ActionVOS: Actions as Prompts for Video Object Segmentation

Liangyang Ouyang¹^(a), Ruicong Liu¹^(b), Yifei Huang¹*^(a), Ryosuke Furuta¹^(b), and Yoichi Sato¹^(b)

The University of Tokyo {oyly, lruicong, hyf, furuta, ysato}@iis.u-tokyo.ac.jp

Abstract. Delving into the realm of egocentric vision, the advancement of referring video object segmentation (RVOS) stands as pivotal in understanding human activities. However, existing RVOS task primarily relies on static attributes such as object names to segment target objects, posing challenges in distinguishing target objects from background objects and in identifying objects undergoing state changes. To address these problems, this work proposes a novel action-aware RVOS setting called ActionVOS, aiming at segmenting only active objects in egocentric videos using human actions as a key language prompt. This is because human actions precisely describe the behavior of humans, thereby helping to identify the objects truly involved in the interaction and to understand possible state changes. We also build a method tailored to work under this specific setting. Specifically, we develop an action-aware labeling module with an efficient action-guided focal loss. Such designs enable ActionVOS model to prioritize active objects with existing readilyavailable annotations. Experimental results on the VISOR dataset reveal that ActionVOS significantly reduces the mis-segmentation of inactive objects, confirming that actions help the ActionVOS model understand objects' involvement. Further evaluations on VOST and VS-COS datasets show that the novel ActionVOS setting enhances segmentation performance when encountering challenging circumstances involving object state changes. We will make our implementation available at https://github.com/ut-vision/ActionVOS.

Keywords: Referring Expression Comprehension · Referring Video Object Segmentation · Active Object Segmentation

1 Introduction

Exploring the domain of egocentric vision (first-person perspective), the development of Referring Video Object Segmentation (RVOS) is critical for comprehending human activities. RVOS aims at segmenting target objects using natural language expressions, serving as a foundation for machines to have a comprehensive understanding of visual-language and temporal information. By

^{*}Corresponding author.



(a) Use static attributes to identify target objects (b) Use human actions to identify target objects

Fig. 1: Human actions as language prompts help to identify active objects.

integrating various modalities, RVOS paves the way for groundbreaking applications in egocentric contexts, such as text-directed object identification and real-time object tracking in videos. This has been exemplified in recent studies, including referring expression comprehension [27, 49], active object localization [68, 75] and intention-driven visual grounding [28, 60]. As highlighted by recent works [8, 9, 16, 32, 78, 80], advancements in egocentric applications have led to a surge in data related to egocentric interactions. This has subsequently increased the demand for RVOS from egocentric perspectives.

In the field of RVOS, existing benchmarks [15,24,53] primarily rely on static attributes, *e.g.*, object names and colors, to describe target objects in the video. In simple scenarios [23,74], such static attributes are adequate to identify the target objects. However, when scenarios become complex, these static attributes fall short in accurately identifying target objects, such as when similar redundant objects coexist or the object state is changing. Fig. 1 (a) illustrates two failure cases of static attributes. In the "carrot&bowl" example, static attributes identify redundant and inactive "carrot&bowl". In the case of "nail", static attributes fail to identify the nail painted from pink to blue.

To address these problems, we employ human actions as a substantial cue for identifying target objects. This is because human actions, as a strong language prompt, precisely describe the behavior of humans. Such action prompts aid in identifying objects truly involved in interactions and comprehending potential object state changes. As illustrated in Fig. 1 (b), when provided with action prompt "put carrot in bowl", the specific carrots and bowl involved in "put" action are accurately identified. Similarly, the specific nail being painted is also correctly identified with the "paint nail" prompt. Therefore, action prompts significantly resolve ambiguity arising from redundant instances and object state changes.

In this work, we propose ActionVOS, a novel action-aware setting for RVOS, segmenting active objects in egocentric videos using action prompts. As shown in Fig. 2, unlike conventional RVOS settings, ActionVOS incorporates an additional language prompt of action narrations. Guided by such action prompts, ActionVOS only segments active objects involved in interactions, regardless of redundancy or state changes.

Unfortunately, existing video object segmentation datasets [9,57,72] lack annotations of identifying active objects, *i.e.*, whether or not they are involved in

interactions. During training, this limitation leads ActionVOS to face difficulty in obtaining annotations that classify whether an object is active or not. To address this issue, we propose an action-aware labeling module to generate pseudo-labels from existing readily-available annotations, including action narrations [7,8,16], semantic segmentation [9,57,72], and hand-object segmentation [9]. This module enables ActionVOS model to obtain training data regarding the objects' involvement in actions without manually annotating their participation. In addition, we design an effective action-guided focal loss working with the action-aware labeling module. This proposed loss reduces the impacts of false positives in generated pseudo-labels, prioritizing the truly active objects.

We evaluate our method on three video object segmentation datasets VI-SOR [9], VOST [57], and VSCOS [72]. Comparing with conventional RVOS setting, ActionVOS significantly reduces the mis-segmentation of inactive objects on VISOR dataset, with a 35.6% mIoU reduction of inactive objects. Evaluation on VOST and VSCOS datasets indicates that the ActionVOS setting enhances the segmentation of objects undergoing state changes, by achieving a 3.0% mIoU increase of state-changed objects. These results confirm that action prompts help ActionVOS model focus on active objects and enhance the understanding of state changes.

The main contributions of this work are as follows:

- We propose a novel action-aware setting for referring video object segmentation, ActionVOS. This setting segments active objects in egocentric videos by employing action narrations as an additional language prompt.
- We develop an action-aware labeling module and an action-guided focal loss for ActionVOS. This design enables ActionVOS models to segment active objects with existing readily-available annotations.
- Extensive evaluation results show that ActionVOS significantly reduces the mis-segmentation of inactive objects, and enhances the segmentation of statechanged objects.

