

# Object-centric Video Representation for Long-term Action Anticipation

Ce Zhang<sup>1\*</sup> Changcheng Fu<sup>1\*</sup> Shijie Wang<sup>1</sup> Nakul Agarwal<sup>2</sup>  
 Kwonjoon Lee<sup>2</sup> Chiho Choi<sup>3†</sup> Chen Sun<sup>1</sup>  
<sup>1</sup> Brown University <sup>2</sup> Honda Research Institute USA <sup>3</sup> Samsung

## Abstract

This paper focuses on building object-centric representations for long-term action anticipation in videos. Our key motivation is that objects provide important cues to recognize and predict human-object interactions, especially when the predictions are longer term, as an observed “background” object could be used by the human actor in the future. We observe that existing object-based video recognition frameworks either assume the existence of in-domain supervised object detectors or follow a fully weakly-supervised pipeline to infer object locations from action labels. We propose to build object-centric video representations by leveraging visual-language pretrained models. This is achieved by “object prompts”, an approach to extract task-specific object-centric representations from general-purpose pretrained models without finetuning. To recognize and predict human-object interactions, we use a Transformer-based neural architecture which allows the “retrieval” of relevant objects for action anticipation at various time scales. We conduct extensive evaluations on the Ego4D, 50Salads, and EGTEA Gaze+ benchmarks. Both quantitative and qualitative results confirm the effectiveness of our proposed method. Our code is available at [github.com/brown-palm/ObjectPrompt](https://github.com/brown-palm/ObjectPrompt).

## 1. Introduction

Given an egocentric video observation, the action anticipation task [55] is defined as generating an action sequence of the camera-wearer in the form of verb and object pairs. Of particular interest is the long-term action anticipation (LTA) task [21], which aims to anticipate future actions over a long time-horizon. A reliable action anticipation algorithm is crucial for building intelligent agents, as it provides important signals for planning in interactive environments.

This paper aims to build effective object-centric video representations for action anticipation. As illustrated in Figure 1, our key motivation is that a detailed, object-centric

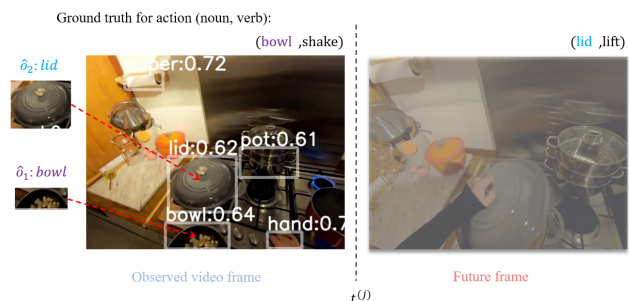


Figure 1. Objects are not only helpful for action recognition (left, *shake a bowl*), they also reveal the possible options for future human-object interactions (right, *lift a lid*). We propose object prompts, which leverage visual-language pretrained models (e.g. GLIP [32]) to build object-centric video representations without dataset-specific finetuning.

understanding of the scene provides visual cues on the goals of the actions and the available tools to be interacted with. The object-centric representation is also inspired and supported by studies on human perception [22, 27, 59]. While objects have been shown to play a crucial role for action understanding in both humans [53, 62] and machines [43, 66], their impacts on video-based action anticipation are yet to be studied. Among the earlier attempts at object-based video action understanding, one common approach is to leverage object detectors trained on in-domain bounding box annotations [8, 25, 43, 66, 68] on the same or similar datasets. While effective, the in-domain bounding box annotation process is time-consuming and labor-intensive, which makes the corresponding frameworks unlikely to scale to visually more diverse videos and complex, cluttered scenes. The other approach is to leverage generic object proposals [58, 61], or to directly work with image patches [17, 48, 57], and rely on the attention mechanism [60] to pick the salient regions with weak supervision from the action labels. Despite being more flexible, this approach does not incorporate prior knowledge on object locations, and often struggles to “detect” the actual objects, especially with limited training data.

We propose to address the limitations of the existing approaches by leveraging visual-language models pretrained

\*The first two authors contribute equally.

†Work done while at Honda Research Institute USA.

on large-scale datasets, such as GLIP [32]. We hypothesize that these pretrained models, whose objective is to (contrastively) associate image regions with text descriptions, learn generic object-centric representation that can be transferred to the action anticipation task, without the need to finetune their weights. We identify two challenges in order to investigate whether visual-language pretrained models help anticipation: First, how to properly “query” the pretrained models to retrieve the most relevant objects in a video observation, based on the domain knowledge; Second, how to associate different objects in an often cluttered scene to predict different human-object interactions in a long time-horizon. For the first challenge, we propose *object prompts*, which incorporate the domain knowledge of the target dataset by mapping the action (*e.g.* verb and object pairs) vocabulary into object prompts. The prompts are used to query the visual-language pretrained models to focus on the objects of interest. For the second challenge, we propose using a *predictive transformer encoder* (PTE), which is a Transformer encoder network that jointly attends to the motion cues (based on pre-trained video encoders) and the object-centric representation (based on the object prompts), and dynamically associates the motion and object evidence in order to predict future actions at different time steps. The proposed framework is trained end-to-end with the future action classification objectives. We refer to our overall framework as **ObjectPrompt**.

We conduct thorough experiments on the Ego4D [21] and the 50Salads [55] long-term action anticipation benchmarks, and the anticipation benchmark EGTEA Gaze+ [33]. Our quantitative experiments confirm that the prompt-based object-centric representations can substantially improve the action anticipation performance. Ablation experiments reveal that it is important to incorporate in-domain knowledge when designing the object prompts, and that object-centric representation significantly outperforms image-level counterpart [10]. In addition, qualitative analysis shows that the model learns to associate the corresponding objects when predicting actions at different time steps.

