

# Segment-to-Act: Label-Noise-Robust Action-Prompted Video Segmentation Towards Embodied Intelligence

Wenxin Li<sup>1</sup>, Kunyu Peng<sup>2,3,\*</sup>, Di Wen<sup>2</sup>, Ruiping Liu<sup>2</sup>, Mengfei Duan<sup>1</sup>, Kai Luo<sup>1</sup>, and Kailun Yang<sup>1,\*</sup>

**Abstract**—Embodied intelligence relies on accurately segmenting objects actively involved in interactions. Action-based video object segmentation addresses this by linking segmentation with action semantics, but it depends on large-scale annotations and prompts that are costly, inconsistent, and prone to multimodal noise such as imprecise masks and referential ambiguity. To date, this challenge remains unexplored. In this work, we take the first step by studying action-based video object segmentation under label noise, focusing on two sources: textual prompt noise (category flips and within-category noun substitutions) and mask annotation noise (perturbed object boundaries to mimic imprecise supervision). Our contributions are threefold. First, we introduce two types of label noises for the action-based video object segmentation task. Second, we build up the first action-based video object segmentation under a label noise benchmark *ActiSeg-NL* and adapt six label-noise learning strategies to this setting, and establish protocols for evaluating them under textual, boundary, and mixed noise. Third, we provide a comprehensive analysis linking noise types to failure modes and robustness gains, and we introduce a Parallel Mask Head Mechanism (*PMHM*) to address mask annotation noise. Qualitative evaluations further reveal characteristic failure modes, including boundary leakage and mislocalization under boundary perturbations, as well as occasional identity substitutions under textual flips. Our comparative analysis reveals that different learning strategies exhibit distinct robustness profiles, governed by a foreground-background trade-off where some achieve balanced performance while others prioritize foreground accuracy at the cost of background precision. These results establish a clear sensitivity profile of action-based video object segmentation to imperfect annotations and set a benchmark for studying noise-robust learning in embodied perception. The established benchmark and source code will be made publicly available at <https://github.com/mylwx/ActiSeg-NL>.

## I. INTRODUCTION

Video Object Segmentation (VOS) is a core perception substrate for embodied intelligence because it anchors language instructions to pixel-level object representations that

This work was supported in part by the National Natural Science Foundation of China (Grant No. 62473139), in part by the Hunan Provincial Research and Development Project (Grant No. 2025QK3019), in part by the State Key Laboratory of Autonomous Intelligent Unmanned Systems (the opening project number ZZKF2025-2-10), and in part by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - SFB 1574 - 471687386.

<sup>1</sup>The authors are with the School of Artificial Intelligence and Robotics and the National Engineering Research Center of Robot Visual Perception and Control Technology, Hunan University, China (email: kailun.yang@hnu.edu.cn).

<sup>2</sup>The authors are with the Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Germany (email: firstname.lastname@kit.edu).

<sup>3</sup>The author is also with INSAIT, Sofia University “St. Kliment Ohridski”, Bulgaria.

\*Corresponding authors: Kailun Yang and Kunyu Peng.

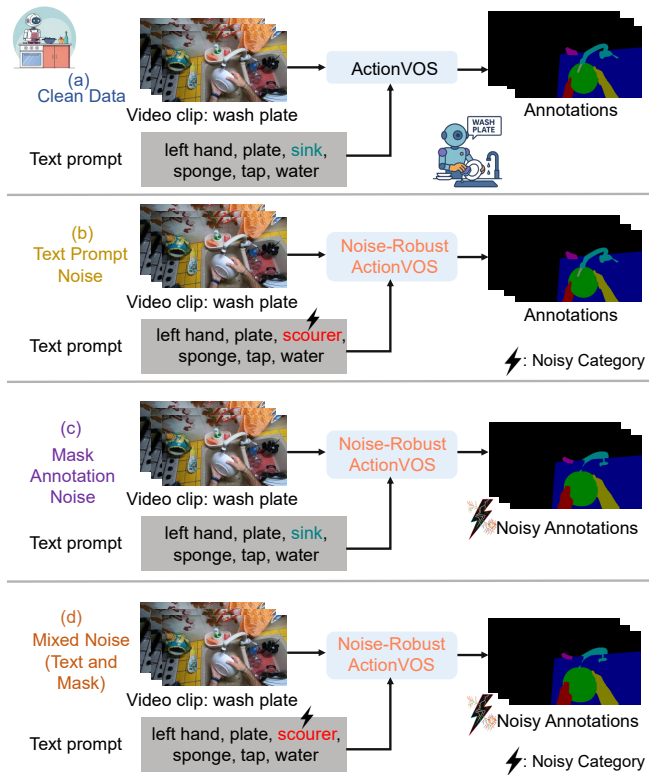


Fig. 1. Training pipeline of action-based video object segmentation (ActionVOS) with one clean reference (a) and three controlled noise scenarios defined by *ActiSeg-NL*, including text prompt noise (b), mask annotation noise (c), and their mixed condition (d). These scenarios approximate perception disturbances observed in egocentric recordings and in robot manipulation, and they test whether segmentation remains stable enough to support downstream action. We then apply and evaluate various adapted noise-robust strategies on each scenario for their mitigation effectiveness.

support planning and control in instruction following, re-arrangement, and manipulation settings [1]. VOS has progressed from classical semi-supervised settings in which the target object is specified by a mask in the first frame to language-guided paradigms [2], [3]. Referring Video Object Segmentation extends VOS by conditioning object segmentation on free-form textual expressions that describe specific targets in a video [4], [5]. As a specialization of R-VOS, action-based video object segmentation (ActionVOS) [6] refines the task by employing structured action narrations as prompts to segment objects actively engaged in human actions. This design facilitates downstream applications in

human-robot interaction and embodied perception, where agents must select task-relevant objects and execute manipulation under egocentric conditions [7].

However, deploying action-based video object segmentation in realistic scenarios presents significant practical challenges. Specifically, large-scale mask annotation in egocentric videos is prohibitively expensive [8], with annotation effort growing proportionally to object density and video length [9]. Additionally, annotations often suffer from subjective inconsistencies across annotators, further complicating data reliability [10]. Meanwhile, linguistic prompts also introduce referential noise, including lexical confusions such as “flour” versus “four” observed in VISOR [11]. Despite its practical importance, such multimodal noise remains largely underexplored, limiting the robustness and generalizability of current action-based video object segmentation models.

Existing noisy-label learning methods have primarily focused on image classification [12], [13], with limited extensions to video segmentation tasks [14]. Moreover, conventional strategies such as sample selection [15] and robust loss functions [16], [17] are ill-suited for action-based video object segmentation, as they fail to simultaneously handle the intertwined pixel-level and language-level label noise. Linguistic noise, in particular, introduces semantic ambiguities and referential mismatches [18], posing unique challenges beyond pixel-level errors.

