

Enhancing Audiovisual Speech Recognition Through Bifocal Preference Optimization

Yihan Wu^{2,1}, Yichen Lu¹, Yifan Peng¹, Xihua Wang², Ruihua Song^{2*}, Shinji Watanabe^{1*}

¹Carnegie Mellon University,

²Gaoling School of Artificial Intelligence, Renmin University of China, {yihanwu, xihuaw, rsong}@ruc.edu.cn, {yichenl5, yifanpen, swatanab}@andrew.cmu.edu

Abstract

Audiovisual Automatic Speech Recognition (AV-ASR) aims to improve speech recognition accuracy by leveraging visual signals. It is particularly challenging in unconstrained real-world scenarios across various domains due to noisy acoustic environments, spontaneous speech, and the uncertain use of visual information. Most previous works fine-tune audio-only ASR models on audiovisual datasets, optimizing them for conventional ASR objectives. However, they often neglect visual features and common errors in unconstrained video scenarios. In this paper, we propose using a preference optimization strategy to improve speech recognition accuracy for real-world videos. First, we create preference data via simulating common errors that occurred in AV-ASR from two focals: manipulating the audio or vision input and rewriting the output transcript. Second, we propose **BPO-AVASR**, a **Bifocal Preference Optimization** method to improve AV-ASR models by leveraging both input-side and output-side preference. Extensive experiments demonstrate that our approach significantly improves speech recognition accuracy across various domains, outperforming previous state-of-the-art models on real-world video speech recognition.

Code — <https://github.com/espnet/espnet>.

1 Introduction

Recently, there has been a growing demand for Automatic Speech Recognition (ASR) systems to evolve into unconstrained audiovisual scenarios, such as YouTube videos, online meetings, and live broadcasts. With both audio and visual streams as input, visual information may help improve speech recognition accuracy, especially in cases where the audio is noisy or unclear. There are usually two scenarios for AudioVisual ASR (AV-ASR), focusing solely on lip motion (Afouras et al. 2022; Ma et al. 2021; Chung et al. 2017) and using full-frame visual features. In this work, we focus on the latter scenario, improving the AV-ASR performance in unconstrained real-world video scenarios. In such settings, the entire visual frame may contribute to ASR performance by providing additional cues on specific objects, background location, or context.

*Corresponding authors: Ruihua Song (rsong@ruc.edu.cn) and Shinji Watanabe (swatanab@cmu.edu).
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

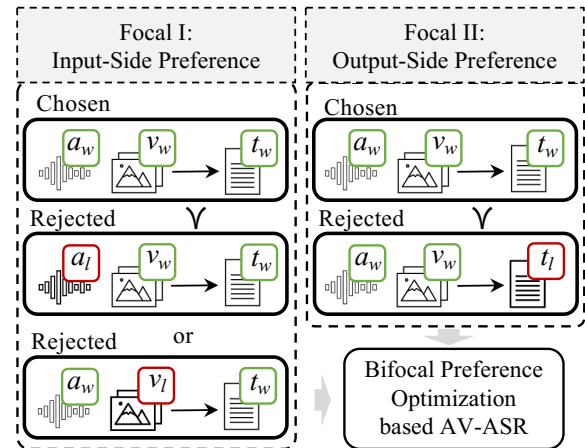


Figure 1: **Overview of our proposed framework.** We construct bifocal preferences by augmenting transcript (T), audio (A) or video (V), resulting in two types of preference pairs: input-side (A and V) and output-side (T). Through bifocal preference optimization, BPO-AVASR learns to generate transcripts that are better aligned with these preferences.

Compared to standard ASR, performing AV-ASR on real-world videos presents the following challenges:

- **Noisy Acoustic Environments.** Homophone issues pose a significant challenge for both standard ASR and AV-ASR. Real-world recordings with variable acoustic environments further severe homophone problems.
- **Spontaneous Speech Scenarios.** Real-world videos often contain spontaneous conversations, which are more variable than read speech or lectures, making accurate recognition more challenging.
- **Uncertain Use of Vision.** The diversity of video domains and unconstrained scenarios leads to the uncertain use of visual information for AV-ASR systems. While some visual cues are irrelevant, others provide crucial localized details or global context for understanding speech, making it challenging to effectively leverage visual information for improved recognition accuracy.

Most current AV-ASR works (Sanabria et al. 2018; Ghorbani et al. 2021; Caglayan et al. 2019; Seo et al. 2023) build

upon ASR models by incorporating visual features as conditions. These models typically adhere to the original ASR training objective, primarily involving the next-token prediction tasks based on the original audiovisual data. However, these supervised fine-tuning (SFT) based methods do not specifically optimize for visual inputs or address scenarios like noisy acoustic environments and spontaneous conversations, which are common in unconstrained AV-ASR tasks. Consequently, this approach may lead the AV-ASR model to neglect visual features or struggle to adapt to real-world videos that feature noisy and spontaneous speech.

