

SynesLM: A Unified Approach for Audio-visual Speech Recognition and Translation via Language Model and Synthetic Data

Yichen Lu^{1*}, Jiaqi Song^{1*}, Xuankai Chang¹, Hengwei Bian¹, Soumi Maiti¹, Shinji Watanabe¹

¹Carnegie Mellon University, Pittsburgh, PA, USA

{yichen15, jiaqison, xuankaic}@andrew.cmu.edu

Abstract

In this work, we present *SynesLM*, an unified model which can perform three multimodal language understanding tasks: audio-visual automatic speech recognition (AV-ASR) and visual-aided speech/machine translation (VST/VMT). Unlike previous research that focused on lip motion as visual cues for speech signals, our work explores more general visual information within entire frames, such as objects and actions. Additionally, we use synthetic image data to enhance the correlation between image and speech data. We benchmark *SynesLM* against the How2 dataset, demonstrating performance on par with state-of-the-art (SOTA) models dedicated to AV-ASR while maintaining our multitasking framework. Remarkably, for zero-shot AV-ASR, *SynesLM* achieved SOTA performance by lowering the Word Error Rate (WER) from 43.4% to 39.4% on the VisSpeech Dataset. Furthermore, our results in VST and VMT outperform the previous results, improving the BLEU score to 43.5 from 37.2 for VST, and to 54.8 from 54.4 for VMT.

Index Terms: Audio-Visual Automatic Speech Recognition, Speech Translation, Multimodal Language Model, Multitask

1. Introduction

Synesthesia is a neurological condition where stimulation of one sensory pathway involuntarily triggers experiences in another, such as perceiving colors when hearing sounds. This phenomenon highlights the complex integration of multisensory inputs in human cognition, essential for comprehending the world through combined audio and visual stimuli [1, 2]. Visual cues, for instance, enhance speech recognition and language understanding, aiding in translation tasks.

Motivated by these insights, we aim to design a unified model for a range of audio-visual tasks that concurrently use audio and visual inputs. Similar to human learning processes, our model benefits from multitask training. Additionally, we found that incorporating pretrained language model weights further improves our model’s performance. Recently, there are several works trying to incorporate visual information with speech to perform automatic speech recognition (ASR). For instance, AV-HuBERT [3] focuses on audio-visual speech representation, utilizing video lip recordings to learn powerful speech representations. AVATAR [4] extends this concept by incorporating full visual frames for unconstrained audio-visual automatic speech recognition (AV-ASR). The AVFormer [5] explores injecting vision into frozen speech models for zero-shot AV-ASR, highlighting the potential of lightweight domain adaptation. However, existing research primarily focuses on AV-ASR, with a noticeable shortage of studies investigating a unified approach for

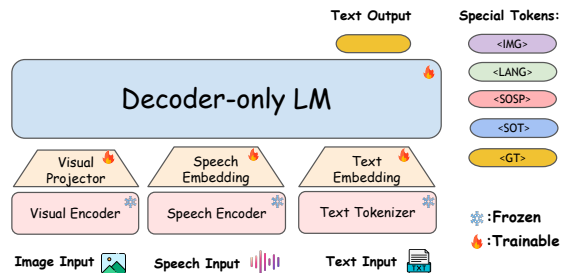


Figure 1: *An overview of SynesLM architecture. The definition of the special tokens will be discussed at the end of the Section 3.*

all audio-visual language understanding tasks, such as visual-aided speech translation.

On the other hand, large language models (LLMs) have surged in popularity due to their advanced capabilities in natural language understanding and generation. Several multimodal language models have been developed for more complex tasks necessitating either visual or audio modalities. LLaVa [6] and BLIP2 [7] incorporate visual modality for visual question answering and image captioning. VoxLM [8] and VioLA [9] utilize speech modality to perform multiple speech tasks, such as speech recognition and text-to-speech. Additionally, works like OneLLM [10] and X-LLM [11] extend the input set to various modalities, including audio and visual. However, none of these models can concurrently process audio and visual inputs, meaning they cannot perform tasks like AV-ASR.

