

MOSE: A New Dataset for Video Object Segmentation in Complex Scenes

Henghui Ding¹✉ Chang Liu¹ Shuting He¹ Xudong Jiang¹ Philip H.S. Torr² Song Bai³
¹Nanyang Technological University ²University of Oxford ³ByteDance

<https://henghuiding.github.io/MOSE>



Figure 1. Examples of video clips from the coMplex video Object SEgmentation (MOSE) dataset. The selected target objects are masked in orange ■. The most notable feature of MOSE is complex scenarios, including the disappearance-reappearance of objects, small/inconspicuous objects, heavy occlusions, crowded scenarios, *etc.* For example, the target player in the 2nd row turns around when reappearing in the 4th and 5th columns after disappearing in the 3rd column, bringing challenges in re-identifying him. Most videos in MOSE contain crowded and occluded objects with the target object seldom being the salient one. The goal of MOSE dataset is to provide a platform that promotes the development of more comprehensive and robust video object segmentation algorithms.

Abstract

Video object segmentation (VOS) aims at segmenting a particular object throughout the entire video clip sequence. The state-of-the-art VOS methods have achieved excellent performance (e.g., 90+% J&F) on existing datasets. However, since the target objects in these existing datasets are usually relatively salient, dominant, and isolated, VOS under complex scenes has rarely been studied. To revisit VOS and make it more applicable in the real world, we collect a new VOS dataset called coMplex video Object SEgmentation (MOSE) to study the tracking and segmenting objects in complex scenarios. MOSE contains 2,149 video clips and 5,200 objects from 36 categories, with 431,725 high-quality object segmentation masks. The most notable feature of MOSE dataset is complex scenes with crowded and occluded objects. The target objects in the videos are commonly occluded by others and disappear in some frames. To analyze the proposed MOSE dataset, we benchmark 18 existing VOS methods under 4 different settings on the proposed MOSE dataset and conduct comprehensive comparisons. The experiments show that

current VOS algorithms cannot well perceive objects in complex scenes. For example, under the semi-supervised VOS setting, the highest J&F by existing state-of-the-art VOS methods is only 59.4% on MOSE, much lower than their ~90% J&F performance on DAVIS. The results reveal that although excellent performance has been achieved on existing benchmarks, there are unresolved challenges under complex scenes and more efforts are desired to explore these challenges in the future.

1. Introduction

Video object segmentation (VOS) [73, 74, 97] aims at segmenting a particular object, e.g., the dominant objects or the objects indicated by users, throughout the entire video sequence. It is one of the most fundamental and challenging computer vision tasks and has a crucial role in many practical applications that involve video analysis and understanding, e.g., self-driving vehicle, augmented reality, video editing, *etc.* There are many different settings for VOS, for example, semi-supervised VOS [4, 71] that gives the first-frame mask of the target object, unsupervised VOS [14, 33] that automatically finds primary objects, and

✉ henghui.ding@gmail.com

interactive VOS [9, 68] that relies on the user’s interactions of the target object. Video object segmentation has been extensively studied in the past using traditional techniques or deep learning methods. Especially, the deep-learning-based methods have greatly improved the performance of video object segmentation and surpassed the traditional techniques by a large margin.

Current state-of-the-art methods have achieved very high performance on two of the most commonly-used VOS datasets DAVIS [73, 74] and YouTube-VOS [97]. For example, XMem [10] achieves **92.0%** $J&F$ on DAVIS 2016 [73], **87.7%** $J&F$ on DAVIS 2017 [74], and **86.1%** G on YouTube-VOS [97]. With such a high performance, it seems that the video object segmentation has been well resolved. However, do we really perceive objects in realistic scenarios? To answer this question, we revisit video object segmentation under realistic and more complex scenes. The target objects in existing datasets [73, 74, 97] are usually salient and dominant. In real-world scenarios, isolated and salient objects rarely appear while complex and occluded scenes happen frequently. To evaluate current state-of-the-art VOS methods under more complex scenarios, we collect 2,149 videos with complex scenes and form a new large-scale challenging video object segmentation benchmark, termed **coMplex video Object SEgmentation (MOSE)**. Specifically, there are 5,200 objects from 36 categories in MOSE, with 431,725 high-quality segmentation masks. As shown in Figure 1, the most notable feature of MOSE is complex scenarios, including disappearance-reappearance of objects, small/inconspicuous objects, heavy occlusions, crowded scenarios, *etc.* For example, the white sedan in the first row of Figure 1 is occluded by the bus, and the heaviest occlusion in 3rd image makes the sedan totally disappear. In the second row of Figure 1, the target player in the crowd is inconspicuous and disappears in the third frame due to the occlusions of the crowd. When the target player reappears, he turns around and shows a different appearance from the first two frames, which makes him very difficult to be tracked. The heavy occlusion and disappearance of objects under complex scenes bring great challenges to video object segmentation. We wish to promote video object segmentation research in complex environments and make VOS applicable in the real world.