In order to fill this gap, we introduce *ActiSeg-NL*, a benchmark designed explicitly to assess the robustness of action-based video object segmentation. Fig. 1 depicts the action-based video object segmentation training pipeline. *ActiSeg-NL* injects synthetic noise into the VISOR training split and defines three controlled scenarios, namely text prompt noise, mask annotation noise, and a mixed condition. We adapt a wide spectrum of noise-robust learning strategies to the action-based video object segmentation setting, including Co-teaching [19], Generalized Cross Entropy (GCE) [16], Symmetric Cross Entropy (SCE) [20], Active Passive Loss (APL) [21], Early Learning Regularization (ELR) [13], and NPN, which integrates partial label learning and negative learning [12]. Our benchmark reveals the limitations of these methods under text prompt noise, mask annotation noise, and the mixed condition within action-based video object segmentation, and underscores the need for solutions tailored to embodied segmentation.

Our contributions are threefold:

- We initiate the study of noisy labels in action-based video object segmentation by formalizing *text prompt noise*, *mask annotation noise*, and *mixed noise*, establishing a noisy-label taxonomy that links semantic ambiguity in language prompts with boundary imprecision in pixel masks.
- We introduce *ActiSeg-NL*, the first benchmark for this setting, and unify adaptations of a broad range of diverse noisy-label learning strategies, including Co-teaching [19], GCE [16], SCE [20], APL [21], ELR [13], and NPN [12], to the pixel-level and language-conditioned setting, providing the first large-

scale comparison of noise-robust training in action-based video object segmentation.

- Our benchmark analysis shows that robustness is not monolithic. Text noise challenges grounding and reduces foreground coverage, while Co-teaching preserves foreground regions, whereas mask noise degrades overlap, and APL improves overlap. Under mixed conditions, pixelwise losses (*e.g.*, GCE, SCE) are more stable than sample filtering, with ELR increasing positive regions at the cost of background precision. Additionally, we propose *Parallel Mask Head Mechanism (PMHM)*, a parallel-head consistency scheme that mitigates mask noise in action-based video object segmentation.

## II. RELATED WORK

### A. Referring Video Object Segmentation

Referring Video Object Segmentation (R-VOS) [4], [5] aims to accurately segment target objects in video sequences using text prompts, enabling applications in human-computer interaction [22] and video analysis [23]. Recent frameworks have considerably pushed forward progress in this domain. MTTR [4] introduces an end-to-end multimodal Transformer that simplifies previously multi-stage workflows. ReferFormer [5] employs language prompts as dynamic queries for efficient mask generation. For handling complex scenes, SgMg [24] incorporates spectrum-guided cross-modal fusion. FindTrack [25] enhances robustness by explicitly decoupling object recognition from mask propagation. Action-based video object segmentation extends R-VOS by incorporating action narrations into text prompts to segment only objects that actively participate in the ongoing action (*e.g.*, narrated target objects, hands, handheld tools, and involved containers and contents), reducing mis-segmentation of inactive objects [6]. While existing datasets lack explicit participation labels and thus depend on potentially erroneous pseudo labels, Action-based video object segmentation mitigates noise in active versus inactive labels [6], yet robustness to text prompt noise and mask annotation noise remains underexplored. We introduce *ActiSeg-NL*, the first benchmark that systematically considers action-based video object segmentation under text prompt noise, mask annotation noise, and their mixed setting.

### B. Learning with Noisy Labels

Noisy labels [26], [27], prevalent in large-scale datasets, significantly impair deep neural network performance due to their reliance on high-quality annotations. Song *et al.* [26] study symmetric and asymmetric noise in image classification. Luo *et al.* [28] leverage high-level spatial structures as supervisory signals to mitigate mislabeled annotations in segmentation tasks. Feng *et al.* [29] enhance weakly supervised segmentation with online pseudo-mask correction. In video object segmentation, Enki *et al.* [14] propose a space-mask-based framework to correct noise using spatial context, whereas Kimhi *et al.* [30] establish benchmarks like COCO-N and VIPER-N [30] to evaluate instance segmentation robustness against class confusion and boundary distortion.

### Algorithm 1 Text Prompt Noise Generation

**Input:** Clean Category  $\mathcal{C}$ , noise rate  $\rho$ , class mapping  $\mathcal{M} : \text{class} \rightarrow \{\mathcal{C}_1, \mathcal{C}_2, \dots\}$ , classes  $\Omega = \{1, \dots, K\}$ .

**Output:** Noisy category  $\mathcal{C}_{\text{noisy}}$ .

**Initialize:**  $\text{class}_c \leftarrow \text{class of } \mathcal{C}$ .

- 1: **if**  $\text{Uniform}(0, 1) < \rho$  **then** ▷ Class flipping
- 2:     Select  $\text{class}_{\text{flip}} \in \Omega \setminus \{\text{class}_c\}$  uniformly at random.
- 3:     Set  $\mathcal{C}_{\text{noisy}} \leftarrow \text{random category from } \mathcal{M}(\text{class}_{\text{flip}})$ .
- 4: **else** ▷ Synonymous replacement
- 5:     Set  $\mathcal{C}_{\text{noisy}} \leftarrow \text{random category from } \mathcal{M}(\text{class}_c)$ .
- 6: **end if**
- 7: **Return:**  $\mathcal{C}_{\text{noisy}}$ .




Fig. 2. Comparison of category distributions before and after applying text prompt noise at 20%, 40%, and 60% noise rates for major categories (proportion  $>1\%$ ) and Others.

However, many of these methods [26], [27] operate at the image level, and most pixel-level approaches [14], [29], [31] do not incorporate text guidance, so they are not directly applicable to action-based video object segmentation, which requires supervision at the pixel level conditioned on language and vision. We fill this gap by formalizing label noise for text prompts and mask annotations, as illustrated in Fig. 1, and by introducing *ActiSeg-NL*, the first benchmark that unifies adaptations of diverse noisy-label learning strategies for action-based video object segmentation.

### III. BENCHMARK

*ActiSeg-NL* is a benchmark for evaluating robustness in action-based video object segmentation under text, mask, and mixed noise, focusing on failure modes in embodied perception. We generate these controlled noises on the VISOR [11] training split for model training and evaluate on the original clean data. This section details our noise construction methods.

#### A. Text Prompt Noise

**Construction.** We generate text prompt noise for the VISOR training split on target objects with two components. First, symmetric noise is built by randomly flipping referent objects in text prompts, such as replacing “container” with “fridge” in the prompt “container used in the action of take celery” with probabilities of 20%, 40%, or 60%, which represent low, medium, and high noise levels commonly used in noisy label learning to assess robustness, following [13], [21]. The target




Fig. 3. Mask annotation noise generation and severity statistics in *ActiSeg-NL*. (a) Generation pipeline: clean instance masks are separated, dilated with different kernel sizes, and recombined to produce noisy annotations. (b) Summary statistics: larger kernel leads to lower mIoU and cIoU, quantifying the decline in annotation quality.

category is randomly selected from a predefined set to ensure equal replacement probability. Second, we substitute nouns with synonymous or related terms within the same category, such as replacing “container” with “food container” or “cheese container” to increase semantic complexity, by uniformly sampling the replacement from  $\mathcal{M}(\text{class}_c)$  when the class flipping condition is not met, as described in Alg. 1. The whole procedure is summarized in Alg. 1.

