
SOC: Semantic-Assisted Object Cluster for Referring Video Object Segmentation

Supplementary Material

Anonymous Author(s)

Affiliation

Address

email

1 Dataset Details

The A2D-Sentences dataset contains 3782 videos and each video has 3-5 annotated segmentation masks and JHMDB-Sentences totally comprises 928 videos, each of which is associated with a text description. For the large-scale dataset, Ref-YouTube-VOS has 3978 videos with about 15K text descriptions. The Ref-DAVIS17 contains 90 videos with 1,544 expressions, including 60 and 30 videos for training and validation respectively.

2 Performance on JHMDB-Sentences

We also compare our SOC with existing methods on JHMDB-Sentences [2] and the results are shown in Table 1. Following ReferFormer [8], we directly report the results utilizing the models trained on A2D-Sentences without finetune. It can be seen that our method achieves new state-of-the-art performance with different backbone and training settings. Compared to other benchmarks, the performance gains on this dataset are relatively small. This can be attributed to JHMDB’s imprecise annotations generated by coarse human puppet model.

Method	Backbone	Precision					mAP	IoU	
		P@0.5	P@0.6	P@0.7	P@0.8	P@0.9		Overall	Mean
Hu <i>et al.</i> [3]	VGG-16	63.3	35.0	8.5	0.2	0.0	17.8	54.6	52.8
Gavrilyuk <i>et al.</i> [2]	I3D	69.9	46.0	17.3	1.4	0.0	23.3	54.1	54.2
CMSA + CFSA [9]	ResNet-101	76.4	62.5	38.9	9.0	0.1	-	62.8	58.1
ACAN [7]	I3D	75.6	56.4	28.7	3.4	0.0	28.9	57.6	58.4
CMPC-V [5]	I3D	81.3	65.7	37.1	7.0	0.0	34.2	61.6	61.7
ClawCraneNet [4]	ResNet-50/101	88.0	79.6	56.6	14.7	0.2	-	64.4	65.6
MTTR [1]	Video-Swin-T	93.9	85.2	61.6	16.6	0.1	39.2	70.1	69.8
ReferFormer [8]	Video-Swin-T	93.3	84.2	61.4	16.4	0.3	39.1	70.0	69.3
SOC(Ours)	Video-Swin-T	94.7	86.4	62.7	17.9	0.1	39.7	70.7	70.1
<i>With Image Pretrain</i>									
ReferFormer	Video-Swin-T	95.8	89.3	66.8	18.9	0.2	42.2	71.9	71.0
ReferFormer	Video-Swin-B	96.2	90.2	70.2	21.0	0.3	43.7	73.0	71.8
SOC(Ours)	Video-Swin-T	96.3	88.7	67.2	19.6	0.1	42.7	72.7	71.6
SOC(Ours)	Video-Swin-B	96.9	91.4	71.1	21.3	0.1	44.6	73.6	72.3

Table 1: Comparison with the state-of-the-art methods on JHMDB-Sentences.