In summary, our contributions in this paper are three-fold: First, we demonstrate the effectiveness of object-centric representation for video action anticipation; Second, we propose “object prompts” to incorporate domain information when querying the pretrained models, and predictive transformer encoder to dynamically associate the object evidence for action anticipation; Finally, we provide extensive quantitative and qualitative analysis of the proposed framework, which achieves competitive performance on three benchmarks. Our implementation along with the pretrained models will be released.

## 2. Related Work

**Object-centric Video Representation** is an active research area for action recognition applications. The motivations include producing a more compact, thus efficient video representation by attending to regions of interest; and enabling compositional generalization to unseen human-object interactions [44] with a structured representation. For example, Wang et al. [61] extract RoI features from an 3D CNN feature maps using off-the-shelf detector, and build a graph neural network on top of RoI features alone. Long-term Feature Banks (LFB) [63] and Object Transformer [64] load off-the-shelf object features from object detectors in video backbones to encode long-term video features. Recently, with the advance of transformer-based architectures, ORViT [25] proposes to crop [24] object regions as a new object tokens and attach them to pixel tokens. ObjectLearner [66] fuses an object-layout stream and a pixels stream using an object-to-pixel transformer. Most of the existing object-based approaches either assume the availability of in-domain bounding box annotations, or apply generic “objectness” criteria (*e.g.* bounding box proposals trained on COCO [35]). ObjectViViT [68] relies on object detectors trained on the target datasets, or ground truth object annotations during evaluation. They observe that a direct application of a general-purpose large-vocabulary object detector [69] does not help video classification. We propose object prompts to leverage a pretrained visual-language model to generate task-specific object representations, which offers a simple yet effective solution to incorporate domain knowledge when they are available.

**Video Transformers.** Transformers [13] are now the predominant architectures for video recognition. The vanilla vision transformer (ViT) [13] evenly divides images into non-overlapping tokens, and run multi-head self-attention [60] over the tokens. TimeSFormer [5] and ViViT [4] extends ViT to videos, by introducing cube tokenization and efficient cross-time attention, *i.e.*, axis-based space-time attention [5] or factorized attention [4]. MotionFormer [48] enhance space-time attention by an implicit trajectory attention module. MViT [14, 34] and VideoSwin Transformer [38] re-introduce resolution-pooling as in convolution nets in video transformers for efficiency.

**Visual-Language Pretrained Models.** We have collectively made huge progress towards building unified learning frameworks for a wide range of tasks, including language understanding [6, 11, 37, 51], visual recognition [16, 29, 31, 65], and multimodal perception [3, 19, 26, 39, 56]. While the earlier visual-language pretrained models work with image-level [39, 50] or video-level [41, 56] representations, more recent models are object grounded [23, 32, 40, 46, 67]. We explore the benefits of both object-level [32] and image-level [50] for action anticipation.

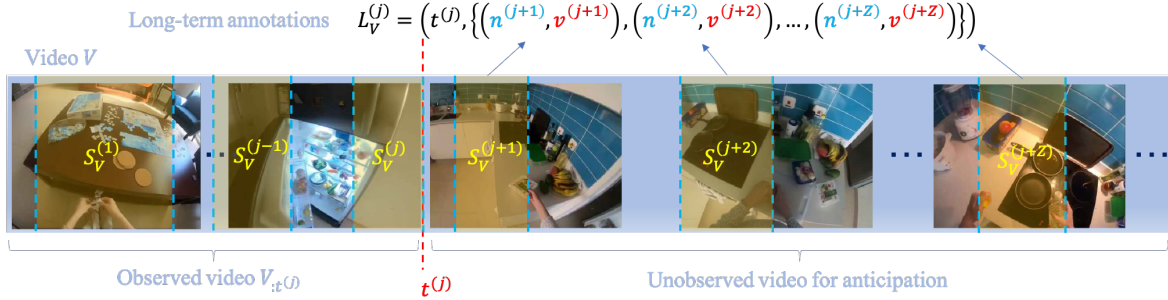


Figure 2. **Illustration of the Long-term Action Anticipation (LTA) task.** The learned model is expected to predict a sequence of  $Z$  actions, in the form of object and verb pairs, given visual observations up to time  $t^{(j)}$  in the video.  $t^{(j)}$  is the end time for the  $j$ -th labeled segment in the original video. During evaluation, the edit distance between a predicted sequence and the ground truth sequence is computed. To account for uncertainty in action anticipation, the model can predict up to  $K$  sequences for each input example.

### 3. Methods

We first introduce the long-term action anticipation (LTA) problem formulation. Then, we describe our overall model architecture. Finally, we describe our choices on object-based video representation in detail.

#### 3.1. The Action Anticipation Task

Given the significant applications of action anticipation, numerous benchmarks have been introduced to evaluate model performance [8, 9, 30, 47, 52, 58]. We focus on the long-term action anticipation (LTA) problem setup introduced by Ego4D [21], whose evaluation metric takes the diversity of future prediction into account. As illustrated in Figure 2, the provided annotations first split a long video  $V$  into smaller segments  $\{S_V^{(i)}\}$ , where  $S_V^{(i)}$  is the  $i$ -th annotated video segment. Each segment is labeled with its starting time, end time, and action label. The action label is represented as one verb, object pair  $(n^{(j)}, v^{(j)})$  for each segment  $S_V^{(j)}$ . The LTA task is specified by a “stop time”  $t^{(j)}$ , which denotes the end time for the last observed video segment  $S_V^{(j)}$ . The learned model is allowed to observe any video frames before  $t^{(j)}$  in order to make future predictions  $\{(n^{(j+1)}, v^{(j+1)}), \dots, (n^{(j+Z)}, v^{(j+Z)})\}$ , where  $Z$  is the number of future steps to predict. For example, suppose a person is frying an egg with a pan in a kitchen where knives, onions, water, and pots are scattered around. In this scenario, the LTA requires the model to predict the person’s upcoming actions sequentially, such as picking up the knife, cutting the onion, and drinking water. To account for the uncertainty of future behaviors, the model is allowed to make up to  $K$  sets of action predictions for each future step. For Ego4D [21], the standard setup and use  $Z = 20$ , and  $K = 5$ . A special case is when  $Z = 1$ , namely next action prediction. More details on the evaluation metric are in Section 4.