Since current multimodal language understanding approaches [4, 5] are limited to handling AV-ASR, and existing MLLMs can only process tasks involving text plus one additional modality, this paper introduces the **Synesthesia Language Model (SynesLM)**. SynesLM is a novel, unified approach capable of performing a variety of audio-visual input tasks within a single model. Drawing inspiration from the phenomenon of synesthesia and leveraging advancements in recent language models, SynesLM is meticulously designed to proficiently handle complex language tasks such as audio-visual automatic speech recognition (AV-ASR), visual speech translation (VST), and visual machine translation (VMT). Additionally, we proposed a data recovery pipeline using LLM and image generative model for multimodal language understanding dataset like How2 [12] to enhanced its multimodal interaction between speech and visual modality.

We make the following contributions in this work:

- **A unified model for multiple audio-visual tasks:** We explore a novel framework capable of comprehending and processing visual, speech, and textual data to perform various of

* Equal contributions

multimodal language understanding tasks.

- **Synthetic visual data recovery pipeline:** To address the poor quality of visual data in multimodal speech datasets like How2 [12], we developed a novel data recovery pipeline. This pipeline significantly improved the integration between visual and speech data.
- **Performance:** We achieved improved performance across all tasks, recording a 4.0% WER absolute improvement in zero-shot AV-ASR and BLEU scores of 43.5 for VST and 54.8 for VMT, showcasing our model’s strong audio processing and visual comprehension capabilities (Table 3).
- **Reproducibility:** For reproducibility, we open-source our code and model checkpoints in the form of ESPNet recipe in <https://github.com/espnet/espnet>.

2. Related Work

Multimodal Language Models (MLMs). Recently, numerous MLMs have been proposed for various modalities. Advances in Visual Language Models (VLMs), such as [6, 13–18], have significantly improved the integration of visual information into pre-trained language models. These models use a pre-trained vision encoder for visual feature extraction, excelling in tasks like image captioning and visual question answering. Similarly, recent studies [8, 19–25] have begun exploring unified models for various speech and text tasks, employing self-supervised learning (SSL) feature extractors as audio encoders alongside pre-trained language models to enhance language comprehension. Some recent research [10, 11, 26–28] has attempted to incorporate both visual and audio modalities into language models. However, these models can only process one specific modality with text and do not explore the interaction between audio and visual modalities. The most related work to ours is Video-SALMONN [29], but it lacks multilingual capabilities, limiting its ability to perform translation tasks.

AV-ASR Methods. Several methods have been designed for AV-ASR tasks. AV-HuBERT [3] uses video recordings to learn robust speech representations through masked multimodal cluster prediction, focusing on lip motion. AVATAR [4] employs a multimodal encoder with a transformer decoder for natural language speech recognition output. AVFormer [5] explores injecting vision into frozen speech models for zero-shot AVSR, showing the potential for lightweight domain adaptation. Additionally, prompting-whisper [30] designs a cascade model that injects visual prompts for AV-ASR. However, these methods are limited to AV-ASR tasks or utilize cascade structures, while we design an end-to-end approach for multiple speech-visual tasks.

3. Method

This section initially outlines the architecture of our model, followed by a description of how we recover the multimodal data and tokenize it.

3.1. Data Representation

Discrete Speech Representation. Recent advances in discrete speech representations for ASR and ST have improved training speed, inference speed, and storage efficiency [31]. These methods utilize SSL’s ability to capture linguistic and acoustic information, surpassing previous techniques like log-Mel filterbanks. By converting continuous audio features into discrete speech tokens, speech input can be handled similarly to text to-

