

# MODELING FINE-GRAINED HAND-OBJECT DYNAMICS FOR EGOCENTRIC VIDEO REPRESENTATION LEARNING

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

In egocentric video understanding, the motion of hands and objects as well as their interactions play a significant role by nature. However, existing egocentric video representation learning methods mainly focus on aligning video representation with high-level narrations, overlooking the intricate dynamics between hands and objects. In this work, we aim to integrate the modeling of fine-grained hand-object dynamics into the video representation learning process. Since no suitable data is available, we introduce **HOD**, a novel pipeline employing a hand-object detector and a large language model to generate high-quality narrations with detailed descriptions of hand-object dynamics. To learn these fine-grained dynamics, we propose **EgoVideo**, a model with a new lightweight motion adapter to capture fine-grained hand-object motion information. Through our co-training strategy, EgoVideo effectively and efficiently leverages the fine-grained hand-object dynamics in the HOD data. Extensive experiments demonstrate that our method achieves state-of-the-art performance across multiple egocentric downstream tasks, including improvements of **6.3%** in EK-100 multi-instance retrieval, **5.7%** in EK-100 classification, and **16.3%** in EGTEA classification in zero-shot settings. Furthermore, our model exhibits robust generalization capabilities in hand-object interaction and robot manipulation tasks.

## 1 INTRODUCTION

Egocentric video understanding has recently garnered increasing attention due to its crucial role in areas such as augmented reality (Pan et al., 2023), embodied AI (Srivastava et al., 2022; Geiger et al., 2013), and personalized assistants (Huang et al., 2018). With the collection of large-scale egocentric video datasets (Damen et al., 2020; Grauman et al., 2022), researchers begin to adopt video-language pretraining (Lin et al., 2022) based on these annotations to learn egocentric video representations. Since the original annotations tend to be highly template-driven and lack diversity, previous works explore using Large Language Models (LLM) to rephrase the narration (Zhao et al., 2023) or introducing new video-language pairs from exocentric datasets (Dou et al., 2024). This scheme has shown its success in a wide range of downstream tasks (Plizzari et al., 2024).

However, as can be seen from the example in Figure 1 right, the original annotations in egocentric video datasets are typically highly condensed, describing only overall actions like “C draws on a book” or “C moves both hands”. Since no additional information is provided, previous works like LaViLa (Zhao et al., 2023) can only rephrase at the same level of abstraction as the original annotations, neglecting a crucial aspect of egocentric videos – the fine-grained dynamics of hands and objects. Most egocentric videos contain a large portion of hand-object interactions, which reflects the camera wearer’s behavior and intentions. As will be seen, integrating this information in vision-language pretraining significantly enhances egocentric video representation learning, resulting in state-of-the-art performance across various benchmarks.

Firstly, to incorporate hand-object dynamics into vision-language pretraining, it is essential to construct data that accurately captures the detailed motion of hands and objects in videos. A recent work directly uses the output of off-the-shelf hand-object detectors (Shan et al., 2020) as the ground truth of auxiliary targets in the pretraining (Zhang et al., 2023). However, this approach only models

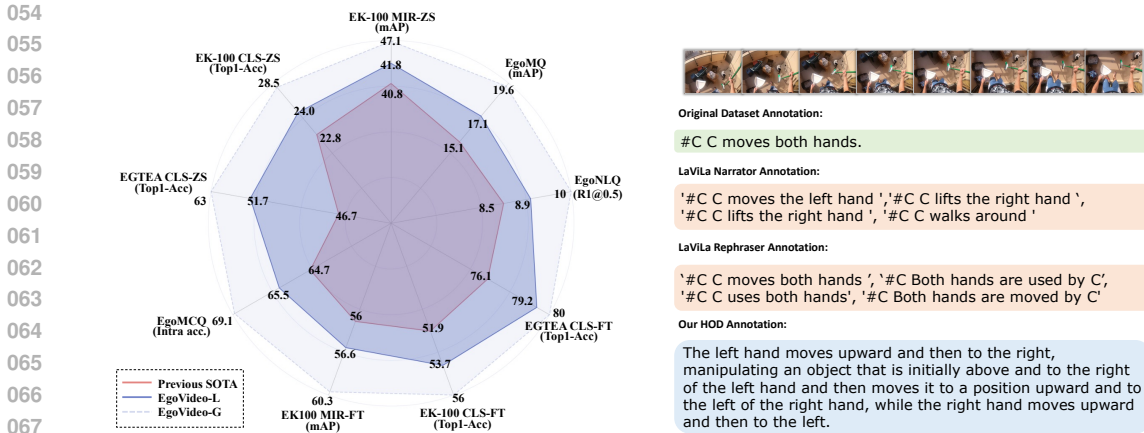


Figure 1: *Left*: Our **EgoVideo** model achieves state-of-the-art performance across multiple video benchmarks by learning fine-grained hand-object dynamics from videos. *Right*: Annotations from different sources: original Ego4D annotation (Grauman et al., 2022), LaViLa (Zhao et al., 2023), and our **HOD**. Our HOD annotations provide a detailed description of hand movements and object manipulation, demonstrating a higher level of detail and context.

the appearance of hands and objects without considering their dynamics. It also fails to learn the semantic connections between hand-object interactions and the original narration. To address this, we introduce **HOD**, a novel framework for generating descriptions with fine-grained **Hand-Object Dynamics** for a given video clip. We begin with using hand-object detectors to obtain bounding boxes of hands and contact objects. Then we design prompts based on these bounding boxes to generate descriptions of the trajectories of the hand and object, as well as their contact states and positions. Finally, using the new prompts and original annotations, we leverage a large language model (LLM) to generate semantically rich captions that encompass the motion states of hands and objects. By utilizing high framerate inputs, we ensure the capture of more detailed motions.

Secondly, to efficiently and effectively exploit the fine-grained spatiotemporal information in HOD, we propose **EgoVideo**, a novel ViT-based model with a lightweight motion adapter. Cooperating with the HOD data, EgoVideo employs a dual-branch design and co-training strategy. The backbone branch is trained normally to learn fundamental video-language alignment, while the adapter branch is trained with a higher framerate to capture detailed hand-object dynamics. The motion adapter has a separable convolution design, allowing for information aggregation from both adjacent frames temporally and from hands and objects at different locations. This design enables EgoVideo to effectively model detailed hand-object dynamics while maintaining low computational costs. Our efficient design also allows us to scale the model size to 1B parameters to fully unlock its potential to comprehend egocentric videos.

We extensively evaluate EgoVideo across multiple pretraining data sources and various egocentric downstream tasks. Experimental results show that our model sets a new state-of-the-art on 9 tasks as partially shown in Figure 1. Notably, our model also achieves the best performance under the same model size in both zero-shot and fine-tuning settings. Further experiments demonstrate that our HOD data is also beneficial for robot manipulation tasks.

Our main contributions are as follows: (1) We develop a HOD data pipeline to generate captions that describe fine-grained hand-object dynamics, which are crucial for egocentric video understanding; (2) We propose EgoVideo, a dual-branch model with a novel lightweight motion adapter and a co-training strategy to leverage the HOD data efficiently and effectively; (3) We demonstrate state-of-the-art performance on 12 downstream tasks, and our approach generalizes well to robot manipulation tasks. All code and data will be made publicly available.

## 2 RELATED WORK

**Egocentric Video Understanding** is receiving increasing research attention. Previous works focus on diverse tasks such as action recognition (Plizzari et al., 2022; Radevski et al., 2023), action antic-

108 ipation (Girdhar & Grauman, 2021), and cross-view understanding (Xue et al., 2022; Huang et al.,  
109 2024; Luo et al., 2024). Recent methods begin to work on egocentric representation learning (Lin  
110 et al., 2022) using the large-scale data from Ego4D (Grauman et al., 2022), or refining the Ego4D  
111 narrations by LLM rephrasing (Zhao & Krähenbühl, 2023). A recent work also searches for addi-  
112 tional data from exocentric datasets to improve the pretraining (Dou et al., 2024). However, since  
113 the Ego4D narrations are highly abstract, these methods fail to learn one critical aspect of egocentric  
114 videos – fine-grained hand-object dynamics of hands and objects. Recently, Helping Hands (Zhang  
115 et al., 2023) utilizes hand and object coordinates as auxiliary targets during pretraining. However,  
116 it only focuses on the spatial information of hands and objects, neglecting their motion dynamics.  
117 Additionally, the provided supervision does not integrate the states of hands and objects with the  
118 video descriptions, limiting the model’s ability to comprehend fine-grained details.

119 Unlike previous works, we propose the first method to integrate the hand-object dynamics into  
120 egocentric representation learning. On the data side, we propose the **HOD** (Hand-Object Dynamics)  
121 pipeline, which generates high-quality video-language pairs. The language in these pairs explicitly  
122 represents the complex states and motions of hands and objects in the videos, enabling the model  
123 to learn detailed information about these dynamics. On the model side, we introduce **EgoVideo**, a  
124 model equipped with a lightweight motion adapter. This adapter is designed to effectively capture  
125 the intricate hand and object dynamics provided by the HOD data, enhancing the model’s ability to  
126 understand and interpret fine-grained dynamics in egocentric videos.

127 **Video-Language Representation Learning** has also attracted researchers after the success of  
128 CLIP (Radford et al., 2021), due to the need for generating robust video representations. Several  
129 large-scale video-language datasets (Kay et al., 2017; Miech et al., 2019; Caba Heilbron et al., 2015)  
130 further fueled the research in this area. However, generating high-quality video-text pairs remains  
131 a challenging task, prompting researchers to develop innovative solutions. LaViLa (Zhao et al.,  
132 2023) leverages Large Language Models (LLMs) to generate dense narrations for videos. Video  
133 Recap (Islam et al., 2024) utilizes a curriculum learning training scheme to generate summaries for  
134 long videos. EMBED (Dou et al., 2024) and EgoInstructor (Xu et al., 2024) use rules or retrieval  
135 models to add additional training data. However, the previous methods can only pretrain their mod-  
136 els at the same abstraction level as the original annotation. In contrast, our approach integrates  
137 finer-level details into the representation learning process.