#### 3.2. Overall Model Architecture

Given a video  $V$  and a stop time  $t^{(j)}$ , our model takes a sequence of video segments  $\{S_V^{(j-N_v+1)}, \dots, S_V^{(j)}\}$  as the input observation, and generate a sequence of actions  $\{(n^{(j+1)}, v^{(j+1)}), \dots, (n^{(j+Z)}, v^{(j+Z)})\}$  as outputs.  $N_v$  is the number of observed video segments, and  $Z$  is the number of future steps. As illustrated in Figure 3, our overall model architecture consists of three modules: (1) a collection of video or object encoders that generate multimodal representations from video segments; (2) an aggregator network which fuses multimodal input representation across space and time; (3) an output decoder which generates action predictions from the aggregated features.

**Video Representation.** We use  $N_v$  clips from the observed video segments, each of which is optionally subsampled (e.g. to have the same number of video frames per clip). We pass the clips to a video encoder (e.g. SlowFast [15]) to compute clip-level representations  $E_{\text{clip}} \in \mathbb{R}^{N_v \times D}$ , where  $D$  is the encoded embedding size. We hypothesize that these clip-level representations capture scene-level appearance and motion information. We then select  $N_o$  objects from the set of detected objects among the  $N_v$  video segments. To compute a visual descriptor for each selected object, we apply an object encoder to generate an object embedding of the same dimension  $D$ . The object representation  $E_{\text{obj}}$  has the shape of  $N_o \times D$ . We discuss how to obtain the objects and their descriptors in Section 3.3.

**Temporal Aggregators.** We now introduce Predictive Transformer Encoder (PTE), a Transformer-based architecture for the action anticipation task. Given  $E_{\text{clip}}$  and  $E_{\text{obj}}$ , PTE is used to generate  $Z$  features  $z_0, z_1, \dots, z_{Z-1}$  for future prediction. PTE has learnable tokens  $E_{\text{POS}} \in \mathbb{R}^{Z \times D}$ . It concatenates  $E_{\text{clip}}$ ,  $E_{\text{obj}}$  and  $E_{\text{POS}}$  to form a sequence of length  $N_v + N_o + Z$ , then adds positional and modality encodings to the entire sequence. PTE uses a bidirectional Transformer encoder [60], so that the learnable to-

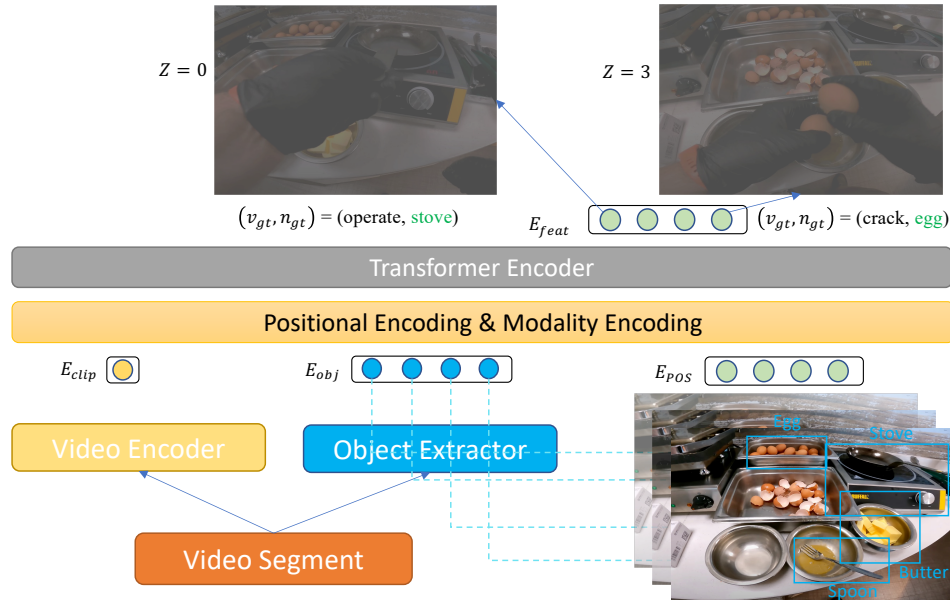


Figure 3. **An illustration of the overall model architecture.** We consider  $N_v = 1$  input video segments, and use  $N_o = 4$  objects for the video input. We consider predicting  $Z = 4$  future actions. Given the features  $E_{clip}, E_{obj}$  extracted by the encoders, Predictive Transformer Encoder (PTE) encodes  $Z$  learnable tokens  $E_{pos}$  to compute the features  $E_{feat}$  for action anticipation. A decoder network is applied on each of the  $Z$  encoded token  $z_i$  from  $E_{feat}$  to compute verb and object probabilities.

kens corresponding to the future actions can be encoded in parallel. We refer to the  $Z$  encoded learnable tokens as  $z_0, z_1, \dots, z_{Z-1}$ .