To overcome the aforementioned challenges, we propose a Bifocal Preference Optimization based AV-ASR (BPO-AVASR) framework, which aims to facilitate the development of AV-ASR models with stronger speech recognition ability in unconstrained real-world video scenarios (as shown in Figure 1). Specifically, we propose a bifocal preference optimization strategy, where one bifocal point is input-side preference and the other focal point is output-side preference. Accordingly, we construct the bifocal preference dataset by simulating common errors associated with the above challenges, including homophone errors, poor performance on spontaneous speech, and inadequate use of visual cues. Subsequently, we leverage these pairs to optimize the AV-ASR model’s preferences. By emphasizing the generation of correct transcripts considering unconstrained audiovisual inputs, BPO-AVASR can better adapt to recognizing speech for real-world videos. The main contributions of this work are summarized as follows:

- We introduce **BPO-AVASR**, a novel framework with **Bifocal Preference Optimization** specifically to align audio-only ASR models with AV-ASR tasks, optimizing AV-ASR as a preference alignment problem.
- We design bifocal preference optimization, a training strategy that optimizes preference for AV-ASR tasks considering both input-side preference and output-side preference. Specifically, we construct the preference data pairs to simulate the common errors in unconstrained video AV-ASR, creating negative samples by augmenting audio, visual, and transcript information.
- Through extensive experiments, we demonstrate the effectiveness of our approach, showing a notable enhancement compared with SFT-based models. Additionally, BPO-AVASR achieves superior performance compared to previous state-of-the-art models across three datasets.

2 Related Works

2.1 Audiovisual Speech Recognition

Recent state-of-the-art ASR models have achieved impressive performance on audio-only benchmarks (Li 2021; Prabhavalkar et al. 2024). Whisper-style ASR models (Radford et al. 2023; Peng et al. 2024), in particular, leverage large-scale supervised learning to achieve robust results across various benchmarks. Building on pre-trained ASR models, several studies explore ASR to audiovisual scenarios (Gabeur et al. 2022; Peng et al. 2023a; Chung et al. 2017; Ghorbani et al. 2021). Most previous AV-ASR works focus

on lip motion (Chung et al. 2017; Afouras et al. 2022; Ma et al. 2021). Moreover, to explore unconstrained AV-ASR for real-world videos, recent works explore different adaptation methods to fine-tune or retrain ASR models with full-frame visual features (Paraskevopoulos et al. 2020; Gabeur et al. 2022; Ghorbani et al. 2021; Peng et al. 2023a; Kumar et al. 2023; Lu et al. 2024). AVFormer (Seo et al. 2023) shows the state-of-the-art performance by integrating visual information into a frozen ASR model and fine-tuning it on the large-scale audiovisual dataset HowTo100M (Miech et al. 2019). However, these models are all trained or supervised fine-tuned toward standard ASR optimization objectives, leading to modality inefficient utilization of audio and visual information, particularly in unconstrained real-world video scenarios. In contrast, this work introduces a novel approach designed to adapt ASR models to unconstrained AV-ASR using bifocal preference optimization.

2.2 Direct Preference Optimization

To align a pre-trained Large Language Model (LLM) with specific preferences of downstream tasks, Direct Preference Optimization (DPO) (Rafailov et al. 2023) is proposed to optimize the LLM with a single-stage policy learning. DPO demonstrates strong performance while eliminating the need for a separate reward model. Intuitively, given the input x and the outputs y_w and y_l ¹, DPO maximizes the difference between the reward for the chosen output $r(x, y_w)$ and that for the rejected output $r(x, y_l)$. Specifically, given a policy model to be optimized π_θ and a reference model π_{ref} , DPO formulates the reward as:

$$r(x, y) = \beta \log \frac{\pi_\theta(y|x)}{\pi_{\text{ref}}(y|x)} + Z(x), \quad (1)$$

where $Z(x)$ is a partition function, and β is the hyperparameter that controls the deviation from the reference model. Based on the Bradley-Terry model (Bradley and Terry 1952), the optimization objective of DPO becomes

$$\mathcal{L}_{\text{DPO}} = -\log \sigma \left(\beta \log \frac{\pi_\theta(y_w|x)}{\pi_{\text{ref}}(y_w|x)} - \beta \log \frac{\pi_\theta(y_l|x)}{\pi_{\text{ref}}(y_l|x)} \right), \quad (2)$$

where σ is the logistic function. Furthermore, the subsequent work of DPO (Rafailov et al. 2024) proves that DPO implicitly learns a token-level reward function, highlighting its potential of using DPO in token-level preference optimization, such as ASR and machine translation (Zeng et al. 2024).

In multimodal scenarios, recent works mainly focus on constructing multimodal preference data to support DPO training in the visual language domain (Wang et al. 2024; Zhang et al. 2024; Zhu et al. 2024). MDPO (Wang et al. 2024) designs a conditional preference optimization strategy to improve multimodal preference alignment in vision language scenarios, while SeVa (Zhu et al. 2024) constructs preference datasets via self-supervised learning. In this work, to the best of our knowledge, we first introduce DPO to AV-ASR tasks, showing a significant improvement

¹ w and l refer to win and loss respectively.