**Analysis.** We diagnose distribution shifts induced by text prompt noise. To validate this noise strategy, we analyze category distributions in the original and noisy training sets. Fig. 2 compares major categories (proportion  $>1\%$  in the original set) and a summarized “Others” category across clean data and 20%, 40%, and 60% noise rates. These statistics show that “left hand” and “right hand” proportions decrease from 14.58% and 15.60% to 6.08% and 6.39% at the 60% noise rate, respectively, while “Others” increases from 44.43% to 75.47%. This shift validates the disruptive effect of our noise strategy, as depicted in Fig. 1, establishing *ActiSeg-NL* as a dedicated benchmark for exploring text prompt noise in video segmentation.

#### B. Mask Annotation Noise

**Construction.** We employ a separate-dilate-combine strategy to simulate mask annotation noise, replicating realistic boundary errors in crowd-sourced annotations for positive objects in mask segmentation tasks. First, we separate clean semantic segmentation annotations into binary masks for each object, such as “right hand” and “fridge” in the “open fridge” action. Next, we apply morphological dilation with a square kernel (kernel sizes 9, 15, or 21) to expand object boundaries, mimicking the boundary blur observed in human annotations. Finally, we combine the dilated per-object masks into noisy annotations using a first-hit rule with a fixed object order, assigning each pixel the label of the first dilated mask that covers it while disregarding subsequent overlaps, thereby ensuring a single label per pixel. The left of Fig. 3 illustrates this process with an example frame. The entire procedure is summarized in Alg. 2.

**Analysis.** By merging the dilated per-object masks with a first hit rule under a fixed object order, overlaps create boundary competition that resembles the blur observed in

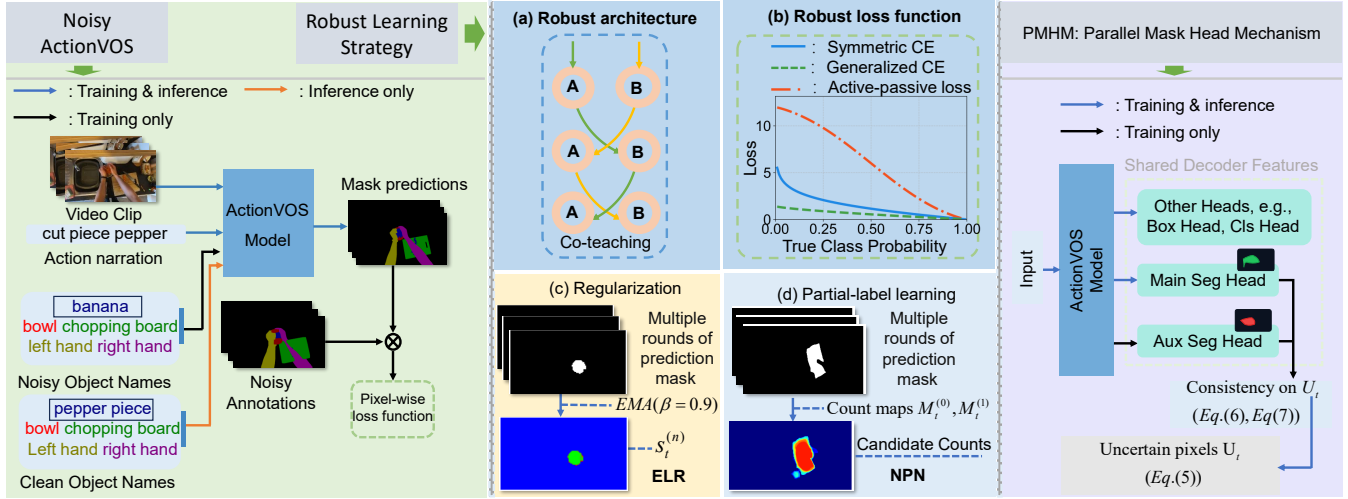


Fig. 4. Overview of the noisy action-based video object segmentation framework and robustness strategies. **Left:** the framework consumes video frames, noisy object names (e.g., “banana”), and action narrations (e.g., “cut piece pepper”), to predict a pixel-level mask. **Middle:** four complementary strategies, (a) *Co-teaching*, where two networks exchange small-loss samples to suppress label noise, (b) *Noise-robust Losses* GCE, SCE, and APL that balance accuracy and robustness, (c) *ELR*, which mitigates overfitting to noisy annotations with an EMA-based regularizer, (d) *NPN*, which integrates candidate-set reasoning with negative learning and with consistency across weak and strong views for pixel supervision. **Right:** *PMHM* architecture. During training, a lightweight auxiliary head runs in parallel with the main head to achieve prediction consistency on uncertain pixels, using symmetric KL divergence across heads and decoder layers.

#### Algorithm 2 Mask Annotation Noise Generation

**Input:** Multi-label annotation mask  $\mathcal{A}$ , positive object IDs  $\Omega = \{\text{obj}_1, \text{obj}_2, \dots, \text{obj}_N\}$ , kernel size  $k \in \{9, 15, 21\}$ .  
**Output:** Noisy annotation mask  $\mathcal{A}_{\text{noise}}$ .  
**Initialize:** Empty list  $\mathcal{R} \leftarrow \emptyset$ ,  $\mathcal{A}_{\text{noise}} \leftarrow \mathbf{0}^{H \times W}$  (zero-initialized mask of same size as  $\mathcal{A}$ ).

- 1: **Separate:** For each  $\text{obj}_i \in \Omega$ , compute binary mask  $\mathcal{B}_i$  where  $\mathcal{B}_i[x, y] = 1$  if  $\mathcal{A}[x, y] = \text{obj}_i$ , else  $\mathcal{B}_i[x, y] = 0$ .
- 2: **for** each  $\text{obj}_i \in \Omega$  **do**
- 3:   Compute binary mask  $\mathcal{B}_i$  from  $\mathcal{A}$  for  $\text{obj}_i$ .
- 4:   Apply morphological dilation to  $\mathcal{B}_i$  with square kernel of size  $k$ , yielding dilated mask  $\mathcal{D}_i$ .
- 5:   Append  $\mathcal{D}_i$  to  $\mathcal{R}$ .
- 6: **end for**
- 7: **Combine:** Initialize  $\mathcal{A}_{\text{noise}}$  as a zero mask of size  $H \times W$ .
- 8: **for** each position  $[x, y]$  in  $\mathcal{A}_{\text{noise}}$  **do**
- 9:   **for** each dilated mask  $\mathcal{D}_i \in \mathcal{R}$  **do**
- 10:     **if**  $\mathcal{A}_{\text{noise}}[x, y] = 0$  and  $\mathcal{D}_i[x, y] \neq 0$  **then**
- 11:       Set  $\mathcal{A}_{\text{noise}}[x, y] \leftarrow \text{obj}_i$ .
- 12:     **end if**
- 13:   **end for**
- 14: **end for**
- 15: **Return:**  $\mathcal{A}_{\text{noise}}$ .

crowdsourced annotations, yet the resulting noise is more challenging than typical crowdsourced boundary errors and grows stronger as the dilation kernel increases. The severity of the generated mask noise is directly controlled by the dilation kernel size. Increasing the kernel size from 9 to 21 monotonically degrades annotation quality (Fig. 3, right).