138 **Hand-Object Interaction Understanding** has long been a key research topic within the field of  
139 egocentric vision. In recent years, several works have made significant strides in modeling estimate  
140 3D hand joints (Brahmbhatt et al., 2020; Cai et al., 2018; Yang & Yao, 2019; Yuan et al., 2018;  
141 Ohkawa et al., 2023) and reconstructing hand-object shape (Cao et al., 2021; Doosti et al., 2020;  
142 Hasson et al., 2019; 2020; Liu et al., 2021). EgoHOS (Zhang et al., 2022) provides a labeled dataset  
143 with fine-grained per-pixel labels of hand and objects and a reliable foundational tool for 2D hand-  
144 object segmentation, 100DOH (Shan et al., 2020) introduces a large-scale video dataset containing  
145 hands and hand-object interactions, providing a rich resource for hand object detector training. In  
146 our work, we utilize existing hand and object detectors in our HOD pipeline to convert information  
147 related to hand/object motion and contact details into natural language descriptions. By integrating  
148 these detailed descriptions with our EgoVideo model, we can integrate this finer level of detail into  
149 the video representation learning process.

## 150 3 METHOD

### 151 3.1 DATA GENERATION PIPELINE: HOD

152  
153  
154 The fine-grained dynamics of hands and objects play a pivotal role in egocentric video understand-  
155 ing (Fathi et al., 2011a). To effectively integrate this information in the video-language pretraining  
156 process, we propose HOD, a novel data generation pipeline to transform hand-object dynamics into  
157 natural languages. An overview of HOD is illustrated in Figure 2 top. First, we utilize an off-the-  
158 shelf hand object detector (Shan et al., 2020) to generate bounding boxes for hands and objects in  
159 each frame of the video clips. Next, we employ a large language model (AI et al., 2024) to enrich  
160 the original video captions. The model is prompted to generate new narrations that integrate the  
161 original captions with hand-object dynamics information, enhancing the semantic richness of the  
162 annotations. Below, we go into the details of the HOD data generation process.

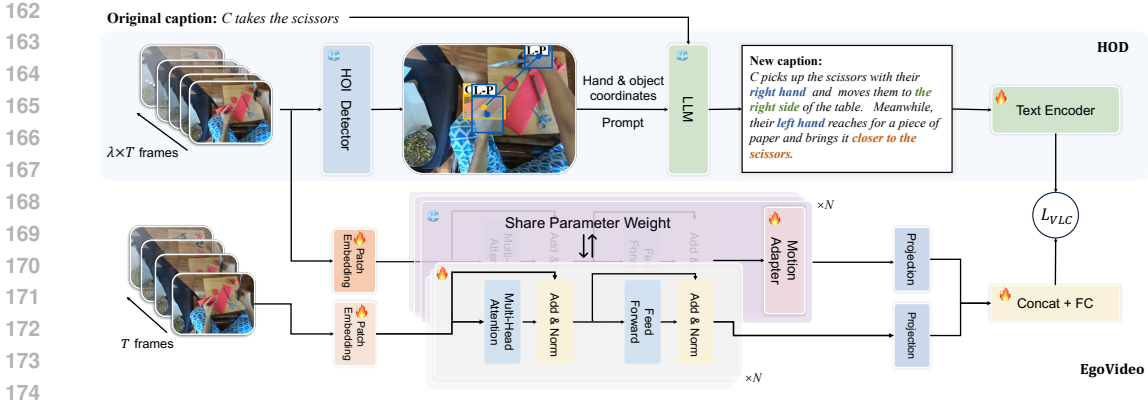


Figure 2: Illustration of our **HOD** pipeline and **EgoVideo** model. In our **Hand-Object Dynamics** data generation pipeline (top), we first use a hand object detector to obtain the spatial coordinates of hands and objects in the clip, then we combine the motion information of hands and objects with the original narrations to generate semantically richer narrations. In our **EgoVideo** model (bottom), the backbone is trained with a lower framerate. We design a lightweight motion adapter to learn fine-grained dynamics efficiently with higher framerate inputs.

### 3.1.1 DATA SELECTION

Before actually going into the generation process, it is essential to select appropriate source data. The basic data component comes from the 4M subset (Lin et al., 2022) of Ego4D (Grauman et al., 2022) which has been proven useful in egocentric video-language pretraining (Pranav et al., 2023). Additionally, we curate data from the large-scale HowTo100M dataset (Miech et al., 2019) since it contains rich hand-object interactions. We specifically choose How2-Interlink7M (Wang et al., 2024a) which contains 7M clips with high-quality GPT-4 (OpenAI, 2023) refined caption. Since the videos are from diverse sources that may include a portion that impedes egocentric representation learning, we employ a filtering technique to retain only clips with egocentric style. To this end, we train a style classifier  $\mathcal{P}$  by manually annotating 10,000 clips as “ego-like” or “non-ego-like”. With this classifier, we obtain 3.4M egocentric-style clips from How2-Interlink and combine these clips with Ego4D, forming our pretraining dataset. More details can be found in Appendix A.

### 3.1.2 GENERATING CAPTIONS WITH HAND OBJECT DYNAMICS

In this section, we introduce our HOD framework. Existing methods of refining descriptions in video-language pretraining (Zhao et al., 2023; Dou et al., 2024) focus on high-level abstracts but overlook the fine-grained details of hand-object dynamics. This oversight is detrimental in egocentric representation learning, where understanding these interactions forms a considerable proportion of egocentric videos by nature. To address this gap, in our HOD framework, we first detect the positions of hands and objects using a hand object detector. With this information, we prompt a large language model to augment the original annotation with detailed descriptions of hand and object movements. Followed by the subsequent video-language pretraining, our EgoVideo model can understand videos at a finer-grained level.

**Hand Object Dynamics Detector.** Thanks to the rapid advancement in the field of hand-object interaction (Jiang et al., 2021; Ohkawa et al., 2023), off-the-shelf hand-object detectors can provide robust hand and object positions. In our framework, we employ 100DOH (Shan et al., 2020) as the detector  $\Phi_{\text{det}}$  for bounding boxes extraction.

For a video clip  $x = (x_1, x_2, \dots, x_T)$ , we uniformly sample  $n = 16$  frames within the clip to obtain fine-grained motion information. Then we use  $\Phi_{\text{det}}$  to acquire the bounding boxes of hands and objects on these frames, which can be represented as

$$LH_i, RH_i, LO_i, RO_i = \Phi_{\text{det}}(x_i) \quad (1)$$

where  $LH_i, RH_i, LO_i, RO_i$  denotes the bounding box of the left hand, right hand, objects in contact with the left hand, and objects in contact with the right hand in the  $i$ -th frame. We use linear interpolation to compensate for missing hand boxes of frame  $t$  if the corresponding hand boxes can be detected for both frames  $x_{t-1}$  and  $x_{t+1}$ .

**Hand Object Dynamics Rephraser.** Current pretraining methods only use high-level language descriptions (e.g., “C takes the scissors” in Figure 2), which lacks important egocentric details like hand and object interaction. In this work, we incorporate these details into the video-language pretraining process. Hand-object dynamics encompass a variety of information, including bounding boxes of hands and objects, hand and object movement directions and trajectories, as well as their contact conditions. To integrate all this information into the video-language pretraining process, we use a LLM as a rephraser to express these dynamics in natural language.

Specifically, we employ Yi-34B (AI et al., 2024) as our LLM. To capture the nuances of hand and object movements, we extract the central points of bounding boxes to derive trajectories for hands and objects. This process yields six essential categories of information: spatial-temporal data for 1) the left hand, 2) the right hand, 3) objects contacted by the left hand, 4) objects contacted by the right hand, 5) objects contacted by both hands, and 6) the original narration. We then prompt the LLM to amalgamate this detailed information, enabling the generation of rich narratives that intricately describe hand-object dynamics. Further details on prompting can be found in Appendix A.

**Analysis of HOD Data.** We conduct additional analyses on our HOD data to evaluate its quality. First, we identify the top 30 most frequent words in HOD captions and the original EgoClip narrations and plot their normalized frequencies in Figure 3. The EgoClip narrations exhibit a more pronounced long-tail distribution, while our HOD captions display a more balanced distribution. Notably, HOD captions include many “dynamic” words, such as “up” and “downwards,” which aligns with the rationale behind our data generation process. To further verify the quality of our HOD, we employ GPT-4o (OpenAI, 2023) for quality assessment. We randomly select 1000 clips and let GPT evaluate the score of the caption data for the video clip in a range from 0 to 10. To ensure GPT does not simply assign high scores based on the length of the captions, we also conduct random gerund replacements on our data for comparison. The results, summarized in Table 1, show that our HOD data have a significantly better GPT-Score. Additional details on the scoring process and evaluations using other metrics are provided in Appendix A.




Figure 3: Normalized frequency of the Top-30 word in EgoClip (green) and our HOD (blue). Our HOD data has a less long-tail distribution, showing its word diversity.

Table 1: Results of narration quality, where HOD-random represents the narration after replacing keywords. The GPT-Score ranges from 0 to 10, with higher values indicating higher quality of narration.

Data	GPT-Score
EgoClip	5.53
HOD-random	3.70
<b>HOD</b>	<b>7.71</b>

### 3.2 EGOCENTRIC REPRESENTATION LEARNING MODEL: EGOVIDEO

The narrations generated by our HOD pipeline are highly detailed. As a result, the previous pre-training scheme struggles to capture the corresponding visual information at this level of detail. In response, we introduce EgoVideo (Figure 2 bottom), a model comprising a backbone and a motion adapter. The motion adapter aids in learning fine-grained hand-object dynamics from densely sampled video frames. Cooperating with a co-training strategy, our EgoVideo model can obtain richer video representations while maintaining computational efficiency.

**Visual and Text Encoder.** Following the standard video-language pretraining setting (Lin et al., 2022), our model includes a visual encoder  $\mathcal{F}_v$  (including our motion adapter) and a text encoder  $\mathcal{F}_t$ . In the visual encoder, for a clip  $x \in \mathbb{R}^{T \times H \times W \times 3}$ , we concatenate image tokens in  $T$  frames with a learnable class token. The output of our visual encoder is  $\mathbf{E}_v \in \mathbb{R}^D$ . For the text encoder,

we employ a 12-layer GPT-like Transformer (Radford et al., 2019) that input tokens after BPE tokenization (Sennrich, 2015). The output of our text encoder is  $\mathbf{E}_t \in \mathbb{R}^D$ .