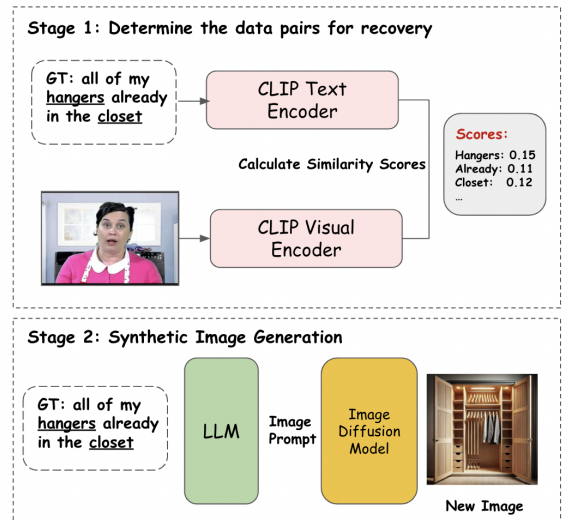


Figure 2: *Synthetic Data Recovery Pipeline.*

kens. This unification allows speech and text tokens, sharing similar semantics, to be processed together in a single LM using a unified character set called discrete speech-text tokens.

Visual Encoders and Features. For the visual modality, we randomly select a single frame from each video clip as the visual input, keeping the image whole rather than dividing it into patches. This simplifies the extraction of object and action information, ensuring one frame provides sufficient details. Using a pre-trained visual encoder from CLIP [32], we extract features from the entire image. To bridge the visual and language modalities, a Vision-Language connector layer (a Multi-Layer Perceptron) maps these visual features into the same embedding space as the discrete audio-text tokens, aligning them with the linguistic components. We also experiment with other CLIP-like pretrained visual encoders [33, 34] to explore their performance in fusing discrete text-speech representations.

Data Format. We use several special tokens to indicate different tasks and different modality. The entire sequence start with a `<SOS>` token, end with a `<EOS>` token. We use a `` token to indicate the position of visual information in input sequence. `<start-of-text>` and `<start-of-speech>` (`<SOT>`, `<SOSP>` in Figure 1) represent the different input modalities. `<generate-text>` (`<GT>` in Figure 1) indicate the output modality of SynesLM. Moreover, the language token `<LANG>` (e.g. `<EN>` and `<PT>`) in the input and output sequence can indicate the source and target language, which enable the translation capacity of our framework. The arrangement of modality and language tokens in SynesLM facilitates its multitasking capability, where each specific combination signifies a distinct task. For example, the token combination `{, <SOSP>, <EN>, <GT>, <PT>}` represents the task of VST from English to Portuguese.

3.2. SynesLM

Figure 1 illustrates the overall architecture of SynesLM. We utilize a transformer-based decoder-only LM as our backbone. We use pre-trained OPT model [35] to initialize our weights to achieve better performance and training efficiency [36].

To process the speech inputs, SynesLM employs a SSL feature extractor like HuBERT [37] and k-means to generate



Figure 3: **Qualitative examples on How2 ASR.** We show that our audio with original visual (A+OV) and audio with synthetic visual (A+SV) method successfully extract and understand the information from the image and incorporate the information with speech representation to perform ASR task.

speech discrete tokens. The integration of visual information is accomplished through a visual encoder and a Vision-Language MLP projector. Details will be explained later in this section.

We concatenate the input text sequence $Y = (y_i \in \mathcal{V}_{\text{txt}} | i = 1, \dots, t_{\text{txt}})$ with the input discrete speech token(dst) sequence $D = (d_i \in \mathcal{V}_{\text{dst}} | i = 1, \dots, t_{\text{dst}})$ as the T -length discrete speech-text token sequence $Z = (z_i \in \mathcal{V} = \mathcal{V}_{\text{txt}} \cup \mathcal{V}_{\text{dst}} \cup \mathcal{V}_{\text{special}} | i = 1, \dots, T)$, where $T = t_{\text{txt}} + t_{\text{dst}}$ and $\mathcal{V}_{\text{special}}$ is the special token set e.g. <LANG> etc. We feed the input discrete speech-text sequence Z into the embedding layer to get the D -dimensional embedding space expressed by $\mathbf{E} = (\mathbf{e}_i \in \mathbb{R}^D | i = 1, \dots, T)$. Additionally, we add a single visual token in the D -dimensional embedding space as \mathbf{V} . We feed the visual embedding \mathbf{V} and discrete speech-text embedding \mathbf{E} in to our auto-regressive LM. The joint probability of the generated text sequence Y can be expressed as:

$$p(Y) = \prod_{i=T+1}^{T+1+L} p(Y_i | \mathbf{V}, \mathbf{E}_{1:i-1}), \quad (1)$$

where L is the generated text length.