#### IV. METHODOLOGY

**Setup and Notation.** We study action-based video object segmentation [6] with noisy supervision. The task is to generate a per-frame binary mask identifying the active object,

given a video clip and a textual prompt. Formally, for each pixel  $(x, y)$  in a frame  $t$ , the model predicts a foreground probability  $p_t(x, y) \in [0, 1]$ . This prediction is trained against a corresponding ground truth label  $y_t(x, y) \in \{0, 1\}$ , which may be noisy. Training minimizes a frame-averaged pixel-wise objective to ensure reliable localization for embodied manipulation.

##### A. Adapted Robust Learners for Pixel-level and Language-conditioned Segmentation

We adapt and integrate a broad set of noisy label strategies into the action-based video object segmentation setting [6]. Each method is independently adapted to operate on pixel probabilities and language-conditioned features. For all methods, the training objective averages per pixel losses over  $\mathcal{I}_t$  and over time. The adapted strategies include Co-teaching [19], GCE [16], SCE [20], APL [21], ELR [13], and NPN [12], enabling the first large-scale comparison of noise-robust training in action-based video object segmentation [6].

*Co-teaching* [19] suppresses heavy prompt mismatch by peer selection of small loss samples between two networks. At epoch  $n$ , each network selects a keep set containing the lowest loss samples in its mini-batch, and its size follows the conservative keep rate  $R_n = 1 - \rho \min(n/T_k, 1)$ . The selected samples are then used to supervise the peer network.

GCE [16] reduces the impact of outlying pixels by interpolating between cross entropy and mean absolute error. Define the correctness score  $r = y_t p_t + (1 - y_t)(1 - p_t)$  and use  $\ell_{\text{GCE}} = (1 - r^q)/q$  with  $q \in (0, 1]$ .

SCE [20] balances the standard cross entropy loss  $\mathcal{L}_{\text{CE}} = \frac{1}{\sum_t |\mathcal{I}_t|} \sum_t \sum_{(x,y) \in \mathcal{I}_t} \ell_{\text{CE}}^{txy}$  and the reverse cross entropy loss  $\mathcal{L}_{\text{RCE}} = \frac{1}{\sum_t |\mathcal{I}_t|} \sum_t \sum_{(x,y) \in \mathcal{I}_t} \ell_{\text{RCE}}^{txy}$  for stability when fore-

TABLE I: Action-based video object segmentation model under text prompt noise (noise-rate sweep).

Noise Rate (%)	p-mIoU↑	n-mIoU↓	p-cIoU↑	n-cIoU↓	gIoU↑	Acc↑
0	65.0	19.7	71.8	34.1	70.7	81.7
20	57.2	13.3	64.4	16.6	69.6	79.3
40	51.2	11.0	57.9	14.2	67.6	75.3
60	49.0	10.0	57.4	12.7	67.1	73.1

ground and background are partially corrupted, with weights  $\alpha$  and  $\beta$  in the combined objective  $\mathcal{L}_{\text{sym}} = \alpha \mathcal{L}_{\text{CE}} + \beta \mathcal{L}_{\text{RCE}}$ .

APL [21] strengthens confident decisions and passively penalizes ambiguous ones. We use the active  $\frac{1-r^q}{q}$  and the passive  $|p_t - y_t|$  in a weighted sum.

ELR [13] prevents noisy label memorization by regularizing its predictions toward their historical average. It combines a primary classification loss with a temporal agreement regularizer. The primary loss is the focal loss [32], which uses weights  $w^+ = \alpha(1-p_t^{(n)})^\gamma$  and  $w^- = (1-\alpha)(p_t^{(n)})^\gamma$  to down-weight well-classified examples:

$$\ell_{\text{focal}}^{txy} = -[y_t w^+ \log p_t^{(n)} + (1-y_t) w^- \log(1-p_t^{(n)})]. \quad (1)$$

To enforce consistency, an exponential moving average (EMA) of predictions,  $s_t^{(n)}$ , is maintained across training epochs  $n$ , updated via  $s_t^{(n)} = (1-\beta)s_t^{(n-1)} + \beta p_t^{(n)}$ . An agreement score  $d_t^{(n)} = p_t^{(n)}s_t^{(n)} + (1-p_t^{(n)})(1-s_t^{(n)})$  measures the similarity between the current prediction and its historical EMA. This score is used to formulate a regularization term  $R^{txy} = -\log(1-d_t^{(n)} + \epsilon)$ , where  $\epsilon$  ensures numerical stability. The final objective is a weighted sum of the focal loss and this regularizer,  $\ell_{\text{focal}}^{txy} + \lambda R^{txy}$ , averaged over all pixels and frames.

NPN [12] integrates partial label learning, negative learning, and dual-view consistency at the pixel level. For each pixel, we form candidate indicators  $A_t^{(1,n)} = y_t + \mathbb{I}[p_t^{(n)} > 0.5]$  and  $A_t^{(0,n)} = (1-y_t) + \mathbb{I}[p_t^{(n)} \leq 0.5]$  and accumulate counts  $M_t^{(1)} \leftarrow M_t^{(1)} + A_t^{(1,n)}$  and  $M_t^{(0)} \leftarrow M_t^{(0)} + A_t^{(0,n)}$ . The proxy label is  $y_t^* = \mathbb{I}[M_t^{(1)} \geq M_t^{(0)}]$  and the reliability is  $w_t = \frac{\max(M_t^{(1)}, M_t^{(0)})}{M_t^{(1)} + M_t^{(0)}}$ . The partial label loss  $\ell_{\text{PLL}}^{txy}$  is

$$\ell_{\text{PLL}}^{txy} = -w_t(y_t^* \log p_t^{(n)} + (1-y_t^*) \log(1-p_t^{(n)})), \quad (2)$$

and the negative learning loss  $\ell_{\text{NL}}^{txy}$  is

$$\begin{aligned} \ell_{\text{NL}}^{txy} = & -(1 - \mathbb{I}[A_t^{(1,n)} \geq 1]) \log(1 - p_t^{(n)}) \\ & - (1 - \mathbb{I}[A_t^{(0,n)} \geq 1]) \log p_t^{(n)}. \end{aligned} \quad (3)$$

Dual-view consistency  $\ell_{\text{CR}}^{txy}$  applies weak-to-strong cross-entropy with hard targets from the weak view. The final loss averages  $\ell_{\text{PLL}}^{txy} + \alpha \ell_{\text{NL}}^{txy} + \beta \ell_{\text{CR}}^{txy}$  over pixels and frames.