**Motion Adapter.** Intuitively, to encode visual representations at the same level of detail as the languages, it is essential to utilize a greater number of frames as input. Since increasing the number of frames in training will result in unacceptable computational overhead, inspired by the PEFT technique in LLMs (Ding et al., 2023), we propose to use a lightweight motion adapter. The motion adapter is injected between the layers of the visual backbone, and is tailored to learn the finer-grained details with a high framerate. Since the hand and object motion forms a spatiotemporal pattern, unlike previous methods (Pan et al., 2022; Xing et al., 2024) that only focus on learning temporal information, our module is designed to learn both spatial and temporal information.

Our motion adapter is attached to the top of each of the  $N$  transformer layers. Without loss of generality, here we illustrate the motion adapter for one transformer layer and illustrate it in Figure 4. Denote  $\mathbf{Y} \in \mathbb{R}^{L \times D}$  as the output of a transformer layer in  $\mathcal{F}_v$  where  $L$  is the number of tokens, we first forward  $\mathbf{Y}$  to a down-projection layer  $W_{\text{down}}$  with ratio  $\gamma$  and followed with a GELU activation function  $\sigma$ . Then, we use a 2D convolution layer  $\text{Conv2D}$  with kernel size  $(k, k)$  to aggregate spatial information from each frame, followed by a 1D temporal convolution layer  $\text{TConv1D}$  and Linear layer  $W_m$  to model the dynamics between adjacent frames. Finally, an up-projection layer  $W_{\text{up}}$  is used to restore the dimension. Formally, the structure can be described as:

$$\begin{aligned} \mathbf{Y}' &= \sigma(\mathbf{Y}W_{\text{down}}), \quad \mathbf{Y}_s = \text{ReLU}(\text{BN}(\text{Conv2D}(\mathbf{Y}'))), \\ \mathbf{Y}_{\text{st}} &= (\text{TConv1D}(\mathbf{Y}_s))W_m, \quad \text{MotionAdapter}(\mathbf{Y}) = \mathbf{Y} + \mathbf{Y}_{\text{st}}W_{\text{up}}, \end{aligned} \quad (2)$$

where  $W_{\text{down}} \in \mathbb{R}^{D \times \gamma D}$ ,  $W_m \in \mathbb{R}^{\gamma D \times \gamma D}$  and  $W_{\text{up}} \in \mathbb{R}^{\gamma D \times D}$ . BN denotes BatchNorm2D.

**Co-training Strategy.** In EgoVideo, the motion adapter receives input at a higher framerate to capture the fine-grained information. Additionally, the backbone must be trained to fully adapt to the egocentric domain. Thus, different from previous PEFT methods that freeze the backbone and only train the adapter part, we need to train both the backbone and adapter parameters. Motivated by the architecture of (Feichtenhofer et al., 2019), we employ a co-training strategy to train the backbone and the motion adapter jointly.

Specifically, we use an upsampling parameter  $\lambda$  to sample the input using two sampling rates. For the input  $x_l \in \mathbb{R}^{T \times H \times W \times C}$  with a low sampling rate, we pass it through the backbone and unfreeze all parameters. As a result, we get the output  $\mathbf{E}_{v1} \in \mathbb{R}^D$ . For the input with a higher sampling rate  $x_h \in \mathbb{R}^{\lambda T \times H \times W \times C}$ , we pass the input through both the backbone and adapter parameters and get the output  $\mathbf{E}_{vh} \in \mathbb{R}^D$ , during which we freeze the parameters of the backbone and only train the adapter. Finally, for the outputs of the two pathways, we concatenate them and pass them through a fully connected layer to obtain the final output  $\mathbf{E}_v \in \mathbb{R}^D$ :

$$\begin{aligned} \mathbf{E}_{v1} &= \mathcal{F}_{\text{backbone}}(x_l), \quad \mathbf{E}_{vh} = \mathcal{F}_v(x_h), \\ \mathbf{E}_v &= [\mathbf{E}_{v1}; \mathbf{E}_{vh}]W_o, \end{aligned} \quad (3)$$

Where “[;]” denotes the concatenation operation,  $\mathcal{F}_{\text{backbone}}$  denotes the visual backbone,  $\mathcal{F}_v$  denotes  $\mathcal{F}_{\text{backbone}}$  with motion adapter and  $W_o \in \mathbb{R}^{2D \times D}$ . With this strategy, we integrate the training of the backbone and adapter into a single stage, reducing the cost of data and computation.

**Vision-Text Alignment.** We follow the standard InfoNCE (Oord et al., 2018) loss as the objective for alignment between visual embedding  $\mathbf{E}_v$  and text embedding  $\mathbf{E}_t$ . For a sampled batch  $\mathcal{B}$ , we have:

$$\mathcal{L} = \frac{1}{|\mathcal{B}|} \sum_{(\mathbf{E}_v^i, \mathbf{E}_t^i) \in \mathcal{B}} \left( \log \frac{e^{s(\mathbf{E}_v^i, \mathbf{E}_t^i)/\tau}}{\sum_{\mathbf{E}_t^j \in \mathcal{B}} e^{s(\mathbf{E}_v^i, \mathbf{E}_t^j)/\tau}} + \log \frac{e^{s(\mathbf{E}_v^i, \mathbf{E}_t^i)/\tau}}{\sum_{\mathbf{E}_v^k \in \mathcal{B}} e^{s(\mathbf{E}_v^k, \mathbf{E}_t^i)/\tau}} \right), \quad (4)$$

where  $s(\mathbf{E}_v^i, \mathbf{E}_t^i)$  denotes dot product operation between the  $i$ -th sample in the batch of  $\mathbf{E}_v$  and  $\mathbf{E}_t$ , and  $\tau$  is a temperature parameter that scales the similarity scores.

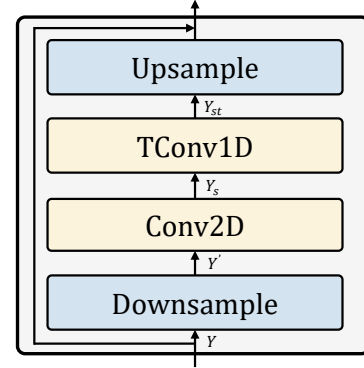


Figure 4: Architecture of our motion adapter. We use a 2D convolution layer and a 1D temporal convolution layer to capture the spatial and temporal dynamics efficiently.



Table 2: Zero-shot performance comparison on 4 tasks between methods with different model sizes ('B' for base, 'L' for large, and 'G' for our 1B parameter backbone). Our **EgoVideo** outperforms previous methods with less but higher-quality pretraining data. Specifically, our **EgoVideo-G** achieves significant performance improvements across all datasets.

Method (ZS)	Data Size	EK-100 MIR		EK-100 CLS		EGTEA		EgoMCQ	
		<i>mAP</i>	<i>nDCG</i>	<i>Top1-acc.</i>	<i>Top5-acc.</i>	<i>Mean-acc.</i>	<i>Top1-acc.</i>	<i>Intra</i>	<i>Inter</i>
EgoVLPv2	4M	26.7	29.1	-	-	-	-	60.9	91.0
LaViLa-B	35M	30.9	32.0	16.4	34.4	28.9	35.4	59.9	93.8
AVION-B	35M	32.9	32.7	-	-	-	-	-	-
EMBED-B	38.3M	36.0	<b>34.9</b>	19.0	39.0	37.0	42.7	61.3	94.5
<b>EgoVideo-B</b>	7.4M	<b>36.5</b>	34.5	<b>22.4</b>	<b>43.3</b>	<b>43.6</b>	<b>51.0</b>	<b>64.6</b>	<b>95.0</b>
LaViLa-L	35M	36.1	34.6	20.8	41.4	34.1	40.1	63.1	94.5
AVION-L	35M	37.6	35.3	-	-	-	-	-	-
Helping Hands	4M	37.5	<b>37.8</b>	-	-	39.1	46.6	63.0	94.5
EMBED-L	38.3M	40.8	37.5	22.8	45.0	40.3	46.7	64.7	95.6
<b>EgoVideo-L</b>	7.4M	<b>41.8</b>	37.0	<b>24.0</b>	<b>46.8</b>	<b>47.1</b>	<b>51.7</b>	<b>65.5</b>	<b>95.9</b>
<b>EgoVideo-G</b>	7.4M	<b>47.1</b>	<b>39.0</b>	<b>28.5</b>	<b>54.3</b>	<b>58.0</b>	<b>63.0</b>	<b>69.1</b>	<b>96.6</b>

## 4 EXPERIMENTS

### 4.1 DATASETS AND EVALUATION PROTOCOLS

**Pretraining Dataset.** As stated in the previous section, the source of our pretraining data comes from Ego4D (Grauman et al., 2022) and How2-Interlink-7M (Wang et al., 2024a). After processing by our HOD pipeline, the total amount of data is 7.4M clips.

**Evaluation Protocols.** We follow previous works (Zhao et al., 2023; Pramanick et al., 2023) and use the following evaluation protocols. (1) Zero-shot (ZS): the pretrained video-text encoders are directly applied to the downstream datasets to perform video-text retrieval tasks without any additional tuning. For classification, we compute the similarity score between the video clip and the textual descriptions of all possible classes. (2) Finetuned (FT): This approach involves taking the pretrained video-text model and performing end-to-end finetuning on the training split of the target downstream dataset. (3) Feature-based: We extract video features using a frozen encoder and only train a task-specific head on the downstream dataset.

**Model Architecture and Hyperparameters.** Our vision-language model follows the initialization of CLIP (Radford et al., 2021), which is composed of a vision encoder and a text encoder. For the base and large models, ViT is used as our vision encoder, and we use a temporal position embedding to learn temporal information, which is randomly initialized. For our giant size model, we use Internvideo2 (Wang et al., 2024b). For hyperparameters, we use  $T = 4$  and  $\lambda = 4$  for frame inputs, and we use downsample ratio  $\gamma = 0.5$  for the motion adapter. During pretraining, we freeze the temperature parameter  $\tau = 0.07$ . More details are placed in Appendix C.

**Downstream Tasks.** We evaluate models on several egocentric downstream tasks: (1) Epic-Kitchens-100 (Damen et al., 2020) (EK-100) tasks. For this dataset, we evaluate our method on multi-instance retrieval (EK-100 MIR) and action recognition (EK-100 CLS) tasks; 2) Ego4D (Grauman et al., 2022) tasks. For Ego4D, we evaluate our model on multiple choice questions (EgoMCQ) (Li et al., 2021), and natural language query (EgoNLQ) and moment query (EgoMQ) tasks; 3) EGTEA (Li et al., 2018) tasks. We evaluate our model on the action recognition task that is focused on fine-grained cooking activities and hand-object interaction. 4) Other tasks. We also evaluate our model on GTEA (Fathi et al., 2011b) and HOI4D (Liu et al., 2022) datasets for the action segmentation task. Meanwhile, to show the generalization ability of our learned video representation, we evaluate the task success rate on Franka Kitchen dataset (Gupta et al., 2019), a simulation environment for embodied AI.

