities always help ASR accuracy?

On the one hand, using multiple modalities could improve accuracy by supplying the model with complementary information. On the other hand, it could conceivably hurt performance, as the additional modalities result in longer input sequences, which might prevent the model from finding relevant information. We thus conduct a modality ablation study to investigate whether fusing more modalities with audio can improve recognition performance. In this experiment we examine the potential benefits of incorporating multiple modalities by averaging across different noise levels.

TABLE I

EVALUATION OF WER (%) AND BENEFIT OF ADDING IMAGE(I), LIP HYPOTHESIS(L), AND DIFFERENT TYPES OF OCR(O) ON THE 3-EQUATIONS 2-NOISE TEST SET, AT DIFFERENT NOISE LEVELS BASED ON SNR. $+\infty$ REPRESENTS CLEAN AUDIO AND $-\infty$ MEANS PURE NOISE.

Task	2-noise, SNR(dB) =									Average
	(clean) $+\infty$	20	10	5	0	-5	-10	-20	(noise) $-\infty$	
	WER ↓ (relative benefit ↑)									
Whisper-base.en	3.88	4.23	4.71	5.32	6.89	17.55	49.78	92.01	123.4	31.4
Whisper-small.en	1.46	1.36	1.48	1.96	2.95	7.96	26.77	55.57	135.2	23.7
Whisper-base.en + OCR prompt	7.79	8.60	8.84	11.84	18.85	44.63	81.58	146.40	293.48	64.0
Whisper-small.en + OCR prompt	6.72	5.48	4.16	4.20	5.70	11.08	34.99	87.04	302.95	46.8
A→T (average of 2 trails)	0.21	0.31	0.57	1.10	3.05	12.73	26.29	37.68	200.75	28.5
I+A→T	0.21 (0.0%)	0.34 (-11.5%)	0.59 (-4.4%)	1.10 (0.0%)	2.97 (+2.6%)	13.05 (-2.5%)	26.78 (-1.9%)	38.07 (-1.0%)	238.1 (-18.6%)	32.3 (-3.5%)
L+A→T	0.27 (-28.6%)	0.45 (-47.5%)	0.76 (-34.5%)	1.24 (-12.7%)	3.12 (-2.3%)	12.43 (+2.4%)	24.71 (+6.0%)	35.14 (+6.7%)	191.8 (+4.5%)	27.2 (-10.5%)
I+L+A→T	0.19 (+9.5%)	0.35 (-14.8%)	0.62 (-9.7%)	1.18 (-7.3%)	3.05 (0.0%)	12.69 (+0.3%)	25.21 (+4.1%)	35.61 (+5.5%)	142.6 (+29.0%)	22.3 (+1.3%)
O+A→T	0.24 (-14.3%)	0.28 (+8.2%)	0.51 (+9.7%)	0.86 (+21.8%)	2.4 (+21.3%)	11.68 (+8.3%)	24.33 (+7.5%)	34.66 (+8.0%)	172.4 (+14.1%)	24.9 (+11.4%)
O+L+A→T	0.22 (-4.8%)	0.29 (+4.9%)	0.51 (+9.7%)	1.01 (+8.2%)	2.57 (+15.7%)	11.72 (+7.9%)	24.13 (+8.2%)	34.03 (+9.7%)	135.9 (+32.3%)	21.2 (+12.5%)
O+I+A→T	0.17 (+19.1%)	0.27 (+11.5%)	0.36 (+36.3%)	0.85 (+22.7%)	2.38 (+22.0%)	10.87 (+14.6%)	23.72 (+9.8%)	34.10 (+9.5%)	296.4 (-47.63%)	37.04 (+12.5%)
O _{oracle,3} +A→T	0.04 (+81.0%)	0.06 (+80.3%)	0.11 (+82.3%)	0.14 (+87.3%)	0.31 (+89.8%)	1.42 (+88.9%)	2.12 (+91.9%)	3.14 (+91.7%)	91.95 (+54.2%)	9.9 (+83.9%)
O _{oracle,10} +A→T	0.07 (+66.7%)	0.08 (+73.8%)	0.13 (+77.0%)	0.32 (+70.9%)	0.79 (+74.1%)	5.78 (+54.6%)	12.32 (+53.1%)	18.79 (+50.1%)	109.4 (+45.5%)	14.8 (+63.9%)

3-Equations: In Table I we list the WER and relative benefit of adding each modality to the audio-only baseline (A→T). To obtain greater statistical reliability, we train the audio-only model twice using the shuffled training data, and calculate the average WERs of these two trails as the baseline. Models are fine-tuned on the 3-Equations 2-noise training set, such that each audio is augmented with random MUSAN noise at an SNR in $[+\infty, 20, 10, 5, 0, -5, -10, -20, -\infty]$. Then the models are evaluated on 2-noise constant SNR test sets in each of $[+\infty, 20, 10, 5, 2.5, 0, -5, -10, -20, -\infty]$.