### B. Proposed PMHM Consistency with Main and Auxiliary Mask Heads

We propose the Parallel Mask Head Mechanism (PMHM), a memory-efficient module for noisy label segmentation that avoids dual models [19] and storing historical predictions for each pixel [12], [32], both of which incur a high GPU

TABLE II: Action-based video object segmentation model with noisy annotations.

kernel size	p-mIoU↑	n-mIoU↓	p-cIoU↑	n-cIoU↓	gIoU↑	Acc↑
0	65.0	19.7	71.8	34.1	70.7	81.7
9×9	55.7	13.7	64.8	20.3	68.2	80.4
15×15	49.1	13.3	61.1	19.7	63.7	80.0
21×21	44.8	13.9	57.4	29.9	60.3	79.2

memory burden. It features a main head and a parallel, lightweight auxiliary head that share decoder features.

During training, we apply mild perturbations (random dropout and decaying-rate freezing) to the auxiliary head to foster prediction diversity. The auxiliary head is discarded at inference, incurring no computational overhead. Our method identifies uncertain pixels  $U_t$  as those near the decision boundary or with high spatial gradients:

$$U_t = \{(x, y) \in \mathcal{I}_t \mid |p_t^{\text{m}}(x, y) - 0.5| < \tau_m \vee \|\nabla p_t^{\text{m}}(x, y)\|_2 > \tau_e\}. \quad (4)$$

The head consistency loss,  $L_{\text{head}}$ , aligns the predictions of the main ( $p_t^{\text{m}}(x, y)$ ) and auxiliary ( $p_t^{\text{a}}(x, y)$ ) heads on these uncertain pixels using a symmetric KL divergence:

$$L_{\text{head}} = \frac{1}{\sum_t |U_t|} \sum_t \sum_{(x,y) \in U_t} \left[ D_{\text{KL}}(p_t^{\text{m}} \| p_t^{\text{a}}) + D_{\text{KL}}(p_t^{\text{a}} \| p_t^{\text{m}}) \right]. \quad (5)$$

We also introduce a layer consistency loss,  $L_{\text{layer}}$ , to align predictions from early decoder stages ( $p_t^{\text{m},\ell}$ ) with the final stage prediction ( $p_t^{\text{m},L}$ ) on the same uncertain pixels  $U_t$ :

$$L_{\text{layer}} = \frac{1}{\sum_t |U_t|} \sum_t \sum_{(x,y) \in U_t} \frac{1}{L-1} \sum_{\ell=1}^{L-1} \left[ D_{\text{KL}}(p_t^{\text{m},L} \| p_t^{\text{m},\ell}) + D_{\text{KL}}(p_t^{\text{m},\ell} \| p_t^{\text{m},L}) \right]. \quad (6)$$

The loss  $L_{\text{seg}}$  combines a hard mask loss  $L_{\text{hard}}$  [6] on confident pixels (outside  $U_t$ ) with our consistency terms:

$$L_{\text{seg}} = L_{\text{hard}} + \lambda_{\text{head}} L_{\text{head}} + \lambda_{\text{layer}} L_{\text{layer}}. \quad (7)$$

Hyperparameters  $\tau_m$ ,  $\tau_e$ ,  $\lambda_{\text{head}}$ , and  $\lambda_{\text{layer}}$  govern uncertainty selection and regularization:  $\tau_m$  defines a margin near 0.5 for low confidence pixels,  $\tau_e$  thresholds  $\|\nabla p_t^{\text{m}}\|$  to capture boundaries, and  $\lambda_{\text{head}}$  with  $\lambda_{\text{layer}}$  weight head and layer consistency against  $L_{\text{hard}}$ ; concrete values are given in Sec. V.

## V. EXPERIMENTS

### A. Datasets and Metrics

*Datasets.* We train on the VISOR training split (13,205 clips; 76,873 objects) with our synthetic noise (Section III), and evaluate on the clean action-based video object segmentation benchmark (294 manually annotated clips) following the protocol in [6]. Details of dataset construction and noise simulation are provided in Section III.

*Metrics.* Following [6], we use several metrics to evaluate performance. We report Mean IoU (mIoU) and cumulative IoU (cIoU) computed separately for active foregrounds (p-mIoU, p-cIoU) and inactive backgrounds (n-mIoU, n-cIoU).

TABLE III: Single-noise benchmark for action-based video object segmentation. Each subtable fixes the noise type and level/kernel: Name-20/40/60 (clean masks); Mask- $k \in \{9, 15, 21\}$  (clean names). PMHM targets mask noise and is therefore evaluated only under noisy mask conditions in this table.

(a) Name-20 (mask clean)							(b) Name-40 (mask clean)						
Method	p-mIoU $\uparrow$	n-mIoU $\downarrow$	p-cIoU $\uparrow$	n-cIoU $\downarrow$	gIoU $\uparrow$	Acc $\uparrow$	Method	p-mIoU $\uparrow$	n-mIoU $\downarrow$	p-cIoU $\uparrow$	n-cIoU $\downarrow$	gIoU $\uparrow$	Acc $\uparrow$
Co-teaching [19]	63.2	18.5	71.1	35.2	70.3	80.9	Co-teaching [19]	63.0	17.4	70.7	29.1	70.6	81.2
SCE [20]	63.8	25.3	67.3	36.9	50.4	67.4	SCE [20]	63.7	25.7	67.7	36.9	47.5	64.7
GCE [16]	59.1	24.1	66.8	37.2	46.9	66.3	GCE [16]	61.3	25.1	65.4	36.9	44.9	63.6
APL [21]	62.4	25.7	67.0	39.0	49.1	66.8	APL [21]	62.7	23.8	65.6	35.4	47.2	65.3
ELR [13]	56.2	12.2	60.6	12.3	69.3	79.0	ELR [13]	54.7	10.1	59.0	7.4	69.5	78.9
NPN [12]	56.1	12.4	60.5	13.8	69.1	78.7	NPN [12]	53.6	10.4	56.9	8.8	68.8	76.9

(c) Name-60 (mask clean)							(d) Mask- $k=9$ (name clean)						
Method	p-mIoU $\uparrow$	n-mIoU $\downarrow$	p-cIoU $\uparrow$	n-cIoU $\downarrow$	gIoU $\uparrow$	Acc $\uparrow$	Method	p-mIoU $\uparrow$	n-mIoU $\downarrow$	p-cIoU $\uparrow$	n-cIoU $\downarrow$	gIoU $\uparrow$	Acc $\uparrow$
Co-teaching [19]	60.2	13.8	68.8	23.0	71.2	80.9	Co-teaching [19]	60.6	23.4	69.8	51.1	65.1	80.1
SCE [20]	58.3	12.8	67.1	18.7	70.7	80.1	SCE [20]	58.7	18.5	67.4	33.1	67.0	81.7
GCE [16]	56.7	11.5	65.3	15.0	70.0	81.1	GCE [16]	58.8	18.5	67.5	34.2	67.1	81.6
APL [21]	57.8	11.6	66.3	15.8	70.9	80.3	APL [21]	58.8	18.7	68.4	33.9	67.0	81.3
ELR [13]	48.5	9.8	51.8	11.8	66.7	72.2	ELR [13]	57.1	18.8	66.0	32.3	64.6	78.0
NPN [12]	47.6	9.2	51.6	7.9	66.3	71.8	NPN [12]	58.8	20.6	67.5	37.1	65.4	78.6
							PMHM	55.6	13.6	64.3	21.5	67.8	81.1