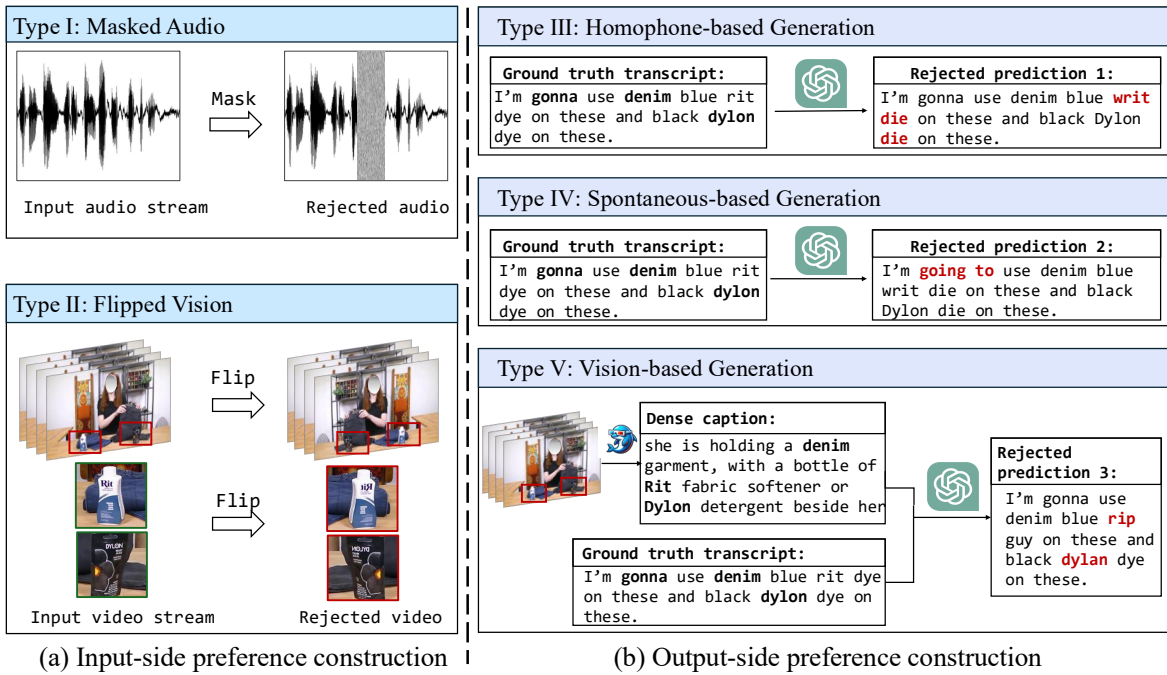


Figure 2: **The illustration of preference dataset constructing strategies:** (a) Input-side preference construction: simulates errors by manipulating input audios and videos; (b) Output-side preference construction: simulates common errors by manipulating ground truth transcripts.

over SFT-based AV-ASR models. Also, we leverage bifocal preference optimization, specifically designed for unconstrained AV-ASR tasks. Unlike DPO, which constructs negative samples by focusing solely on the output-side, BPO optimizes the AV-ASR model using both input-side preferences and output-side preferences.

3 Proposed Method

We describe our proposed BPO-AVASR framework in this section. To support bifocal alignment optimization, we design two key strategies to construct the bifocal preference dataset, constructing both input-side preferences and output-side preferences (Section 3.1). Then we outline the training procedure, including the supervised fine-tuning stage (Section 3.2) and the preference optimization stage (Section 3.3).

3.1 Creation of Bifocal Preference Dataset

Suppose we have an audiovisual speech recognition dataset $\mathcal{D} = \{(a^i, v^i, t^i)\}_{i=1}^K$, which contains K elements. Each element $\tau^i = (a^i, v^i, t^i)$ consists of a speech a^i , a video v^i , and the corresponding transcript t^i . For AV-ASR tasks, both a^i and v^i served as the inputs (denoted collectively as c^i), while the corresponding transcript t^i is the output. Thus each element can be better denoted as $\tau^i = (t^i | c^i) = (t^i | a^i, v^i)$. Given the challenges of performing AV-ASR on real-world videos (as analyzed in Section 1), we construct preference pairs by manipulating three modalities, resulting in input-side preferences (audio, video) and output-side preferences (text). By simulating common errors in AV-ASR and constructing the preference dataset accordingly, the model can

be optimized by maximizing the distinction between chosen and rejected pairs. This approach enables the model to learn to avoid these kinds of errors during inference.

Input-side Preference. To simulate insufficient use of input information, we create rejected samples by independently manipulating the audio and video inputs.

- **Masked Audio.** We mask certain audio frames to simulate situations where speech is noisy or ambiguous. In detail, we mask content words and add Gaussian noise to the corresponding audio input, as shown in Type I of Figure 2. Therefore, we obtain rejected sample $\tau_l = (t_w^i | c_l^i) = (t_w^i | a_l^i, v_w^i)$, which serves as a hard negative example to a chosen sample $\tau_w = (t_w^i | a_w^i, v_w^i)$.
- **Flipped Vision.** To mimic the situation where the AV-ASR model uses visual information inadequately, we intentionally construct samples with rejected visual inputs. As shown in Type II of Figure 2, we create the rejected sample v_l^i by flipping the video v_w^i to simulate the situation where the ASR model lacks sufficient visual information, especially for the detailed object information. Flipped images can make optical character recognition (OCR) harder and induce potential rejected responses. Therefore, we obtain the rejected sample $\tau_l = (t_w^i | c_l^i) = (t_w^i | a_w^i, v_l^i)$ for a chosen sample $\tau_w = (t_w^i | a_w^i, v_w^i)$.

As a result, we obtain the preference dataset $\tilde{\mathcal{D}}_c$.

Output-side Preference. To construct a preference dataset that includes rejected samples simulating common errors in recognized text, we use ChatGPT as a proxy

for the human generator. This approach allows us to build a cost-effective yet efficient output-side preference dataset. More specifically, for each ground truth element $\tau_w = (t_w^i | c_w^i) = (t_w^i | a_w^i, v_w^i)$, we generate a rejected text t_l^i to construct the corresponding rejected element $\tau_l = (t_l^i | c_w^i) = (t_l^i | a_w^i, v_w^i)$. To account for the different causes of text recognition errors, we employ three distinct prompts for the ChatGPT, each designed to generate rejected text based on specific error types.