During training and inference process, given the previous information \mathbf{V} , $\mathbf{E}_{1:i-1}$, we can use this formula to predict the probability distribution in the next time step \hat{p}_i by $\hat{p}_i = \text{SynesLM}(Y_i | \mathbf{V}, \mathbf{E}_{1:i-1})$. Then the cross entropy (CE) loss of the model can be expressed as:

$$\mathcal{L}_{\text{CE}}(p_i, \hat{p}_i) = - \sum_{c=1}^{|\mathcal{V}|} p_i(c) \log \hat{p}_i(c), \quad (2)$$

where p_i is the reference probability distribution.

3.3. Data Recovery via Synthetic Image Data

We mainly utilize two different multimodal datasets in our study:

- **How2** [12] is a rich multimodal collection of instructional videos with English subtitles and Portuguese translations. We use the 300 hours subset of video content in this project. We utilize the training set contains 182,167 clips, the validation set includes 1,939 clips, and the test set consists of 2,298 clips.
- **VisSpeech** [4] is constructed by selecting 482 video clips from the extensive HowTo100M dataset [38], aimed at creating a robust test set. The selected clips are characterized

by their strong visual-audio correlation, providing an ideal platform to evaluate the effectiveness of our visual modality.

We use the How2 dataset [12] for training and testing. To ensure the correlation between visual and textual data, we initially calculate the similarity score between the ground truth transcription and the selected frame. As illustrated in Figure 2, for a dataset containing image-text pairs, we input the images and text into CLIP [32] to calculate the cosine similarity score for each word in the ground truth. The distribution of similarity scores indicates that over 62% of visual data have a similarity score below 20%, reflecting the poor quality of the How2 visual data. To address this issue, we designed a data recovery pipeline to enhance modality interaction within the How2 dataset. If the maximum score for each words in the sentence is below a given threshold ($\tau = 0.2$ in our experiment), we generate new visual data based on the ground truth. For image data generation, we first input the ground truth data into a Large Language Model (LLM) to extract object or action information and generate an image generation prompt based on that information. Finally, we feed the image generation prompt into our image diffusion model [39].

4. Experiments

4.1. Experiment Setup

According to the survey of discrete representation [31], we fixed our BPE size to 3K and k-mean cluster number to 2K since How2 only contains about 300 hours of speech data. We did all of our experiment on the open-source E2E speech processing toolkit ESPNet [40]. For the speech encoder, we select XLS-R [41] as our SSL feature extraction method. In all experiments, we use $2 \times \text{V100}$ GPUs for training our models. Our core model is based on a Decoder-only Transformer architecture, featuring 12 layers with a 768-dimensional feature space and 12 attention heads, culminating in a total of 125 million trainable parameters. For multitask training, we ensured an equal distribution of data across the different tasks. We froze the vision encoder to concentrate training on the vision-language projector and the decoder-only language model, achieving an end-to-end training process.

4.2. Results

Table 1 shows the experimental results evaluating the impact of visual features in both single-task and multi-task settings. Our

Table 1: *Experimental results comparing on single-task and multi-task on different visual encoders. The relative improvement rate represents the performance influence of incorporating visual modality to the original model under the certain tasks, the visual influence in single-task or multi-task scenarios are calculated independently.*