### 4.2 COMPARISON TO STATE-OF-THE-ART

**Zero-shot Evaluation.** Table 2 shows the results on 4 tasks in the zero-shot setting. We compare our method against previous egocentric representation learning methods including EgoVLPv2 (Pramanick et al., 2023), LaViLa (Zhao et al., 2023), AVION (Zhao & Krähenbühl, 2023), Helping

Table 3: Fine-tuning performance of models of different sizes on 5 tasks. Compared to the previous SOTA EMBED-L, our EgoVideo achieves new state-of-the-art performances on all datasets.

Method (FT)	EK-100 MIR		EK-100 CLS	EGTEA	EgoNLQ	EgoMQ	
	<i>mAP</i>	<i>nDCG</i>	<i>Top1-acc.</i>	<i>Top1-acc.</i>	<i>R1@0.5</i>	<i>R1@0.5</i>	<i>mAP</i>
EgoVLPv2-B	47.3	61.9	-	-	7.9	31.1	12.2
EgoVideo-B	<b>52.7</b>	<b>65.3</b>	<b>49.8</b>	<b>74.6</b>	<b>8.1</b>	<b>34.7</b>	<b>14.7</b>
Helping Hands-L	-	-	-	-	7.9	33.4	16.0
LaViLa-L	50.9	66.5	51.0	76.0	7.3	32.5	13.4
EMBED-L	56.0	67.9	51.9	76.1	8.5	33.9	15.1
EgoVideo-L	<b>56.6</b>	<b>69.0</b>	<b>53.7</b>	<b>79.2</b>	<b>8.9</b>	<b>36.6</b>	<b>17.1</b>
EgoVideo-G	<b>60.3</b>	<b>70.0</b>	<b>56.0</b>	<b>80.0</b>	<b>10.0</b>	<b>38.7</b>	<b>19.6</b>

Hands (Zhang et al., 2023) and EMBED (Dou et al., 2024). Notably, despite the use of refined captions and significantly larger training datasets, LaViLa, AVION, and EMBED fail to achieve results as our EgoVideo. In the following experiments we will demonstrate that both our high-quality HOD data and our design of the EgoVideo model play important roles in achieving good performance. Helping Hands uses a stronger backbone TimeSformer (Bertasius et al., 2021) and adds additional decoders for auxiliary object-oriented tasks. However, our method can still outperform Helping Hands, demonstrating the superiority of our representation learning scheme.

Specifically, in the EK-100 MIR task, our EgoVideo outperforms EMBED by 0.5%, 1.0%, 6.3% in mAP and significantly outperforms the LaViLa at the same model size. In the EK-100 CLS task, our EgoVideo-B model demonstrates superior performance with a top-1 accuracy of 22.4% and a top-5 accuracy of 43.3%, significantly outperforming LaViLa-B and EMBED-B.

On the EGTEA dataset known for its focus on hand-object interactions, our EgoVideo-B achieves a mean accuracy of 43.6% and a top-1 accuracy of 51.0%, surpassing EMBED-B and even EMBED-L. This underscores the importance of learning hand object dynamics and shows the strong generalization capability of our model. The EgoMCQ task further highlights the efficacy of our method, with EgoVideo-B outperforming LaViLa-B by 4.7%, 1.2% and EMBED-B’s by 3.3% and 0.5% on the inter-class and intra-class accuracy, respectively. Our EgoVideo-L model also shows significant improvements with an inter-class accuracy of 65.5% and an intra-class accuracy of 95.9%. These results demonstrate the superior performance and generalization capability of our method without any additional supervision. We take a step forward to explore the scaling law in egocentric representation learning, finding that EgoVideo-G has elevated performance to the next level.

**Fine-tuning Evaluation.** Table 3 shows the result of the fine-tuning evaluation. Our EgoVideo method outperforms previous approaches across all tasks and datasets. Our EgoVideo-B demonstrates significant performance enhancements compared to EgoVLPv2-B, with improvements of 5.4% and 2.5% in mAP for the EK-100 MIR and EgoMCQ tasks, respectively. This performance is even comparable to the larger LaViLa-L. For our EgoVideo-L, we observe consistent improvements across all tasks, including a substantial enhancement by 1.8% and 3.1% in EK-100 CLS and EGTEA action recognition tasks, highlighting the superior performance of our model in fine-grained action understanding. Moreover, we achieve improvements of 0.4% in R1@0.5 in the EgoNLQ task and 2.7% and 2.0% in R1@0.5 and mAP in the EgoMQ task, confirming the richness of representations learned by our model and its capacity to capture intricate hand-object interaction information.

### 4.3 ABLATION STUDIES

**Pretraining Data.** We first conduct experiments by fixing the models and varying the pretraining data. Here we choose to use AVION for fair comparison since both AVION and EgoVideo use ViT as the backbone. As shown in Table 4, both our EgoVideo and AVION achieve the best performance when the combination of Ego4D-HOD data and How2-HOD data is used, and EgoVideo consistently outperforms AVION when trained on the same data, emphasizing the effectiveness of the model design. Comparing models trained with EgoClip and Ego4D-HOD (rows 1,2 and 5,6), it is clear that significant improvements can be observed in the EK-100 MIR and EGTEA tasks. Adding additional data from How2-HOD can improve both models substantially (rows 1,3 and 5,7). Furthermore, when



432 using only Ego4D-HOD, the performance on EGTEA surpasses EgoClip and How2-HOD together,  
 433 indicating the beneficial impact of our data on fine-grained dynamics understanding.  
 434

435 Table 4: Ablations on different pretrain datasets, include original EgoClip (Zhao & Krähenbühl,  
 436 2023), Ego4D-HOD and How2-HOD selected by our classifier from How2Interlink-7M.  
 437

ID	Model	Ego4D-	Ego4D-	How2-	EK-100 MIR		EGTEA	
		EgoClip	HOD	HOD	<i>mAP</i>	<i>nDCG</i>	<i>Mean-acc.</i>	<i>Top1-acc.</i>
1		✓			27.3	29.3	26.2	30.5
2	AVION-B		✓		31.0 <sub>(+3.7)</sub>	31.3 <sub>(+2.0)</sub>	32.3 <sub>(+6.1)</sub>	37.0 <sub>(+6.5)</sub>
3		✓		✓	33.2 <sub>(+5.9)</sub>	32.5 <sub>(+3.2)</sub>	31.6 <sub>(+5.4)</sub>	35.6 <sub>(+5.1)</sub>
4			✓	✓	<b>34.4<sub>(+7.1)</sub></b>	<b>33.7<sub>(+4.4)</sub></b>	<b>39.4<sub>(+13.2)</sub></b>	<b>46.4<sub>(+15.9)</sub></b>
5		✓			31.1	32.0	30.8	36.0
6	EgoVideo-B		✓		34.4 <sub>(+3.3)</sub>	33.9 <sub>(+1.9)</sub>	41.1 <sub>(+10.3)</sub>	47.9 <sub>(+11.9)</sub>
7		✓		✓	35.5 <sub>(+4.4)</sub>	34.1 <sub>(+2.1)</sub>	40.8 <sub>(+10.0)</sub>	47.1 <sub>(+11.1)</sub>
8			✓	✓	<b>36.5<sub>(+5.4)</sub></b>	<b>34.5<sub>(+2.5)</sub></b>	<b>43.6<sub>(+12.8)</sub></b>	<b>51.0<sub>(+15.0)</sub></b>

449  
 450 Table 5: Comparison of the number of pa-  
 451 rameters.  
 452

Method	Backbone	Params	EK100 mAP
LaViLa-B	TSF-B	121M	30.9
AVION-B	ViT-B	86M	32.9
EMBED-B	TSF-B	121M	36.0
EgoVideo-B	ViT-B	112M	<b>36.5</b>
LaViLa-L	TSF-L	438M	36.1
AVION-L	ViT-L	307M	37.6
EMBED-L	TSF-L	438M	40.8
EgoVideo-L	ViT-L	375M	<b>41.8</b>
EgoVideo-G	ViT-G	1050M	<b>47.1</b>

453  
 454 Table 6: The computational cost during infer-  
 455 ence. Views = #frames × #spatial crops ×  
 456 #temporal clips. “Extra GFLOPs” means extra  
 457 computation compared to ViT under the same  
 458 number of views.  
 459

Method	Views	GFLOPs	Extra GFLOPs
ViT-B	4 × 1 × 3	201	-
ViT-B	16 × 1 × 3	804	-
LaViLa-B	16 × 1 × 3	1432	
EgoVideo-B	16 × 1 × 3	1092	288
ViT-L	4 × 1 × 3	1047	-
ViT-L	16 × 1 × 3	4188	-
LaViLa-L	16 × 1 × 3	4956	
EgoVideo-L	16 × 1 × 3	5350	1162

467 **Model Size and Inference Computational Cost.** In Table 5, we compare the number of param-  
 468 eters of our model to other methods. Our EgoVideo model maintains a relatively small parameter  
 469 count, and even with the addition of the motion adapter, the total remains lower than that of LaViLa  
 470 and EMBED, highlighting the efficiency of our approach. Meanwhile, in Table 6 we compare the  
 471 inference computational cost of our EgoVideo with ViT and LaViLa. Thanks to our MotionAdapter,  
 472 the increase in inference time for our model compared to ViT at 16 frames, is only similar to ViT’s  
 473 inference time at 4 frames.

474 **Training Efficiency.** In Table 7, we compare the performance and computational speed of our  
 475 EgoVideo-B with AVION-B, where AVION-B is trained under two different parameter settings:  
 476 pre-training with 16 frames and pre-training with 4 frames. Our EgoVideo is trained in a mixed 16  
 477 and 4 frame fashion, thus being faster than directly using all 16 frames to train the whole backbone.  
 478 Meanwhile, EgoVideo achieves the best performance on the EK-100 MIR and EGTEA datasets.  
 479 These results strongly demonstrate the effectiveness of our training strategy and motion adapter  
 480 design in the EgoVideo model.