**Decoders and Training Objectives.** For each  $z_i$  encoded by PTE, we apply one set of linear classifiers on top as the decoder to compute the logits for the actions. We use two linear classifiers, one for verb, the other for noun. We apply Softmax Cross-Entropy as the loss function, and assign equal weights to all  $Z$  future steps to be predicted.

**Fusion strategies.** PTE takes video representations video encoders that capture motion, and object encoders to extract object-centric representation. These input representations can be concatenated together to form the input of PTE, which are then fused with bidirectional Transformer encoders. We refer to this approach of combining multiple video representations as early fusion. To better understand the best strategy to combine motion and object information, we also explore late fusion, where we use the two encoders to generate logits for future actions independently, and then take the average of the logits to form final predictions.

### 3.3. Object-centric Video Representation

We now describe how to extract object-centric representations from videos. As illustrated in Figure 4, we propose to leverage a general purpose grounded vision-language model, such as GLIP [32], to recognize objects of interest and compute their corresponding visual descriptors. We hypothesize that models pretrained on diverse object detec-

tion and phrase grounding datasets like GLIP already encode transferrable object information, but it is crucial to *guide* the model to focus on the objects of interest based on domain-knowledge (when they are available). We thus design an object prompt strategy that effectively incorporates the prior domain knowledge.

#### 3.3.1 Leveraging Pre-trained Grounding Models

Our proposed framework is able to leverage any recent image-language grounding models, such as region-text contrastive models [67], or open-vocabulary object detectors [23, 46]. To demonstrate the effectiveness of our proposed framework, we choose one particular model Grounded Language-Image Pre-training (GLIP) [32], and we believe advances in grounding models would lead to more performant object-centric representations. GLIP [32] proposes a pre-training strategy to build vision-language foundation models that can be further adopted for “zero-shot” object detection. During pre-training, GLIP is agnostic to the choice of object detectors to generate region proposals. In our work, we follow GLIP’s setting to use Dynamic Head [7] object detector for images, and the BERT [12] encoder for text inputs. Similar to applying CLIP [50] for zero-shot image classification, the zero-shot object detection can be achieved by querying with “object prompts” (e.g. “an image region with a cat”) to the pre-trained language-image grounding model.

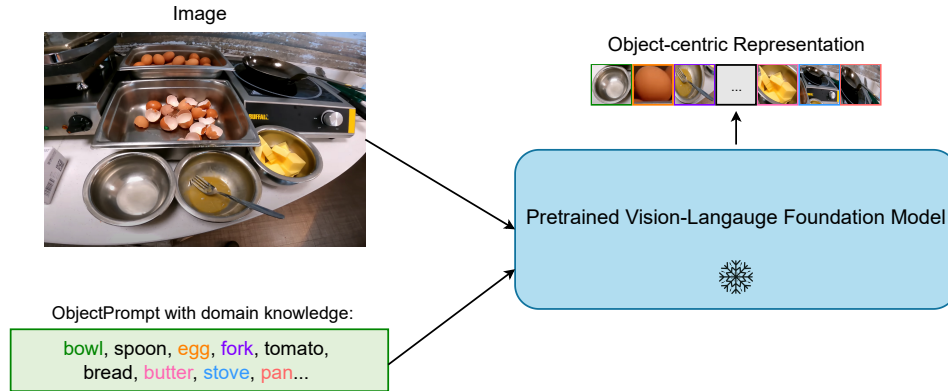


Figure 4. An illustration of **ObjectPrompt** to construct object-centric representation. On the left we show the input image and the list of objects of interest selected based on domain knowledge (e.g. kitchen and cooking) using ObjectPrompt. The prompts are then provided to a frozen vision-language model, such as GLIP [32], to detect the objects of interest and extract object-centric visual representations.

To retrieve the object descriptors computed by GLIP for the finally selected object proposals, we explore two complementary approaches. First, we inject a special identifier inside each region proposal data structure in order to gain original object features for each proposal. Second, we crop and feed the object proposal region into a pretrained image encoder to generate object features. Additionally, we also extract the object category level alignment scores and the box location to append to the object-level representation. The second approach allows us to leverage additional image-language pretrained model, such as CLIP [50].

### 3.3.2 ObjectPrompt with Domain Knowledge

Intuitively, domain knowledge is helpful to guide a grounding model to focus on objects of *interest*. For example, for humans performing cooking activities in a kitchen (e.g. Figure 4), it is important to focus on the food and cookware to model human-object interactions. Grounded pretrained models naturally support such guidance in the form of text prompts, namely a list of objects that the model is expected to ground. We thus hypothesize that the desirable object prompts should incorporate the domain knowledge (e.g. common objects appearing in the dataset), and explore different strategies to design the object prompts.

In order to obtain the object prompts, we first define the object vocabulary that contains the object classes to be detected. One simple and intuitive solution is to reuse the object vocabulary as provided by the task: For LTA, this refers to the list of objects (nouns) to be interacted with. To refine the vocabulary, we explore two complementary strategies: First, pick the most common object categories based on their frequency in the training data of the target task; Second, use word embedding (e.g. word2vec [45]) and K-Means clustering to group similar categories. In addition, we also explore standard object vocabulary, such as the one

used by the COCO [35] dataset for comparison. Figure 5 shows the actual detections by GLIP with different object prompt strategies. Compared to COCO, both *most-common* and *kmeans* are able to better detect the important objects (e.g. keyboard, music instrument, and hand) thanks to the in-domain knowledge when designing object prompts.