	ASR (How2) WER(↓)	ASR (VisSpeech) WER(↓)	ST BLEU(↑)	MT BLEU(↑)
How2 Baseline w/ visual [12]	18.0	-	37.2	54.4
Single w/o visual	17.6	41.6	40.5	55.2
Single w/ visual (CLIP [32])	17.0 (+3.41%)	41.7 (-0.24%)	40.7 (+0.49%)	55.6 (+0.72%)
Single w/ visual (EVA-CLIP [34])	17.6 (+0.00%)	40.4 (+2.88%)	41.3 (+1.98%)	54.7 (-0.91%)
Single w/ visual (SigLIP [33])	17.3 (+1.70%)	42.3 (-1.68%)	41.4 (+2.22%)	54.7 (-0.91%)
Multi w/o visual	16.4	40.8	42.9	54.7
Multi w/ visual (CLIP [32])	16.1 (+1.83%)	40.1 (+1.72%)	43.0 (+0.23%)	54.0 (-1.30%)
Multi w/ visual (EVA-CLIP [34])	15.9 (+3.05%)	40.2 (+1.47%)	43.4 (+1.17%)	53.9 (-1.46%)
Multi w/ visual (SigLIP [33])	15.7 (+4.27%)	39.4 (+3.43%)	43.5 (+1.40%)	54.8 (+0.18%)

Table 2: **Ablation Study.** *Experimental AV-ASR task results comparing the visual influence on multi-task SigLIP [33] encoder. Random visual means randomly select the visual features from other video clips.*

	ASR (How2) WER(↓)	ASR (Visspeech) WER(↓)
Multi w/o visual	16.4	40.8
Multi w random visual	16.4	41.0
Multi w visual	15.9	39.4
Multi w visual + synthetic	15.7	39.4

methodology is evaluated across three distinct tasks: ASR, ST, and MT. Furthermore, we explore the effectiveness of three different visual encoders [32–34] to determine which encoder best aligns with speech-text discrete representations.

There are some notable instances where the visual modality enhances ASR task performance. As demonstrated in Figure 3, visual cues are instrumental in recognizing seldom-used vocabulary, especially when there is a strong correlation between the visual content and these specific words. This highlights the potential of language models to comprehend visual information and merge it with speech data for a multimodal understanding. The findings indicate that incorporating SynesLM’s visual modality consistently enhances performance across all the tasks when compared to an audio-only baseline (e.g., improving the WER from 16.4% to 15.7% in a multitask setting using SigLIP [33]). Among the visual encoders tested, CLIP [32] achieves the highest performance in single-task experiments. On the other hand, in the multitask framework, SigLIP [33] demonstrates superior efficacy, notably achieving a 3.43% relative performance increase in the zero-shot AV-ASR on the Vis-Speech [4] dataset. Furthermore, to delve deeper into the impact of visual features, we conduct an ablation study where the visual input is replaced with a random image. As shown in Table 2, the WER for the random visual input scenario increases from 15.7% to 16.4%, which is the same as the performance without visual input. In addition, our synthetic data recovery technique further improves the performance from 15.9% to 15.7%, indicating that better audio-visual correlation could further benefit model performance. This outcome underscores the robustness of our model in audio-visual tasks.

Compare with the SOTA methods. In Table 3, we com-

Table 3: *Comparison with the state-of-the-art on AV-ASR task. The train set section for last two rows indicates that those methods use additional dataset other than How2 [12] for pre-training. † denotes initialization with OPT. Results are reported as WER (% , lower is better).*

Method	Train Set	How2	VisSpeech
How2 Base [12]	300hrs	18.0	-
LLD [42]	300hrs	16.7	-
VAT [43]	300hrs	18.0	-
MultiRes [44]	300hrs	20.5	-
AVATAR [4]	300hrs	15.6	43.4
Ours [†]	300hrs	15.7	39.4
AVFormer [5]	960hrs + 6500hrs	13.6	16.6
AVATAR [4]	300hrs + 131k hrs	9.1	11.3
Prompt-whisper [30]	680k hrs	-	7.16

pare our model with state-of-the-art methods on the AV-ASR task. The results demonstrate that our approach surpasses most of the methods when utilizing only the How2 dataset. Note that, both AVATAR [4] and AVFormer [5] pretrain on the vast HowTo100M [38] dataset, which makes them perform well explicitly on single AV-ASR task. Contrary to the design purposes of these models, we aim to explore an unified model architecture for different audio-visual related tasks. Under the multitasking scenario, our model not only retains high performance in the AV-ASR task but also surpasses the How2 baseline in ST and MT tasks. Specifically, as shown in Table 1, we observed a significant improvement in the BLEU score for ST, increasing from 37.2 to 43.5. Similarly, for MT, there was a modest enhancement from 54.4 to 54.8 in the BLEU score.