481 **Motion Adapter vs Other Adapters.** To verify the effectiveness of our motion adapter, we com-  
 482 pare it with the standard Adapter (Houlsby et al., 2019) and the ST-adapter (Pan et al., 2022). In  
 483 the standard adapter, we only use a downsample MLP and upsample MLP, while in the ST-adapter,  
 484 we perform convolution operations solely along the temporal dimension. As shown in Table 8, the  
 485 results on the EK-100 MIR task demonstrate that both the ST-adapter and our motion adapter out-  
 perform the standard adapter. This improvement can be attributed to the limited parameters of the

standard adapter, which restrict its ability to capture complex, fine-grained information. Compared to the ST-adapter, our Motion Adapter achieves the best performance by adding a spatial convolution operation, suggesting that both spatial and temporal information is crucial for egocentric video representation learning, especially for hand and object dynamics.

Table 7: Ablations on training strategy. Models are all trained on our Ego4D-HOD dataset with 10 epochs.

Method	GPU Hours	EK-100 MIR		EGTEA
		<i>mAP</i>	<i>nDCG</i>	<i>Top1-acc.</i>
AVION-4f	<b>95.5</b>	34.4	33.7	46.4
AVION-16f	395.5	36.2	34.3	47.4
EgoVideo	180.6	<b>36.5</b>	<b>34.5</b>	<b>51.0</b>

Table 8: Experiment with different adapters. Our motion adapter achieves the best performance with a small increase in parameters.

Design	Param Size	EK-100 MIR	
		<i>mAP</i>	<i>nDCG</i>
Adapter	8.28M	34.7	33.0
ST-adapter	10.08M	35.9	34.1
Motion Adapter	26.01M	<b>36.5</b>	<b>34.5</b>

#### 4.4 FEATURE-BASED EVALUATION ON OTHER TASKS

With the knowledge of hand-object dynamics, EgoVideo features can well generalize to other human behavior understanding tasks and robot manipulation tasks. Table 9 shows the results of Action Segmentation on the HOI4D (Liu et al., 2022) and GTEA (Fathi et al., 2011a) datasets, using features extracted from I3D (Carreira & Zisserman, 2017), AVION, and our EgoVideo. The results demonstrate our EgoVideo is also effective in the action segmentation task, especially for HOI4D which requires differentiating fine-grained hand-object interaction.

Also, we test the generalization capability of EgoVideo on the robot manipulation task on the Franka Kitchen dataset (Gupta et al., 2019). We follow the same setting and compare with previous robotic representation learning works MVP (Radosavovic et al., 2023), Voltron (Karamcheti et al., 2023) and MPI (Zeng et al., 2024). For MPI we compare both MPI with and without additional detection supervision. From Table 10, our EgoVideo can consistently surpass MVP and Voltron on the “Turn knob (TK)”, “Open Microwave (OM)” and “Open door (OD)” tasks. While MPI uses additional detection and prediction transformers and performs better than EgoVideo on two tasks, EgoVideo still performs comparably in overall success rate. Complete results with more details and analyses can be seen in Appendix E. The results strongly prove the delicacy and generalization ability of our EgoVideo learned representations.

Table 9: Experiment on the action segmentation task. We report results of ASFormer (Yi et al., 2021) with different input features.

Feature	HOI4D		GTEA	
	<i>F1@50</i>	<i>Edit</i>	<i>F1@50</i>	<i>Edit</i>
I3D	35.0	80.3	79.2	84.6
AVION	70.2	89.1	84.5	89.4
EgoVideo	<b>74.8</b>	<b>90.1</b>	<b>87.1</b>	<b>90.1</b>

Table 10: Results on Franka Kitchen. We report the success rate (%) on 50 sampled trajectories.

Method	TK	OM	OD	Avg.
MVP	79.0	41.0	48.0	56.0
Voltron	76.0	41.0	45.3	54.1
MPI	<b>85.5</b>	49.0	52.5	62.3
EgoVideo	80.1	<b>65.0</b>	<b>52.7</b>	<b>66.0</b>
MPI+Det	89.0	54.0	57.7	66.9

## 5 CONCLUSION

In this work, we inject fine-grained hand-object dynamics into egocentric video representation learning. Our method addresses the drawbacks of the existing method from two perspectives. On the data side, we propose **HOD**, a novel framework to generate new paired video-language data, where the language contains intricately depicted hand-object dynamics. On the model side, we propose **EgoVideo**, where we use a model with a motion adapter combined with a co-training technique, to fully exploit the fine-grained dynamics provided by HOD data in the representation learning process. Experimental results demonstrate that our method achieves state-of-the-art performance across multiple downstream tasks, and can generalize in the embodied manipulation environment.

## REFERENCES

01. AI, :, Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, Kaidong Yu, Peng Liu, Qiang Liu, Shawn Yue, Senbin Yang, Shiming Yang, Tao Yu, Wen Xie, Wenhao Huang, Xiaohui Hu, Xiaoyi Ren, Xinyao Niu, Pengcheng Nie, Yuchi Xu, Yudong Liu, Yue Wang, Yuxuan Cai, Zhenyu Gu, Zhiyuan Liu, and Zonghong Dai. Yi: Open foundation models by 01.ai, 2024.
- Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, volume 2, pp. 4, 2021.
- Samarth Brahmabhatt, Chengcheng Tang, Christopher D Twigg, Charles C Kemp, and James Hays. Contactpose: A dataset of grasps with object contact and hand pose. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIII 16*, pp. 361–378. Springer, 2020.
- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 961–970, 2015.
- Yujun Cai, Lihao Ge, Jianfei Cai, and Junsong Yuan. Weakly-supervised 3d hand pose estimation from monocular rgb images. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 666–682, 2018.
- Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12417–12426, 2021.
- Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6299–6308, 2017.
- Guo Chen, Yifei Huang, Jilan Xu, Baoqi Pei, Zhe Chen, Zhiqi Li, Jiahao Wang, Kunchang Li, Tong Lu, and Limin Wang. Video mamba suite: State space model as a versatile alternative for video understanding. *arXiv preprint arXiv:2403.09626*, 2024.
- Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. Scaling egocentric vision: The epic-kitchens dataset. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 720–736, 2018.
- Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, et al. The epic-kitchens dataset: Collection, challenges and baselines. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4125–4141, 2020.
- Ning Ding, Yujia Qin, Guang Yang, Fuchao Wei, Zonghan Yang, Yusheng Su, Shengding Hu, Yulin Chen, Chi-Min Chan, Weize Chen, et al. Parameter-efficient fine-tuning of large-scale pre-trained language models. *Nature Machine Intelligence*, 5(3):220–235, 2023.
- Bardia Doosti, Shujon Naha, Majid Mirbagheri, and David J Crandall. Hope-net: A graph-based model for hand-object pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6608–6617, 2020.
- Zi-Yi Dou, Xitong Yang, Tushar Nagarajan, Huiyu Wang, Jing Huang, Nanyun Peng, Kris Kitani, and Fu-Jen Chu. Unlocking exocentric video-language data for egocentric video representation learning. *arXiv preprint arXiv:2408.03567*, 2024.
- Yazan Abu Farha and Jurgen Gall. Ms-tcn: Multi-stage temporal convolutional network for action segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3575–3584, 2019.
- Alireza Fathi, Ali Farhadi, and James M Rehg. Understanding egocentric activities. In *2011 international conference on computer vision*, pp. 407–414. IEEE, 2011a.

- 594 Alireza Fathi, Xiaofeng Ren, and James M Rehg. Learning to recognize objects in egocentric activ-  
595 ities. In *CVPR 2011*, pp. 3281–3288. IEEE, 2011b.
- 596
- 597 Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video  
598 recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp.  
599 6202–6211, 2019.
- 600
- 601 Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The  
602 kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- 603
- 604 Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *Proceedings of the  
605 IEEE/CVF international conference on computer vision*, pp. 13505–13515, 2021.
- 606
- 607 Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Gird-  
608 har, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in  
609 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision  
and Pattern Recognition*, pp. 18995–19012, 2022.
- 610
- 611 Abhishek Gupta, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy  
612 learning: Solving long-horizon tasks via imitation and reinforcement learning. *arXiv preprint  
arXiv:1910.11956*, 2019.
- 613
- 614 Yana Hasson, Gul Varol, Dimitrios Tzionas, Igor Kalevatykh, Michael J Black, Ivan Laptev, and  
615 Cordelia Schmid. Learning joint reconstruction of hands and manipulated objects. In *Proceedings  
of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11807–11816, 2019.
- 616
- 617 Yana Hasson, Bugra Tekin, Federica Bogo, Ivan Laptev, Marc Pollefeys, and Cordelia Schmid.  
618 Leveraging photometric consistency over time for sparsely supervised hand-object reconstruction.  
619 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 571–  
620 580, 2020.
- 621
- 622 Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, An-  
623 drea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp.  
624 In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.
- 625
- 626 Yifei Huang, Minjie Cai, Zhenqiang Li, and Yoichi Sato. Predicting gaze in egocentric video by  
627 learning task-dependent attention transition. In *European Conference on Computer Vision*, 2018.
- 628
- 629 Yifei Huang, Guo Chen, Jilan Xu, Mingfang Zhang, Lijin Yang, Baoqi Pei, Hongjie Zhang,  
630 Lu Dong, Yali Wang, Limin Wang, et al. Egoxolearn: A dataset for bridging asynchronous  
631 ego-and exo-centric view of procedural activities in real world. In *Proceedings of the IEEE/CVF  
Conference on Computer Vision and Pattern Recognition*, pp. 22072–22086, 2024.
- 632
- 633 Md Mohaiminul Islam, Ngan Ho, Xitong Yang, Tushar Nagarajan, Lorenzo Torresani, and Gedas  
634 Bertasius. Video recap: Recursive captioning of hour-long videos. In *Proceedings of the  
IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18198–18208, 2024.
- 635
- 636 Hanwen Jiang, Shaowei Liu, Jiashun Wang, and Xiaolong Wang. Hand-object contact consistency  
637 reasoning for human grasps generation. In *Proceedings of the IEEE/CVF international conference  
on computer vision*, pp. 11107–11116, 2021.
- 638
- 639 Siddharth Karamcheti, Suraj Nair, Annie S Chen, Thomas Kollar, Chelsea Finn, Dorsa Sadigh,  
640 and Percy Liang. Language-driven representation learning for robotics. *arXiv preprint  
arXiv:2302.12766*, 2023.
- 641
- 642 Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijaya-  
643 narasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action  
644 video dataset. *arXiv preprint arXiv:1705.06950*, 2017.
- 645
- 646 Shijie Li, Yazan Abu Farha, Yun Liu, Ming-Ming Cheng, and Juergen Gall. Ms-tcn++: Multi-stage  
647 temporal convolutional network for action segmentation. *IEEE transactions on pattern analysis  
and machine intelligence*, 45(6):6647–6658, 2020.