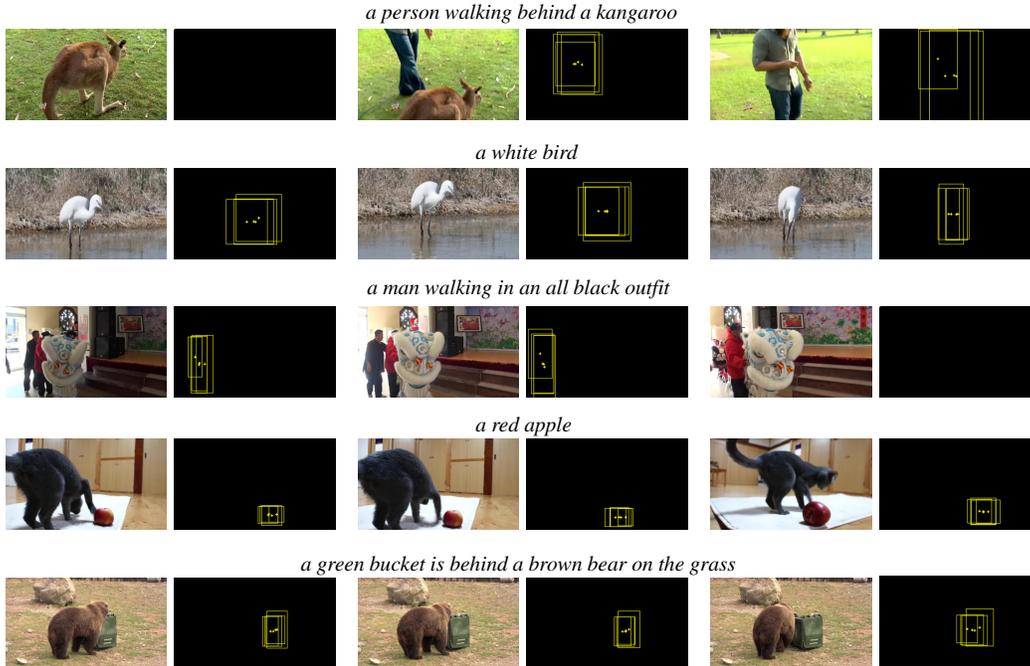


Figure 1: Visualization of the frame-level object query

14 3 More Implementation Details

15 **Training Settings** Our models are trained with the AdamW optimizer using Pytorch. The weight
 16 decay is 1×10^{-4} . The batch size is set to 56 for pretraining and 8 for main training. The models
 17 are trained for 30 epochs. The initial learning rate is set to 1×10^{-4} for Ref-YouTube-VOS and
 18 RefCOCO+/g, 5×10^{-5} for A2D-Sentences. The learning rate decays by 10 for the backbone
 19 network. During training, we apply RandomResize and Horizontal Flip for data augmentation.
 20 Specifically, all frames are downsampled to 360×640 for Ref-YouTube-VOS and RefCOCO+/g,
 21 320×576 for A2D-Sentences.

22 **Inference Settings** During inference, the input videos are downsampled to 320×576 for A2D-
 23 Sentences dataset and 360p for other datasets. We directly output the segmentation masks without
 24 any post-process.

25 4 Additional Qualitative Results

26 4.1 Query Visualization

27 To demonstrate that the frame-level query embeddings can represent the referred object in a specific
 28 frame, we visualize the predicted bounding boxes corresponding to the query embeddings. As
 29 illustrated in Fig. 1, the majority of queries focus on regions of the referred object as expected.
 30 This indicates that the compact frame-level query embeddings indeed reflect object information and
 31 subsequent video-level object cluster is performing temporal interaction for referred objects.

32 4.2 Segmentation Stability Visualization

33 The benchmark performance and IoU variance analysis in the main paper have proven the effective-
 34 ness and stability of our method. Here we incorporate visual comparisons to further validate the
 35 segmentation stability of our model. In Fig. 2 (a), benefiting from the global object view, SOC is
 36 capable of tracking the referred object across frames in coherence. On the contrary, ReferFormer [8],
 37 the existing state-of-the-art method, may generate segmentation masks with high degree of variance,
 38 indicating that the frame-based paradigm fails to accurately understand the state of the object in the
 39 context of the entire video (see in Fig. 2 (b)).