2 Related Works

2.1 Referring Expression Comprehension

Referring expression comprehension (REC) aims to localize target objects described by a referring expression in natural language. Established REC benchmarks [10, 21, 23, 37, 43, 64, 74] and REC methods [22, 29, 35, 42, 58, 59, 71, 73, 76] contribute to this fundamental yet challenging task. A new benchmark GREC [19, 34] introduces generalized referring expression comprehension, extending REC by permitting expressions to describe any number of target objects.

In addition to REC in images, there has been a growing interest in videobased REC [5,11,31,61,65,70], which requires both temporal and spatial localization of text-referred objects in video frames. Recent works [27,68,75] introduce REC to track and localize active objects in egocentric videos. However, in these works, the number of target active objects is typically limited to one or two

in each video. In this work, we not only extend localization to segmentation, but also aim to identify a broader range of active objects, e.g., hands, tools, containers and other entities.

2.2 Referring Video Object Segmentation

Referring video object segmentation (RVOS) aims to segment the target object indicated by a given expression across the entire video clip. Conventional RVOS datasets [15, 24, 48, 53, 69] are constructed by adding language expressions to existing video object segmentation datasets. These datasets often provide an expression for a single object, which usually describes the static attributes of the target object. A recent dataset MeViS [12] focuses on segmenting objects in video content based on a sentence describing their motions. Existing RVOS methods [2, 6, 13, 30, 44, 53, 62, 66, 67, 71] employ various approaches to address the RVOS task. Among these works, SLVP [44] is the first to adopt RVOS to VISOR [9] dataset. Comparing to SLVP, our work incorporates an additional action narration in the language prompt to describe and segment only active objects in egocentric videos.

2.3 Action-object Relation

The relations between human actions and objects have been extensively studied over time. Previous works [3, 9, 14, 20, 36, 54, 75, 77] focus on hand-object interactions as a basic for understanding active objects. Besides hand-objects, many works have introduced different representations to model action-object relations across various applications, such as graphical models [17], object-action complexes (OAC) [26], object affordances [25], action-objects [1], active entities [9], objects undergoing change with tools [68], and action scene graphs [52]. In comparison to prior works, our work broadens the range of active objects in Sec. 3. This includes not only objects described by action narrations but also treats hands, hand tools, containers, and contents as active objects in human actions.

3 Problem Setting

Input. ActionVOS task is a Referring Video Object Segmentation (RVOS) task focused on active objects involved in human actions. Its input contains three parts: 1) A video clip $\mathcal{V} = \{V_t\}_{t=1}^T$, where $V_t \in \mathbb{R}^{H \times W \times 3}$ is an RGB image from T frames of the video clip. H and W stand for height and width, respectively. 2) An action narration \mathcal{A} , which describes the human action in \mathcal{V} . 3) A set of N object names $\mathcal{O} = \{O_i\}_{i=1}^N$, where O_i is a noun of an object. Note that N is arbitrary and any object name can be in \mathcal{O} , making ActionVOS an open-vocabulary setting.

Output. ActionVOS aims to predict T segmentation masks $\mathcal{M} = \{M_t\}_{t=1}^T$, where $M_t \in \mathbb{R}^{N \times H \times W}$ for N objects. We use $M_t(O_i)$ to represent the binary



Fig. 2: Comparison between ActionVOS and conventional RVOS settings.



Fig. 3: Examples of positive objects in ActionVOS.

segmentation mask for O_i in frame t, where each pixel belongs to one single object or background.

Comparing with conventional RVOS tasks, ActionVOS focuses on whether the referring object is interacted in the ongoing human action. We define the objects interacted in the action as positive $\mathcal{O}_{\mathcal{P}}$, while other objects are negative $\mathcal{O}_{\mathcal{N}}$. For the positive objects, all of them should be segmented through all frames. For the negative objects, their mask predictions should be all-zero since they do not participate in the action. Fig. 2 compares the inputs and outputs of ActionVOS with conventional RVOS settings. Compared to RVOS, ActionVOS incorporates additional action prompts as input, expressed as an action narration "open tofu container". In this example, only active objects, *i.e.*, hands, tofu, and tofu container are segmented in ActionVOS outputs.

Definition of "positive". One of the most important concepts of ActionVOS is the definition of positive objects. According to the action prompt, we define positive objects as follows:

- 1) Objects described by the action narration.
- 2) Hands and hand-tools used for the action.
- 3) Containers and contents interacted in the action.

1) The objects described by action narration are unquestionably defined as positive. 2) Hand-tools, being operated by human during the action, are thus defined as positive. 3) Taking the action "put pan" in Fig. 3 as an example, we need to segment the objects inside the pan as target objects (*e.g.*, meat, spoon). As the objects inside the pan are also put down with the "pan" through the action "put", they are subjected to the action of "put pan". Similarly, if an object moving with

a container is mentioned in the action narration while the container itself is not, the container should also be subjected to the action. Therefore, we define containers and contents interacted in the action as positive. As shown in Fig. 3, the "cutting board" is defined as the positive object for the action "take eggplants". On the other hand, if the "cutting board" as a moving vessel of "eggplants" is not defined as positive, this action becomes non-existent.

Data & annotation. Existing video object segmentation datasets such as VI-SOR [9], VOST [57] and VSCOS [72] are collected on egocentric videos, providing both semantic segmentation labels and human action narrations [7,8,16]. As VOST and VSCOS focus on objects undergoing state changes, they have only one active object being annotated for each action. We only use the validation sets of these two datasets to evaluate ActionVOS performance on state-changed objects. VISOR annotates a set of objects masks for each action, but lack precise indication of objects' involvement in actions, *i.e.*, positive and negative classification labels. To address this issue, we propose a labeling module in Sec. 4.2 to generates such classification labels with existing annotations.



Fig. 4: Overview of the proposed method.