(e) Mask- $k=15$ (name clean)							(f) Mask- $k=21$ (name clean)						
Method	p-mIoU $\uparrow$	n-mIoU $\downarrow$	p-cIoU $\uparrow$	n-cIoU $\downarrow$	gIoU $\uparrow$	Acc $\uparrow$	Method	p-mIoU $\uparrow$	n-mIoU $\downarrow$	p-cIoU $\uparrow$	n-cIoU $\downarrow$	gIoU $\uparrow$	Acc $\uparrow$
Co-teaching [19]	53.4	21.3	65.4	47.0	60.9	79.5	Co-teaching [19]	46.7	18.7	59.5	45.2	58.1	79.5
SCE [20]	51.4	14.7	63.4	28.4	64.2	82.4	SCE [20]	44.1	12.3	55.8	24.8	60.6	80.0
GCE [16]	51.5	13.9	62.1	28.4	64.6	82.1	GCE [16]	44.7	12.5	57.6	22.1	60.9	79.9
APL [21]	51.6	13.6	61.5	24.0	64.7	82.5	APL [21]	45.7	14.3	58.4	30.3	60.8	82.0
ELR [13]	51.2	17.4	62.0	30.2	62.0	79.1	ELR [13]	45.2	16.3	57.4	34.2	58.4	78.0
NPN [12]	52.2	18.2	62.2	34.9	61.9	78.6	NPN [12]	46.2	17.2	58.5	37.0	58.9	78.0
PMHM	48.7	11.8	59.5	19.2	64.6	81.2	PMHM	43.8	11.0	56.6	20.9	61.7	80.1

Additionally, we use generalized IoU (gIoU) [33] for overall segmentation quality and Accuracy (Acc) to evaluate the classification of active versus inactive objects.

### B. Implementation Details

We use ResNet-101 [34] as the visual encoder and RoBERTa [35] as the text encoder (frozen). All models are initialized from Ref-YouTube-VOS checkpoints [36]. The optimizer and loss coefficients follow [6]. Data augmentations follow [37]; for ELR [13] and NPN [12], we disable spatial transforms and keep only pixel-wise photometric distortion and normalization to preserve per-pixel alignment required by historical prediction updates. Method-specific settings are: Co-teaching keeps samples with a loss rank  $\leq 0.95$ , uses gradient accumulation of  $128\times$ , and sets  $T_k=6$ ; GCE and APL use  $q=0.7$  [16]; ELR and NPN both use focal BCE and serialize per-pixel histories to disk, with ELR adding temporal regularization ( $\beta=0.9$ ,  $\epsilon=10^{-6}$ ,  $\alpha=0.5$ ,  $\gamma=2$ ,  $\lambda=1$ ) and NPN applying thresholding ( $\alpha=0.1$ ,  $\beta=0.2$ ). PMHM uses thresholds  $\tau_m=0.20$ ,  $\tau_e=0.85$ ; loss weights are  $\lambda_{\text{head}}=0.1$ ,  $\lambda_{\text{layer}}=0.1$ . Training runs on  $4\times$  RTX 3090 for 6 epochs, except that Co-teaching uses 10 epochs to ensure convergence of its progressive sample selection mechanism under noisy labels.

### C. Influence of Text Prompt Noise

Under a name-perturbation sweep (Table I), we observe a clear pattern: as text noise increases from 0% to 60%, positive-region segmentation and accuracy decline while negative-region metrics ( $\downarrow$  better) improve, leaving the global score only mildly affected. Most of the damage occurs

early at 20~40% noise (p-mIoU  $65.0\rightarrow 57.2\rightarrow 51.2$ ; p-cIoU  $71.8\rightarrow 64.4\rightarrow 57.9$ ), and positive cumulative IoU then largely plateaus from 40% to 60% (p-cIoU  $57.9\rightarrow 57.4$ ). This asymmetry indicates that mismatched prompts drive the model to back off to conservative masks, reducing spurious activations on background (n-mIoU  $19.7\rightarrow 10.0$ ; n-cIoU  $34.1\rightarrow 12.7$ ) while hurting recall on active regions, consistent with a fallback to visual cues when language becomes unreliable. Consequently, gIoU changes only slightly ( $70.7\rightarrow 67.1$ ) and can obscure failure modes, making p/n-partitioned metrics the more informative diagnostics.

### D. Effects of Mask Annotation Noise

With  $k\times k$  morphological dilation applied to ground-truth masks (Table II), performance degrades chiefly on positive-region metrics and the aggregate score: p-mIoU drops steeply at moderate noise ( $65.0\rightarrow 55.7$  at  $9\times 9$ ) and then continues to fall with diminishing increments ( $49.1$  at  $15\times 15$ ,  $44.8$  at  $21\times 21$ ), while gIoU declines from  $70.7\rightarrow 68.2\rightarrow 63.7\rightarrow 60.3$ . Negative-region metrics ( $\downarrow$  better) improve at moderate noise but show a non-monotone rebound at the highest level: n-mIoU  $19.7\rightarrow 13.7\rightarrow 13.3\rightarrow 13.9$  and n-cIoU  $34.1\rightarrow 20.3\rightarrow 19.7\rightarrow 29.9$ . This pattern indicates that boundary dilation first discourages spurious background activations, pushing the model toward conservative masks with fewer false positives. Yet, once the noisy boundary band becomes too wide, supervision near edges turns conflicting, inducing boundary drift and shape jitter that hurt positive overlap and increase large area background overlap (hence the n-cIoU rebound). Acc changes little ( $81.7\rightarrow 79.2$ ), suggesting the classifier is relatively insensitive compared with

TABLE V: Name noise 60% + Mask kernel  $21 \times 21$  on action-based video object segmentation (mixed setting).  $\uparrow/\downarrow$  denotes that higher/lower is better.

Method	p-mIoU $\uparrow$	n-mIoU $\downarrow$	p-cIoU $\uparrow$	n-cIoU $\downarrow$	gIoU $\uparrow$	Acc $\uparrow$
FLac [6]	37.8	7.6	49.7	11.2	60.5	76.6
Co-teaching [19]	36.2	10.9	48.7	24.3	57.9	70.4
SCE [20]	36.4	6.5	47.8	6.0	60.9	72.9
GCE [16]	38.7	7.9	49.7	11.8	61.1	75.4
APL [21]	36.5	7.1	48.1	9.9	60.5	73.6
ELR [13]	41.5	15.9	51.2	33.3	48.8	72.4
NPN [12]	37.6	8.0	49.7	14.0	60.2	75.2
PMHM	36.8	8.4	47.1	14.6	59.9	74.1

segmentation overlap metrics. At matched severities, mask noise is also more damaging than text noise (e.g., gIoU 60.3 at  $21 \times 21$  vs. 67.1 at Name-60; cf. Table I), underscoring the need for boundary-focused robustness and for reporting p/n-partitioned metrics over a single aggregate score.