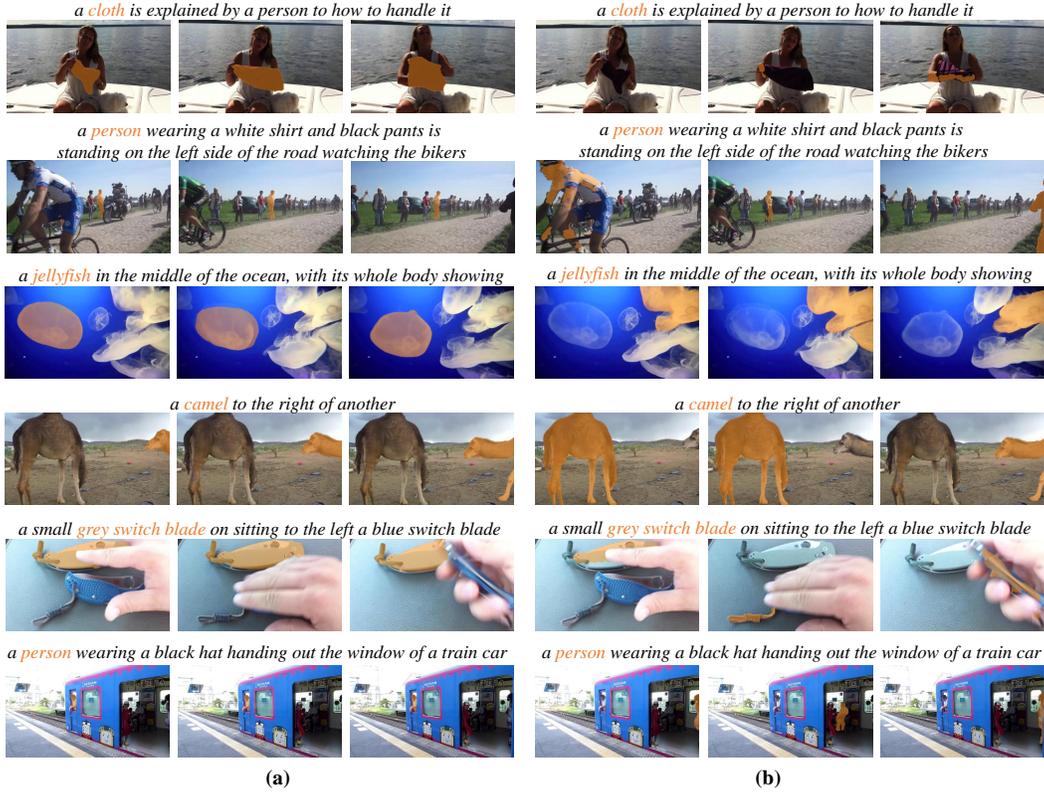


Figure 2: Visualization comparisons of segmentation stability between our SOC and existing state-of-the-art method ReferFormer [8]. (a) and (b) denote our SOC and ReferFormer, respectively.

4.3 Adaptability for Texts Describing Temporal Variation

Figure 6 in the main paper has shown some results to demonstrate that our SOC can better handle descriptions that focus on temporal variation. Here we provide more cases to demonstrate the adaptability of our method to such text descriptions. Fig. 4 and Fig. 5 show the segmentation results of our SOC and ReferFormer, where (a) indicates the segmentation results by SOC and (b) represents the results by ReferFormer[8].

5 Comprehensive Evaluation

We comprehensively measure our method by different perspectives, *e.g.*, performance, inference speed and computation cost under fair comparison. It is noted that the horizontal axis of Fig. 3 denotes performance on Ref-YouTube-VOS, vertical axis is FPS and the radius of the circle represents the relative FLOPs. Compared with ReferFormer [8] (blue ●), Our method (red ●) achieves superior performance with faster inference speed and less computation cost. Although MTTR [1] (orange ●) [1] has the lowest FLOPs, the lack of elaborate multi-modal fusion and temporal interaction significantly degrade the segmentation accuracy. In contrast, our method leverages video-level multi-modal understanding, which brings a significant increase in performance with only minimal computational costs.

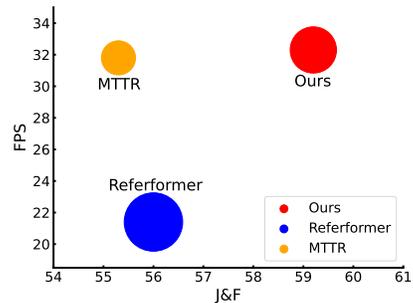


Figure 3: Performance vs Inference Speed vs Computation Cost.