4 Proposed Method

As illustrated in Fig. 4, we propose a method for the ActionVOS setting. In Sec. 4.1, we develop an ActionVOS model S, which is constructed by adding an extra classification head to an RVOS model. In Sec. 4.2, we propose an action-aware labeling module Φ . This module generates pseudo-labels of active objects, addressing the problem that existing datasets lack indication of objects' involvement in actions. In Sec. 4.3, we propose an action-guided focal loss to reduce the impact of false positives from the generated pseudo-labels.

Dataflow. During training, our method takes video \mathcal{V} , action narration \mathcal{A} , object names \mathcal{O} , object masks \mathcal{M} , and hand-object masks \mathcal{M}_{h-obj} as input, which are all from existing annotations. An ActionVOS model \mathcal{S} outputs classification of objects' involvement $\hat{Cls} \in [0,1]$ and mask predictions $\hat{\mathcal{M}}$, *i.e.*, $\hat{Cls}, \hat{\mathcal{M}} = \mathcal{S}(\mathcal{V}, \mathcal{A}, \mathcal{O})$. Given the input, the action-aware labeling module Φ generates pseudo-labels correspondingly, *i.e.*, $Cls, \mathcal{M}_{act} = \Phi(\mathcal{A}, \mathcal{O}, \mathcal{M}, \mathcal{M}_{h-obj})$. Along with the labeling module Φ , a generating function g generates pixel-wise weights for segmentation loss, *i.e.*, $\mathcal{W} = g(\mathcal{A}, \mathcal{O}, \mathcal{M}, \mathcal{M}_{h-obj})$.

4.1 ActionVOS Model

We construct an ActionVOS model S by adding an extra classification head to an RVOS model. Following state-of-the-art RVOS models [66, 71], which use a classification head to enhance segmentation performance, we add this classification head to distinguish positive and negative objects. This head predicts $\hat{Cls}(O_i) \in [0, 1]$ indicating the probability of object O_i being positive in action \mathcal{A} . During inference, we set a threshold θ to determine an object's positivity. Specifically, $\hat{Cls}(O_i)$ is used to adjust segmentation results $\hat{\mathcal{M}}(O_i)$ as follows:

$$\hat{\mathcal{M}}(O_i) = \begin{cases} \hat{\mathcal{M}}(O_i), & \hat{Cls}(O_i) \ge \theta\\ 0, & \hat{Cls}(O_i) < \theta, \end{cases}$$
(1)

where θ is set as 0.75 according to experiments in Sec. 5.4.

4.2 Action-aware Labeling Module

To address the problem that existing datasets lack indication of objects' involvement in actions, we propose an action-aware labeling module Φ to generate pseudo-labels of objects' involvement. By using annotations of action narrations, semantic segmentation and hand-object segmentation, we label three types of objects as positive based on the guidance of action narrations and hand-object masks as follows:

- 1) Objects mentioned in the action narrations.
- 2) Objects inside hand-object masks.
- 3) Objects that intersect with hand-object bounding boxes.

These three types correspond to the three definitions of "positive" in Sec. 3. For type 3), such a design identifies a large number of objects that are potentially positive, because these objects within close reach are highly likely to be relevant to the action, such as containers and contents.

Pseudo-labels generated by the labeling module Φ contains two parts: classification labels *Cls* and action-aware object masks \mathcal{M}_{act} . For each object O_i , its pseudo-label is formulated as:

$$Cls(O_i) = \begin{cases} 1, & O_i \in \mathcal{A} \\ 1, & \mathcal{M}(O_i) \in M_{h-obj} \\ 1, & \mathcal{M}(O_i) \cap B_{h-obj} \neq \emptyset \end{cases} \quad \mathcal{M}_{act}(O_i) = \begin{cases} \mathcal{M}(O_i), & Cls(O_i) = 1 \\ 0, & Cls(O_i) = 0, \end{cases}$$
(2)

where B_{h-obj} stands for the minimal bounding box of hand-object mask \mathcal{M}_{h-obj} . The generated pseudo-labels Cls, \mathcal{M}_{act} are used to train the ActionVOS model.

However, since Eq. (2) provides a more relaxed definition of positive compared to Sec. 3, practical issues may arise. While it correctly identifies many potential positives, it also introduces multiple false positives simultaneously. To address this issue, we introduce an action-guided focal loss in Sec. 4.3 aimed at reducing the impact of false positives.

4.3 Action-guided Focal Loss

To reduce the impact of false positives, action-guided focal loss FL_{act} is proposed by adding pixel-wise action-guided weights W to segmentation focal loss FL [33].

Fig. 5 (a) analyzes typical mistakes in action-aware object masks generated from Φ . In the "take container" example, although the object is in contact with left hand, it is not involved in this action. In the "put down pan" example, redundant instances of "pan" are not active even though they are mentioned by the narration.



Fig. 5: Action-aware object masks and action-guided weights. In action-guided weights, λ_{pos} is in red, λ_{nar} and λ_{h-obj} are in blue, λ_{neg} is in yellow.

To address these false positives, we adjust the weight when calculating the pixel-level segmentation loss, to make the segmentation model aware of interacted positive objects. We establish three rules:

- 1) Objects in both action narration and hand-object bounding box > those solely in either.
- 2) Objects mentioned in action narration or in contact with hands > those only intersecting with hand-object bounding boxes.
- 3) Objects labeled as negative are assigned a high weight within their masks.

Rules 1,2 aim to prioritize objects more likely to be actively involved among all potential positives, while the last rule penalizes the model for mis-segment negative objects as positive. As shown in Fig. 5 (b), the "container" (1st column), both mentioned in narration and in contact with right hand, receives higher weights than the left-hand object, which is only in contact with left hand (according to rule 1). The "pot" (3rd column), held by right hand, receives higher weight than the "pan" under left hand (according to rule 2).