**Discussion.** Though well supported by research on human cognition [22, 27, 53, 59, 62], how to effectively construct object-centric video representations remains an open problem. For example, ObjectViViT [68] observes that general-purpose object detectors, even capable of detecting over twenty-thousand object categories [69], do not lead to better video representation for the classification task. They instead have to resort to object detectors trained on the target video datasets, with their collected object bounding box annotations. We hypothesize that one important reason for the challenges of applying general object detectors for representing videos is the lack of attention to the objects of interest: As a result, the irrelevant objects may distract or mislead the learning framework. Instead of relying on in-domain object annotations, which are time-consuming to collect, our ObjectPrompt framework only relies on domain knowledge, which are extremely easy to collect (e.g. always detect hands since it is important to model human-object interactions). In our experiments, we simply use the list of actions in a target dataset to compute the object prompts.

## 4. Experiment

We now report experimental observations on three benchmarks, Ego4D, EGTEA Gaze+, and 50Salads to demonstrate the effectiveness of ObjectPrompt on action anticipation task. We perform extensive ablation study and qualitative / quantitative result analysis to support our findings. Additional implementation details can be found in the supplementary material.




Figure 5. Example GLIP detection results with different object prompts on randomly sampled video frames of Ego4D. We observe that the K-Means clustering strategy (last column) offers the best precision and recall for the objects of interest in all four cases.

#### 4.1. Experiment Setup

**Ego4D** [21] contains 3,670 hours of daily life activity egocentric video spanning hundreds of scenarios. We focus on videos under Forecasting subset designed for the LTA benchmark. It contains 1723 clips with 53 scenarios and around 116 hours. The label space includes 115 verb as interaction behaviors and 478 noun as objects. We follow the standard train, validation, and test splits from [21] annotations for evaluation. Test split performance is obtained via the official evaluation server.

**EGTEA Gaze+** [33] contains 28 hours egocentric videos of cooking activities from 86 unique sessions of 32 subjects. Each video is annotated with 19 verbs and 53 nouns. We follow the same train and test splits.

**50Salads** [55] contains 25 people preparing salads in 50 videos with actions falling into 17 classes. The average video length is 6 minutes and on average the number of instances per video is around 20. We follow the dataset standard and perform 5-fold cross-validation in our evaluation.

**Evaluation Metrics.** For Ego4D, we follow the standard evaluation protocol from [21], and report edit distance (ED). ED is computed as the Damerau-Levenshtein distance over sequences of predictions of verbs, nouns and actions. For

the  $K$  possible sequences a model predicts, the metric chooses the smallest edit distance between the ground truth and any of the  $K$  sequences. Following [21], we set  $K = 5$  in our experiments, and report Edit Distance at  $Z = 20$  (ED@20) on the test set and Average Edit Distance (AUED) on the validation set.

For EGTEA Gaze+, we report top1 accuracy and class-mean top1 accuracy. We report next action prediction performance which is the standard practice on this dataset. For 50Salads, we follow the LTA evaluation metric in [2] and report the mean over classes accuracy (MoC@( $O, F$ )), where  $O$  is the percentage of observed video frames, and  $F$  is the percentage of future video frames to be labeled.

#### 4.2. Ablation Study

We focus our ablation study on the Ego4D dataset. To better understand the effectiveness of our object-centric representation, we only provide the object features  $E_{obj}$  as the input to the PTE model.

**Object prompts.** We first study how to properly incorporate domain knowledge into object prompt design. We explore three strategies: “most-common”, “kmeans”, “coco”. We fix the number of prompts to contain 80 object prompts, and empirically observe that the performance is robust even with fewer (e.g. 20 and 40) prompts. As a recap, the “most-common” strategy picks the top 80 most frequent nouns from the Ego4D vocabulary. The “kmeans” strategy contains the top 80 most frequent prompts after the entire list of objects is clustered to reduce redundancy. The “coco” baseline uses the 80 object categories from the COCO [35] dataset. We also include “random” as an additional baseline. It randomly picks regions in images as bounding boxes. From Table 1a, we can see that both “most-common” and “kmeans” outperform the baselines “coco” and “random” significantly, demonstrating the effectiveness and importance of our designed object prompts.

**Object locations and categories.** In Table 1c, we observe that both location and category features provide useful information when included in the object-centric representation, especially for the noun ED. We also observe that object location information appears to be less important than category information. We conjecture that the category information may provide the model with high-level scene information (e.g. *person*, *laptop*, and *book* is very likely to indicate a library), which helps the model infer long-term future actions.

**Object quantity and quality.** Each grounded object by the GLIP model comes with a confidence score. Intuitively, higher scores correspond to higher quality objects, at the cost of lower recall. In Table 1d, we use a threshold on confidence scores to filter objects. We can see that the model performs the best when we set threshold to 0.3. In Table 1b, we vary the maximum number of objects to be selected per

Vocabulary	Verb ↓	Noun ↓	#obj	Verb ↓	Noun ↓
random bbox	0.745	0.945	1	0.735	0.834
coco	0.737 (-1.0%)	0.789 (-16.5%)	3	0.727	0.800
most common	0.734 (-1.5%)	0.776 (-17.8%)	5	<b>0.728</b>	<b>0.771</b>
kmeans	<b>0.728 (-2.3%)</b>	<b>0.771 (-18.4%)</b>	10	0.731	0.781

(a) Object Prompt Strategy				(b) Number of Objects		
Loc.	Cate.	Verb ↓	Noun ↓	Threshold	Verb ↓	Noun ↓
✗	✗	0.730	0.789	0.00	0.731	0.781
✓	✗	0.730	0.786	0.30	<b>0.728</b>	<b>0.771</b>
✗	✓	0.731	0.774	0.45	0.728	0.773
✓	✓	<b>0.728</b>	<b>0.771</b>	0.55	0.731	0.787

(c) Location and Category				(d) Threshold		
---------------------------	--	--	--	---------------	--	--

Table 1. **Ablation study of object-centric representation on the Ego4D validation set.** We conduct detailed ablation on (1) object vocabulary, (2) number of object per frame, (3) object location and category, (4) detection threshold. We only provide object-centric representation  $E_{obj}$  as the video representation.