5. Conclusion

We introduce SynesLM, a novel multimodal language model designed for multiple audio-visual tasks. SynesLM outperforms existing single-task methods, showcasing the effective integration of audio and visual data. The experiments also reveal SynesLM’s ability to synergize auditory, textual and visual information effectively. The results underscore SynesLM’s proficiency across all evaluated tasks and highlight its potential for broader applications in audio-visual processing.

6. Acknowledgements

Experiments of this work used the Bridges2 system at PSC and Delta system at NCSA through allocations CIS210014 and IRI120008P from the Advanced Cyber infrastructure Coordination Ecosystem: Services & Support (ACCESS) program, supported by National Science Foundation grants #2138259, #tel:2138286, #tel:2138307, #tel:2137603, and #tel:2138296.

7. References

- [1] F. N. Newell and K. J. Mitchell, "Multisensory integration and cross-modal learning in synaesthesia: a unifying model," *Neuropsychologia*, vol. 88, pp. 140–150, 2016.
- [2] R. Chiou, M. Stelter, and A. N. Rich, "Beyond colour perception: Auditory–visual synaesthesia induces experiences of geometric objects in specific locations," *Cortex*, vol. 49, no. 6, pp. 1750–1763, 2013.
- [3] B. Shi, W.-N. Hsu, K. Lakhotia, and A. Mohamed, "Learning audio-visual speech representation by masked multimodal cluster prediction," in *International Conference on Learning Representations*, 2022.
- [4] V. Gabeur *et al.*, "AVATAR: Unconstrained audiovisual speech recognition," in *Interspeech*, 2022.
- [5] P. H. Seo, A. Nagrani, and C. Schmid, "AVFormer: Injecting vision into frozen speech models for zero-shot av-asr," in *CVPR*, 2023.
- [6] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," in *Advances in Neural Information Processing Systems*, 2023.
- [7] J. Li, D. Li, S. Savarese, and S. Hoi, "Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models," *arXiv preprint arXiv:2301.12597*, 2023.
- [8] S. Maiti *et al.*, "Voxlm: unified decoder-only models for consolidating speech recognition/synthesis and speech/text continuation tasks," in *ICASSP*, 2024.
- [9] T. Wang *et al.*, "Viola: Unified codec language models for speech recognition, synthesis, and translation," *arXiv preprint arXiv:2305.16107*, 2023.
- [10] J. Han *et al.*, "Onellm: One framework to align all modalities with language," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [11] F. Chen *et al.*, "X-llm: Bootstrapping advanced large language models by treating multi-modalities as foreign languages," 2023.
- [12] R. Sanabria *et al.*, "How2: A large-scale dataset for multimodal language understanding," 2018.
- [13] W. Dai *et al.*, "Instructblip: Towards general-purpose vision-language models with instruction tuning," *arXiv preprint arXiv:2305.06500*, 2023.
- [14] J.-B. Alayrac *et al.*, "Flamingo: a visual language model for few-shot learning," in *Advances in Neural Information Processing Systems*, 2022.
- [15] J. Bai *et al.*, "Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond," 2023.
- [16] K. Chen *et al.*, "Shikra: Unleashing multimodal llm's referential dialogue magic," *arXiv preprint arXiv:2306.15195*, 2023.
- [17] J. Chen *et al.*, "Minigpt-v2: large language model as a unified interface for vision-language multi-task learning," *arXiv preprint arXiv:2310.09478*, 2023.
- [18] D. Driess *et al.*, "Palm-e: An embodied multimodal language model," *arXiv preprint arXiv:2303.03378*, 2023.
- [19] Z. Borsos *et al.*, "Audiollm: a language modeling approach to audio generation," in *ICASSP*. IEEE, 2023.