Following these rules, we propose a generating function of action-guided weights \mathcal{W} . For a pixel (h, w) inside object O_i 's region, *i.e.*, $M_t(O_i, h, w) = 1$, its action-guided weight is:

$$W_t(O_i, h, w) = \begin{cases} \lambda_{pos}, & O_i \in \mathcal{A} \text{ and } B_{h-obj,t}(h, w) = 1\\ \lambda_{nar}, & O_i \in \mathcal{A} \text{ and } B_{h-obj,t}(h, w) = 0\\ \lambda_{h-obj}, & O_i \notin \mathcal{A} \text{ and } M_{h-obj,t}(h, w) = 1\\ \lambda_{neg}, & O_i \notin \mathcal{A} \text{ and } B_{h-obj,t}(h, w) = 0\\ 1, & \text{otherwise.} \end{cases}$$
(3)

The relationship of these weights are defined as $\lambda > 1$, $\lambda_{pos} > \lambda_{nar}$, and $\lambda_{pos} > \lambda_{h-obj}$. Empirically, we set $\lambda_{pos} = 5$, $\lambda_{nar} = 2$, $\lambda_{h-obj} = 2$, $\lambda_{neg} = 5$. For pixels outside O_i 's region, where $M_t(O_i, h, w) = 0$, the weight is set as 1.

The action-guided weights are added to the pixel-level segmentation focal loss [33]. Let $p_{h,w} \in [0,1]$ be the probability for pixel (h,w)'s positivity, which is predicted from the model's output $\hat{\mathcal{M}}$. $y_{h,w} \in \{0,1\}$ is the generated pseudolabels of the pixel on the same location, coming from \mathcal{M}_{act} . The action-guided focal loss FL_{act} for each frame is expressed as:

$$p_{t} = p_{h,w} \cdot y_{h,w} + (1 - p_{h,w}) \cdot (1 - y_{h,w}),$$

$$\alpha_{t} = \alpha \cdot y_{h,w} + (1 - \alpha) \cdot (1 - y_{h,w}),$$

$$FL_{act} = -\frac{1}{H * W} \sum_{h=1,w=1}^{H,W} W_{t}(O_{i}, h, w) \cdot \alpha_{t} \cdot (1 - p_{t})^{\gamma} \cdot \log(p_{t}),$$
(4)

where α is a balanced parameter, γ is a focusing parameter. Following previous work [66], we set $\alpha = 0.25$ and $\gamma = 2$.

As Eqs. (2) and (4) indicates, both the action-guided focal loss and actionaware labeling module have no trainable parameters. This parameter-free design allows our method to segment active objects with existing readily-available data. In addition, as our method's inputs and outputs align with the conventional RVOS task, it is compatible with existing RVOS models.

Hand-object masks \mathcal{M}_{h-obj} used in Eqs. (2) and (3) are obtained from the human annotation in VISOR [9] dataset. Hand-object masks are only used in training for generating pseudo-labels. During inference, ActionVOS model does not take hand-object masks as input and has no need to estimate the contact with hands.

5 Experiments

5.1 Experiment Settings

Datasets. We use three egocentric video object segmentation datasets, VISOR [9], VOST [57] and VSCOS [72] in experiments. Note that only VISOR is used for training the models, and we evaluate the models on three datasets. Details of the datasets are provided in the supplemental material.

Evaluation Metrics. Following [34], we employ mean IoU (mIoU), cumulative IoU (cIoU), generalized IoU (gIoU), and accuracy (Acc) as evaluation metrics. We report mIoU and cIoU separately for positive and negative objects, *i.e.*, p-mIoU, n-mIoU, p-cIoU, n-cIoU.

Model Settings. We apply ReferFormer (RF) [66] with different visual backbones as baseline RVOS models in our experiments. The backbone of ReferFormer can be replaced with ResNet-101 [18], Swin-L [39], or Video-Swin-Base [40]. RoBERTa [38] is employed as the text encoder, where its parameters are retrained in our experiment. The extra classification head is a linear layer, which receives averaged features from the last output layer [4, 81] to predict binary classification, defined by nn.Linear(256,1) in pytorch [47] implementation.

Training Details. All models are trained from best checkpoints on Refer-YouTube-VOS [53] benchmarks. We follow training settings of ReferFormer [66], including epochs, optimizer [41], loss coefficients [33, 46, 51] and data augmentations [63]. We replace the segmentation focal loss [33] with our proposed actionguided focal loss, and introduce a binary cross-entropy loss to train the additional classification head. The weight for the extra classification loss is set to 2.

5.2 ActionVOS Results

Quantitative results on VISOR. We analyze the segmentation performance of ActionVOS models. Here, the results of RVOS is provided as the upper bound of p-mIoU/p-cIoU, as it treats all objects as positive. As shown in Tab. 1, compared to RVOS, ActionVOS offers a significant n-mIoU/n-cIoU decrease while ensuring there is only a slight decrease in p-mIoU/p-cIoU. This indicates that ActionVOS significantly reduces the mis-segmentation of non-interacted objects while keeping the ability of segmenting active objects. We also evaluate our method by removing action prompts input and training with only object names. For example, we replace the language input "knife used in the action of cut apple" with "knife", while keeping the same pseudo-labels and loss weights. Experimental results show that the models trained with action prompts achieve much better performance, confirming that action prompts help ActionVOS models focus on active objects.