Aggregator	Modality	Fusion	Verb ↓	Noun ↓
Transformer	video	-	0.751	0.766
PTE	video	-	0.713 (-5.1%)	0.753 (-1.7%)
PTE	video+object	early	<b>0.707(-5.9%)</b>	<b>0.743 (-3.0%)</b>
PTE	video+object	late	0.709 (-5.6%)	0.748 (-2.3%)

Table 2. **Temporal modeling and modality fusion on Ego4D LTA validation set.** We observe that PTE outperforms the baseline temporal aggregator, and object-centric representation can effectively improve the performance, especially with early fusion.

frame, sorted via confidence scores. We can see the model performs the best when we use 5 objects per frame. We conjecture that a trade-off exists between the quantity and quality of objects.

**Temporal modeling.** We compare our temporal modeling methods with Ego4D LTA Baseline [21] in the first two rows of Table 2. We observe that with the same video representation  $E_{clip}$  encoded by a pre-trained SlowFast network, PTE achieves significant improvement over the baseline temporal aggregator with a vanilla Transformer with no learnable tokens for future prediction.

**Effectiveness of the object-centric representation.** The next two rows of Table 2 study the impact of object-centric representation. We observe that the object modality helps reduce the edit distance substantially, especially for noun prediction. We also observe that early fusion works better than late fusion.

**Enhancing object-centric representation.** While we use GLIP [32] to detect and represent objects, there are many other pre-trained models we can use to obtain object representation. In Table 3, we explore additional pretrained models for object representation. We still use the same detected

Model	Modality	Verb ↓	Noun ↓
-	video	0.713	0.753
GLIP	video+object	0.707 (-5.9%)	0.743 (-3.0%)
CLIP	video+object	<b>0.700 (-6.8%)</b>	<b>0.717 (-4.8%)</b>

Table 3. **Impact of object-centric representation on Ego4D LTA validation set.** Incorporating CLIP brings additional performance gain over GLIP object embeddings.

Model	Verb ↓	Noun ↓	Action ↓
ICVAE [42]	0.7410	0.7396	0.9304
SlowFast [21]	0.7389	0.7800	0.9432
VCLIP [10]	0.7389	0.7688	0.9412
ObjectPrompt (ours)	<b>0.7265</b>	<b>0.7396</b>	<b>0.9290</b>

Table 4. Comparison with the state-of-the-art on the Ego4D LTA benchmark. We report test split results from the evaluation server.

objects, and use CLIP [50] (ViT-L/14@336px) to represent detected objects. We observe that incorporating CLIP representation further reduces the edit distance. An interesting future work is to explore other pre-trained visual encoders.

**Per-step analysis.** We further analyze the per-step edit distance for verbs and nouns on the Ego4D validation set in Figure 6. “video (baseline)” uses the vanilla Transformer Encoder as temporal aggregator. The other methods use PTE as the temporal aggregator. We can see that our object-centric representation leads to consistent and significant improvements across all 20 steps, especially for the nouns.




Figure 6. **Per-step ED** on Ego4D validation (lower is better). **Top:** verb. **Bottom:** noun. ED@Z is the edit distance at  $Z^{th}$  step.

Observation → Future →	20%				30%			
	10%	20%	30%	50%	10%	20%	30%	50%
RNN [2]	30.1	25.4	18.7	13.5	30.8	17.2	14.8	9.8
CNN [2]	21.2	19.0	15.9	9.8	29.1	20.1	17.5	10.9
Sener <i>et al.</i> [54]	34.7	26.3	23.7	15.7	34.5	26.1	22.7	17.1
Qi <i>et al.</i> [49]	37.9	28.8	21.3	11.1	<b>37.5</b>	24.1	17.1	9.1
FUTR [20]	<b>39.6</b>	27.5	23.3	17.8	35.2	<b>24.9</b>	24.2	15.3
ObjectPrompt (ours)	37.4	<b>28.9</b>	<b>24.2</b>	<b>18.1</b>	28.0	24.0	<b>24.3</b>	<b>19.3</b>

Table 5. Comparison with state-of-the-art on 50Salads LTA task.

Model	Top-1 Acc. ↑	Class-mean Acc. ↑
I3D-Res50 [28]	34.8	23.2
FHOI [36]	36.6	25.3
AVT [18]	43.0	35.2
ObjectPrompt (ours)	<b>44.6</b>	<b>36.4</b>

Table 6. Comparison with state-of-the-art on EGTEA Gaze+.

### 4.3. Comparison to the State-of-the-art

Finally, we report the action anticipation performance on Ego4D, 50Salads, and EGTEA Gaze+ benchmarks, and compare with previously published results.

**Ego4D.** Table 4 compares our best model (Slowfast + PTE (video + CLIP object, early)) with the recent state-of-the-art methods under comparable settings. We report verb, noun and action ED on the test set. Our model is able to outperform competitive baseline methods.

**50Salads.** We conduct experiments on long-term action anticipation benchmarks and report class-mean top1 action accuracy in Table 5. Following common practice [2], we predict 10%, 20%, 30%, and 50%, respectively, of the video after observing the previous 20% or 30%. We use PTE as a temporal modeling method in video-only models and CLIP object representation in video+object models. Our method

achieves overall on par performance with FUTR [20] and achieves the best performance for longer-term anticipation.