- **Homophone-based Generation.** Given the ground truth text t_w^i , we prompt ChatGPT to replace the word with their homophone, e.g., “dye” is replaced by “die” (See Type III of Figure 2). We specifically focus on content words, as ASR errors frequently involve homophone and near-homophone errors in these words.
- **Spontaneous-based Generation.** For real-world videos, in addition to read speech, there is a significant amount of spontaneous speech. To address this error, we use ChatGPT to generate rejected text t_l^i from the ground truth text t_w^i . It simulates typical ASR errors in spontaneous scenarios, such as generating “going to” instead of “gonna” (See Type IV of Figure 2).
- **Vision-based Generation.** Neglecting or misusing visual information is another common error in AV-ASR. To simulate errors where visual information is ignored during recognition, we remove parts of the ground truth text that are related to the video. Specifically, we transfer the video stream v_w^i into dense and high-quality caption via ShareGPT4Video (Chen et al. 2024). Then, we use this caption and ground truth text t_w^i as a prompt to replace the words that appear in the dense caption, simulating the omission of visual cues. For example, “Rit” in the frame is re-written as “rip” (See Type V of Figure 2). Rather than using video directly, we observe that using these detailed video captions allows for better control in the construction process.

As a result, we obtain the output-side preference dataset \tilde{D}_p by rewriting the transcripts systematically.

3.2 Supervised Fine-tuning

Pre-trained ASR Model. We use the audio-only ASR model OWSM v3.1² as our backbone (Peng et al. 2024). OWSM v3.1 is an open-source pre-trained speech recognition model that achieves robust performance on standard ASR benchmarks (Panayotov et al. 2015; Bu et al. 2017; Pratap et al. 2020; Hernandez et al. 2018). It utilizes an encoder-decoder architecture (Vaswani et al. 2017), with the stack of E-Branchformer encoders (Kim et al. 2022) and Transformer decoders (Vaswani et al. 2017). Being trained on large amounts of ASR data, OWSM v3.1 has a good generalization ability for AV-ASR tasks.

Visual Encoder. We leverage a pre-trained visual encoder to extract visual tokens as conditions to the audio-only ASR

²We opt for OWSM over Whisper due to potential data contamination concerns. Whisper’s training data, sourced through web crawling, might include our test set.

model. Specifically, we use CLIP (Radford et al. 2021) to ensure that all relevant information within the video frames is effectively encoded for AV-ASR. We sample M frames from video v^i uniformly, then use the CLIP to extract visual tokens from these frames. A visual projection layer is subsequently applied to map visual tokens into speech space.

Fine-tuning Objective Given M visual tokens and N speech tokens, we concatenate them along the sequence dimension, resulting in a combined sequence of length $N+M$. Then, we feed them into the ASR model and fine-tune the model π_{ref} in Equation 1 with attention loss \mathcal{L}_{ATT} and the CTC loss \mathcal{L}_{CTC} :

$$\begin{aligned} \mathcal{L}_{\text{SFT}}(\pi_{\text{ref}}, \mathcal{D}) &= \mathcal{L}_{\text{ATT}}(\pi_{\text{ref}}, \mathcal{D}) + \alpha \cdot \mathcal{L}_{\text{CTC}}(\pi_{\text{ref}}, \mathcal{D}), \\ \mathcal{L}_{\text{ATT}}(\pi_{\text{ref}}, \mathcal{D}) &= - \sum_u \ln P_{\text{ATT}}(t_u^* | a, v, t_{1:u-1}^*), \\ \mathcal{L}_{\text{CTC}}(\pi_{\text{ref}}, \mathcal{D}) &= - \ln P_{\text{CTC}}(t^* | a, v), \end{aligned} \quad (3)$$

where $t_{1:u-1}^*$ is the preceding tokens of the ground truth character sequence y^* .

3.3 Bifocal Preference Optimization

To align the fine-tuned ASR model π_{ref} (Equation 3) with the AV-ASR task, we use input-side preference dataset \tilde{D}_c and the output-side preference dataset \tilde{D}_p (as described in Section 3.1) to train the BPO-AVASR model π_θ through bifocal preference optimization.

Analogous to the Equation 2, the bifocal preference optimization objective for AV-ASR is formulated as

$$\begin{aligned} \mathcal{L}_{\text{BPO}}(\pi_\theta; \pi_{\text{ref}}, \tilde{D}) &= \mathcal{L}_{\text{input}}(\pi_\theta; \pi_{\text{ref}}, \tilde{D}_c) + \mathcal{L}_{\text{output}}(\pi_\theta; \pi_{\text{ref}}, \tilde{D}_p), \\ \mathcal{L}_{\text{input}}(\pi_\theta; \pi_{\text{ref}}, \tilde{D}_c) &= - \log \sigma \left(\beta \log \frac{\pi_\theta(t_w | c_w)}{\pi_{\text{ref}}(t_w | c_w)} - \beta \log \frac{\pi_\theta(t_w | c_l)}{\pi_{\text{ref}}(t_w | c_l)} \right), \\ \mathcal{L}_{\text{output}}(\pi_\theta; \pi_{\text{ref}}, \tilde{D}_p) &= - \log \sigma \left(\beta \log \frac{\pi_\theta(t_w | c_w)}{\pi_{\text{ref}}(t_w | c_w)} - \beta \log \frac{\pi_\theta(t_l | c_w)}{\pi_{\text{ref}}(t_l | c_w)} \right), \end{aligned} \quad (4)$$

where $\mathcal{L}_{\text{input}}$ and $\mathcal{L}_{\text{output}}$ are preference optimization loss for input-side preference and output-side preference respectively. $(t_l | c_w)$ and $(t_w | c_l)$ serve as hard negatives to $(t_w | c_w)$, sampled from the preference datasets \tilde{D}_p and \tilde{D}_c respectively. Through this approach, we train the BPO-AVASR model π_θ to avoid generating the above errors in the application, thereby improving the ability of the AV-ASR model to take good advantage of complex audiovisual information, especially for unconstrained real-world videos.

