

Retrieval-Augmented Egocentric Video Captioning

Jilan Xu^{1,5} Yifei Huang^{2,5} Junlin Hou¹ Guo Chen^{3,5} Yuejie Zhang^{1*} Rui Feng^{1*} Weidi Xie^{4,5}
¹Fudan University ²The University of Tokyo ³Nanjing University
⁴CMIC, Shanghai Jiao Tong University ⁵Shanghai AI Laboratory

Abstract

Understanding human actions from videos of first-person view poses significant challenges. Most prior approaches explore representation learning on egocentric videos only, while overlooking the potential benefit of exploiting existing large-scale third-person videos. In this paper, (1) we develop **EgoInstructor**, a retrieval-augmented multimodal captioning model that automatically retrieves semantically relevant third-person instructional videos to enhance the video captioning of egocentric videos, (2) for training the cross-view retrieval module, we devise an automatic pipeline to discover ego-exo video pairs from distinct large-scale egocentric and exocentric datasets, (3) we train the cross-view retrieval module with a novel **EgoExoNCE** loss that pulls egocentric and exocentric video features closer, by aligning them to shared text features that describe similar actions, (4) through extensive experiments, our cross-view retrieval module demonstrates superior performance across seven benchmarks. Regarding egocentric video captioning, EgoInstructor exhibits significant improvements by leveraging third-person videos as references.

1. Introduction

Recently, egocentric video understanding is getting increasing attention in the vision community, for example, on action recognition [5, 18, 40, 59], detection [66, 77], video-text retrieval [43, 78], grounding [14, 43, 54] and gaze estimation [35, 39]. As egocentric videos are normally recorded from a “first-person” view, reflecting activities engaging with the environment, it plays a paramount role in deploying vision models into real-world scenarios, such as robotics [33] and augmented reality [47].

As preliminary baseline solutions for egocentric video understanding, early works have directly trained models on the egocentric videos [18, 39, 40, 73], however, such approaches are limited by dataset scale and have overlooked the benefit of exploiting the large corpus of third-person (ex-

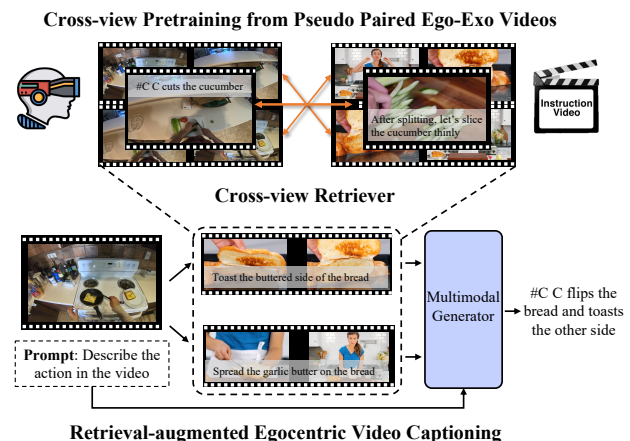


Figure 1. EgoInstructor is a retrieval-augmented multimodal captioning model that retrieves relevant exocentric videos as references to generate the caption of egocentric videos. The cross-view retrieval ability is enabled by training on automatically constructed large-scale pseudo paired ego-exo videos.

ocentric) videos. To leverage the information in exocentric videos, some researchers have tried to directly transfer models learnt from third-person videos [31, 41, 42, 60]. Due to the discrepancy in terms of recording perspective, camera motion, video continuity, and domain shift, the representations learnt from exocentric videos may not be optimal for egocentric videos. Another line of research involves the joint learning of egocentric and exocentric video representations [3, 24, 59, 71, 72]. However, training these models would require time-synchronised ego- and exo-video pairs collected in the same environment [19, 34, 58, 59].

In this paper, we take inspiration from the fact that humans are born with the ability of learning by observation, *i.e.*, the ability of observing third-view demonstrations and subsequently generalising towards egocentric perspective in different environment [25, 51, 56]. For a more specific example, people often seek answers to “How To” questions and watch third-person instructional videos before attempting tasks they are unfamiliar with, *e.g.* how to make a cheesebeef sandwich. When engaging in the task, people also recall key procedures, visual demonstrations and

*Corresponding author

detailed instructions from instructional videos. We explore retrieval-augmented egocentric video captioning, an alternative way for transferring knowledge from exocentric videos to enhance egocentric video captioning.

Specifically, as shown in Fig. 1, our goal is to learn a unified ego-exo representation space that allows for explicitly retrieving relevant third-view videos available on the Internet (*e.g.* from HowTo100M [46]), to assist video captioning in the egocentric perspective. This is achieved by first introducing an automated pipeline, generating pseudo ego-exo video pairs at scale by aligning the captions that describe similar actions. Subsequently, the paired videos are employed to train a cross-view retrieval module via a novel EgoExoNCE loss, which aligns both egocentric and exocentric video features with shared text features describing similar action semantics. With this cross-view retrieval capability, we develop a retrieval-augmented egocentric video captioning model by leveraging the retrieved exocentric videos as references for caption generation.

As a consequence, we evaluate the cross-view retrieval module on seven benchmarks, *e.g.*, EK100 Multi-instance Retrieval [18], Ego4d Multiple Choice Questions [21], YouCook2 video-text retrieval [79] and CharadesEgo cross-view video retrieval [59], the model demonstrates superior performance on all datasets. Regarding egocentric video captioning on Ego4d [21], we show the benefits of leveraging semantically related exocentric instructional videos.

2. Related Work

Egocentric video understanding. The unique viewpoint of egocentric videos poses a broad range of challenges in human activity analysis [11], including action recognition and detection [18, 31, 42, 66, 77], human pose estimation [29, 49], gaze estimation [26, 39, 73], and captioning [30, 48]. Recently, the introduction of the Ego4d [21] dataset has prompted a series of studies for representation learning in egocentric videos [5, 43, 76, 78]. For instance, EgoVLP [43] performed contrastive learning over paired egocentric videos and narrations. LaViLa [78] learnt egocentric video-language representations by using pseudo-labelled egocentric videos. Despite the progress in egocentric video understanding, these models lack enough generalisation ability on exocentric videos. In contrast, our cross-view retrieval module is designed to learn a unified ego-exo video-language representation space.

Ego-exo video understanding. Ego-exo video understanding models involve both egocentric and exocentric videos to solve a variety of vision tasks [2, 3, 20, 24, 42, 59, 69, 71]. The majority of these approaches explore shared visual clues from synchronous egocentric and exocentric data recorded simultaneously [20, 58, 59], which poses additional difficulty and cost in data collection. Subsequently,

versatile methods are developed without relying on paired data. EgoExo [42] transferred exocentric pre-trained model to egocentric videos by mining egocentric pseudo labels. AE2 [72] temporally aligned the representations of the unpaired egocentric and exocentric videos via Dynamic Time Warping [9]. Sum-L [67] created cross-view pseudo pairs from egocentric dataset [18, 59] and exocentric dataset [28] using the ground-truth semantic category labels. In contrast, our method learns joint ego-exo representations on unpaired egocentric [21] and exocentric [46] datasets by leveraging captions/narrations and moves beyond vision-only [42, 67, 72] tasks to video-language understanding.