Tab. 1 also shows the results of replacing the backbone network with different structures, including ResNet-101 [18], Swin-L [39], and Video-Swin-Base [40]. With our proposed setting, all backbones achieve significant improvements on these metrics. Quantitative results demonstrate that the proposed ActionVOS is compatible with various existing network structures.

| Model | Setting | AP | p-mIoU \uparrow | n-mIoU \downarrow | p-cIoU \uparrow | n-cIoU \downarrow | $\mathrm{gIoU}\uparrow$ | Acc \uparrow |
|-----------|------------------------|----|---------------------|---------------------|---------------------|---------------------|-------------------------|---------------------|
| | RVOS* | | 67.7 | 54.2 | 73.2 | 67.5 | 43.8 | 59.1 |
| RF-R101 | ActionVOS ActionVOS | × | 56.3 65.4 | 19.9 19.0 | 61.2 72.4 | 32.8 32.7 | 66.8 70.9 | 72.9 82.4 |
| | RVOS* | | 71.8 | 59.7 | 79.9 | 73.0 | 46.8 | 59.4 |
| RF-SwinL | ActionVOS ActionVOS | × | 64.4 69.1 | 28.2 24.6 | 71.8 75.7 | 47.3 46.5 | 65.1 70.3 | 72.8 80.7 |
| | RVOS* | | 70.5 | 58.5 | 78.1 | 71.7 | 45.6 | 59.2 |
| RF-VSwinB | ActionVOS ActionVOS | × | 61.6 68.2 | 25.2 22.0 | 66.8 75.0 | 44.5 41.5 | 65.7 70.6 | 72.5 81.2 |

Table 1: Results of ActionVOS on VISOR. "AP" indicates whether action prompts are used. * indicates RVOS is the upper bound of p-IoUs under this experimental setting.

Table 2: Comparison with RVOS model under scenarios with object state changes.

| Setting | AP | n mIoII A | ST alou * | VSCOS | | |
|-----------|----|-------------|-----------|-------|------|--|
| RVOS | ×× | 29.3 | 17.5 | 46.4 | 44.9 | |
| ActionVOS | | 9.0 | 8.2 | 22.5 | 33.0 | |
| ActionVOS | | 32.3 | 22.8 | 49.4 | 49.6 | |

Quantitative results under object state changes. We evaluate models on VOST and VSCOS datasets, which consist of object state changes. As shown in Tab. 2, ActionVOS model outperforms RVOS across both datasets. This suggests that our method effectively handles scenarios involving object state changes, with action prompts providing enhanced understanding of state changes.

Table 3: Comparion with HOS [9] and RVOS [66]+S4.2. RVOS model is RF-R101.

| Method | p-mIoU ↑ | n-mIoU↓ | p-cIoU ↑ | n-cIoU ↓ | gIoU \uparrow | Acc \uparrow | p-mIoU | <u>ST</u> p-cIoU | p-mIoU | DS p-cIoU |
|-------------------------------|------------------------|------------------------|------------------------|---|-----------------------------|-----------------------------|-----------------------------|------------------------|------------------------|---------------------------------|
| HOS RVOS+S4.2 ActionVOS | $56.2 \\ 65.3 \\ 65.4$ | $11.4 \\ 35.2 \\ 19.0$ | $58.1 \\ 71.5 \\ 72.4$ | $ \begin{array}{r} {f 16.8} \\ 56.4 \\ 32.7 \end{array} $ | 68.8 60.4 70.9 | 77.0 75.1 82.4 | 19.4 29.3 32.3 | $13.1 \\ 17.5 \\ 22.8$ | $34.4 \\ 46.4 \\ 49.4$ | $24.1 \\ 44.9 \\ \textbf{49.6}$ |

Comparison with baseline methods. We compare ActionVOS model with baseline methods on three datasets. The baseline methods are: 1)Hand-object segmentation (HOS) model. We take the best HOS model in [9] as a baseline model, which segments hands and hand-objects as positive object masks. 2)RVOS+S4.2. We use Eq. (2) in Sec. 4.2 as a post-process of RVOS model outputs. Note that Eq. (2) take ground-truth hand-object masks as input. The comparisons are shown in Tab. 3. ActionVOS model outperforms other baselines in terms of positive IoUs, gIoU, and accuracy. HOS model has lower negative IoUs, because it only segment hands and hand-objects. RVOS+S4.2 shows worse results since S4.2 brings amounts of false positives.

Action Vocabulary. To evaluate the open-vocabulary setting of ActionVOS, we conduct experiments on unseen action categories. Refer to the supplemental material for statistics, evaluations and discussions on action vocabulary.



Fig. 6: Visualization results of ActionVOS models trained w/ and w/o action prompts.



Fig. 7: Segmentation results of ActionVOS in the same scenes. All frames share the same input object names.

5.3 Qualitative results

Action prompts. Fig. 6 compares segmentation predictions from models trained with and without action prompts. Given the same input object names, model trained with action prompts correctly identifies objects involved in the action. In contrast, model trained with only object names tend to segment inactive objects, *e.g.*, redundant instances.

Effect of action prompts in identical scenes. Fig. 7 shows ActionVOS's segmentation results in the same scenes. Even though those actions occur within



Fig. 8: Segmentation results of ActionVOS under Fig. 9: The effect of the proposed scenarios with object state changes. action-guided focal loss FL_{act} .

Table 4: Ablation Table 5: The impact of language for the classi- prompts and fine-tuning text en- Table 6. The effect of study fication head and its coders. threshold θ .

| Table 0. | The effect of |
|-------------|---------------|
| proposed | action-guided |
| focal loss. | |
| | |

| | | | | Prompt | Tuned | p-mIoU | n-mIoU | gIoU | Acc |
|--|----------------------|-----------------------------|--|------------------------------|-------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|
| | p-mIoU | n-mIoU | gIoU Acc | NoAction | × | 54.5 | 17.7 | 66.1 | 71.3 |
| No head | 67.4 | 31.2 | $63.2 \ 76.0$ | $^{+,Action}_{+sAction}$ | x | 60.4 61.7 | 17.0 16.3 | 70.0 70.6 | 79.5 80.2 |
| $ \begin{aligned} \theta &= 0.25 \\ \theta &= 0.50 \\ \theta &= 0.75 \end{aligned} $ | 67.6 67.5 65.4 | 32.7 30.6 19.0 | 61.8 75.3 63.5 76.7 70.9 82.4 | NoAction +, Action + sAction | 1 | 56.3 65.1 65.4 | 19.9 19.6 19.0 | 66.8 70.8 70.9 | 72.9 82.4 82.4 |

| 1 | p-mIoU | n-mIoU | JgIoU Acc |
|-----------------------|---|---------------------|-----------------------|
| $FL [33] \\ FL_{act}$ | $\begin{array}{c} 65.4 \\ 65.4 \end{array}$ | 20.5 19.0 | 70.6 82.1 70.982.4 |

the same scene and share identical input object names, our method still correctly segments active objects. This underscores the model's comprehension of humanobject interaction, facilitated by action narrations.