4 Experimental Settings

4.1 Implementation Details

As described in Section 3.2, we use OWSM v3.1 as our speech recognition backbone and develop two versions of BPO-AVASR: BPO-AVASR small and BPO-AVASR medium respectively. Specifically, BPO-AVASR-small consists of 9 E-Branchformer (Kim et al. 2022) encoder blocks and 9 Transformer decoder blocks. The hidden size is 768. BPO-AVASR-medium comprises 18 E-Branchformer encoder blocks and 18 Transformer decoder blocks with a hidden size of 1024. Following AVFormer (Seo et al. 2023),










When vision provides content cues:		
		
GT: We use the cherry tomatoes . w/o BPO: We use the chair to mirror . BPO-AVASR: We use the cherry tomatoes .	GT: That's how you tune a ukulele . w/o BPO: That's how you tuna ucal aile . BPO-AVASR: That's how you tune a ukulele .	GT: Salad and I really like this dole . w/o BPO: Spell it and I really like this dull . BPO-AVASR: Salad and I really like this dole .
When vision offers context clues:		
		
GT: how well it's chopping off . w/o BPO: how well it's chap enough . BPO-AVASR: how well it's chopping off .	GT: We slowly ferried over in a small boat. w/o BPO: We slowly varied over in a small boat. BPO-AVASR: We slowly ferried over in a small boat.	GT: A seething cosmic maelstrom. w/o BPO: A seeding cosmic maelstrom. BPO-AVASR: A seething cosmic maelstrom.
When speech is under spontaneous scenarios :		
		
GT: One five . w/o BPO: Voa Voa . BPO-AVASR: One five .	GT: Cuz we take our game serious . w/o BPO: Because we take our game **** . BPO-AVASR: Cuz we take our game serious .	GT: I started drinking before I like . w/o BPO: I start drinking **** **** . BPO-AVASR: I started drinking before I like .

Figure 3: **Qualitative Results.** We show the ground truth text (GT), and predictions from the OWSM-visual small(w/o DPO) and BPO-AVASR small. We show the enhanced performance in three scenarios: when vision provides content cues (top), when vision offers context clues (middle), and when speech is under spontaneous scenarios (bottom). Errors in the predicted words compared to the GT are highlighted in red. Faces are blurred for privacy.

we utilize CLIP-Large (Radford et al. 2021) as the visual encoder. For each video segment, we sample 4 frames uniformly (i.e. $M = 4$ as described in Section 3.2). During the SFT stage, we fine-tune OWSM v3.1 for 10 epochs with a batch size of 64. We set α to 0.3 in Equation 3. In the BPO fine-tuning stage, we set β to 0.1 in Equation 4. We conduct the training using 1 V100 GPU, with a batch size of 512 and a learning rate of $2e-6$. We use word error rate (WER) as the evaluation metric for all experiments, with lower values indicating better performance.

4.2 Datasets

In the supervised fine-tuning stage, we apply How2 as the fine-tuning dataset. In the bifocal preference optimization stage, we construct the preference dataset based on the How2. BPO-AVASR models are evaluated on three datasets, How2, VisSpeech, and Ego4D.

- **How2** (Sanabria et al. 2018) is an instructional video dataset designed for multimodal understanding. Following AVFormer, we use the 300-hour version of How2. These videos are segmented into short clips, averaging 5.8 seconds each, with a user-uploaded transcript.
- **VisSpeech** (Gabeur et al. 2022) is an AV-ASR benchmark sampled from the HowTo100M dataset (Miech et al. 2019). It contains 508 video clips with manually

annotated transcripts. VisSpeech uses a video-text similarity model to ensure high audiovisual correspondence.

- **Ego4D** (Grauman et al. 2022) is an egocentric daily-life activity video dataset. We use the audiovisual diarization benchmark and evaluate our model on the validation set with ground truth annotations. Videos are segmented into shorter clips for analysis. Unlike other datasets, Ego4D contains noisier and more spontaneous videos across different domains, increasing the difficulty of AV-ASR.

4.3 Baselines

We compare BPO-AVASR models to strong baselines, including robust audio-only ASR models (see the upper section of Table 1) and state-of-the-art AV-ASR models (see the middle section of Table 1).

- **Audio-only ASR Models.** We compare BPO-AVASR series with (i) **BEST-RQ**, a robust ASR model pre-trained on LibriLight (Kahn et al. 2020), LibriSpeech (Panayotov et al. 2015) and HowTo100M (Miech et al. 2019) datasets ($\sim 190K$ hours data in total); (ii) **OWSM-ft small**, the OWSM v3.1 small model pre-trained on public speech datasets and fine-tuned on the How2 using audio-only data ($\sim 180K$ hours data in total).
- **AV-ASR Models.** We compare the BPO-AVASR series with the following AV-ASR baselines: **How2**