- 648 Yanghao Li, Tushar Nagarajan, Bo Xiong, and Kristen Grauman. Ego-exo: Transferring visual rep-  
649 resentations from third-person to first-person videos. In *Proceedings of the IEEE/CVF Conference*  
650 *on Computer Vision and Pattern Recognition*, pp. 6943–6953, 2021.
- 651 Yin Li, Miao Liu, and James M Rehg. In the eye of beholder: Joint learning of gaze and actions in  
652 first person video. In *Proceedings of the European conference on computer vision (ECCV)*, pp.  
653 619–635, 2018.
- 654 Kevin Qinghong Lin, Jinpeng Wang, Mattia Soldan, Michael Wray, Rui Yan, Eric Z Xu, Difei  
655 Gao, Rong-Cheng Tu, Wenzhe Zhao, Weijie Kong, et al. Egocentric video-language pretraining.  
656 *Advances in Neural Information Processing Systems*, 35:7575–7586, 2022.
- 657 Shaowei Liu, Hanwen Jiang, Jiarui Xu, Sifei Liu, and Xiaolong Wang. Semi-supervised 3d hand-  
658 object poses estimation with interactions in time. In *Proceedings of the IEEE/CVF Conference*  
659 *on Computer Vision and Pattern Recognition*, pp. 14687–14697, 2021.
- 660 Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu,  
661 He Wang, and Li Yi. Hoi4d: A 4d egocentric dataset for category-level human-object interaction.  
662 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
663 21013–21022, 2022.
- 664 Mi Luo, Zihui Xue, Alex Dimakis, and Kristen Grauman. Put myself in your shoes: Lifting the  
665 egocentric perspective from exocentric videos. *arXiv preprint arXiv:2403.06351*, 2024.
- 666 Antoine Miech, Dimitri Zhukov, Jean-Baptiste Alayrac, Makarand Tapaswi, Ivan Laptev, and Josef  
667 Sivic. Howto100m: Learning a text-video embedding by watching hundred million narrated  
668 video clips. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp.  
669 2630–2640, 2019.
- 670 Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A univer-  
671 sal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022.
- 672 Takehiko Ohkawa, Kun He, Fadime Sener, Tomas Hodan, Luan Tran, and Cem Keskin. Assembly-  
673 hands: Towards egocentric activity understanding via 3d hand pose estimation. In *Proceedings of*  
674 *the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12999–13008, 2023.
- 675 Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predic-  
676 tive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- 677 R OpenAI. Gpt-4 technical report. arxiv 2303.08774. *View in Article*, 2(5), 2023.
- 678 Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient  
679 image-to-video transfer learning. *Advances in Neural Information Processing Systems*, 35:26462–  
680 26477, 2022.
- 681 Xiaqing Pan, Nicholas Charron, Yongqian Yang, Scott Peters, Thomas Whelan, Chen Kong, Omkar  
682 Parkhi, Richard Newcombe, and Yuheng Carl Ren. Aria digital twin: A new benchmark dataset  
683 for egocentric 3d machine perception. In *Proceedings of the IEEE/CVF International Conference*  
684 *on Computer Vision*, pp. 20133–20143, 2023.
- 685 Chiara Plizzari, Mirco Planamente, Gabriele Goletto, Marco Cannici, Emanuele Gusso, Matteo  
686 Matteucci, and Barbara Caputo. E2 (go) motion: Motion augmented event stream for egocentric  
687 action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern*  
688 *recognition*, pp. 19935–19947, 2022.
- 689 Chiara Plizzari, Gabriele Goletto, Antonino Furnari, Siddhant Bansal, Francesco Ragusa, Gio-  
690 vanni Maria Farinella, Dima Damen, and Tatiana Tommasi. An outlook into the future of egocen-  
691 tric vision. *International Journal of Computer Vision*, pp. 1–57, 2024.
- 692 Shraman Pramanick, Yale Song, Sayan Nag, Kevin Qinghong Lin, Hardik Shah, Mike Zheng Shou,  
693 Rama Chellappa, and Pengchuan Zhang. Egovlpv2: Egocentric video-language pre-training with  
694 fusion in the backbone. In *Proceedings of the IEEE/CVF International Conference on Computer*  
695 *Vision*, pp. 5285–5297, 2023.

- 702 Gorjan Radevski, Dusan Grujicic, Matthew Blaschko, Marie-Francine Moens, and Tinne Tuyte-  
703 laars. Multimodal distillation for egocentric action recognition. In *Proceedings of the IEEE/CVF*  
704 *International Conference on Computer Vision*, pp. 5213–5224, 2023.
- 705 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language  
706 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 707 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
708 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
709 models from natural language supervision. In *International conference on machine learning*, pp.  
710 8748–8763. PMLR, 2021.
- 711 Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell.  
712 Real-world robot learning with masked visual pre-training. In *Conference on Robot Learning*, pp.  
713 416–426. PMLR, 2023.
- 714 Snehanshu Saha, Archana Mathur, Aditya Pandey, and Harshith Arun Kumar. Diffact: A unifying  
715 framework for activation functions. In *2021 International Joint Conference on Neural Networks*  
716 *(IJCNN)*, pp. 1–8. IEEE, 2021.
- 717 Rico Sennrich. Neural machine translation of rare words with subword units. *arXiv preprint*  
718 *arXiv:1508.07909*, 2015.
- 719 Dandan Shan, Jiaqi Geng, Michelle Shu, and David F Fouhey. Understanding human hands in  
720 contact at internet scale. In *Proceedings of the IEEE/CVF conference on computer vision and*  
721 *pattern recognition*, pp. 9869–9878, 2020.
- 722 Sanjana Srivastava, Chengshu Li, Michael Lingelbach, Roberto Martín-Martín, Fei Xia, Kent Elliott  
723 Vainio, Zheng Lian, Cem Gokmen, Shyamal Buch, Karen Liu, et al. Behavior: Benchmark for  
724 everyday household activities in virtual, interactive, and ecological environments. In *Conference*  
725 *on robot learning*, pp. 477–490. PMLR, 2022.
- 726 Alex Jinpeng Wang, Linjie Li, Kevin Qinghong Lin, Jianfeng Wang, Kevin Lin, Zhengyuan Yang,  
727 Lijuan Wang, and Mike Zheng Shou. Cosmo: Contrastive streamlined multimodal model with  
728 interleaved pre-training. *arXiv preprint arXiv:2401.00849*, 2024a.
- 729 Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yanan He, Guo Chen, Baoqi Pei, Rongkun Zheng,  
730 Jilan Xu, Zun Wang, et al. Internvideo2: Scaling video foundation models for multimodal video  
731 understanding. *arXiv preprint arXiv:2403.15377*, 2024b.
- 732 Zhen Xing, Qi Dai, Zejia Weng, Zuxuan Wu, and Yu-Gang Jiang. Aid: Adapting image2video  
733 diffusion models for instruction-guided video prediction. *arXiv preprint arXiv:2406.06465*, 2024.
- 734 Jilan Xu, Yifei Huang, Junlin Hou, Guo Chen, Yuejie Zhang, Rui Feng, and Weidi Xie. Retrieval-  
735 augmented egocentric video captioning. In *Proceedings of the IEEE/CVF Conference on Com-*  
736 *puter Vision and Pattern Recognition*, pp. 13525–13536, 2024.
- 737 Hongwei Xue, Tiankai Hang, Yanhong Zeng, Yuchong Sun, Bei Liu, Huan Yang, Jianlong Fu, and  
738 Baining Guo. Advancing high-resolution video-language representation with large-scale video  
739 transcriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
740 *Recognition*, pp. 5036–5045, 2022.
- 741 Linlin Yang and Angela Yao. Disentangling latent hands for image synthesis and pose estimation.  
742 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.  
743 9877–9886, 2019.
- 744 Fangqiu Yi, Hongyu Wen, and Tingting Jiang. Asformer: Transformer for action segmentation.  
745 *arXiv preprint arXiv:2110.08568*, 2021.
- 746 Shanxin Yuan, Guillermo Garcia-Hernando, Björn Stenger, Gyeongsik Moon, Ju Yong Chang, Ky-  
747 oung Mu Lee, Pavlo Molchanov, Jan Kautz, Sina Honari, Lihao Ge, et al. Depth-based 3d  
748 hand pose estimation: From current achievements to future goals. In *Proceedings of the IEEE*  
749 *conference on computer vision and pattern recognition*, pp. 2636–2645, 2018.

Jia Zeng, Qingwen Bu, Bangjun Wang, Wenke Xia, Li Chen, Hao Dong, Haoming Song, Dong Wang, Di Hu, Ping Luo, et al. Learning manipulation by predicting interaction. *arXiv preprint arXiv:2406.00439*, 2024.

Chuhan Zhang, Ankush Gupta, and Andrew Zisserman. Helping hands: An object-aware egocentric video recognition model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13901–13912, 2023.

Hao Zhang, Aixin Sun, Wei Jing, and Joey Tianyi Zhou. Span-based localizing network for natural language video localization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 6543–6554, Online, July 2020. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.acl-main.585>.

Lingzhi Zhang, Shenghao Zhou, Simon Stent, and Jianbo Shi. Fine-grained egocentric hand-object segmentation: Dataset, model, and applications. In *European Conference on Computer Vision*, pp. 127–145. Springer, 2022.

Chen Zhao, Ali K Thabet, and Bernard Ghanem. Video self-stitching graph network for temporal action localization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13658–13667, 2021.

Yue Zhao and Philipp Krähenbühl. Training a large video model on a single machine in a day. *arXiv preprint arXiv:2309.16669*, 2023.

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*, pp. 6586–6597, 2023.

## A DETAILS ABOUT HOD



Figure 5: Examples of egocentric video/ego-like video and non-ego video.

**Data Selection** Since our HOD involves data from not only Ego4D but also How2link-7M, we use a style classifier  $\mathcal{P}$  to filter egocentric style videos from the How2link-7M dataset. Specifically, our style classifier employs a simple two-layer MLP architecture. We utilize InternVideo2 (Wang et al., 2024b) to extract video features from all videos of the How2-Interlink7M dataset. After that, we manually annotate 10,000 clips with positive and negative labels, where the positive label indicates this video is an egocentric video (or ego-like video). Examples of positive and negative labeled videos can be found in Figure 5. We randomly select 10% of these clips to form the validation set. After training our classifier on the train set, we get 89% accuracy on the validation set.