Effect of action prompts under state changes. Fig. 8 visualizes the segmentation results of ActionVOS in comparison with RVOS under scenarios with object state change. RVOS method fails to segment objects after a change of state, such as broken eggshells and yolks, and sliced lemons. In contrast, our method successfully identifies these state-changed objects, confirming that action prompts help to enhance understanding of state changing. More visualization results are provided in the supplemental material.

$\mathbf{5.4}$ Ablations

We perform extensive ablation studies to analyze the impact of components of ActionVOS. All ablations are conducted on VISOR dataset and RF-R101.

Classification Head. As illustrated in Sec. 4.1, there is an extra classification head in ActionVOS model to predict objects' positivity. We compare the segmentation results with and without this classification head, and we also test different threshold θ of this binary classification during inference (Eq. (1)). As shown in Tab. 4, using the classification head brings improvements to all metrics, indicating that the head has a strong ability to distinguish active objects, which is a simple yet efficient modification to the model. For the threshold θ , using a higher threshold significantly reduces the mis-segmentation of negative objects (decrease in n-mIoU) while ensuring there is only a slight decrease in

p-mIoU. Considering the best trade-off between positive and negative IoUs, we set $\theta = 0.75$ for all the experiments.

Text Prompts. We compare three types of language prompts, as the design of language prompts is important in language and vision-language tasks [45,50,55, 56,79]. These three types are as follows:

- NoAction. The text prompt is the object class name. e.g., "knife".
- +, Action. The text prompt is the object class name and the action narration combined with a comma. e.g., "knife, cut apple".
- +sAction. The text prompt is a natrual sentence consisting of the object class name and action narration, e.g., "knife used in the action of cut apple".

We compare these three types with a text decoder frozen and tuned. As Tab. 5 shows, no matter if the text encoder is tuned, +sAction enhances the segmentation results most. This indicates that using a natural language description helps the segmentation model better understand the action. In other experiments, the text encoder is tuned and the language prompt is +sAction.

Action-guided Focal Loss. In Sec. 4.3, the action-guided focal loss FL_{act} is proposed to reduce the impact of false positives, prioritize truly active objects. We compare the proposed FL_{act} with focal loss FL [33] in Tab. 6. The proposed loss function offers improvements to the gIoU and Acc, and decrease in n-mIoU. This indicates that the impact of false positives have been reduced. Visualization in Fig. 9 shows the segmentation results when there are objects in both hands. The model trained with the proposed action-guided focal loss well prioritizes truly active objects and ignores those false positives. In the "put olive oil" example, a bottle of olive oil is in the left hand while a bottle of salt is in the right hand. The model trained without action-guided focal loss failed to segment the olive oil and predicted the salt as positive. In contrast, the model trained with our proposed loss successfully segments the olive oil as the only active object.

6 Conclusion

In this paper, we propose ActionVOS, a novel action-aware setting for referring video object segmentation. This setting segments active objects in egocentric videos by employing action narrations as an additional language prompt. Specifically, we develop an action-aware labeling module and an action-guided focal loss for ActionVOS. This design enables ActionVOS models to segment active objects with existing readily-available annotations. As for future work, we consider extending ActionVOS by incorporating various action-object relations, reducing the heavy reliance on the availability of dense annotations, and adapting ActionVOS in open-world applications.

Acknowledgements

This work was supported by JST ASPIRE Grant Number JPMJAP2303, JST SPRING Grant Number JPMJSP2108, JSPS KAKENHI Grant Numbers JP22KF0119, JP23H00488, JP24K02956.

References

- Bertasius, G., Park, H.S., Stella, X.Y., Shi, J.: First-person action-object detection with egonet. In: Robotics: Science and Systems (2017)
- Botach, A., Zheltonozhskii, E., Baskin, C.: End-to-end referring video object segmentation with multimodal transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4985–4995 (2022)
- 3. Cai, M., Kitani, K.M., Sato, Y.: Understanding hand-object manipulation with grasp types and object attributes. In: Robotics: Science and Systems. vol. 3 (2016)
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: Proceedings of the European Conference on Computer Vision. pp. 213–229 (2020)
- Chen, Z., Ma, L., Luo, W., Wong, K.Y.K.: Weakly-supervised spatio-temporally grounding natural sentence in video. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 1884–1894 (2019)
- Cheng, H.K., Oh, S.W., Price, B., Schwing, A., Lee, J.Y.: Tracking anything with decoupled video segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1316–1326 (2023)
- Damen, D., Doughty, H., Farinella, G.M., Fidler, S., Furnari, A., Kazakos, E., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Scaling egocentric vision: The epic-kitchens dataset. In: Proceedings of the European Conference on Computer Vision. pp. 720–736 (2018)
- Damen, D., Doughty, H., Farinella, G.M., Furnari, A., Kazakos, E., Ma, J., Moltisanti, D., Munro, J., Perrett, T., Price, W., et al.: Rescaling egocentric vision: Collection, pipeline and challenges for epic-kitchens-100. International Journal of Computer Vision pp. 1–23 (2022)
- Darkhalil, A., Shan, D., Zhu, B., Ma, J., Kar, A., Higgins, R., Fidler, S., Fouhey, D., Damen, D.: Epic-kitchens visor benchmark: Video segmentations and object relations. Advances in Neural Information Processing Systems 35, 13745–13758 (2022)
- De Vries, H., Strub, F., Chandar, S., Pietquin, O., Larochelle, H., Courville, A.: Guesswhat?! visual object discovery through multi-modal dialogue. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5503– 5512 (2017)
- Deruyttere, T., Vandenhende, S., Grujicic, D., Van Gool, L., Moens, M.F.: Talk2car: Taking control of your self-driving car. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing. pp. 2088–2098 (2019)
- Ding, H., Liu, C., He, S., Jiang, X., Loy, C.C.: Mevis: A large-scale benchmark for video segmentation with motion expressions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2694–2703 (2023)
- Ding, H., Liu, C., Wang, S., Jiang, X.: Vlt: Vision-language transformer and query generation for referring segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022)
- Fu, Q., Liu, X., Kitani, K.: Sequential voting with relational box fields for active object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2374–2383 (2022)
- Gavrilyuk, K., Ghodrati, A., Li, Z., Snoek, C.G.: Actor and action video segmentation from a sentence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 5958–5966 (2018)