Model	Modality	How2	VisSpeech	Ego4D
OWSM-ft small (Peng et al. 2023b)	A	10.8	16.6	70.6
BEST-RQ (Chiu et al. 2022)	A	15.3	16.7	68.3
How2 base (Sanabria et al. 2018)	A+V	18.0	–	–
VAT (Caglayan et al. 2019)	A+V	18.0	–	–
MultiRes (Paraskevopoulos et al. 2020)	A+V	20.5	–	–
LLD (Ghorbani et al. 2021)	A+V	16.7	–	–
AVATAR (Gabeur et al. 2022)	A+V	15.6	43.4	–
SynesLM (Lu et al. 2024)	A+V	15.7	39.4	–
AVATAR [†] (Gabeur et al. 2022)	A+V	9.1	35.7	92.0
AVFormer [†] (Seo et al. 2023)	A+V	13.6	16.4	64.6
BPO-AVASR small	A+V	9.3 (31.6%)	<u>15.6</u> (4.9%)	<u>59.2</u> (8.4%)
BPO-AVASR medium	A+V	<u>9.2</u> (32.4%)	14.5 (11.6%)	56.5 (12.5%)

Table 1: **Comparison to baseline methods across different datasets.** Results are reported as WER%. For the VisSpeech and Ego4D datasets, models are evaluated without any fine-tuning. [†] denotes these models are trained on a much larger additional audiovisual dataset, HowTo100M. “A” and “V” refer to audio and visual, respectively. Bolded values indicate the best results, while underlined values indicate the second-best results. The numbers in parentheses indicate the improvement in model performance (relative reduction in WER) compared to the state-of-the-art model considering three test sets, AVFormer[†].

base (Sanabria et al. 2018) re-trains an ASR model by learning a video-specific bias from speech features. **VAT** (Caglayan et al. 2019) integrates different visual features into the ASR model with adaptive training. **MultiRes** (Paraskevopoulos et al. 2020) fuses video features with an additional cross-modal multi-head attention layer. **LLD** (Ghorbani et al. 2021) uses a deliberation model to leverage video and text representations extracted from a self-supervised text-video embedding model. **AVATAR** (Gabeur et al. 2022) is an encoder-decoder based AV-ASR model which fuses visual and speech information through a multimodal encoder. **AVFormer** (Seo et al. 2023) injects visual information into the frozen ASR model, BEST-RQ, with lightweight trainable adaptors, which shows good performance by training on 131k hours of audiovisual dataset. **SynesLM** (Lu et al. 2024) is a unified speech language model capable of performing audiovisual speech recognition. It explores decoder-only architecture within a multitask learning objective. **OWSM-visual small** is the OWSM v3.1 small model fine-tuned on the How2 dataset that incorporates visual information (the first row of Table 2).

5 Experimental Results

5.1 Comparison with SOTA Models

We compare BPO-AVASR with baseline models on three test sets in Table 1. The BPO-AVASR models outperform all baselines trained on the same audiovisual dataset (How2 base, VAT, MultiRes, LLD, SynesLM, and AVATAR), particularly on the noisy and spontaneous Ego4D dataset, highlighting the effectiveness of bifocal preference optimization in adapting ASR to unconstrained AV-ASR tasks. Furthermore, BPO-AVASR medium achieves better results than BPO-AVASR small, demonstrating the benefit of using larger models. While AVATAR[†] achieves the best per-

formance on How2, it shows poor generalization ability on the out-of-domain dataset Ego4D. In contrast, the BPO-AVASR series shows consistently better results across different domain datasets, demonstrating the robustness and generalization capability of BPO-AVASR. Compared with the previous state-of-the-art work AVFormer, the BPO-AVASR series shows significant improvement, especially on How2 (31.6% and 32.4%, respectively). Notably, both AVFormer[†] and BPO-AVASR are based on pre-trained ASR models with over 100K hours of training data. However, while AVFormer[†] is further fine-tuned on 131k hours of audiovisual data from HowTo100M, BPO-AVASR achieves superior results using only 300 hours of audiovisual data.

5.2 Qualitative Analysis

Examples in Figure 3 illustrate the effectiveness of BPO-AVASR. Qualitative results show that OWSM-visual (w/o BPO) still struggles to obtain accurate transcripts in real-world video scenarios. In contrast, qualitative examples demonstrate how BPO-AVASR improves ASR performance. As shown in Figure 3, BPO-AVASR is effective in three scenarios: when vision provides content cues, when vision offers context clues, and when speech is under spontaneous scenarios. For instance, visual content helps the model recognize objects directly in the video, such as brand names (row 1 column 3) and object names (row 1 column 1, row 1 column 2). Additionally, by using visual information as contextual clues, the model can correctly distinguish between similarly pronounced words (row 2). Furthermore, by constructing rejected samples considering spontaneous words, BPO-AVASR improves the accuracy of recognizing filler words in spontaneous speech scenarios (row 3).

5.3 Ablation Studies

Analysis of Rejected Samples. To evaluate the effects of different preference construction strategies, we create pref-

Rejected Data Construction Strategy		How2	VisSpeech	Ego4D
N/A (OWSM-visual small)	–	10.5	15.8	59.9
Input-side preference	Masked audio	9.4	15.7	59.2
	Random cropped vision	12.1	15.7	61.5
	Flipped vision	9.7	15.5	59.3
Output-side preference	Rule-based replacement	10.4	16.0	62.4
	Homophone-based generation	9.5	15.8	59.8
	Spontaneous-based generation	10.8	16.1	59.0
	Vision-based generation	9.7	15.6	59.3
Mixture (BPO-AVASR small)	–	9.3 (11.4%)	15.6 (1.3%)	59.2 (1.7%)

Table 2: **Comparison of preference data constructing strategies.** Results are reported as WER%. N/A indicates that preference data and preference optimization are not used. The digits in parentheses indicate the improvement in model performance (relative reduction in WER) compared to the model without preference optimization, i.e., OWSM-visual small.