Retrieval-augmented models. Retrieval-augmented language models in the NLP community aim at retrieving external knowledge to enhance performance in various NLP tasks [4, 12, 22, 36, 53]. Recent progresses in vision-language domain also retrieve semantically related samples to improve the performance, including image recognition [27, 45], captioning [55, 57], visual question answering [15, 74], image generation [16, 75] and vision-language pre-training [70]. While most methods primarily focus on image domain and employ a pre-trained CLIP [52] model for cross-modal retrieval, such methods are sub-optimal for video data, especially in the case between egocentric and exocentric views. In contrast, our approach involves training a cross-view retrieval module tailored for retrieving exocentric videos for egocentric video captioning.

3. Methodology

The overall architecture of our proposed EgoInstructor is shown in Fig. 2. Given an egocentric video v^{ego} , our model first retrieves K semantically relevant third-view instructional videos and associated texts with the cross-view retrieval, *i.e.*, $\{x_1^{\text{exo}}, \dots, x_K^{\text{exo}}\} = \Phi_{\text{retrieve}}(v^{\text{ego}})$, where $x_i^{\text{exo}} = (v_i^{\text{exo}}, t_i^{\text{exo}})$ includes both video and associated text narrations. Then, the multimodal captioning model takes both the ego-video and retrieved samples to generate the text description, $t^{\text{ego}} = \Psi_{\text{generate}}(v^{\text{ego}}, \{x_1^{\text{exo}}, \dots, x_K^{\text{exo}}\})$. In the following sections, we start by describing the procedure for training the cross-view retrieval module with unpaired data (Sec. 3.1), then followed by the full retrieval-augmented multimodal captioning model (Sec. 3.2).

3.1. Cross-view Visual Representation Alignment

Our goal in this step is to train a cross-view retrieval module that associates the ego- and exo-videos sharing the same semantics, using videos with only caption or narrations, *i.e.*, without manual annotation for pairing ego-exo videos. Specifically, we adopt **egocentric** videos from Ego4D [21] with 4M manually labelled captions, and **exocentric** videos from HowTo100M [46], containing detailed instructions for performing daily procedural activities.




Figure 2. An overview of our EgoInstructor. Given an egocentric video, we first retrieve relevant exocentric instructional videos using a frozen cross-view retrieval module pre-trained on pseudo ego-exo pairs generated automatically. The multimodal captioning model (consisting of a visual encoder, a perceiver resampler, and a text decoder.) takes the egocentric video and the retrieved videos and captions as references, and generates the caption of the ego-video.

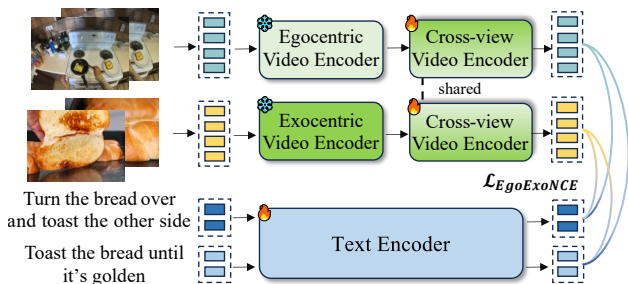


Figure 3. Our cross-view retrieval module trained via EgoExoNCE loss. We keep the egocentric and exocentric video encoders frozen and train the cross-view video encoder and text encoder.

3.1.1 Architecture Detail

Egocentric visual-language representation. For egocentric video clips, we extract representations via a CLIP-style dual encoder architecture, adopting a frozen VideoMAE [62] as the visual encoder, $z_i^{\text{ego}} = f^{\text{ego}}(v_i^{\text{ego}}) \in \mathbb{R}^{T \times d}$ with T frames; and a BERT model [32] as the text encoder $u_i^{\text{ego}} = f^{\text{text}}(t_i^{\text{ego}}) \in \mathbb{R}^d$. Both visual and text encoders have been pre-trained on Ego4d video-text pairs.

Exocentric visual-language representation. For exocentric video clip and its instruction from HowTo100M dataset, the visual representation is computed using a frozen InternVideo [68] video encoder. This encoder is a Uniformer [38] model pre-trained on large corpus of exocentric video-text pairs [68], $z_i^{\text{exo}} = f^{\text{exo}}(v_i^{\text{exo}}) \in \mathbb{R}^{T \times d}$. The text features are also computed with the same BERT encoder used for ego-text processing, $u_i^{\text{exo}} = f^{\text{text}}(t_i^{\text{exo}}) \in \mathbb{R}^d$.

Cross-view video encoder. With computed visual-language representations, we train a cross-view video encoder $\phi(\cdot)$ that takes either egocentric or exocentric video

feature as input, and maps them into an embedding space, such that the videos with similar semantics are close together regardless of the shooting perspective. We obtain $z_i^{\text{ego}} = \phi(z_i^{\text{ego}}) \in \mathbb{R}^d$ and $z_i^{\text{exo}} = \phi(z_i^{\text{exo}}) \in \mathbb{R}^d$ with the temporal dimension being aggregated with average pooling. The cross-view retrieval module is shown in Fig. 3.

3.1.2 Automatic Ego-Exo Pair Generation

For training the described cross-view retrieval module, we propose a scalable approach that exploits natural languages as a bridge to pair videos from different viewpoints, *i.e.*, grouping egocentric and exocentric videos if their corresponding language narrations and instructions are highly relevant. In practice, unlike egocentric videos from Ego4D that have manually annotated captions, instructions for HowTo100M videos are normally acquired from ASR transcripts [8]. This incurs two issues: (i) the transcript may not be of descriptive style, containing redundancy or ambiguity, as depicted in Fig. 4 (left); (ii) the transcript may not be well aligned with the visual signals. There might be greetings from the speaker or inaccuracies in timing, *e.g.*, describing the action after performing it, as discussed in [23].

Caption refinement via large language model. We transform the ASR transcript into similar style as those of Ego4D captions with a large language model (LLM). Specifically, given an instructional video from HowTo100M, *i.e.*, $x^{\text{exo}} = \{(v_j, t_j, s_j, e_j)\}_{j=1}^M$, where each narration t_j is annotated with a start time s_j and end time e_j , we refine the ASR transcript by prompting the LLM to only capture the main actions in the transcript and output descriptive captions:

$$[t'_1, \dots, t'_M] = \text{LLM}([\text{Prompt}, (t_1, s_1, e_1), \dots, (t_M, s_M, e_M)]),$$

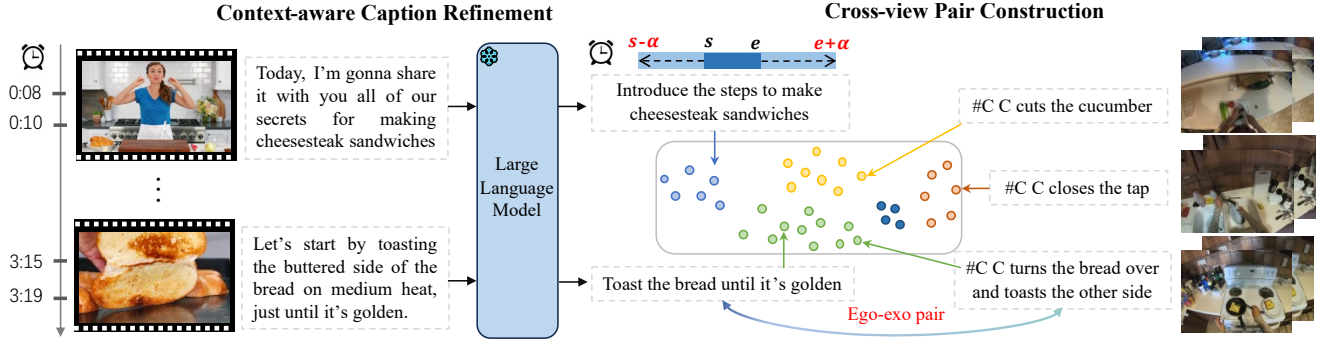


Figure 4. An illustration of context-aware caption refinement (left) and cross-view pair construction (right). The ASR transcripts of instructional videos are concatenated and refined by a LLM to match the descriptive style of manually labelled captions in Ego4d. We construct the ego-exo pairs by choosing the ego and exo captions that describe the similar action (e.g., toast the bread).