- 16 L. Ouyang et al.
- Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al.: Ego4d: Around the world in 3,000 hours of egocentric video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18995–19012 (2022)
- Gupta, A., Davis, L.S.: Objects in action: An approach for combining action understanding and object perception. In: 2007 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–8 (2007)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 770–778 (2016)
- He, S., Ding, H., Liu, C., Jiang, X.: Grec: Generalized referring expression comprehension. arXiv preprint arXiv:2308.16182 (2023)
- Higgins, R.E.L., Fouhey, D.F.: Moves: Manipulated objects in video enable segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6334–6343 (2023)
- Hu, R., Rohrbach, M., Darrell, T.: Segmentation from natural language expressions. In: Proceedings of the European Conference on Computer Vision. pp. 108–124. Springer (2016)
- Kamath, A., Singh, M., LeCun, Y., Synnaeve, G., Misra, I., Carion, N.: Mdetrmodulated detection for end-to-end multi-modal understanding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1780–1790 (2021)
- Kazemzadeh, S., Ordonez, V., Matten, M., Berg, T.: Referitgame: Referring to objects in photographs of natural scenes. In: Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing. pp. 787–798 (2014)
- Khoreva, A., Rohrbach, A., Schiele, B.: Video object segmentation with language referring expressions. In: Asian Conference on Computer Vision. pp. 123–141 (2018)
- Kjellström, H., Romero, J., Kragić, D.: Visual object-action recognition: Inferring object affordances from human demonstration. Computer Vision and Image Understanding 115(1), 81–90 (2011)
- Krüger, N., Geib, C., Piater, J., Petrick, R., Steedman, M., Wörgötter, F., Ude, A., Asfour, T., Kraft, D., Omrčen, D., et al.: Object–action complexes: Grounded abstractions of sensory–motor processes. Robotics and Autonomous Systems 59(10), 740–757 (2011)
- Kurita, S., Katsura, N., Onami, E.: Refego: Referring expression comprehension dataset from first-person perception of ego4d. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15214–15224 (2023)
- Lee, C., Kumar, M.G., Tan, C.: Determinet: A large-scale diagnostic dataset for complex visually-grounded referencing using determiners. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 20019–20028 (2023)
- Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., et al.: Grounded language-image pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10965–10975 (2022)
- Li, X., Wang, J., Xu, X., Li, X., Raj, B., Lu, Y.: Robust referring video object segmentation with cyclic structural consensus. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 22236–22245 (2023)
- Li, Z., Tao, R., Gavves, E., Snoek, C.G., Smeulders, A.W.: Tracking by natural language specification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6495–6503 (2017)

- 32. Lin, K.Q., Wang, J., Soldan, M., Wray, M., Yan, R., XU, E.Z., Gao, D., Tu, R.C., Zhao, W., Kong, W., et al.: Egocentric video-language pretraining. Advances in Neural Information Processing Systems 35, 7575–7586 (2022)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2980–2988 (2017)
- Liu, C., Ding, H., Jiang, X.: Gres: Generalized referring expression segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23592–23601 (2023)
- Liu, J., Ding, H., Cai, Z., Zhang, Y., Satzoda, R.K., Mahadevan, V., Manmatha, R.: Polyformer: Referring image segmentation as sequential polygon generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18653–18663 (2023)
- Liu, R., Ohkawa, T., Zhang, M., Sato, Y.: Single-to-dual-view adaptation for egocentric 3d hand pose estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 677–686 (June 2024)
- Liu, R., Liu, C., Bai, Y., Yuille, A.L.: Clevr-ref+: Diagnosing visual reasoning with referring expressions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4185–4194 (2019)
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692 (2019)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10012–10022 (2021)
- Liu, Z., Ning, J., Cao, Y., Wei, Y., Zhang, Z., Lin, S., Hu, H.: Video swin transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3202–3211 (2022)
- 41. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations (2018)
- 42. Lüddecke, T., Ecker, A.: Image segmentation using text and image prompts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7086–7096 (2022)
- Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A.L., Murphy, K.: Generation and comprehension of unambiguous object descriptions. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 11–20 (2016)
- 44. Mei, J., Piergiovanni, A., Hwang, J.N., Li, W.: Slvp: Self-supervised languagevideo pre-training for referring video object segmentation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 507–517 (2024)
- Miao, Z., Zhao, K., Tsuruoka, Y.: Improving arithmetic reasoning ability of large language models through relation tuples, verification and dynamic feedback. arXiv preprint arXiv:2406.17873 (2024)
- Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth International Conference on 3D Vision. pp. 565–571 (2016)
- 47. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: an imperative style, high-performance deep learning library. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems. pp. 8026–8037 (2019)