**HOD Rephraser** We use Yi-34B model to generate hand object dynamics narrations. The Yi-34B model is trained on a corpus of over 150,000 high-quality texts and its model weights are open-source, which has a high ranking among all existing open-source Large Language Models. We directly use the model without finetuning.

To generate reliable narrations, we need to convert the obtained hand-object information into appropriate texts. For the movement trajectories of hands and objects, we directly calculate the center points of the bounding boxes and perform normalization to get a sequence  $L =$



810  $((w_0, h_0), (w_1, h_1), \dots, (w_{15}, h_{15}))$ . To determine whether the object is contacted by the left hand or  
 811 right hand separately, or contacted by both hands simultaneously, we use a generalized IoU function  
 812 on the left-contact object and right-contact object. For IoU value greater than 0.9, we classify the  
 813 object as being contacted by both hands.

814 Subsequently, we prompt the LLM with a system prompt of:

```
815 ## System Prompt
816 Now you are a captioning assistant, you need to generate hand object interaction
817 caption and combine them with the origin narration. Given the origin narration
818 of the video clip and spatial localization ([x, y]) of hands and objects in the
819 clip, please help me describe the direction of motion of the left and right hands,
820 their relative relationship to objects and whether they are touching or not. Do not
821 mention the pixel info. Two_hand.object means objects with two hands in contact,
822 left_hand.object means objects with left hand in contact, right_hand.object means
823 objects with right hand in contact.
824
825
826 ## Hand Object Dynamics
827 left_hand: ((w_0, h_0), (w_1, h_1), ..., (w_15, h_15))
828 right_hand: ((w_0, h_0), (w_1, h_1), ..., (w_15, h_15))
829 left_hand.object: ((w_0, h_0), (w_1, h_1), ..., (w_15, h_15))
830 right_hand.object: ((w_0, h_0), (w_1, h_1), ..., (w_15, h_15))
831 two_hand.object: ((w_0, h_0), (w_1, h_1), ..., (w_15, h_15))
832 origin narration: C takes a scissors.
```

```
834 ## System Prompt
835 Please help me summarize the direction of movement of the left hand, right hand,
836 and objects, and generate a new caption based on the original caption. It is
837 strictly forbidden to mention the frame number and spatial position coordinates in
838 the description.
```

839 For the computational cost, it takes around 2 days to extract bounding boxes from all vision-language  
 840 clips and 3 days to generate narrations with LLM using 32 A100 GPUs, resulting in a total of around  
 841 4000 GPU-hours.

842 **Data Evaluation.** Here, we provide a detailed explanation of our evaluation process. First, we  
 843 prompted the LLM to generate new narrations with different verbs/nouns using the prompt:

```
845 Please help me modify the key verbs and nouns in this sentence to slightly alter
846 its meaning while keeping the sentence structure largely unchanged. Just return
847 the modified sentence to me. Ensure the semantic shift is minimal, such as
848 changing one or two verbs and nouns.
```

849 Then we utilized GPT-4o as a judge to determine the quality of the narrations with the prompt:

```
851 You are a judge. There are 16 frames in the video, I have three captions and
852 need your help to score the three captions based on three criteria: relevance,
853 accuracy, and level of detail. The score ranges from 0 to 10, with a higher score
854 indicating better quality of the caption. You can just answer me in the following
855 format: First: score1, Second: score2, Third: score3. First caption: text1
856 Second caption: text2 Third caption: text3
```

857 As mentioned in the main manuscript, to further validate the quality of our HOD dataset, we utilize  
 858 two standard unsupervised automatic metrics to evaluate the quality of narrations. We use the human  
 859 narration as the ground truth and compare our HOD data with LaViLa-Narrator on METEOR and  
 860 CIDEr scores. The results in Table 11 reveal that while our HOD data achieves a slightly lower  
 861 METEOR score, it outperforms LaViLa-Narrator in CIDEr. This discrepancy arises because many  
 862 LaViLa narrations closely mirror the original text, whereas our narrations incorporate additional  
 863 dynamic information. Although our performance does not drastically exceed that of LaViLa, the  
 results demonstrate that our narrations successfully retain the original semantic content.

Table 11: Comparison with LaViLa-Narrator on the narration quality.

Text	METEOR	CIDEr
LaViLa-Narrator	0.45	0.34
HOD	0.39	0.40

**Limitations and future work.** Our model relies on the quality of hand-object detection and the rephrasing of LLM, which may include error accumulation. In addition to reducing the error in the data construction, exploring how the hand-object dynamics can be better involved into language or in other formats is a promising direction for our future work.

## B DATASET DETAILS

**Ego4D** Ego4D (Grauman et al., 2022) contains 3,670 hours of egocentric videos with temporally dense narrations. Each narration has a timestamp and an associated free-form sentence. We follow previous works ((Zhao et al., 2023),(Lin et al., 2022)) to prepare the Ego4D dataset for vision-language pretraining. Specifically, we drop the narrations that either contain “#unsure”/“#Unsure” tags or are shorter than 4 words. This results in 4M video-text clip pairs.

**Howto-Interlink7M** Howto-Interlink7M (Wang et al., 2024a) contains 1M videos and 7M clips, which is part of the broader Howto100M dataset. Diverging from the original dataset, clips in Howto-Interlink7M have concise descriptions, and dense region captions and leverage GPT-4 to generate comprehensive summaries from detailed annotation. We use a classifier to select 3.3M vision-text pairs from the dataset.

**EpicKitchens-100** The Epic-Kitchens-100 (EK-100) dataset (Damen et al., 2020; 2018) contains 100 hours of egocentric cooking videos. Each clip is annotated with a start and end timestamp, a short textual narration, and a verb and noun class that the narration belongs to. The action class can also be uniquely determined by combining the verb and the noun. In EpicKitchens-MIR, we use Mean Average Precision and normalized Discounted Cumulative Gain (nDCG) as evaluation metrics. In EpicKitchens-CLS, we use top-1 action accuracy and top-5 action accuracy as evaluation metrics.

**EGTEA** EGTEA (Li et al., 2018) contains 28 hours of cooking activities from 86 unique sessions of 32 subjects. In zero-shot evaluation, we compute the similarity score between every video embedding and the 106 text embeddings, and take the text embedding with the highest similarity score as the predicted class. In fine-tuning evaluation, we finetune the video encoder for action classification. using the linear probing protocol.

**GTEA** The Georgia Tech Egocentric Activities (GTEA) dataset (Fathi et al., 2011b) consists of seven distinct types of everyday activities, including making a sandwich, preparing tea, and brewing coffee. Each of these activities is demonstrated by four different individuals, resulting in a total of 28 unique video recordings. Each video captures around 20 fine-grained action instances, such as “take bread” or “pour ketchup,” all occurring within approximately one minute. This dataset provides a comprehensive look at egocentric perspectives, making it an invaluable resource for research in activity recognition and human-computer interaction.

**HOI4D** The HOI4D dataset (Liu et al., 2022) represents a significant advancement in the study of category-level human-object interaction, offering a large-scale 4D egocentric resource enriched with detailed annotations. Comprising 2.4 million RGB-D egocentric video frames across more than 4,000 sequences, the dataset captures interactions performed by nine participants with 800 unique object instances spanning 16 categories within 610 diverse indoor environments. To foster advancements in category-level human-object interaction, HOI4D introduces three benchmarking tasks: semantic segmentation of 4D dynamic point cloud sequences, category-level object pose tracking, and egocentric action segmentation involving a variety of interaction targets.

**Franka Kitchen** The Franka Kitchen dataset (Gupta et al., 2019) is a comprehensive resource designed to facilitate research in robotic manipulation and human-robot interaction within a kitchen environment. This dataset comprises a diverse collection of videos showcasing a humanoid robot, Franka Emika Panda, performing various cooking tasks. The setup features a Franka robot with 9

degrees of freedom positioned within a kitchen environment equipped with various common household items, including a microwave, kettle, overhead light, cabinets, and an oven. This environment is designed for multitask objectives, requiring the robot to interact with these items to achieve specific goal configurations.

## C IMPLEMENTATION DETAILS

**Pretraining Details** We pre-train on the video-narration pairs generated by our HOD from Ego4D and How-InterLink7M. We use AdamW optimizer with betas = (0.9,0.999) for 15 epochs. We use different settings for different size models. For EgoVideo-B, we adopt a batch size of 128 over 16 GPUs with a fixed learning rate of 5e-5, For EgoVideo-L, we use a batch size of 32 over 16 GPUs with a fixed learning rate of 3e-5. For EgoVideo-G, we choose to use a batch size of 16 over 16 GPUs with a fixed learning rate of 1e-5. For input frames, we preprocess the frames by resizing the shorter side to 320 pixels, which accelerates the data loading speed. Subsequently, we applied a standard RandomResizedCrop function (Zhao & Krähenbühl, 2023) with a scale parameter of (0.5, 1.0) to obtain the corresponding input frames.

**Finetuning Details** We finetune the downstream tasks using AdamW with  $(\beta_1, \beta_2) = (0.9, 0.999)$  and weight decay of 0.05 with cosine annealing. Table 12 shows the hyperparameters details and in all tasks we use 8 GPUs for finetuning. During the training phase, we resize the shorter side of the video to 256 pixels and subsequently extract a 224×224 crop. During the testing phase, we scale the shorter side to 224 pixels and take a central 224×224 crop.

Table 12: Hyperparameters for Different Downstream Tasks

Task	Model Size	Epochs	LR_start	LR_end	Batch Size
EK100-MIR	EgoVideo-B	100	1e-6	1e-5	256
	EgoVideo-L	70	5e-7	5e-6	64
	EgoVideo-G	50	4e-7	4e-6	32
EK100-CLS	EgoVideo-B	100	1e-6	1e-5	256
	EgoVideo-L	70	5e-7	5e-6	64
	EgoVideo-G	60	4e-7	4e-6	32
EGTEA	EgoVideo-B	100	1e-6	1e-5	256
	EgoVideo-L	70	7e-7	7e-6	64
	EgoVideo-G	50	4e-7	4e-6	32

For the EgoNLQ task (Grauman et al., 2022), we build on the methodologies introduced by EgoVLP (Lin et al., 2022) and LAVILA (Zhao et al., 2023) for fairness. We adopt VSLNet (Zhang et al., 2020) as the task head. We train the task head for 50 epochs, using a learning rate of 3e-3, dropout 0.3, batch size 32 on a single A100 GPU. For the EgoMQ task, we use VSGN (Zhao et al., 2021) as our task head for training. We set batch size as 16, learning rate as 2e-4, gamma as 0.6, and train the task head on a single A100 GPU.