where **Prompt** is composed of task instructions, a list of 10 caption rules followed by example cases. Please refer to **Supplementary** for the exact prompt. Compared with clip-wise processing of each transcript individually, video-level refinement enables the LLM to infer the long-term goal of the activity within the context, and manages to replace the pronouns with specific objects deduced from other transcripts for better visual-text alignment. Additionally, we extend the temporal boundary of the narration to $[s - \alpha, e + \alpha]$ by α seconds, to mitigate the temporal shift problem. As shown in Fig. 4 (middle), the ASR transcripts are effectively transformed into shorter yet more informative sentences.

Pairing ego-exo videos via language alignment. In addition to refining the captions of HowTo100M videos, we also remove the character indicator in Ego4D narrations, e.g., the '#C' in captions. Till this point, the egocentric and exocentric narrations share similar textual formats, we then establish pairings via a two-step procedure. *First*, we conduct an initial categorisation based on the daily activity scenarios, such as food and entertaining, crafts, gardening, etc. Egocentric and exocentric videos belonging to the same scenario are grouped together; *Second*, we extract nouns and verbs using Spacy [64] from each sentence as they collaboratively determine the semantics of the activity in the video. For each egocentric video, the exocentric video(s) with the highest overlaps in terms of nouns and verbs are selected as the paired sample(s). As both egocentric and refined exo-narrations focus on the main actions/objects, using nouns and verbs is sufficient to represent the sentence semantics while being computationally efficient. Leveraging sentence-level similarity is also expected to work well.

Long-form video-text pairs construction. Till here, we have obtained paired ego-exo short video clips, where the duration of video clips are determined by the timestamp of the captions. As both Ego4d and HowTo100M consist of long videos lasting minutes, training on short video-text pairs may not fully exploit the information in long videos. To address this, we propose a simple strategy to

better utilise long-form information. In specific, as the narrations/captions are dense, we create 'fake' long-form video-text pairs by cropping longer video clips and summarising their accompanied narrations into one sentence via the LLM. These pseudo long-form pairs are integrated into the original data for training, and pairing long-form ego-exo video pairs can be executed similarly as before.

3.1.3 Retrieval Module Training and Inference

For training cross-view retrieval, we introduce a novel contrastive loss, termed as EgoExoNCE loss. Specifically, the proposed loss incorporates three variants of positive pairs: (1) the original video-text pairs $(\hat{z}_i^{\text{ego}}, u_i^{\text{ego}})$ and $(\hat{z}_i^{\text{exo}}, u_i^{\text{exo}})$ as adopted in the InfoNCE loss [52]; (2) cross-view pairs mined via language alignment, i.e., $(\hat{z}_i^{\text{ego}}, u_i^{\text{exo}})$ and $(\hat{z}_i^{\text{exo}}, u_i^{\text{ego}})$, as both u_i^{ego} and u_i^{exo} represent similar actions in the video; (3) all pairs that share at least one noun or verb *regardless of views*. The remaining pairs in each mini-batch are considered as negative pairs.

At training time, assume we have \mathcal{B} samples in each mini-batch, the positive pair set for the i -th sample is thus defined as:

$$\mathcal{P}_i = \{(\hat{z}_i^{\text{ego}}, u_i^{\text{ego}}), (\hat{z}_i^{\text{exo}}, u_i^{\text{exo}})\} \cup \{(\hat{z}_i^{\text{ego}}, u_i^{\text{exo}}), (\hat{z}_i^{\text{exo}}, u_i^{\text{ego}})\} \cup \{(\hat{z}_i^{\text{ego}}, u_j^c), (\hat{z}_i^{\text{exo}}, u_j^c) | \mathcal{N}(t_i^c) \cap \mathcal{N}(t_j^c) \neq \emptyset, \mathcal{V}(t_i^c) \cap \mathcal{V}(t_j^c) \neq \emptyset\},$$

where $j \in \mathcal{B}$, $c \in \{\text{ego}, \text{exo}\}$, $\mathcal{N}(t)$ and $\mathcal{V}(t)$ represent the sets of nouns and verbs in text t . The video-to-text EgoExoNCE loss $\mathcal{L}_{\text{EgoExoNCE}}^{\text{v2t}}$ can be calculated as:

$$\mathcal{L}_{\text{EgoExoNCE}}^{\text{v2t}} = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\sum_{(\hat{z}_i, u) \in \mathcal{P}_i} \langle \hat{z}_i, u \rangle}{\sum_{j \in \mathcal{B}} \langle \hat{z}_i^{\text{ego}}, u_j \rangle + \langle \hat{z}_i^{\text{exo}}, u_j \rangle}, \quad (1)$$

where $\langle z, u \rangle = \exp(z^T u / \tau)$ and τ is the temperature. The total loss $\mathcal{L}_{\text{EgoExoNCE}}$ is the sum of video-to-text loss and symmetrical text-to-video loss. Despite the training procedure only encourages cross-view, cross-modal retrieval, the ego-exo video embeddings are also implicitly pulled closer, by aligning to the shared semantics of the actor's action.

At inference time, given an egocentric video and a candidate set of exocentric videos with associated texts, we first calculate egocentric video representation (\hat{z}^{ego}) and exocentric video and text representations ($\hat{z}^{\text{exo}}, u^{\text{exo}}$). Then, for the egocentric video, we choose the exocentric sample with the highest averaged similarity, considering both (1) ego-video to exo-video similarity and (2) ego-video to exo-text similarity, which is defined as:

$$\arg \max_{(v^{\text{exo}}, t^{\text{exo}}) \in \mathcal{C}} \frac{1}{2} (\langle \hat{z}^{\text{ego}}, \hat{z}^{\text{exo}} \rangle + \langle \hat{z}^{\text{ego}}, u^{\text{exo}} \rangle). \quad (2)$$

3.2. Retrieval-augmented Captioning

In this section, we detail the proposed retrieval-augmented model for egocentric video captioning. In addition to the target video as input, retrieved instructional videos and associated texts are also incorporated as conditions for generating captions. Next, we present three key modules in our captioning model, as shown in Fig. 2.

Visual encoder. We adopt a frozen CLIP vision transformer encoder that takes the concatenation of exo- and ego-videos $v = [v_1^{\text{exo}}; \dots; v_K^{\text{exo}}; v^{\text{ego}}] \in \mathbb{R}^{(K+1) \times T \times C \times H \times W}$ as input, tokenises each frame into N image patches, and compute dense visual representation per frame, $z = \Phi_v(v) \in \mathbb{R}^{(K+1) \times T \times N \times d}$. Here, $[\cdot]$ refers to concatenation.