- 18 L. Ouyang et al.
- Perazzi, F., Pont-Tuset, J., McWilliams, B., Van Gool, L., Gross, M., Sorkine-Hornung, A.: A benchmark dataset and evaluation methodology for video object segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 724–732 (2016)
- Qi, Y., Wu, Q., Anderson, P., Wang, X., Wang, W.Y., Shen, C., Hengel, A.v.d.: Reverie: Remote embodied visual referring expression in real indoor environments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9982–9991 (2020)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763 (2021)
- Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 658–666 (2019)
- Rodin, I., Furnari, A., Min, K., Tripathi, S., Farinella, G.M.: Action scene graphs for long-form understanding of egocentric videos. arXiv preprint arXiv:2312.03391 (2023)
- Seo, S., Lee, J.Y., Han, B.: Urvos: Unified referring video object segmentation network with a large-scale benchmark. In: Proceedings of the European Conference on Computer Vision. pp. 208–223 (2020)
- Shan, D., Geng, J., Shu, M., Fouhey, D.F.: Understanding human hands in contact at internet scale. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9869–9878 (2020)
- 55. Shin, T., Razeghi, Y., Logan IV, R.L., Wallace, E., Singh, S.: Autoprompt: Eliciting knowledge from language models with automatically generated prompts. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing. pp. 4222–4235 (2020)
- Tateno, M., Yagi, T., Furuta, R., Sato, Y.: Learning object states from actions via large language models. arXiv preprint arXiv:2405.01090 (2024)
- 57. Tokmakov, P., Li, J., Gaidon, A.: Breaking the" object" in video object segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22836–22845 (2023)
- Wang, P., Wang, S., Lin, J., Bai, S., Zhou, X., Zhou, J., Wang, X., Zhou, C.: Onepeace: Exploring one general representation model toward unlimited modalities. arXiv preprint arXiv:2305.11172 (2023)
- Wang, P., Yang, A., Men, R., Lin, J., Bai, S., Li, Z., Ma, J., Zhou, C., Zhou, J., Yang, H.: Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In: International Conference on Machine Learning. pp. 23318–23340. PMLR (2022)
- Wang, W., Zhang, Y., He, X., Yan, Y., Zhao, Z., Wang, X., Liu, J.: Beyond literal descriptions: Understanding and locating open-world objects aligned with human intentions. arXiv preprint arXiv:2402.11265 (2024)
- 61. Wang, X., Shu, X., Zhang, Z., Jiang, B., Wang, Y., Tian, Y., Wu, F.: Towards more flexible and accurate object tracking with natural language: Algorithms and benchmark. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13763–13773 (2021)
- Wang, X., Zhang, X., Cao, Y., Wang, W., Shen, C., Huang, T.: Seggpt: Towards segmenting everything in context. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1130–1140 (2023)

- 63. Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-end video instance segmentation with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8741–8750 (2021)
- Wu, C., Lin, Z., Cohen, S., Bui, T., Maji, S.: Phrasecut: Language-based image segmentation in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10216–10225 (2020)
- Wu, D., Han, W., Wang, T., Dong, X., Zhang, X., Shen, J.: Referring multi-object tracking. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14633–14642 (2023)
- Wu, J., Jiang, Y., Sun, P., Yuan, Z., Luo, P.: Language as queries for referring video object segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4974–4984 (2022)
- Wu, J., Jiang, Y., Yan, B., Lu, H., Yuan, Z., Luo, P.: Segment every reference object in spatial and temporal spaces. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2538–2550 (2023)
- Wu, T.L., Zhou, Y., Peng, N.: Localizing active objects from egocentric vision with symbolic world knowledge. In: Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing. pp. 4991–5006 (2023)
- Xu, N., Yang, L., Fan, Y., Yang, J., Yue, D., Liang, Y., Price, B., Cohen, S., Huang, T.: Youtube-vos: Sequence-to-sequence video object segmentation. In: Proceedings of the European Conference on Computer Vision. pp. 585–601 (2018)
- Yamaguchi, M., Saito, K., Ushiku, Y., Harada, T.: Spatio-temporal person retrieval via natural language queries. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 1453–1462 (2017)
- 71. Yan, B., Jiang, Y., Wu, J., Wang, D., Luo, P., Yuan, Z., Lu, H.: Universal instance perception as object discovery and retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15325–15336 (2023)
- Yu, J., Li, X., Zhao, X., Zhang, H., Wang, Y.X.: Video state-changing object segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 20439–20448 (2023)
- Yu, L., Lin, Z., Shen, X., Yang, J., Lu, X., Bansal, M., Berg, T.L.: Mattnet: Modular attention network for referring expression comprehension. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1307–1315 (2018)
- Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: Proceedings of the European Conference on Computer Vision. pp. 69–85. Springer (2016)
- Zhang, C., Gupta, A., Zisserman, A.: Helping hands: An object-aware ego-centric video recognition model. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13901–13912 (2023)
- Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: Glipv2: Unifying localization and vision-language understanding. Advances in Neural Information Processing Systems 35, 36067–36080 (2022)
- Zhang, L., Zhou, S., Stent, S., Shi, J.: Fine-grained egocentric hand-object segmentation: Dataset, model, and applications. In: Proceedings of the European Conference on Computer Vision. pp. 127–145 (2022)
- 78. Zhang, M., Huang, Y., Liu, R., Sato, Y.: Masked video and body-worn imu autoencoder for egocentric action recognition. arXiv preprint arXiv:2407.06628 (2024)
- Zhou, K., Yang, J., Loy, C.C., Liu, Z.: Learning to prompt for vision-language models. International Journal of Computer Vision 130(9), 2337–2348 (2022)

- 20 L. Ouyang et al.
- Zhu, C., Xiao, F., Alvarado, A., Babaei, Y., Hu, J., El-Mohri, H., Chang, S., Sumbaly, R., Yan, Z.: Egoobjects: A large-scale egocentric dataset for fine-grained object understanding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2023)
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: International Conference on Learning Representations (2020)