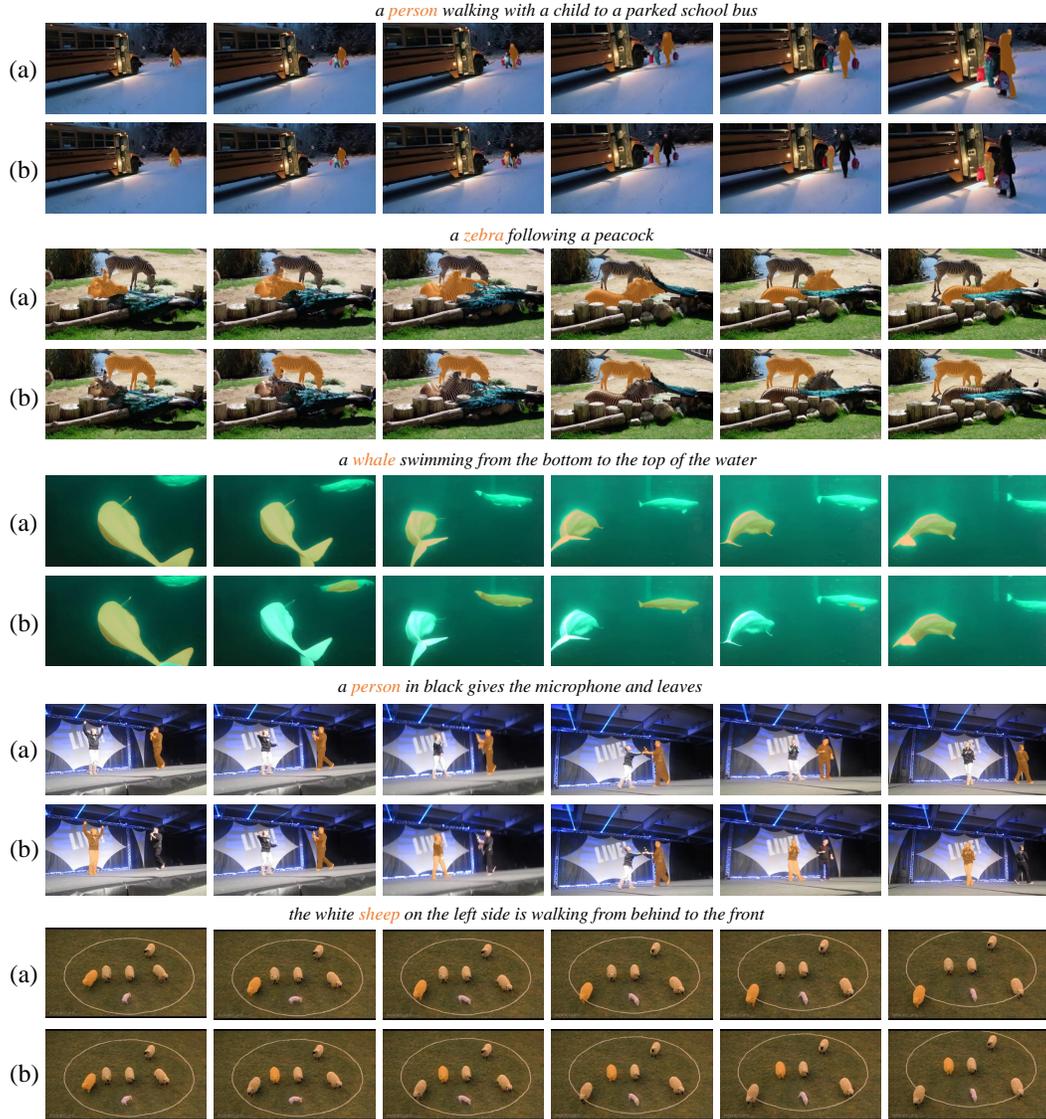


Figure 4: Visualization comparison using text expressions about temporal variation. (a) and (b) are segmentation results of our SOC and ReferFormer [8], respectively.

62 6 Error Bar

63 We have retrained our model several times on Ref-YouTube-VOS [6] dataset. The results demonstrate
 64 that the randomness of the model has little effect on the performance, *i.e.*, the max deviation is about
 65 0.5% $J&F$.

66 7 Broader Impact

67 Malicious use of the RVOS model may lead to potential negative societal impacts, including but not
 68 limited to unauthorized surveillance or privacy-infringing tracking. However, we firmly believe that
 69 the task itself is neutral with positive implications, such as video editing and human-robot interaction.
 70