Perceiver resampler serves as a bottleneck module, that models the correspondence among the visual features and squeezes the visual features to a bottleneck feature with fixed length L , typically $L \ll T \times N$. This is achieved with Transformer Decoder using fixed-length of learnable query tokens $q \in \mathbb{R}^{L \times d}$, that cross-attend to the visual features, i.e., $\hat{q} = \Phi_p(q, [z + \mathbf{pe}; q]) \in \mathbb{R}^{(K+1) \times L \times d}$. Here, \mathbf{pe} denotes temporal positional embedding.

Text decoder is a visually conditioned LLM. It takes captions of the retrieved samples as input prompt and cross-attends to visual features for caption generation. The text input prompt \mathbf{t}_p is formed by stacking captions for all exocentric videos, which is denoted as:

$$\mathbf{t}_p = [\langle \text{video} \rangle t_1^{\text{exo}} \langle \text{eoc} \rangle; \dots; \langle \text{video} \rangle t_K^{\text{exo}} \langle \text{eoc} \rangle; \langle \text{video} \rangle],$$

where $\langle \text{video} \rangle$ and $\langle \text{eoc} \rangle$ are special media tokens representing the video and end-of-chunk in text format, respectively. In particular, several gated cross-attention modules [1, 6] are inserted between the original LLM blocks. Compared with the standard LLM block, the gated cross-attention module exhibits two key variations: (1) the cross-attention layer takes text features as query and visual features as keys and values, enabling visually conditioned generation. Note that, each text token only cross-attends to the visual feature of its corresponding video. (2) A tanh gating mechanism is added. The output of cross-attention layer and feed-forward-network in each gated cross-attention block is multiplied by a tanh layer with zero-initialisation. This ensures

Dataset	Evaluation Metrics	#Len(s)	#Num
<i>Egocentric benchmark</i>			
EK100 MIR [18]	mAP, nDCG	3.7	9668
EgoMCQ [21]	inter-/intra-video acc.	34.2	39751
SummMCQ [21]	inter-video acc.	181.6	1614
<i>Exocentric benchmark</i>			
YouCook2-Clip [79]	R@1, R@5, R@10	19.7	3350
YouCook2-Video [79]	R@1, R@5, R@10	212.4	436
<i>Ego-Exo benchmark</i>			
CharadesEgo [59]	Ego2Exo/Exo2Ego Avg. R@1/@5/@10	22.7	145
EgoLearner-MCQ	Ego2Exo/Exo2Ego acc.	5.1	1951
<i>Video Captioning</i>			
Ego4d cooking [21]	BLEU-4, METEOR, ROUGE-L, CIDER	0.7	7161
EgoLearner	BLEU-4, METEOR, ROUGE-L, CIDER	7.4	1089

Table 1. Datasets for evaluation. We list evaluation metrics, average clip/video length (#Len), and the number of test samples.

that the decoder is initially equivalent to the original LLM, thereby improving the training stability. The captioning model is trained by optimising the standard cross-entropy loss for each predicted token.

4. Experiments

4.1. Experimental Setups

Training datasets. For egocentric videos, we use the public Ego4d dataset [21], consisting of 3670 hours of daily activity videos. Following prior works [43, 78], videos belonging to the validation and test sets are excluded during training, resulting in 4.3M video-text pairs for training [78]. For exocentric videos, we adopt the large-scale instructional video dataset HowTo100M [46], containing more than 1M videos from YouTube. We use the annotation from TAN [23] including 30M video and ASR transcript pairs with better temporally aligned narration timestamps. For training the captioning model, we use the Ego4d [21] cooking subset, containing 0.55M video clips with manually labelled captions. All video clips are obtained based on the timestamps of the narrations following [43].

Downstream benchmarks. We evaluate the zero-shot transfer ability of our **retrieval module** on the validation/test set using benchmark datasets outlined in Table 1. Our evaluation includes 7 retrieval benchmarks that cover both egocentric [18, 21, 59] and exocentric videos [59, 79], ranging from short-term to long-term. Notably, CharadesEgo and an in-house dataset, termed as EgoLearner, with ground-truth ego-exo pairs are adopted to evaluate model’s video-to-video retrieval ability. The retrieval evaluations include video-text, video-video retrieval, and multiple choice questions (MCQ). Specifically, MCQ involves retrieving the matched video from 5 candidates from a query. The query is the associated text in EgoMCQ and SummMCQ. While in EgoLearner, the video from the other view serves as the query. We report the R@1, R@5 and R@10 for retrieval tasks and Top-1 accuracy for MCQ tasks. For **video captioning**, we sample 7161 cooking video clips

ID	Loss	Caption	View	Ego-centric Benchmark				Exo-centric Benchmark			Ego-Exo Benchmark			
				EK100-MIR		EgoMCQ		YouCook2 Clip→Text			EgoLearner		CharadesEgo	
				mAP	nDCG	Inter Acc.	Intra Acc.	R@1	R@5	R@10	Ego2Exo	Exo2Ego	Ego2Exo	Exo2Ego
1	InfoNCE	-	Ego	29.9	33.9	91.5	54.6	0.1	0.5	0.6	29.4	22.6	6.1	7.4
2	InfoNCE	ASR	Exo	6.7	12.3	35.0	22.6	21.7	48.1	61.2	24.6	21.9	1.8	4.1
3	InfoNCE	ASR	Ego+Exo	28.1	32.1	91.5	51.8	21.3	45.4	57.3	38.7	29.9	18.3	7.6
4	InfoNCE	Refined Cap.	Ego+Exo	28.4	32.0	92.3	53.2	23.3	49.2	60.7	46.8	30.1	17.8	9.8
5	EgoNCE	ASR	Ego+Exo	29.0	33.1	91.5	52.5	21.4	45.0	57.6	39.2	30.9	19.0	8.7
6	EgoExoNCE	ASR	Ego+Exo	28.6	33.2	91.9	51.4	21.1	45.6	57.7	41.3	31.3	20.8	13.6
7	EgoExoNCE	Refined Cap.	Ego+Exo	29.2	33.1	92.8	54.0	23.7	49.4	61.8	46.4	30.5	22.6	15.2

Table 2. Ablation study on the exocentric instructional video textual descriptions (original ASR transcripts vs. Refined Caption) and the cross-view alignment loss (InfoNCE [52] vs. EgoNCE [43] vs. EgoExoNCE). We report metrics on short-term benchmarks. The baseline model trained on ego-only/exo-only data (ID-1, ID-2) fails on exo/ego benchmarks due to the discrepancy between views.

from the Ego4d [21] validation set. We also choose videos from EgoLearner which contains 1089 video clips with fine-grained video captions. Please refer to the **Supplementary** for details on each dataset.

Training and inference pipeline. To begin with, we take egocentric videos from Ego4d and exocentric videos from HowTo100M to construct pseudo ego-exo pairs offline, which are subsequently used to train the cross-view retrieval module. Next, we take the videos from the cooking subset of Ego4d, along with the paired exocentric videos and captions, to train the retrieval-augmented captioning model. At inference stage, given an egocentric video, we retrieve relevant exocentric videos via the cross-view retrieval module, and use them as references to generate the ego-caption.