### E. Evaluation of Robustness Strategies under Combined Noise Conditions

Under the mixed setting (Name-60 + Mask- $21 \times 21$ ; Table V), overall accuracy and positive-region IoUs drop across the board, and the ranking shifts relative to the single-noise sweeps (Table III): methods that filter sample-level noise (e.g., Co-teaching [19]) no longer dominate, whereas losses that damp per-pixel label noise (GCE [16], SCE [20], APL [21]) and consistency-based training (NPN [12]) tend to rank higher on gIoU and the p-/n-partitioned IoUs. This aligns with the structure of multimodal corruption: name mismatches introduce instance-level grounding errors, whereas dilated masks introduce dense boundary bias; sample selection helps when labels are globally wrong but is brittle when many edge pixels are systematically corrupted. We also observe the familiar positive–negative trade-off; recovering active-region IoUs often increases background errors, so gIoU alone can misorder methods, making the p-/n-partitioned scores the more informative diagnostics.

*Observation 1* is the non-transitive robustness across noise types. Methods that prevail under text-only noise (e.g., Co-teaching at Name-40/60 in Table III) do not necessarily remain superior under mixed noise (Table V). This is consistent with the premise of Co-teaching [19], which mitigates instance-level label contradictions but offers limited protection against dense, pixel-localized boundary corruption.

*Observation 2* is a gradient-conflict effect under mixed noise. Prompt mis-grounding encourages conservative masks (shrinking foreground to avoid false activations), whereas boundary dilation encourages expansion near edges; pixel-wise robust losses (GCE [16], SCE [20], APL [21]) and dual-view consistency (NPN [12]) better reconcile these opposing update directions than sample filtering alone. This interpretation matches the design goals of GCE/SCE/APL and NPN, now evidenced in a multimodal, pixel-level setting.

*Observation 3* concerns metric sensitivity and benchmark coverage. Since positive and negative regions shift oppositely under mixed noise, aggregate gIoU can mask failures; reporting p-/n-partitioned mIoU/cIoU surfaces boundary-specific

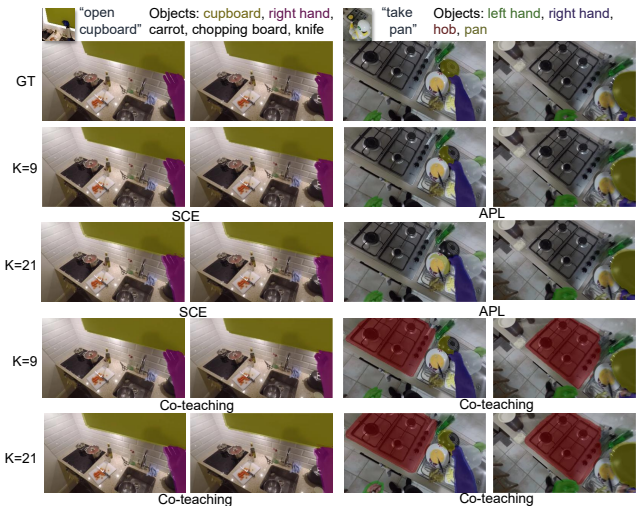


Fig. 5. Mask annotation noise qualitative results on *ActiSeg-NL*. Larger kernels thicken boundaries, coarsen edges, and introduce redundant regions, revealing trade-offs between foreground coverage and background precision.

errors and background precision. The consistent re-ranking between Table III and Table V validates the need for our multimodal protocols and partitioned metrics to fully characterize robustness in action-based video object segmentation.

### F. PMHM under Boundary and Mixed Noise

As *PMHM* is designed to address corrupted mask labels, we evaluated its performance under boundary and mixed noise conditions. In the boundary-only noise setting (kernel sizes 9, 15, and 21), *PMHM* consistently reduces background errors and improves global scores over the baseline. At the highest noise level ( $21 \times 21$ ), it achieves a gIoU of 61.7, outperforming the baseline’s 60.3 (Table II). In the mixed-noise setting (Name-60 + Mask-21), where severe text noise is introduced, *PMHM*’s performance (gIoU 59.9) does not surpass the baseline FLac (gIoU 60.5), as shown in Table V. This confirms that while *PMHM* resists boundary corruption, its gains fade when combined with significant text noise.

### G. Implications for Embodied Perception

For tasks that require strong background suppression, the more balanced profiles of GCE and SCE are preferable. They reduce background leakage while maintaining active-region coverage, consistent with the re-ranking between Table III and Table V. When language prompts are unreliable, Co-teaching with extended training preserves foreground under text noise, yet it provides limited protection against dense boundary corruption. It should therefore be paired with boundary-aware checks before execution. APL shows cleaner object edges at small kernel sizes but may overextend coverage at large kernel sizes, which warrants caution in tight workspaces. ELR increases positive-region IoUs and can support robust pick-up under occlusion, although the reduction in background precision calls for uncertainty gating. In practice, the partitioned metrics are task-aligned proxies:

p-cIoU and p-mIoU reflect contact placement accuracy, and n-cIoU and n-mIoU reflect collision risk. Their joint trend provides a more reliable go/no-go signal than a single aggregate score. The qualitative patterns in Fig. 5 translate into concrete risks and design choices for robotic manipulation and closed-loop interaction. Boundary thickening and edge coarsening can cause grasp-pose drift and increase collision risk near contact surfaces. Spurious activation on inactive objects, as seen on the hob in the “take pan” sequence, can mislead target selection.

## VI. CONCLUSION

We introduce *ActiSeg-NL*, a benchmark to assess action-based video object segmentation under text and mask noise, towards advancing embodied intelligence. Our analysis shows that model performance degrades significantly with noise, especially from masks. We find that robust learning strategies present clear trade-offs: some preserve foregrounds under text noise (Co-teaching), others improve overlap (APL), or achieve a superior balance in mixed conditions (GCE, SCE). Future work seeks to solve this trade-off by integrating vision language foundation models to drive boundary-focused objectives and uncertainty-guided consistency under real-world noise, together with systematic validation in closed-loop robotic systems.