Method	WER%
AVATAR (Gabeur et al. 2022)	55.3
AVFormer (Seo et al. 2023)	55.2
OWSM-visual small	52.3
BPO-AVASR small	50.0 (4.4%)

Table 3: **Generalization to Ego4D dataset.** All models are fine-tuned on the Ego4D training set. The digits in parentheses indicate the performance improvement (relative reduction in WER) compared to the OWSM-visual small.

erence datasets with eight different strategies and then use them to optimize the BPO-AVASR small model. From the results in Table 2, we have the following observations:

- **Overall.** The bifocal preference dataset constructed by mixing all strategies leads to the best and most balanced performance across all test sets. Moreover, when comparing BPO-AVASR small with OWSM-visual, the significant improvement (especially the 11.4% improvement on How2) highlights the effectiveness of our proposed bifocal preference optimization over SFT optimization.
- **Importance of Hard Negatives.** Constructing hard negative samples is crucial for providing effective preference optimization signals. For input-side preference, frame flipping shows better results than random cropping. We attribute this to the fact that random cropping often fails to remove useful information due to the redundancy of visual data in unconstrained AV-ASR, making it difficult to create truly hard negatives. For output-side preference construction, we compare the rule-based construction strategy with the LLM-based strategy described in Section 3.1. Specifically, the rule-based strategy uses a homophone dictionary to replace content words in the ground truth transcripts. The LLM-based preference construction shows greater improvement, due to the LLM’s strong text understanding and generation capabilities, which are more effective in constructing hard negatives.
- **Effectiveness of Different Strategies.** Preference opti-

mization strategies show different improvements on three test sets. Masking audio improves speech recognition performance across all three datasets. The Ego4D dataset benefits significantly from the strategy that addresses spontaneous speech (from 59.9 to 59.0).

Generalization to Other Datasets. To demonstrate the generalization ability of BPO-AVASR, we further construct a bifocal preference dataset using the Ego4D training set (Grauman et al. 2022), following the same strategies described in Section 3.1. We then fine-tune BPO-AVASR small on this constructed preference dataset. As shown in Table 3, BPO-AVASR outperforms all previous works, demonstrating the effectiveness of bifocal preference optimization on real-world videos, including instructional (How2) and egocentric (Ego4D) videos. Moreover, fine-tuning on Ego4D with preference optimization results in a significant improvement in WER compared to zero-shot testing on Ego4D (Table 1). This highlights that constructing preferences within the same domain further enhances the model’s performance through preference optimization.

6 Conclusion

In this paper, we first formulate the AV-ASR task as a preference optimization problem. Accordingly, we develop BPO-AVASR, an AV-ASR system optimized by bifocal preference optimization to improve speech recognition accuracy for real-world videos. First, we introduce a simple yet effective method to create preference data by simulating common errors related to different modalities in AV-ASR. Second, we propose a bifocal preference optimization strategy to optimize AV-ASR models by emphasizing the distinction between correct and incorrect answers. BPO-AVASR outperforms previous state-of-the-art models, demonstrating the effectiveness of using preference optimization to align the audio-only ASR model to real-world video scenarios. For future work, we plan to build a high-quality open-domain AV-ASR dataset to facilitate future research.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (No. 62276268) and the Outstanding Innovative Talents Cultivation Funded Programs 2022 of Renmin University of China. Experiments of this work used the Bridges2 system at PSC and Delta system at NCSA through allocations CIS210014 and IRI120008P from the Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support (ACCESS) program.