Implementation details. In the retrieval module, the egocentric video encoder is a VideoMAE-L [62] model pre-trained on Ego4d. The exocentric video encoder is a UniFormerV2 [38] model with InternVideo [68] pre-trained weights on third-person video datasets. Both video encoders output a 768-d feature vector for an 8-frame input. Following [23, 46], we get one feature per second by decoding video at 8 fps. To match the average duration of Ego4d and HowTo100M, we randomly sample 4 frame features as input to the cross-view video encoder by default. In terms of pseudo long-form video-text pairs, we crop long video clips that last between 60 and 300 seconds, covering 20 narrations. We employ Stable-Vicuna-v1.5 [17], an advanced instruction-tuned LLM for both long-form text summarisation and caption refinement. The cross-view video encoder consists of four Transformer encoder layers [65]. The text encoder is a BERT model [32]. The retrieval module is trained for 5 epochs using the AdamW optimizer with a learning rate of 3×10^{-5} , and the batch size is set to 4096.

For retrieval-augmented captioning model, the visual encoder is CLIP-L/14 and the perceiver resampler consists of 6-layer transformer blocks with 64 learnable query tokens. We adopt two variants of language models, LLaMA-7B [63] and MPT-1B [61]. For fast convergence, we initialise the captioning model with pre-trained weights [37].

We train the model for 3 epochs with an initial learning rate of 1×10^{-5} and cosine learning rate decay. We set batch size to 8/32 for LLaMA-/MPT-based models.

4.2. Experimental Results

4.2.1 Results on Cross-view Retrieval

We evaluate the effect of caption refinement, EgoExoNCE loss, and different egocentric video encoders in this section.

Effect of caption refinement. We show comparison between our refined exocentric video caption with the original ASR transcript across multiple losses in Table 2 (ID-3 vs. ID-4, ID-6 vs. ID-7). The refined caption leads to consistent improvements on YouCook2 exocentric retrieval, indicating better visual-text alignment. Surprisingly, we observe 1%~2% improvement on egocentric benchmark, EgoMCQ. We conjecture that the performance gain is from the alignment between egocentric and exocentric caption formats achieved through caption refinement. This alignment allows egocentric videos to align with broader semantically rich captions, thereby enhancing the representation of egocentric video-text relations. On cross-view benchmark, *i.e.*, EgoLearner video retrieval, the accuracy consistently improves in the Ego2Exo setting regardless of the used losses, as semantically similar ego- and exo-videos are better aligned via similar captions. In the Exo2Ego setting, however, the model exhibits minimal improvement, probably due to the fact that the egocentric videos in EgoLearner are primarily collected in a similar environment, making it challenging to distinguish these ego-videos.

Effect of EgoExoNCE loss. To investigate the impact of different loss functions on aligning egocentric and exocentric videos, we compared our EgoExoNCE with two commonly used loss functions, namely, InfoNCE [52, 78] and EgoNCE [43]. InfoNCE aligns each video with its paired caption without considering cross-view information. EgoNCE extends this by incorporating videos performing similar actions as positive samples, while different actions in the same video as negative samples.

Method	View	Backbone	Ego-centric Benchmark				Exo-centric Benchmark				Ego-Exo Benchmark				
			EK100-MIR		EgoMCQ		SummMCQ	YC2-Clip		YC2-Video		EgoLearner		CharadesEgo	
			mAP	nDCG	Inter Acc.	Intra Acc.	Acc	R@1	R@5	R@1	R@5	Ego2Exo	Exo2Ego	Ego2Exo	Exo2Ego
HierVL [5]	Ego	FIT	18.9	24.7	90.5	52.4	95.4 [†]	-	-	-	-	-	-	-	-
CLIP [52]	Ego+Exo	ViT-L	18.4	24.0	88.7	34.6	92.0	20.2	44.0	36.0	76.4	48.9	41.5	50.1	49.0
CLIP + Ours	Ego+Exo	ViT-L	19.4	24.5	88.3	35.0	93.5	23.6	52.8	58.0	90.4	49.0	42.3	59.5	59.1
EgoVLP [43]	Ego+Exo	TSF-B	24.7	28.4	91.2	45.7	88.8	20.2	44.8	39.6	78.0	52.0	43.8	21.7	22.9
EgoVLP + Ours	Ego+Exo	TSF-B	25.7	29.1	91.9	47.2	91.6	22.7	50.5	59.4	91.3	53.2	46.8	21.2	29.2
LaViLa [78]	Ego+Exo	TSF-L	24.3	28.2	90.8	45.3	89.9	19.7	44.5	37.6	76.6	52.6	43.3	37.1	25.9
LaViLa + Ours	Ego+Exo	TSF-L	25.7	29.0	91.2	46.5	94.2	23.2	51.2	61.0	89.7	53.7	41.8	35.5	33.6
InternVideo [68]	Ego+Exo	UF-L	21.4	25.1	88.9	38.7	88.6	20.4	45.4	34.8	73.8	50.1	42.2	63.9	62.0
InternVideo + Ours	Ego+Exo	UF-L	25.7	27.6	90.3	39.6	93.4	23.6	52.1	59.6	87.6	52.8	48.3	79.8	70.3
VideoMAE [62]	Ego+Exo	ViT-L	28.1	32.1	91.5	51.8	87.1	21.3	45.4	37.6	75.4	38.7	29.9	18.3	7.6
VideoMAE + Ours	Ego+Exo	ViT-L	31.6	34.9	92.7	54.2	90.3	25.5	51.8	73.4	96.8	44.0	30.4	25.1	19.3

Table 3. Effect of our cross-view visual representation alignment over various egocentric video encoders, while fixing the exocentric video encoder as InternVideo [68]. We report clip/video to text retrieval results on YouCook-Clip/YouCook-Video. † indicates different test data (still unavailable), which is not fairly comparable. FIT, TSF and UF refer to Frozen in Time [7], TimeSformer [10] and UniformerV2 [38].

As shown in Table 2, EgoNCE improves InfoNCE on egocentric benchmarks (ID-3 vs. ID-5), *i.e.*, EK100-MIR and EgoMCQ, however, only marginal improvement is observed on cross-view benchmarks. In contrast, EgoExoNCE boosts the accuracy of InfoNCE by an average of 2.0% and 4.3% on cross-view benchmarks (ID-3 vs. ID-6), *i.e.*, EgoLearner and CharadesEgo, respectively. The results can be further improved by using refined captions (ID-4 vs. ID-7). This indicates the superiority of our EgoExoNCE loss in leveraging captions to mine cross-view samples, thereby enhancing the model’s video alignment ability. In summary, our cross-view visual alignment outperforms the baseline on all benchmarks (ID-3 vs. ID-7).

Analysis on egocentric video encoder. We extract egocentric video features using various pre-trained encoders, including CLIP [52], EgoVLP [43], LaViLa [78], InternVideo [68] and VideoMAE [62], while ensuring the consistency of exocentric video features extracted via InternVideo [68]. The results are shown in Table 3.

All the baseline models are trained on both original egocentric and exocentric data via InfoNCE. By adding our cross-view visual representation alignment strategies, *i.e.*, training on refined paired data with EgoExoNCE loss, all models exhibit consistent improvements on most performance metrics. Specifically, on cross-view evaluation benchmarks, decent improvement can be seen on CLIP and InternVideo. For instance, our model improves the baseline InternVideo model by 15.9% and 7.7% on CharadesEgo.