- [20] P. K. Rubenstein *et al.*, "Audiopalm: A large language model that can speak and listen," *arXiv preprint arXiv:2306.12925*, 2023.
- [21] Q. Dong *et al.*, "Polyvoice: Language models for speech to speech translation," 2023.
- [22] K. Lakhotia *et al.*, "Generative spoken language modeling from raw audio," 2021.
- [23] C. Tang *et al.*, "Salmonn: Towards generic hearing abilities for large language models," *arXiv preprint arXiv:2310.13289*, 2023.
- [24] J. Wu *et al.*, "On decoder-only architecture for speech-to-text and large language model integration," in *Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2023, pp. 1–8.
- [25] Z. Chen *et al.*, "Salm: Speech-augmented language model with in-context learning for speech recognition and translation," *arXiv preprint arXiv:2310.09424*, 2023.
- [26] C. Lyu *et al.*, "Macaw-llm: Multi-modal language modeling with image, audio, video, and text integration," *arXiv preprint arXiv:2306.09093*, 2023.
- [27] H. Zhang, X. Li, and L. Bing, "Video-llama: An instruction-tuned audio-visual language model for video understanding," in *EMNLP*, 2023.
- [28] M. Kim *et al.*, "Tmt: Tri-modal translation between speech, image, and text by processing different modalities as different languages," 2024.
- [29] G. Sun *et al.*, "video-SALMONN: Speech-enhanced audio-visual large language models," in *Forty-first International Conference on Machine Learning*, 2024. [Online]. Available: <https://openreview.net/forum?id=nYsh5GF1qX>
- [30] P. Peng, B. Yan, S. Watanabe, and D. Harwath, "Prompting the hidden talent of web-scale speech models for zero-shot task generalization," in *Interspeech*, 2023.
- [31] X. Chang *et al.*, "Exploring speech recognition, translation, and understanding with discrete speech units: A comparative study," in *ICASSP*, 2024.
- [32] A. Radford *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*. PMLR, 2021, pp. 8748–8763.
- [33] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer, "Sigmoid loss for language image pre-training," 2023.
- [34] Q. Sun *et al.*, "Eva-clip: Improved training techniques for clip at scale," *arXiv preprint arXiv:2303.15389*, 2023.
- [35] S. Zhang *et al.*, "Opt: Open pre-trained transformer language models," *arXiv preprint arXiv:2205.01068*, 2022.
- [36] M. Hassid *et al.*, "Textually pretrained speech language models," in *Advances in Neural Information Processing Systems*, 2024.
- [37] W.-N. Hsu *et al.*, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2021.
- [38] A. Miech *et al.*, "HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips," in *ICCV*, 2019.
- [39] R. Rombach *et al.*, "High-resolution image synthesis with latent diffusion models," 2022.
- [40] S. Watanabe, H. Hori, S. Karita *et al.*, "ESPnet: End-to-end speech processing toolkit," in *Interspeech*, 2018.
- [41] A. Babu *et al.*, "Xls-r: Self-supervised cross-lingual speech representation learning at scale," *arXiv preprint arXiv:2111.09296*, 2021.
- [42] S. Ghorbani, Y. Gaur, Y. Shi, and J. Li, "Listen, look and deliberate: Visual context-aware speech recognition using pre-trained text-video representations," in *IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2021, pp. 621–628.
- [43] O. Caglayan *et al.*, "Multimodal grounding for sequence-to-sequence speech recognition," in *ICASSP*. IEEE, 2019, pp. 8648–8652.
- [44] G. Paraskevopoulos, S. Parthasarathy, A. Khare, and S. Sundaram, "Multimodal and multiresolution speech recognition with transformers," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 2381–2387.