**EGTEA Gaze+.** Table 6 shows results on the next action prediction benchmarks for EGTEA+ Split 1 following the standard practice. We use CLIP object representation in video+object models, PTE as temporal modeling, and finetune the video backbone. We report top1 accuracy and class-mean top1 accuracy.

### 4.4. Qualitative Analysis

In Fig. 7, we show several qualitative examples of object attention weights produced by PTE. We use attention rollout [1] to compute attention weights from  $Z$  output action (noun, verb) pairs to previous visual observations and choose the top 10 objects which have the relatively highest weight. Comparing with the ground truth label shows the model learns to associate the corresponding objects when predicting actions at different time steps. More results can be found in the supplementary materials.

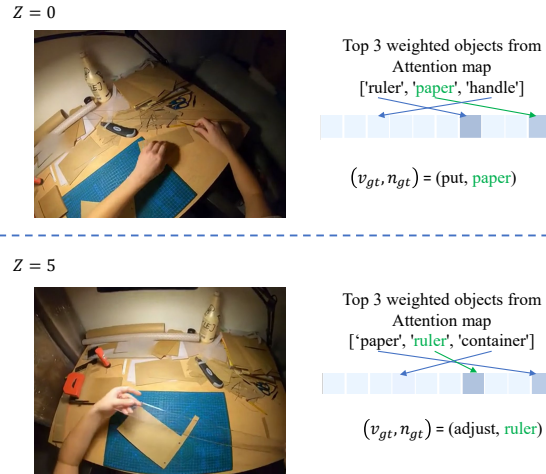


Figure 7. **Object attention visualizations.** We show the top 3 predicted objects generated with attention rollout [1].

## 5. Conclusion

We propose ObjectPrompt, a framework to effectively construct object-centric video representation for long-term action anticipation. Our framework has two modules, an object prompt strategy to effectively leverage domain knowledge and construct object-centric representation with pre-trained grounding models, and predictive transformer encoder, which dynamically associates the object evidence for long time-horizon action prediction. Experiments results confirm that both modules improve the action anticipation performance. We report competitive results on Ego4D LTA, 50Salads, and EGTEA Gaze+ benchmarks.

**Acknowledgements.** This work is in part supported by Honda Research Institute and Meta AI.



## References

- [1] Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. In *ACL*, 2020. 8
- [2] Yazan Abu Farha, Alexander Richard, and Juergen Gall. When will you do what? - anticipating temporal occurrences of activities. In *CVPR*, 2018. 6, 8
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, et al. Flamingo: a visual language model for few-shot learning. In *NeurIPS*, 2022. 2
- [4] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lucic, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, 2021. 2
- [5] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 2
- [6] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. Language models are few-shot learners. In *NeurIPS*, 2020. 2
- [7] Xiyang Dai, Yinpeng Chen, Bin Xiao, Dongdong Chen, Mengchen Liu, Lu Yuan, and Lei Zhang. Dynamic head: Unifying object detection heads with attentions. In *CVPR*, 2021. 4
- [8] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The EPIC-KITCHENS dataset. In *ECCV*, 2018. 1, 3
- [9] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Antonino Furnari, Evangelos Kazakos, Jian Ma, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Rescaling egocentric vision: collection, pipeline and challenges for epic-kitchens-100. *International Journal of Computer Vision*, 2022. 3
- [10] Srijan Das and Michael S Ryoo. Video+ clip baseline for ego4d long-term action anticipation. *arXiv preprint arXiv:2207.00579*, 2022. 2, 7
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 4
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 2
- [14] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 2
- [15] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. *CoRR*, abs/1812.03982, 2018. 3
- [16] Golnaz Ghiasi, Barret Zoph, Ekin D Cubuk, Quoc V Le, and Tsung-Yi Lin. Multi-task self-training for learning general representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8856–8865, 2021. 2
- [17] Rohit Girdhar, Joao Carreira, Carl Doersch, and Andrew Zisserman. Video action transformer network. In *CVPR*, 2019. 1
- [18] Rohit Girdhar and Kristen Grauman. Anticipative video transformer. In *ICCV*, 2021. 8
- [19] Rohit Girdhar, Mannat Singh, Nikhila Ravi, Laurens van der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A single model for many visual modalities. In *CVPR*, 2022. 2
- [20] Dayoung Gong, Joonseok Lee, Manjin Kim, Seong Jong Ha, and Minsu Cho. Future transformer for long-term action anticipation. In *CVPR*, 2022. 8
- [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 *CVPR*, 2022. 1, 2, 3, 6, 7
- [22] Kalanit Grill-Spector and Nancy Kanwisher. Visual recognition: As soon as you know it is there, you know what it is. *Psychological Science*, 2005. 1, 5
- [23] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. 2, 4
- [24] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 2
- [25] Roei Herzig, Elad Ben-Avraham, Karttikeya Mangalam, Amir Bar, Gal Chechik, Anna Rohrbach, Trevor Darrell, and Amir Globerson. Object-region video transformers. In *CVPR*, 2022. 1, 2
- [26] Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021. 2
- [27] Daniel Kahneman, Anne Treisman, and Brian J Gibbs. The reviewing of object files: Object-specific integration of information. *Cognitive psychology*, 1992. 1, 5
- [28] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 8
- [29] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *CVPR*, 2018. 2
- [30] Kris M Kitani, Brian D Ziebart, James Andrew Bagnell, and Martial Hebert. Activity forecasting. In *ECCV*, 2012. 3
- [31] Iasonas Kokkinos. Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *CVPR*, 2017. 2
- [32] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, Kai-Wei Chang, and