4.2.2 Results on Retrieval-augmented Captioning

Effect of retrieved exocentric videos. Table 4 shows the comparison results on the impact of retrieved samples. We vary the number of exocentric videos K from 0 to 8, which is referred to as K -shot evaluation [1]. In the 0-shot scenario, where no exocentric videos are used, both LLaMA- and MPT-based captioning models achieve relatively low captioning performance, with the MPT-based model achiev-

Model	Shot	BLEU-4	METEOR	ROUGE-L	CIDER
LaViLA-GPT2 [78]	0	13.1	28.7	47.8	75.1
LaViLA-GPT2-XL	0	17.0	30.5	51.4	103.0
LaViLA-LLaMA7B	0	17.4	31.3	51.6	108.9
<i>Ours</i>					
LLaMA-7B	0	12.5	27.6	45.8	54.1
LLaMA-7B-ASR	1	27.7	38.4	58.3	166.5
LLaMA-7B	1	28.4	38.7	58.9	173.8
MPT-1B	0	12.0	27.9	44.7	46.8
MPT-1B	1	32.2	40.0	61.2	197.3
MPT-1B	2	32.5	40.2	61.8	205.4
MPT-1B	4	32.6	40.4	62.0	207.0
MPT-1B	8	32.0	40.7	62.2	210.7

Table 4. Video captioning results on Ego4d validation set (cooking subset). Shot refers to the number of exocentric videos.

ing only a 0.468 CIDER score. In contrast, by leveraging only one exocentric video as reference, the model demonstrates a significant improvement across all metrics. The captioning performance gradually improves as the number of retrieved samples increases, indicating that the multi-modal captioning model successfully leverages relevant visual and textual information from exocentric videos. This observation aligns with prior studies on few-shot learning [1, 13, 44, 74]. Compared with refined caption, using the original ASR transcript as the text input prompt leads to a slight decrease in performance for LLaMA-based captioning model (CIDER 1.665 vs. 1.768). We also re-implement a prior egocentric video captioning model on our training set, *i.e.*, LaViLA-GPT2 [78]. Our retrieval-augmented captioning model outperforms LaViLA by a large margin.

Performance on EgoLearner. We evaluate the generalisation ability of EgoInstructor by directly transferring the model trained on Ego4d cooking subset to EgoLearner cooking videos. As seen from Table 5, in the 0-shot scenario, our captioning models fail to achieve satisfactory results, which can be attributed to the discrepancy between the video and caption distribution of two datasets. In the 1-shot scenario, a significant boost can be observed, demonstrat-

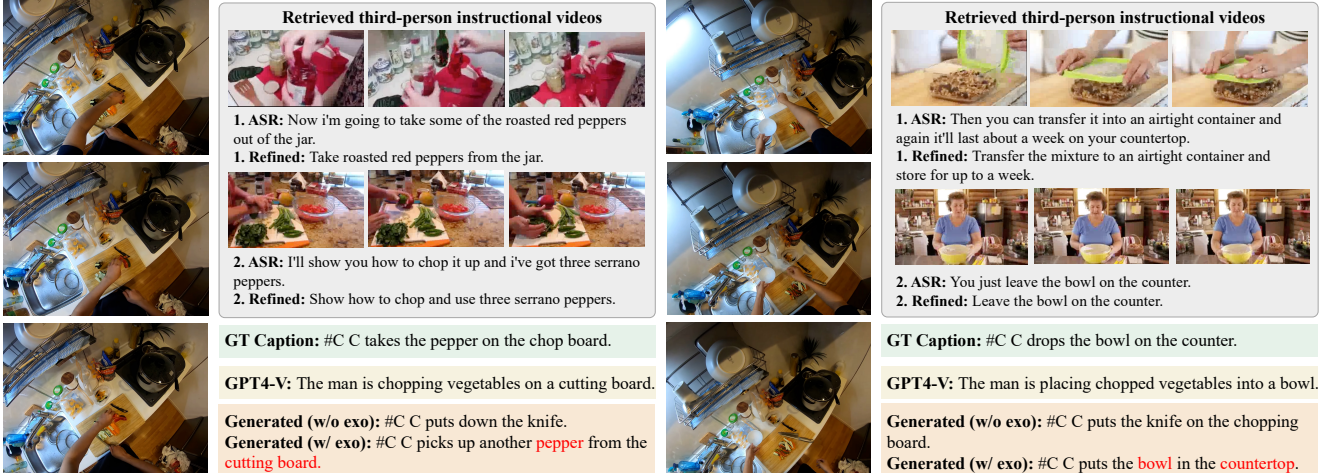


Figure 5. **Visualisation results.** For each egocentric video, we show two retrieved third-person instructional videos and their original ASR/refined captions. By leveraging retrieved exocentric samples, the generated captions capture correct actions and interacting objects.

ing that EgoInstructor effectively transfers knowledge from retrieved exocentric videos to assist ego-video captioning, achieving fast adaptation from Ego4d to EgoLearner. In addition, we study two variants of retrieved samples: (1) random selection and (2) ground-truth paired exocentric videos. Table 5 reveals that using randomly selected videos yields similar performance to 0-shot captioning, indicating that random samples do not provide useful semantics, and the model could be misled by irrelevant referenced videos. In contrast, samples retrieved by our cross-view retriever lead to a large gain of captioning performance, and using GT paired exo-videos performs even better.

4.2.3 Qualitative Results

As shown in Fig. 5, we showcase the retrieved third-person videos relevant to the given egocentric videos. As can be seen from examples, the retrieved samples are generally associated with actions or objects present in the egocentric videos. With the assistance of these third-person videos and corresponding captions, the captioning model accurately describes fine-grained objects in the video, such as “pepper” or “bowl”. In contrast, the captioning model without third-person video references generates less accurate descriptions. These results demonstrate the effectiveness of our EgoInstructor in leveraging retrieved third-person videos to improve egocentric video captioning. We have also experimented with the GPT4-Vision model [50]. Though GPT4-V excels at detecting objects in the video, the generated caption is not precise due to the incorrect perception of interacting objects. We include failure cases in the supplementary.

5. Conclusion

In this paper, we consider the problem of egocentric video captioning, and present a retrieval-augmented model,

Model	Shot	BLEU-4	METEOR	ROUGE-L	CIDER
LaViLA-GPT2 [78]	0	1.6	6.9	17.5	8.5
LaViLA-GPT2-XL	0	1.7	6.9	17.5	7.8
LaViLA-LLaMA7B	0	1.5	8.2	17.3	8.5
<i>Ours</i>					
LLaMA-7B	0	0.5	7.1	14.0	8.2
LLaMA-7B	1	2.4	12.0	22.6	20.4
LLaMA-7B (random)	1	0.8	7.1	17.3	5.2
LLaMA-7B (GT)	1	2.9	13.4	24.0	25.1
MPT-1B	0	0.6	7.0	15.6	6.5
MPT-1B	1	2.4	13.0	24.1	24.5
MPT-1B (random)	1	1.0	7.1	18.2	5.4
MPT-1B (GT)	1	3.2	15.3	26.1	32.7

Table 5. Zero-shot and one-shot video captioning results on EgoLearner. All the models are trained on Ego4d cooking subset and directly transferred to EgoLearner.

named EgoInstructor. It aims to retrieve semantically similar third-person instructional videos as references and generate captions for egocentric videos. We design a cross-view visual representation alignment for retrieval, where we train a cross-view retrieval module on automatically curated ego- and exo-video pairs. We propose a novel EgoExoNCE loss for training retrieval module, which aligns ego- and exo-video features to shared text features representing similar action semantics. Across seven benchmarks, our retrieval module demonstrates strong zero-shot cross-view and cross-modal retrieval ability. With the capability of cross-view retrieval, quantitative and qualitative results show the benefits of leveraging exocentric videos to enhance egocentric video captioning. In the future, we will extend the model to solve diverse egocentric video tasks.