## REFERENCES

- [1] K. Grauman *et al.*, “Ego4D: Around the world in 3,000 hours of egocentric video,” in *Proc. CVPR*, 2022, pp. 18 973–18 990.
- [2] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixe, D. Cremers, and L. Van Gool, “One-shot video object segmentation,” in *Proc. CVPR*, 2017, pp. 5320–5329.
- [3] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim, “Video object segmentation using space-time memory networks,” in *Proc. ICCV*, 2019, pp. 9225–9234.
- [4] A. Botach, E. Zheltonozhskii, and C. Baskin, “End-to-end referring video object segmentation with multimodal transformers,” in *Proc. CVPR*, 2022, pp. 4975–4985.
- [5] J. Wu, Y. Jiang, P. Sun, Z. Yuan, and P. Luo, “Language as queries for referring video object segmentation,” in *Proc. CVPR*, 2022, pp. 4974–4984.
- [6] L. Ouyang, R. Liu, Y. Huang, R. Furuta, and Y. Sato, “ActionVOS: Actions as prompts for video object segmentation,” in *Proc. ECCV*, vol. 15068, 2024, pp. 216–235.
- [7] S. Yenamandra *et al.*, “HomeRobot: Open-vocabulary mobile manipulation,” in *Proc. CoRL*, vol. 229, 2023, pp. 1975–2011.
- [8] Z. Chen, L. Zhang, P. Hu, H. Lu, and Y. He, “MaskTrack: Auto-labeling and stable tracking for video object segmentation,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 36, no. 7, pp. 12 052–12 065, 2025.
- [9] D. Damen *et al.*, “Rescaling egocentric vision: Collection, pipeline and challenges for EPIC-KITCHENS-100,” *International Journal of Computer Vision*, vol. 130, no. 1, pp. 33–55, 2022.
- [10] S. Ye, Y. Xu, D. Chen, S. Han, and J. Liao, “Learning a single network for robust medical image segmentation with noisy labels,” *IEEE Transactions on Medical Imaging*, vol. 43, no. 9, pp. 3188–3199, 2024.
- [11] A. Darkhalil *et al.*, “EPIC-KITCHENS VISOR benchmark: Video segmentations and object relations,” in *Proc. NeurIPS*, vol. 35, 2022, pp. 13 745–13 758.
- [12] M. Sheng, Z. Sun, Z. Cai, T. Chen, Y. Zhou, and Y. Yao, “Adaptive integration of partial label learning and negative learning for enhanced noisy label learning,” in *Proc. AAAI*, vol. 38, no. 5, 2024, pp. 4820–4828.
- [13] S. Liu, J. Niles-Weed, N. Razavian, and C. Fernandez-Granda, “Early-learning regularization prevents memorization of noisy labels,” in *Proc. NeurIPS*, vol. 33, 2020, pp. 20 331–20 342.
- [14] E. Cho, J. U. Kim, and S. T. Kim, “Spatial mask-based adaptive robust training for video object segmentation with noisy labels,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 35, no. 5, pp. 4976–4990, 2025.
- [15] Q. Miao *et al.*, “Learning with noisy labels using collaborative sample selection and contrastive semi-supervised learning,” *Knowledge-Based Systems*, vol. 296, p. 111860, 2024.
- [16] Z. Zhang and M. Sabuncu, “Generalized cross entropy loss for training deep neural networks with noisy labels,” in *Proc. NeurIPS*, 2018, pp. 8792–8802.
- [17] X. Zhou, X. Liu, D. Zhai, J. Jiang, and X. Ji, “Asymmetric loss functions for noise-tolerant learning: Theory and applications,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 7, pp. 8094–8109, 2023.
- [18] S. Yan *et al.*, “Referred by multi-modality: A unified temporal transformer for video object segmentation,” in *Proc. AAAI*, vol. 38, no. 6, 2024, pp. 6449–6457.
- [19] B. Han *et al.*, “Co-teaching: Robust training of deep neural networks with extremely noisy labels,” in *Proc. NeurIPS*, 2018, pp. 8536–8546.
- [20] Y. Wang, X. Ma, Z. Chen, Y. Luo, J. Yi, and J. Bailey, “Symmetric cross entropy for robust learning with noisy labels,” in *Proc. ICCV*, 2019, pp. 322–330.
- [21] X. Ye, X. Li, S. Dai, T. Liu, Y. Sun, and W. Tong, “Active negative loss functions for learning with noisy labels,” in *Proc. NeurIPS*, 2023, pp. 6917–6940.
- [22] H.-Y. Joo-Won, “Hand-based technologies in first-person view: Bridging the gap between human interaction and intelligent systems,” *Available at SSRN 5032628*, 2024.
- [23] T. Liang, K.-Y. Lin, C. Tan, J. Zhang, W.-S. Zheng, and J.-F. Hu, “ReferDINO: Referring video object segmentation with visual grounding foundations,” in *Proc. ICCV*, 2025, pp. 20 009–20 019.
- [24] B. Miao, M. Bennamoun, Y. Gao, and A. Mian, “Spectrum-guided multi-granularity referring video object segmentation,” in *Proc. ICCV*, 2023, pp. 920–930.
- [25] S. Cho, S. Lee, M. Lee, J. Lee, and S. Lee, “Find first, track next: Decoupling identification and propagation in referring video object segmentation,” *arXiv preprint arXiv:2503.03492*, 2025.
- [26] W. Song, Z. Cao, J. Zhang, C. Xu, and A. Lim, “Learning variable ordering heuristics for solving constraint satisfaction problems,” *Engineering Applications of Artificial Intelligence*, vol. 109, p. 104603, 2022.
- [27] S. Liu, K. Liu, W. Zhu, Y. Shen, and C. Fernandez-Granda, “Adaptive early-learning correction for segmentation from noisy annotations,” in *Proc. CVPR*, 2022, pp. 2596–2606.
- [28] Y. Luo, G. Liu, Y. Guo, and G. Yang, “Deep neural networks learn meta-structures from noisy labels in semantic segmentation,” in *Proc. AAAI*, vol. 36, no. 2, 2022, pp. 1908–1916.
- [29] J. Feng, X. Wang, T. Li, S. Ji, and W. Liu, “Weakly-supervised semantic segmentation via online pseudo-mask correcting,” *Pattern Recognition Letters*, vol. 165, pp. 33–38, 2023.
- [30] M. Kimhi, O. Kerem, E. Grad, E. Rivlin, and C. Baskin, “Noisy annotations in semantic segmentation,” *arXiv preprint arXiv:2406.10891*, 2024.
- [31] J. Yao, Y. Zhang, S. Zheng, M. Goswami, P. Prasanna, and C. Chen, “Learning to segment from noisy annotations: A spatial correction approach,” in *Proc. ICLR*, 2023.
- [32] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, “Focal loss for dense object detection,” in *Proc. ICCV*, 2017, pp. 2999–3007.
- [33] C. Liu, H. Ding, and X. Jiang, “GRES: Generalized referring expression segmentation,” in *Proc. CVPR*, 2023, pp. 23 592–23 601.
- [34] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proc. CVPR*, 2016, pp. 770–778.
- [35] Y. Liu *et al.*, “RoBERTa: A robustly optimized BERT pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [36] S. Seo, J.-Y. Lee, and B. Han, “URVOS: Unified referring video object segmentation network with a large-scale benchmark,” in *Proc. ECCV*, 2020, pp. 208–223.
- [37] Y. Wang *et al.*, “End-to-end video instance segmentation with transformers,” in *Proc. CVPR*, 2021, pp. 8741–8750.