Compared to the audio-only baseline, the average benefits of adding a single modality of either image (I+A→T), lip (L+A→T), or OCR (O+A→T) are -3.5%, -10.5%, and +11.4%, respectively. Hence, only the addition of OCR brings a consistent benefit across noise levels. However, when considering 3-modality combinations, we observe more consistent benefits: adding both image and lip (I+L+A→T), the model surpasses the audio-only model by +1.3%; adding both OCR and lip (O+L+A→T), the model achieves an average benefit of +12.5%; and adding both OCR and image (O+I+A→T), the model also achieves an average benefit of +12.5%.

To ensure that visual cues are not sufficient to complete this task, we also include experiments for OCR-only (O→T), lip-only (L→T), and image-only (I→T). The WERs of O→T, L→T, and I→T are 77.3, 41.8, and 79.7, respectively. When comparing O→T to O+A→T, it's easy to figure out that the model needs audio information at all noise levels except the pure noise scenes. The overall WER increases by 52.4 in the absence of audio information. Same conclusion holds for the lip-only model and the image-only model, with an overall WER increment of 47.4 and 14.6. The performance gap between the single modality models and the 2-modalities models reflects the importance of integrating more modal information for such tasks.

In addition to varying the input modalities, we explore different representations of visual modality, from implicit to explicit. We consider tokens generated by the image encoder (DALL-E) as *implicit* visual representation, OCR as *explicit*, and oracle OCR as *accurate explicit*. As summarized in Table I, the average benefits of these models follow the order: I+A<O+A<O_{oracle,3}+A. This can be attributed to the visual representation becoming more explicit and accurate, making it easier for the model to use, which suggests that better visual

TABLE II

EVALUATION OF WER (%) AND BENEFIT OF ADDING OCR(O) ON THE SLIDEAVSR DATASET, AT DIFFERENT NOISE LEVELS BASED ON SNR AND FQ RANKER K^a VALUES. O_{ALL} MEANS USING ALL OCR WORDS.

Task	SNR(dB) =				Average
	(clean) $+\infty$	10	0	-10	
	WER ↓ (relative benefit ↑)				
A→T	33.8	42.5	44.8	70.6	47.9
O _{ALL} +A→T	31.4 (+7.0%)	37.2 (+12.4%)	47.7 (-6.4%)	75.9 (-7.6%)	48.0 (+1.4%)
O _{K=30} +A→T	30.5 (+9.7%)	35.7 (+15.8%)	46.9 (-4.7%)	75.5 (-6.9%)	47.2 (+3.5%)
O _{K=10} +A→T	30.6 (+9.5%)	34.6 (+18.5%)	46.2 (-3.1%)	77.9 (-10.3%)	47.3 (+3.6%)

^aIndicating maximum word counts for OCR.

representation can lead to better supplementary performance.

We also compare our model to two Whisper models that are of similar size as ours. Although our model is only trained on about 4k hours of audio (per-training and fine-tuning combined), which is less than 1% of Whisper's training set, it shows better performance by harnessing multi-modal capabilities. This suggests leveraging more modalities can optimize performance even with a limited amount of data. Since Whisper models are able to accept prompts, we also consider comparing to the prompted Whisper (as shown in Table I, Whisper + OCR prompt). However, the performance degraded when prompted with OCR texts, we suspect this is because Whisper models suffer more from hallucination problem when provided with a prompt.

SlideAVSR: Based on previous 3-Equations experiments, image data with more text tends to perform better with explicit representation. Therefore, we similarly conduct OCR-based experiments on SlideAVSR. As shown in Table II, adding extra OCR modality (O_{ALL}+A→T) shows an average relative benefit of +1.4%, which confirms that adding OCR improves overall recognition performance.