## D ADDITIONAL ABLATIONS


**Pretraining Data** To further demonstrate the effectiveness of our HOD, we fix the amount of data and conduct experiments using different model sizes. As shown in Table 13, with the same model size and the same data size, using our Ego4D-HOD data can consistently achieve improvement. Since one sample in EgoClip corresponds strictly to one sample in Ego4D-HOD, this table strongly demonstrates the high quality of our HOD data.

**Adapter Downsampling Ratio** We test the design of our motion adapter by studying the effect of adapter downsampling ratio  $\gamma$ , and show the result in Figure 6. It can be observed that as the value of  $\gamma$  increases, the model’s performance continues to improve. This indicates that our generated

Table 13: Comparison of performance across different model sizes and vision-language datasets.

Model	Pretrain Data	Data Size	EK-100 MIR	
			mAP	nDCG
EgoVideo-B	EgoClip	4M	31.1	32.0
EgoVideo-B	Ego4D-HOD	4M	34.4	33.9
EgoVideo-L	EgoClip	4M	35.3	34.6
EgoVideo-L	Ego4D-HOD	4M	38.3	35.9
EgoVideo-G	EgoClip	4M	42.1	37.5
EgoVideo-G	Ego4D-HOD	4M	44.8	38.2

Frame count	$\lambda$	EK MIR mAP
4	1	34.2
8	2	35.3
12	3	35.9
16	4	<b>36.5</b>
32	8	<b>36.5</b>

Table 14: Comparison of performance across different frame up-sampling rate  $\lambda$ .Figure 6: Ablations on Downsampling Ratio  $\gamma$ .

narrations contain rich semantic information and further validates the effectiveness of our motion adapter. To reduce computational overhead, we ultimately decide to set  $\gamma = 0.5$ .

**Frame number** We further study the effect of the number of sampled frames as input. We consistently use 4 frames as the sampling rate for the backbone part. The results in Table 14 indicate that as the number of frames increases from 4 to 16, the model’s performance improves continuously from 34.2% to 36.5%. However, when the frame count reaches to 32, the performance plateaus, showing no significant improvement with further increases in frame count. Besides, increasing the number of frames beyond this point incurs substantial computational cost. As a result, we choose to use  $\lambda = 4$  as the default value in our EgoVideo, balancing the speed and accuracy.

## E ADDITIONAL RESULTS

Table 15: Comparison with the state-of-the-art methods on 50Salads, GTEA and HOI4D dataset.

Feature	Method	GTEA			HOI4D		
		<i>F1@10, 25, 50</i>	<i>Edit</i>	<i>Acc</i>	<i>F1@10, 25, 50</i>	<i>Edit</i>	<i>Acc</i>
I3D	MS-TCN	85.8 / 83.4 / 69.8	79.0	76.3	55.6 / 47.8 / 31.8	74.7	44.2
I3D	MS-TCN++	88.8 / 85.7 / 76.0	83.5	80.1	54.7 / 46.5 / 30.3	75.2	42.2
I3D	ASFormer	90.1 / 88.8 / 79.2	84.6	79.7	-	-	-
I3D	DiffAct	92.5 / 91.5 / 84.7	89.6	82.2	-	-	-
AVION	ASFormer	92.5 / 91.0 / 84.5	89.4	81.4	84.4 / 81.1 / 70.2	89.2	74.2
EgoVideo	ASFormer	<b>92.7 / 92.2 / 87.1</b>	<b>90.1</b>	<b>82.7</b>	<b>88.9 / 85.3 / 74.8</b>	<b>90.1</b>	<b>76.2</b>

### E.1 DETAILS ON ACTION SEGMENTATION TASKS.

We evaluate our model on two benchmark datasets: GTEA (Fathi et al., 2011b), and HOI4D (Liu et al., 2022). We follow the previous work to use four-fold cross-validations on both datasets. We use accuracy (Acc), the edit distance (Edit), and the F1 scores at overlap thresholds 10%, 25%, 50% (F1@10, 25, 50) as metrics for evaluation.


We use ASFormer (Yi et al., 2021) as the task head, with input features extracted by our EgoVideo, I3D (Carreira & Zisserman, 2017), and AVION (Zhao & Krähenbühl, 2023). We follow (Chen et al., 2024), using learning rate =  $5e-4$ , drop rate = 0.3, epoch = 100 for training. Table 15 presents the experimental results of our method and other recent approaches, including MS-TCN (Farha & Gall, 2019), MS-TCN++ (Li et al., 2020), ASFormer (Yi et al., 2021), and DiffAct (Saha et al., 2021). The results clearly show the high quality of our EgoVideo feature. With the same task head ASFormer, the EgoVideo feature can outperform the AVION feature consistently. EgoVideo can even help ASFormer to beat the stronger task head DiffAct.



Turn on the knob



Open microwave



Open door

Figure 7: Qualitative results on the Franka Kitchen dataset. We show the tasks of turning the stove top knob, opening the microwave and opening the door. All tasks are trained with 25 demonstrations.

### E.2 DETAILS ON FRANKA KITCHEN DATASET.

Here we introduce the details of experiments on the Franka Kitchen dataset (Gupta et al., 2019). In this dataset, we adopt 3 tasks, including “Turn the stove top knob (TK)”, “Open the microwave (OM)” and “Open the left door (OD)”. The goal is to predict 9-DoF joint velocities (7 joints and 2 grippers) based on the visual representations and proprioceptive states (*i.e.*, joint velocities). We follow the MPI mode (Zeng et al., 2024), which trains a shallow MLP policy network. For evaluation, we follow the R3M method (Nair et al., 2022) and (Karamcheti et al., 2023) and calculate the average success rates for each setting across the 3 tasks.

We compare our EgoVideo with MVP (Radosavovic et al., 2023), Voltron (Karamcheti et al., 2023) and MPI (Zeng et al., 2024). MVP learns representation for robot manipulation by masked image modeling. Voltron takes a step forward to combine masked image modeling with vision text align-

1080 ment. MPI designs detection and prediction transformers to use object detection signals as additional  
1081 guidance. Notably, these works also use Ego4D as training data.

1082 The experimental results in Table 10 indicate that our model significantly outperforms both the  
1083 MVP and Voltron models by more than 10% in average success rate, and exhibits performance  
1084 comparable to the more advanced MPI model, which integrates multiple pre-training tasks related to  
1085 robot learning. When the MPI model is solely trained using contrastive learning and masked signal  
1086 modeling as supervision, we achieve 3.7% improvements in average success rate than MPI model.  
1087 When MPI incorporates the video prediction task, which has been proven crucial for robot learning,  
1088 our average success rate is only 0.9% lower. This demonstrates the robust generalization capabilities  
1089 of our model and highlights the contribution of our hand-object dynamics learning scheme to fine-  
1090 grained hand operations. Figure 7 shows the qualitative results on turning on the knob, and opening  
1091 the microwave and opening the door tasks.

## 1092 F QUALITATIVE RESULTS

1093 In Figure 8, we show more examples to compare narrations generated by our HOD with LaViLa  
1094 Rephraser and the original narrations. We can observe that the narrations generated by our HOD  
1095 model can well describe the hand-object dynamics (*e.g.*, ‘The left hand moves downwards to touch  
1096 the bicycle tire’). Moreover, compared to the LaViLa rephraser, which often merely changes word  
1097 order or modifies nouns/verbs, our model can combine original actions to generate more semanti-  
1098 cally rich descriptions of actions and scenes, resulting in significantly higher quality narrations. (See  
1099 the first example: our HOD generates ‘Person C picks a card with their right hand, which is then  
1100 handed to their left hand.’ while LaViLa yields ‘#C C chooses a card/#C C selects a card/#C C picks  
1101 a card’).



1134

1135

1136

1137

1138



ORIGIN: #C C picks a card.

LAVILA REPHRASER:

#C C picks a card. #C C chooses a card.

#C C selects a card.

#C C chooses a card from the deck.

Our HOD:

Person C picks a card with their right hand, which is then handed to their left hand.

1139

1140

1141

1142

1143

1144

1145

1146



ORIGIN: #c c presses piano while tapping table.

LAVILA REPHRASER:

#c c presses piano while tapping table.

When tapping the table c presses piano.

Our HOD:

The left hand moves up and down on the piano while the right hand taps on the table.

1147

1148

1149

1150

1151

1152

1153



ORIGIN: #C C touches the bicycle tire.

LAVILA REPHRASER:

#C C is referring to C bicycle tire.

#C C addresses the bicycle tire.

#C C mentions the bicycle tire

Our HOD:

The left hand moves downwards to touch the bicycle tire, while the right hand remains still.

1154

1155

1156

1157

1158

1159

1160

1161



ORIGIN: #C C carries mold.

LAVILA REPHRASER:

#C C is the product of mold.

#C C is a fan of the mold.

Our HOD:

A person is carrying a mold with their right hand, while their left hand is empty.

1162

1163

1164


1165

1166

1167

1168

1169



ORIGIN:

#C C gathers the short bamboo with both hands.

LAVILA REPHRASER:

#C C gathers the short bamboo with both hands.

#C C holds the short bamboo with both hands.

Our HOD:

The left hand moves from the left to the top center, gathering short bamboo. The right hand assists from the right, helping to collect the bamboo. Together, the hands work to gather the bamboo into a compact bundle.

1170

1171

1172

1173


1174

1175

1176

1177

1178



ORIGIN: #C C takes a piece of cloth.

LAVILA REPHRASER:

#C C takes a piece of cloth.

#C C is given a piece of cloth.

A piece of cloth is taken from C bag.

Our HOD:

The left hand moves from the left to the center, while the right hand remains still. The object, a piece of cloth, is taken by the left hand and brought towards the center of the frame.

1179

1180

1181

1182

1183

1184

1185

1186

1187

Figure 8: Comparison between: 1) the original Ego4D narrations; 2) LaViLa narrations; and 3) Narrations generated by our HOD. Our narration can describe the dynamic motion information of hands and objects, enhancing the semantic richness of the original narration.