- Jianfeng Gao. Grounded language-image pre-training. In *CVPR*, 2022. 1, 2, 4, 5, 7
- [33] Yin Li, Miao Liu, and James M Rehg. In the eye of beholder: Joint learning of gaze and actions in first person video. In *ECCV*, 2018. 2, 6
- [34] Yanghao Li, Chao-Yuan Wu, Haoqi Fan, Karttikeya Mangalam, Bo Xiong, Jitendra Malik, and Christoph Feichtenhofer. Mvitv2: Improved multiscale vision transformers for classification and detection. In *CVPR*, 2022. 2
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 5, 6
- [36] Miao Liu, Siyu Tang, Yin Li, and James Rehg. Forecasting human-object interaction: Joint prediction of motor attention and actions in first person video, 2019. 8
- [37] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 2
- [38] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *CVPR*, 2022. 2
- [39] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, 2019. 2
- [40] Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Motlaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. *arXiv preprint arXiv:2206.08916*, 2022. 2
- [41] Huaishao Luo, Lei Ji, Botian Shi, Haoyang Huang, Nan Duan, Tianrui Li, Jason Li, Taroon Bharti, and Ming Zhou. Univl: A unified video and language pre-training model for multimodal understanding and generation. *arXiv preprint arXiv:2002.06353*, 2020. 2
- [42] Esteve Valls Mascaro, Hyemin Ahn, and Dongheui Lee. Intention-conditioned long-term human egocentric action forecasting @ ego4d challenge 2022. 7
- [43] Joanna Materzynska, Tete Xiao, Roei Herzig, Huijuan Xu, Xiaolong Wang, and Trevor Darrell. Something-else: Compositional action recognition with spatial-temporal interaction networks. In *CVPR*, 2020. 1
- [44] Joanna Materzynska, Tete Xiao, Roei Herzig, Huijuan Xu, Xiaolong Wang, and Trevor Darrell. Something-else: Compositional action recognition with spatial-temporal interaction networks. In *CVPR*, 2020. 2
- [45] Tomas Mikolov, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. In *ICLR*, 2013. 5
- [46] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, Xiao Wang, Xiaohua Zhai, Thomas Kipf, and Neil Houlsby. Simple open-vocabulary object detection with vision transformers. In *ECCV*, 2022. 2, 4
- [47] Tushar Nagarajan, Yanghao Li, Christoph Feichtenhofer, and Kristen Grauman. Ego-topo: Environment affordances from egocentric video. In *CVPR*, 2020. 3
- [48] Mandela Patrick, Dylan Campbell, Yuki Asano, Ishan Misra, Florian Metz, Christoph Feichtenhofer, Andrea Vedaldi, and João F Henriques. Keeping your eye on the ball: Trajectory attention in video transformers. *NeurIPS*, 2021. 1, 2
- [49] Zhaobo Qi, Shuhui Wang, Chi Su, Li Su, Qingming Huang, and Qi Tian. Self-regulated learning for egocentric video activity anticipation. *TPAMI*, 2021. 8
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020*, 2021. 2, 4, 5, 7
- [51] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019. 2
- [52] Nicholas Rhinehart and Kris M Kitani. Online semantic activity forecasting with darko. *arXiv preprint arXiv:1612.07796*, 2016. 3
- [53] Scott J. Robson and Valerie A. Kuhlmeier. Infants’ understanding of object-directed action: An interdisciplinary synthesis. *Frontiers in Psychology*, 7, 2016. 1, 5
- [54] Fadime Sener, Dipika Singhania, and Angela Yao. Temporal aggregate representations for long-range video understanding. In *ECCV*, 2020. 8
- [55] Sebastian Stein and Stephen J. McKenna. Combining embedded accelerometers with computer vision for recognizing food preparation activities. In *UbiComp*, 2013. 1, 2, 6
- [56] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *ICCV*, 2019. 2
- [57] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Kevin Murphy, Rahul Sukthankar, and Cordelia Schmid. Actor-centric relation network. In *ECCV*, 2018. 1
- [58] Chen Sun, Abhinav Shrivastava, Carl Vondrick, Rahul Sukthankar, Kevin Murphy, and Cordelia Schmid. Relational action forecasting. In *CVPR*, 2019. 1, 3
- [59] Joshua B Tenenbaum, Charles Kemp, Thomas L Griffiths, and Noah D Goodman. How to grow a mind: Statistics, structure, and abstraction. *Science*, 2011. 1, 5
- [60] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017. 1, 2, 3
- [61] Xiaolong Wang and Abhinav Gupta. Videos as space-time region graphs. In *ECCV*, 2018. 1, 2
- [62] A. Woodward. Infants selectively encode the goal object of an actor’s reach. *Cognition*, 1998. 1, 5
- [63] Chao-Yuan Wu, Christoph Feichtenhofer, Haoqi Fan, Kaiming He, Philipp Krahenbuhl, and Ross Girshick. Long-term feature banks for detailed video understanding. In *CVPR*, 2019. 2

- [64] Chao-Yuan Wu and Philipp Krahenbuhl. Towards long-form video understanding. In *CVPR*, 2021. [2](#)
- [65] Amir R Zamir, Alexander Sax, William Shen, Leonidas J Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *CVPR*, 2018. [2](#)
- [66] Chuhan Zhang, Ankush Gupta, and Andrew Zisserman. Is an object-centric video representation beneficial for transfer? In *ACCV*, 2022. [1](#), [2](#)
- [67] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, and Jianfeng Gao. Regionclip: Region-based language-image pretraining. In *CVPR*, 2022. [2](#), [4](#)
- [68] Xingyi Zhou, Anurag Arnab, Chen Sun, and Cordelia Schmid. How can objects help action recognition? In *CVPR*, 2023. [1](#), [2](#), [5](#)
- [69] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Phillip Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. *ECCV*, 2022. [2](#), [5](#)