B. Does each modality provide a uniform accuracy boost?

We explored the hypothesis that the benefits of multimodality would be largest when there is a "medium" amount of noise because (a) if the audio is very clean, the other modalities are unnecessary and (b) if the audio is very noisy, the model cannot find a correspondence between the auditory and visual channels. Results on both datasets are described below.

3-Equations: In Table I, when audio is clean, neither image nor lip helps improve recognition performance. However, when looking into higher noise levels, the relative benefit of adding images exhibits a very different trend compared to adding lips. As shown in Fig. 4, with increasing noise

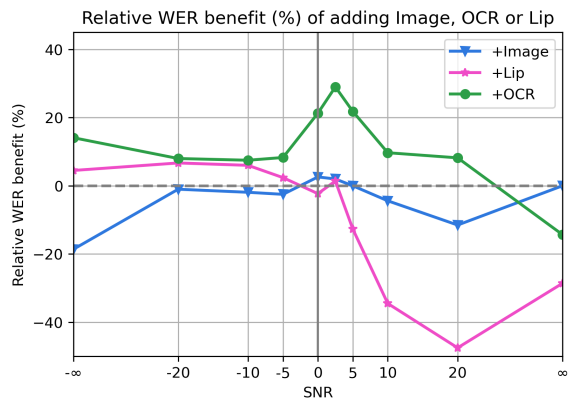


Fig. 4. Relative WER benefit (%) of adding single image, OCR, or lip modalities on 3-Equations 2-noise test set.

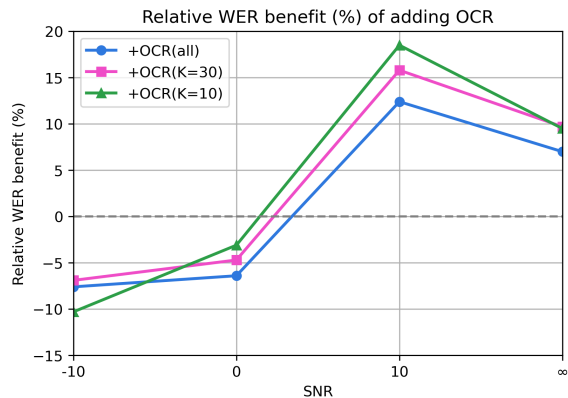


Fig. 5. Relative WER benefit (%) of adding OCR with different K values for FQ Ranker on SlideAVSR test set.

levels, the benefit of lips in enhancing accuracy becomes more amplified, showing the same trend as discussed in previous works [15]. Conversely, the benefit of images follows a trend of first increasing and then decreasing, peaking in the middle when SNR=0dB. This discrepancy could be attributed to image modality not being inherently synchronized with speech, unlike lip movements. Therefore, when above a “sweet spot”, the audio is too noisy for the image to establish a reliable correspondence, thus the relative benefit begins to decline. The benefit of OCR shows a similar trend as image modality, but the sweet spot is likely different from the image due to visual information quality (in this case, SNR=2.5dB).

SlideAVSR: The results in Table II and Fig. 5 show a similar phenomenon: Including the OCR modality ($O_{ALL}+A \rightarrow T$) outperforms the audio-only baseline at low noise levels (∞ , 10dB), but performs worse at higher noise levels. The relative benefit of adding OCR initially increases, then decreases, achieving the greatest benefit of +12.4% at SNR=10dB.

C. How does irrelevant visual info. affect ASR performance?

Since 3-Equations dataset has 3 written equations but only 2 spoken equations in each example, the OCR inputs inherently contain 1/3 irrelevant information. Therefore, we are also interested in how the proportion of irrelevant visual inputs affects performance. On the 3-Equations dataset, we thus add 7 extra irrelevant oracle OCR sentences in addition to the 3

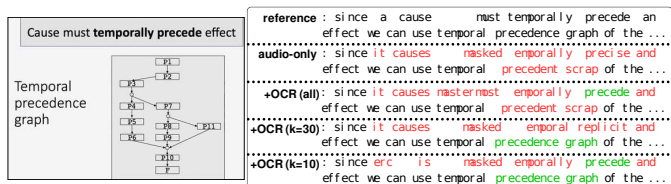


Fig. 6. An example on SlideAVSR of how filtering relevant information helps the accuracy. The corresponding slide is on the left. Red words are wrong predictions, and green words are those that were wrongly predicted by the audio-only model but corrected by the OCR plus audio model.

relevant sentences, making a dataset with 4/5 irrelevant inputs. The result is included in Table I ($O_{oracle,10}+A \rightarrow T$). We observe that the overall benefit is much worse than using accurate oracle OCR. This trend confirms that adding more irrelevant visual information will hinder the model from finding the correct information, especially in noisy environments.