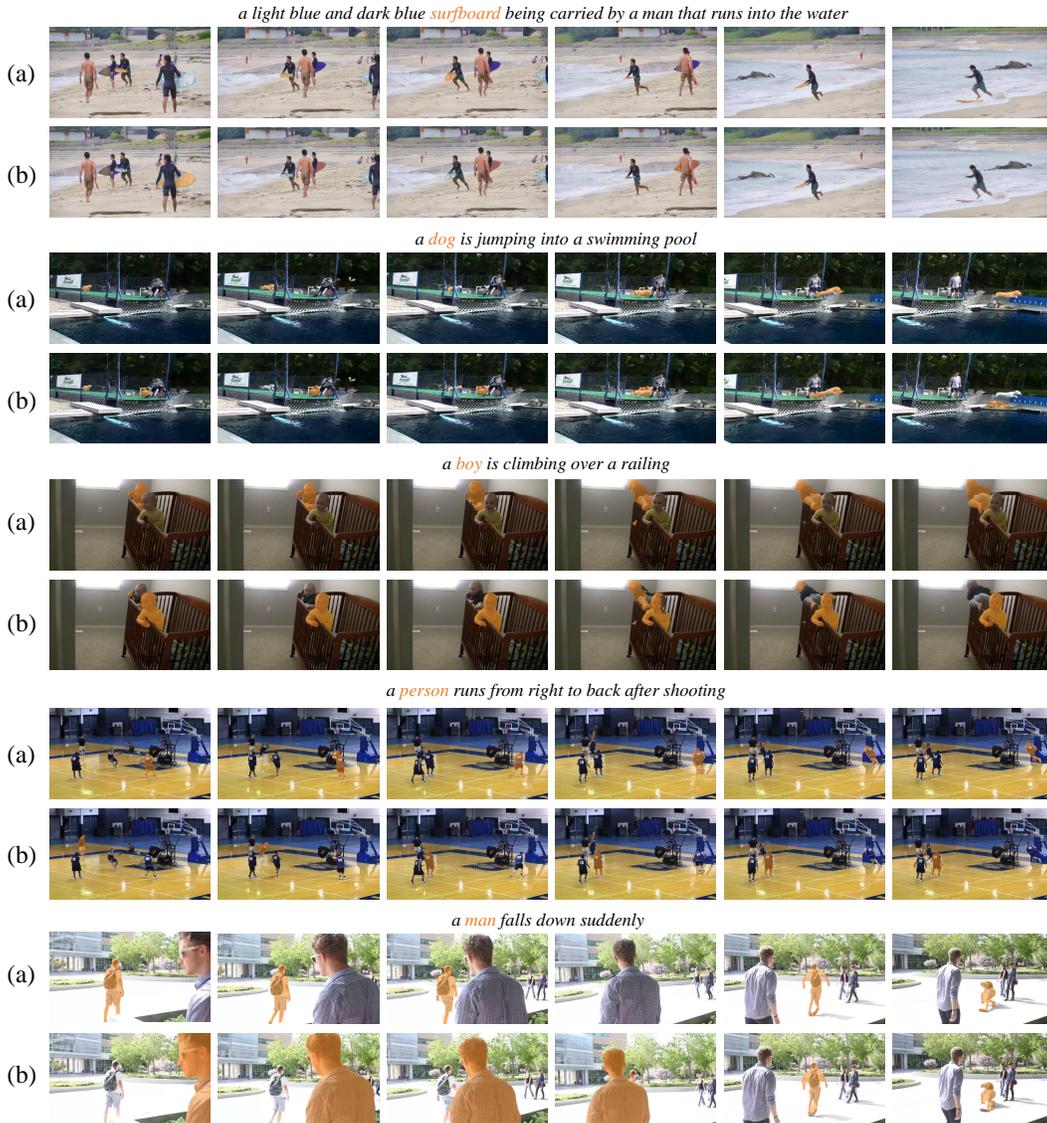


Figure 5: Visualization comparison using text expressions about temporal variation. (a) and (b) are segmentation results of our SOC and ReferFormer [8], respectively.

71 References

- 72 [1] Botach, A., Zheltonozhskii, E., Baskin, C.: End-to-end referring video object segmentation with
73 multimodal transformers. In: CVPR. pp. 4975–4985 (2022) 1, 3
- 74 [2] Gavriluk, K., Ghodrati, A., Li, Z., Snoek, C.G.M.: Actor and action video segmentation from a
75 sentence. In: CVPR. pp. 5958–5966 (2018) 1
- 76 [3] Hu, R., Rohrbach, M., Darrell, T.: Segmentation from natural language expressions. In: Leibe,
77 B., Matas, J., Sebe, N., Welling, M. (eds.) ECCV. pp. 108–124 (2016) 1
- 78 [4] Liang, C., Wu, Y., Luo, Y., Yang, Y.: Clawcranenet: Leveraging object-level relation for text-
79 based video segmentation. arXiv preprint arXiv:2103.10702 (2021) 1
- 80 [5] Liu, S., Hui, T., Huang, S., Wei, Y., Li, B., Li, G.: Cross-modal progressive comprehension for
81 referring segmentation. TPAMI pp. 4761–4775 (2022) 1

- 82 [6] Seo, S., Lee, J., Han, B.: URVOS: unified referring video object segmentation network with a
83 large-scale benchmark. In: ECCV. pp. 208–223 (2020) [4](#)
- 84 [7] Wang, H., Deng, C., Yan, J., Tao, D.: Asymmetric cross-guided attention network for actor and
85 action video segmentation from natural language query. In: ICCV. pp. 3938–3947 (2019) [1](#)
- 86 [8] Wu, J., Jiang, Y., Sun, P., Yuan, Z., Luo, P.: Language as queries for referring video object
87 segmentation. In: CVPR. pp. 4964–4974 (2022) [1](#), [2](#), [3](#), [4](#), [5](#)
- 88 [9] Ye, L., Rochan, M., Liu, Z., Zhang, X., Wang, Y.: Referring segmentation in images and videos
89 with cross-modal self-attention network. TPAMI pp. 3719–3732 (2022) [1](#)