References

- Afouras, T.; et al. 2022. Deep Audio-Visual Speech Recognition. *IEEE Trans. Pattern Anal.*, 44(12): 8717–8727.
- Bradley, R. A.; and Terry, M. E. 1952. Rank analysis of incomplete block designs: I. The method of paired comparisons. *Biometrika*, 39(3/4): 324–345.
- Bu, H.; et al. 2017. AISHELL-1: An open-source Mandarin speech corpus and a speech recognition baseline. In *O-COCOSDA*, 1–5.
- Caglayan, O.; et al. 2019. Multimodal Grounding for Sequence-to-sequence Speech Recognition. In *ICASSP*, 8648–8652.
- Chen, L.; Wei, X.; Li, J.; Dong, X.; Zhang, P.; Zang, Y.; Chen, Z.; Duan, H.; Lin, B.; Tang, Z.; et al. 2024. ShareGPT4Video: Improving Video Understanding and Generation with Better Captions. *arXiv preprint arXiv:2406.04325*.
- Chiu, C.; Qin, J.; Zhang, Y.; Yu, J.; and Wu, Y. 2022. Self-supervised learning with random-projection quantizer for speech recognition. In Chaudhuri, K.; Jegelka, S.; Song, L.; Szepesvári, C.; Niu, G.; and Sabato, S., eds., *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, 3915–3924. PMLR.
- Chung, J. S.; et al. 2017. Lip Reading Sentences in the Wild. In *CVPR*, 3444–3453.
- Gabeur, V.; et al. 2022. AVATAR: Unconstrained Audiovisual Speech Recognition. In *Interspeech*, 2818–2822.
- Ghorbani, S.; et al. 2021. Listen, Look and Deliberate: Visual Context-Aware Speech Recognition Using Pre-Trained Text-Video Representations. In *IEEE SLT Workshop*, 621–628.
- Grauman, K.; et al. 2022. Ego4D: Around the World in 3, 000 Hours of Egocentric Video. In *CVPR*, 18973–18990.
- Hernandez, F.; Nguyen, V.; Ghannay, S.; Tomashenko, N. A.; and Estève, Y. 2018. TED-LIUM 3: Twice as Much Data and Corpus Repartition for Experiments on Speaker Adaptation. In Karpov, A.; Jokisch, O.; and Potapova, R., eds., *Speech and Computer - 20th International Conference, SPECOM 2018, Leipzig, Germany, September 18-22, 2018, Proceedings*, volume 11096 of *Lecture Notes in Computer Science*, 198–208. Springer.
- Kahn, J.; Rivièrè, M.; Zheng, W.; Kharitonov, E.; Xu, Q.; Mazaré, P.; Karadayi, J.; Liptchinsky, V.; Collobert, R.; Fuegen, C.; Likhomanenko, T.; Synnaeve, G.; Joulin, A.; Mohamed, A.; and Dupoux, E. 2020. Libri-Light: A Benchmark for ASR with Limited or No Supervision. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP 2020, Barcelona, Spain, May 4-8, 2020*, 7669–7673. IEEE.
- Kim, K.; et al. 2022. E-Branchformer: Branchformer with Enhanced Merging for Speech Recognition. In *IEEE SLT Workshop*, 84–91.
- Kumar, V. B.; et al. 2023. Visual Information Matters for ASR Error Correction. In *ICASSP*, 1–5.
- Li, J. 2021. Recent Advances in End-to-End Automatic Speech Recognition. *arXiv:2111.01690*.
- Lu, Y.; Song, J.; Chang, X.; Bian, H.; Maiti, S.; and Watanabe, S. 2024. SynesLM: A Unified Approach for Audio-visual Speech Recognition and Translation via Language Model and Synthetic Data. *arXiv:2408.00624*.
- Ma, P.; et al. 2021. End-To-End Audio-Visual Speech Recognition with Conformers. In *ICASSP*, 7613–7617.
- Miech, A.; et al. 2019. HowTo100M: Learning a Text-Video Embedding by Watching Hundred Million Narrated Video Clips. In *2019 IEEE/CVF ICCV*, 2630–2640.
- Panayotov, V.; et al. 2015. Librispeech: An ASR corpus based on public domain audio books. In *ICASSP*, 5206–5210.
- Paraskevopoulos, G.; et al. 2020. Multimodal and Multiresolution Speech Recognition with Transformers. In *ACL*, 2381–2387.
- Peng, P.; et al. 2023a. Prompting the Hidden Talent of Web-Scale Speech Models for Zero-Shot Task Generalization. In *Interspeech*, 396–400.
- Peng, Y.; et al. 2023b. Reproducing Whisper-Style Training Using An Open-Source Toolkit And Publicly Available Data. In *ASRU*, 1–8.
- Peng, Y.; et al. 2024. OWSM v3.1: Better and Faster Open Whisper-Style Speech Models based on E-Branchformer. *arXiv:2401.16658*.
- Prabhavalkar, R.; et al. 2024. End-to-End Speech Recognition: A Survey. *IEEE ACM Trans. Audio Speech Lang. Process.*, 32: 325–351.
- Pratap, V.; Xu, Q.; Sriram, A.; Synnaeve, G.; and Collobert, R. 2020. MLS: A Large-Scale Multilingual Dataset for Speech Research. In Meng, H.; Xu, B.; and Zheng, T. F., eds., *21st Annual Conference of the International Speech Communication Association, Interspeech 2020, Virtual Event, Shanghai, China, October 25-29, 2020*, 2757–2761. ISCA.
- Radford, A.; et al. 2021. Learning Transferable Visual Models From Natural Language Supervision. In *ICML*, volume 139, 8748–8763.
- Radford, A.; et al. 2023. Robust Speech Recognition via Large-Scale Weak Supervision. In *ICML*, volume 202, 28492–28518.
- Rafailov, R.; Hejna, J.; Park, R.; and Finn, C. 2024. From r to Q^* : Your Language Model is Secretly a Q-Function. *CoRR*, abs/2404.12358.
- Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; and Finn, C. 2023. Direct Preference Optimization:

Your Language Model is Secretly a Reward Model. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

Sanabria, R.; et al. 2018. How2: A Large-scale Dataset for Multimodal Language Understanding. [abs/1811.00347](https://arxiv.org/abs/1811.00347).

Seo, P. H.; et al. 2023. AVFormer: Injecting Vision into Frozen Speech Models for Zero-Shot AV-ASR. In *CVPR*, 22922–22931.

Vaswani, A.; et al. 2017. Attention is All you Need. In *NeurIPS*, 5998–6008.

Wang, F.; Zhou, W.; Huang, J. Y.; Xu, N.; Zhang, S.; Poon, H.; and Chen, M. 2024. mDPO: Conditional Preference Optimization for Multimodal Large Language Models. *CoRR*, [abs/2406.11839](https://arxiv.org/abs/2406.11839).

Zeng, J.; Meng, F.; Yin, Y.; and Zhou, J. 2024. Teaching Large Language Models to Translate with Comparison. In Wooldridge, M. J.; Dy, J. G.; and Natarajan, S., eds., *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024, Vancouver, Canada*, 19488–19496. AAAI Press.

Zhang, R.; Gui, L.; Sun, Z.; Feng, Y.; Xu, K.; Zhang, Y.; Fu, D.; Li, C.; Hauptmann, A.; Bisk, Y.; and Yang, Y. 2024. Direct Preference Optimization of Video Large Multimodal Models from Language Model Reward. [arXiv:2404.01258](https://arxiv.org/abs/2404.01258).

Zhu, K.; Zhao, L.; Ge, Z.; and Zhang, X. 2024. Self-Supervised Visual Preference Alignment. *CoRR*, [abs/2404.10501](https://arxiv.org/abs/2404.10501).