In SlideAVSR, one slide sample can contain hundreds of OCR words, but only a small proportion is relevant to the speech. Following the work in [18], we used FQ Ranker that calculates word ranks based on the frequency of word occurrences in English Wikipedia, and filters the OCR words based on word frequency. We use 10 and 30 as the maximum word count (K) for prompts. This preprocessing step helps us filter the most relevant or long-tail words in OCR words, and conceivably helps the recognition performance. An example of how this filtering process helps the performance is shown in Fig. 6. As shown in Table II and Fig. 5, although the performance of adding OCR is worse than that of audio-only at some noise levels, the relative benefit of OCR generally increases as we filter more stringently (i.e., smaller K).

VI. CONCLUSIONS

We investigated how multiple modalities impact the accuracy of speech recognition performed by decoder-only discrete speech models. Our experiments suggest that fusing multiple modalities generally enhances recognition performance, but with caveats: Image information exhibits a different trend from lip movements. Typically, as the noise level increases, the accuracy benefit of the lip information grows larger, whereas images provide the greatest benefit at moderate noise levels. Also, we observe a steady performance improvement when relevant visual information is filtered in preprocessing. To our knowledge, this paper is the first to show the benefit of combining audio, images, and lip movements in one model.

However, a limitation of our work is that using a synthesized dataset may introduce bias due to inherent synthesis errors. In future research, other modality backbone models, such as VQ-Wav2Vec and ViViT, and more real-world datasets should be explored to assess the generalizability of our findings.

Acknowledgment: This research was supported by the NSF National AI Institute for Student-AI Teaming (iSAT) under grant DRL #2019805, and also from an NSF CAREER grant #2046505. The opinions expressed are those of the authors and do not represent views of the NSF.