Acknowledgement This work is supported by National Science and Technology Major Project (No. 2021ZD0114001), Natural Science Foundation of China (No. 62172101), the Science and Technology Commission of Shanghai Municipality (No. 21511101000; No. 23511100602).

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. [5](#), [7](#)
- [2] Shervin Ardeshir and Ali Borji. Ego2top: Matching viewers in egocentric and top-view videos. In *Proceedings of the European Conference on Computer Vision*, pages 253–268. Springer, 2016. [2](#)
- [3] Shervin Ardeshir and Ali Borji. An exocentric look at egocentric actions and vice versa. *Computer Vision and Image Understanding*, 171:61–68, 2018. [1](#), [2](#)
- [4] Akari Asai, Sewon Min, Zexuan Zhong, and Danqi Chen. Retrieval-based language models and applications. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts)*, pages 41–46, 2023. [2](#)
- [5] Kumar Ashutosh, Rohit Girdhar, Lorenzo Torresani, and Kristen Grauman. Hiervl: Learning hierarchical video-language embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23066–23078, 2023. [1](#), [2](#), [7](#)
- [6] Anas Awadalla, Irena Gao, Josh Gardner, Jack Hessel, Yusuf Hanafy, Wanrong Zhu, Kalyani Marathe, Yonatan Bitton, Samir Gadre, Shiori Sagawa, et al. Openflamingo: An open-source framework for training large autoregressive vision-language models. *arXiv preprint arXiv:2308.01390*, 2023. [5](#)
- [7] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1728–1738, 2021. [7](#)
- [8] Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. Whisperx: Time-accurate speech transcription of long-form audio. *INTERSPEECH 2023*, 2023. [3](#)
- [9] Donald J Berndt and James Clifford. Using dynamic time warping to find patterns in time series. In *Proceedings of the 3rd International Conference on Knowledge Discovery and Data Mining*, pages 359–370, 1994. [2](#)
- [10] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *International Conference on Machine Learning*, page 4, 2021. [7](#)
- [11] Alejandro Betancourt, Pietro Morerio, Carlo S Regazzoni, and Matthias Rauterberg. The evolution of first person vision methods: A survey. *IEEE Transactions on Circuits and Systems for Video Technology*, 25(5):744–760, 2015. [2](#)
- [12] Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George Bm Van Den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, et al. Improving language models by retrieving from trillions of tokens. In *International Conference on Machine Learning*, pages 2206–2240. PMLR, 2022. [2](#)
- [13] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020. [7](#)
- [14] Guo Chen, Sen Xing, Zhe Chen, Yi Wang, Kunchang Li, Yizhuo Li, Yi Liu, Jiahao Wang, Yin-Dong Zheng, Bingkun Huang, et al. Internvideo-ego4d: A pack of champion solutions to ego4d challenges. *arXiv preprint arXiv:2211.09529*, 2022. [1](#)
- [15] Wenhui Chen, Hexiang Hu, Xi Chen, Pat Verga, and William W Cohen. Murag: Multimodal retrieval-augmented generator for open question answering over images and text. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2022. [2](#)
- [16] Wenhui Chen, Hexiang Hu, Chitwan Saharia, and William W Cohen. Re-imagen: Retrieval-augmented text-to-image generator. *arXiv preprint arXiv:2209.14491*, 2022. [2](#)
- [17] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90 [6](#)
- [18] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision*, pages 1–23, 2022. [1](#), [2](#), [5](#)
- [19] Fernando De la Torre, Jessica Hodgins, Adam Bargteil, Xavier Martin, Justin Macey, Alex Collado, and Pep Beltran. Guide to the carnegie mellon university multimodal activity (cmu-mmact) database. 2009. [1](#)
- [20] Chenyou Fan, Jangwon Lee, Mingze Xu, Krishna Kumar Singh, Yong Jae Lee, David J Crandall, and Michael S Ryoo. Identifying first-person camera wearers in third-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5125–5133, 2017. [2](#)
- [21] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022. [2](#), [5](#), [6](#)
- [22] Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. Retrieval augmented language model pre-training. In *International Conference on Machine Learning*, pages 3929–3938. PMLR, 2020. [2](#)
- [23] Tengda Han, Weidi Xie, and Andrew Zisserman. Temporal alignment networks for long-term video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2906–2916, 2022. [3](#), [5](#), [6](#)
- [24] Hsuan-I Ho, Wei-Chen Chiu, and Yu-Chiang Frank Wang. Summarizing first-person videos from third persons’ points of view. In *Proceedings of the European Conference on Computer Vision*, pages 70–85, 2018. [1](#), [2](#)
- [25] Nicola J Hodges, A Mark Williams, Spencer J Hayes, and