REFERENCES

- [1] X. Chang, B. Yan *et al.*, “Exploring speech recognition, translation, and understanding with discrete speech units: A comparative study,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 11 481–11 485.
- [2] G. Team, R. Anil *et al.*, “Gemini: a family of highly capable multimodal models,” *arXiv preprint arXiv:2312.11805*, 2023.
- [3] V. A. Trinh, R. Southwell *et al.*, “Discrete multimodal transformers with a pretrained large language model for mixed-supervision speech processing,” *arXiv preprint arXiv:2406.06582*, 2024.
- [4] A. Radford, J. W. Kim *et al.*, “Robust speech recognition via large-scale weak supervision,” in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [5] T. Afouras, J. S. Chung *et al.*, “Deep audio-visual speech recognition,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 12, pp. 8717–8727, 2018.
- [6] B. Shillingford, Y. Assael *et al.*, “Large-scale visual speech recognition,” *arXiv preprint arXiv:1807.05162*, 2018.
- [7] J. Achiam, S. Adler *et al.*, “Gpt-4 technical report,” *arXiv preprint arXiv:2303.08774*, 2023.
- [8] D. Ivanko, D. Ryumin, and A. Karpov, “A review of recent advances on deep learning methods for audio-visual speech recognition,” *Mathematics*, vol. 11, no. 12, 2023. [Online]. Available: <https://www.mdpi.com/2227-7390/11/12/2665>
- [9] D. Ryumin, D. Ivanko, and E. Ryumina, “Audio-visual speech and gesture recognition by sensors of mobile devices,” *Sensors*, vol. 23, no. 4, p. 2284, 2023.
- [10] P. Ma, S. Petridis, and M. Pantic, “End-to-end audio-visual speech recognition with conformers,” in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 7613–7617.
- [11] Q. Song, B. Sun, and S. Li, “Multimodal sparse transformer network for audio-visual speech recognition,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 34, no. 12, pp. 10 028–10 038, 2022.
- [12] M. Burchi and R. Timofte, “Audio-visual efficient conformer for robust speech recognition,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 2258–2267.
- [13] B. Shi, W.-N. Hsu *et al.*, “Learning audio-visual speech representation by masked multimodal cluster prediction,” in *International Conference on Learning Representations*.
- [14] B. Shi, W.-N. Hsu, and A. Mohamed, “Robust self-supervised audio-visual speech recognition,” *arXiv preprint arXiv:2201.01763*, 2022.
- [15] P. Ma, A. Haliassos *et al.*, “Auto-avsr: Audio-visual speech recognition with automatic labels,” in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.
- [16] J. Hong, M. Kim *et al.*, “Watch or listen: Robust audio-visual speech recognition with visual corruption modeling and reliability scoring,” in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2023, pp. 18 783–18 794.
- [17] H. Wang, F. Yu *et al.*, “Slidespeech: A large scale slide-enriched audio-visual corpus,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 11 076–11 080.
- [18] H. Wang, S. Kurita *et al.*, “Slideavsr: A dataset of paper explanation videos for audio-visual speech recognition,” *arXiv preprint arXiv:2401.09759*, 2024.
- [19] D. Zhang, Y. Yu *et al.*, “Mm-llms: Recent advances in multimodal large language models,” *arXiv preprint arXiv:2401.13601*, 2024.
- [20] Y. Chu, J. Xu *et al.*, “Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models,” *arXiv preprint arXiv:2311.07919*, 2023.
- [21] H. Liu, C. Li *et al.*, “Visual instruction tuning,” *Advances in neural information processing systems*, vol. 36, 2024.
- [22] K. Zheng, X. He, and X. E. Wang, “Minigt-5: Interleaved vision-and-language generation via generative vokens,” *arXiv preprint arXiv:2310.02239*, 2023.
- [23] S. Wu, H. Fei *et al.*, “Next-gpt: Any-to-any multimodal llm,” in *Forty-first International Conference on Machine Learning*.
- [24] J. H. Yeo, S. Han *et al.*, “Where visual speech meets language: Vsp-llm framework for efficient and context-aware visual speech processing,” *arXiv preprint arXiv:2402.15151*, 2024.
- [25] U. Cappellazzo, M. Kim *et al.*, “Large language models are strong audio-visual speech recognition learners,” *arXiv preprint arXiv:2409.12319*, 2024.
- [26] S. Zhang, S. Roller *et al.*, “Opt: Open pre-trained transformer language models,” *arXiv preprint arXiv:2205.01068*, 2022.
- [27] L. Barrault, Y.-A. Chung *et al.*, “Seamlessm4t-massively multilingual & multimodal machine translation,” *arXiv preprint arXiv:2308.11596*, 2023.
- [28] A. Ramesh, M. Pavlov *et al.*, “Zero-shot text-to-image generation,” in *International conference on machine learning*. Pmlr, 2021, pp. 8821–8831.
- [29] V. Panayotov, G. Chen *et al.*, “Librispeech: an asr corpus based on public domain audio books,” in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [30] Y. Jia, M. T. Ramanovich *et al.*, “Cvss corpus and massively multilingual speech-to-speech translation,” *arXiv preprint arXiv:2201.03713*, 2022.
- [31] C. Wang, A. Wu, and J. Pino, “Covost 2 and massively multilingual speech-to-text translation,” *arXiv preprint arXiv:2007.10310*, 2020.
- [32] T.-Y. Lin, M. Maire *et al.*, “Microsoft coco: Common objects in context,” in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. Springer, 2014, pp. 740–755.
- [33] I. Loshchilov and F. Hutter, “Decoupled weight decay regularization,” in *International Conference on Learning Representations*.
- [34] “pytsx3: Offline text to speech (tts) converter for python,” <https://github.com/nateshmbhat/pytsx3>.
- [35] “Easyocr: Ready-to-use ocr with 80+ supported languages,” <https://github.com/JaidedAI/EasyOCR>.
- [36] K. Prajwal, R. Mukhopadhyay *et al.*, “A lip sync expert is all you need for speech to lip generation in the wild,” in *Proceedings of the 28th ACM international conference on multimedia*, 2020, pp. 484–492.
- [37] D. Snyder, G. Chen, and D. Povey, “Musan: A music, speech, and noise corpus,” *arXiv preprint arXiv:1510.08484*, 2015.
- [38] C. Cohn, C. Snyder *et al.*, “A multimodal approach to support teacher, researcher and ai collaboration in stem+c learning environments,” *British Journal of Educational Technology*.