- Gavin Breslin. What is modelled during observational learning? *Journal of Sports Sciences*, 25(5):531–545, 2007. 1
- [26] Yifei Huang, Minjie Cai, Zhenqiang Li, and Yoichi Sato. Predicting gaze in egocentric video by learning task-dependent attention transition. In *Proceedings of the European Conference on Computer Vision*, pages 754–769, 2018. 2
- [27] Ahmet Iscen, Mathilde Caron, Alireza Fathi, and Cordelia Schmid. Retrieval-enhanced contrastive vision-text models. *arXiv preprint arXiv:2306.07196*, 2023. 2
- [28] Baoxiong Jia, Yixin Chen, Siyuan Huang, Yixin Zhu, and Song-Chun Zhu. Lemma: A multi-view dataset for learning multi-agent multi-task activities. In *Proceedings of the European Conference on Computer Vision*, pages 767–786. Springer, 2020. 2
- [29] Hao Jiang and Kristen Grauman. Seeing invisible poses: Estimating 3d body pose from egocentric video. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3501–3509. IEEE, 2017. 2
- [30] Soo-Han Kang and Ji-Hyeong Han. Video captioning based on both egocentric and exocentric views of robot vision for human-robot interaction. *International Journal of Social Robotics*, pages 1–11, 2021. 2
- [31] Evangelos Kazakos, Arsha Nagrani, Andrew Zisserman, and Dima Damen. Epic-fusion: Audio-visual temporal binding for egocentric action recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5492–5501, 2019. 1, 2
- [32] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of North American Chapter of the Association for Computational Linguistics-HLT*, page 2, 2019. 3, 6
- [33] Ashish Kumar, Saurabh Gupta, and Jitendra Malik. Learning navigation subroutines from egocentric videos. In *Conference on Robot Learning*, pages 617–626. PMLR, 2020. 1
- [34] Taein Kwon, Bugra Tekin, Jan Stühmer, Federica Bogo, and Marc Pollefeys. H2o: Two hands manipulating objects for first person interaction recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10138–10148, 2021. 1
- [35] Bolin Lai, Miao Liu, Fiona Ryan, and James M Rehg. In the eye of transformer: Global–local correlation for egocentric gaze estimation and beyond. *International Journal of Computer Vision*, pages 1–18, 2023. 1
- [36] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474, 2020. 2
- [37] Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. Otter: A multi-modal model with in-context instruction tuning. *arXiv preprint arXiv:2305.03726*, 2023. 6
- [38] Kunchang Li, Yali Wang, Yanan He, Yizhuo Li, Yi Wang, Limin Wang, and Yu Qiao. Uniformerv2: Unlocking the potential of image vits for video understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1632–1643, 2023. 3, 6, 7
- [39] Yin Li, Alireza Fathi, and James M Rehg. Learning to predict gaze in egocentric video. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3216–3223, 2013. 1, 2
- [40] Yin Li, Zhefan Ye, and James M Rehg. Delving into egocentric actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 287–295, 2015. 1
- [41] Yin Li, Miao Liu, and James M Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *Proceedings of the European Conference on Computer Vision*, pages 619–635, 2018. 1
- [42] Yanghao Li, Tushar Nagarajan, Bo Xiong, and Kristen Grauman. Ego-exo: Transferring visual representations from third-person to first-person videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6943–6953, 2021. 1, 2
- [43] Kevin Qinghong Lin, Jimpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Z XU, Difei Gao, Rong-Cheng Tu, Wenzhe Zhao, Weijie Kong, et al. Egocentric video-language pretraining. *Advances in Neural Information Processing Systems*, 35:7575–7586, 2022. 1, 2, 5, 6, 7
- [44] Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. What makes good in-context examples for gpt-3? *arXiv preprint arXiv:2101.06804*, 2021. 7
- [45] Alexander Long, Wei Yin, Thalaisyasingam Ajanthan, Vu Nguyen, Pulak Purkait, Ravi Garg, Alan Blair, Chunhua Shen, and Anton van den Hengel. Retrieval augmented classification for long-tail visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6959–6969, 2022. 2
- [46] Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2630–2640, 2019. 2, 5, 6
- [47] Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, and Kristen Grauman. Ego-topo: Environment affordances from egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 163–172, 2020. 1
- [48] Katsuyuki Nakamura, Hiroki Ohashi, and Mitsuhiro Okada. Sensor-augmented egocentric-video captioning with dynamic modal attention. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 4220–4229, 2021. 2
- [49] Evonne Ng, Donglai Xiang, Hanbyul Joo, and Kristen Grauman. You2me: Inferring body pose in egocentric video via first and second person interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9890–9900, 2020. 2

- [50] OpenAI. Gpt-4 technical report, 2023. 8
- [51] David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4):515–526, 1978. 1
- [52] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 2, 4, 6, 7
- [53] Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. In-context retrieval-augmented language models. *arXiv preprint arXiv:2302.00083*, 2023. 2
- [54] Santhosh Kumar Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Naq: Leveraging narrations as queries to supervise episodic memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6694–6703, 2023. 1
- [55] Rita Ramos, Bruno Martins, Desmond Elliott, and Yova Kementchedjhiava. Smallcap: lightweight image captioning prompted with retrieval augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2840–2849, 2023. 2
- [56] Giacomo Rizzolatti and Laila Craighero. The mirror-neuron system. *Annu. Rev. Neurosci.*, 27:169–192, 2004. 1
- [57] Sara Sarto, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Retrieval-augmented transformer for image captioning. In *Proceedings of the 19th International Conference on Content-based Multimedia Indexing*, pages 1–7, 2022. 2
- [58] Fadime Sener, Dibyadip Chatterjee, Daniel Sheleпов, Kun He, Dipika Singhania, Robert Wang, and Angela Yao. Assembly101: A large-scale multi-view video dataset for understanding procedural activities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21096–21106, 2022. 1, 2
- [59] Gunnar A Sigurdsson, Abhinav Gupta, Cordelia Schmid, Ali Farhadi, and Karteek Alahari. Actor and observer: Joint modeling of first and third-person videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7396–7404, 2018. 1, 2, 5
- [60] Swathikiran Sudhakaran, Sergio Escalera, and Oswald Lanz. Lsta: Long short-term attention for egocentric action recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9954–9963, 2019. 1
- [61] MosaicML NLP Team. Introducing mpt-7b: A new standard for open-source, commercially usable llms, 2023. Accessed: 2023-05-05. 6
- [62] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *Advances in Neural Information Processing Systems*, 35:10078–10093, 2022. 3, 6, 7
- [63] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 6
- [64] Yuli Vasiliev. *Natural language processing with Python and spaCy: A practical introduction*. No Starch Press, 2020. 4
- [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017. 6
- [66] Huiyu Wang, Mitesh Kumar Singh, and Lorenzo Torresani. Ego-only: Egocentric action detection without exocentric transferring. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5250–5261, 2023. 1, 2
- [67] Qitong Wang, Long Zhao, Liangzhe Yuan, Ting Liu, and Xi Peng. Learning from semantic alignment between unpaired multiviews for egocentric video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3307–3317, 2023. 2
- [68] Yi Wang, Kunchang Li, Yizhuo Li, Yanan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, et al. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. 3, 6, 7
- [69] Xinxiao Wu, Han Wang, Cuiwei Liu, and Yunde Jia. Cross-view action recognition over heterogeneous feature spaces. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 609–616, 2013. 2
- [70] Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metzke, Luke Zettlemoyer, and Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 2021. 2
- [71] Mingze Xu, Chenyou Fan, Yuchen Wang, Michael S Ryoo, and David J Crandall. Joint person segmentation and identification in synchronized first-and third-person videos. In *Proceedings of the European Conference on Computer Vision*, pages 637–652, 2018. 1, 2
- [72] Zihui Xue and Kristen Grauman. Learning fine-grained view-invariant representations from unpaired ego-exo videos via temporal alignment. *Advances in Neural Information Processing Systems*, 2023. 1, 2
- [73] Kentaro Yamada, Yusuke Sugano, Takahiro Okabe, Yoichi Sato, Akihiro Sugimoto, and Kazuo Hiraki. Attention prediction in egocentric video using motion and visual saliency. In *Advances in Image and Video Technology: 5th Pacific Rim Symposium, PSIVT 2011, Gwangju, South Korea, November 20-23, 2011, Proceedings, Part I 5*, pages 277–288. Springer, 2012. 1, 2
- [74] Zhengyuan Yang, Zhe Gan, Jianfeng Wang, Xiaowei Hu, Yumao Lu, Zicheng Liu, and Lijuan Wang. An empirical study of gpt-3 for few-shot knowledge-based vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3081–3089, 2022. 2, 7
- [75] Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Richard James, Jure Leskovec, Percy Liang, Mike Lewis,

- Luke Zettlemoyer, and Wen-tau Yih. Retrieval-augmented multimodal language modeling. 2023. [2](#)
- [76] Chuhan Zhang, Ankush Gputa, and Andrew Zisserman. Helping hands: An object-aware ego-centric video recognition model. In *International Conference on Computer Vision*, 2023. [2](#)
- [77] Chen-Lin Zhang, Jianxin Wu, and Yin Li. Actionformer: Localizing moments of actions with transformers. In *Proceedings of the European Conference on Computer Vision*, pages 492–510. Springer, 2022. [1](#), [2](#)
- [78] Yue Zhao, Ishan Misra, Philipp Krähenbühl, and Rohit Girdhar. Learning video representations from large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6586–6597, 2023. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [79] Luowei Zhou, Chenliang Xu, and Jason Corso. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. [2](#), [5](#)