To analyze the proposed MOSE dataset, we retrain and evaluate some existing VOS methods on MOSE. Specifically, we retrain 6 state-of-the-art methods under the semi-supervised setting using mask as the first-frame reference, 2 methods under the semi-supervised setting using bounding box as the first-frame reference, 3 methods under the multi-object zero-shot video object segmentation setting, and 7 methods under interactive setting. The experimental results show that videos of complex scenes make the current

state-of-the-art VOS methods less pronounced, especially in terms of tracking objects that disappear for a while due to occlusions. For example, the $J&F$ performance of XMem [10] on DAVIS 2016 is **92.0%** but drop to **57.6%** on MOSE, the $J&F$ performance of DeAOT [104] on DAVIS 2016 is **92.9%** but drop to **59.4%** on MOSE, which consistently reveal the difficulties brought by complex scenes.

The poor performance on MOSE is due to not only occlusions/crowds/small-scale in static images but also objects’ disappearance-reappearance and flickering across the temporal domain. While the heavy occlusions, crowds, and small objects bring challenges to the segmentation of objects in images, the disappearance-reappearance of objects makes it extremely difficult to track an occluded object, increasing the challenge of association.

In a summary, our main contributions are as follows:

- We build a new video object segmentation benchmark dataset termed MOSE (**coMplex video Object SEgmentation**). MOSE focuses on understanding objects in videos under complex environments.
- We conduct comprehensive comparison and evaluation of state-of-the-art VOS methods on the MOSE dataset under 4 different settings, including mask-initialization semi-supervised, box-initialization semi-supervised, unsupervised, and interactive settings.
- Taking a close look at MOSE, we analyze the challenges and potential directions for future video understanding research in complex scenes.

2. Related Work

2.1. Video Object Segmentation (VOS)

Video object segmentation (VOS) aims at segmenting a particular object throughout the entire video clip sequence. According to how to indicate the particular object, there are mainly four different settings, *i.e.*, semi-supervised VOS (or semi-automatic VOS [89] or one-shot VOS), unsupervised VOS (or automatic VOS [89] or zero-shot VOS), interactive VOS, and referring VOS.

• **Semi-supervised VOS.** Semi-supervised video object segmentation (or one-shot video object segmentation) [4] gives the first frame object mask and target to segment the target object throughout the remaining video frames. Most existing works can be categorized into propagation-based methods [8, 13, 25, 30–32, 34, 35, 51, 72, 94, 95, 97, 98, 109] and matching-based methods [13, 21, 22, 29, 40, 69, 85, 91, 102, 107, 111]. The propagation-based methods utilize the mask of previous frame to guide the mask generation of the current frame, propagating the clues of target object frame by frame. The matching-based methods memorize the embedding of target object and then conduct per-pixel

classification to measure the similarity of each pixel’s feature to the embedding of target object.

Since pixel-wise masks are hard to be obtained, some semi-supervised VOS works propose utilize **bounding box** as the first-frame clue to indicate the target object [50, 81, 86]. For example, SiamMask [86] applies a mask prediction branch on fully-convolutional Siamese object tracker to generate binary segmentation masks.

- **Interactive VOS.** Interactive video object segmentation is a special form of semi-supervised video object segmentation [6, 9, 11, 13, 28, 66, 68, 106], it aims at segmenting the target object in a video indicated by user’s interaction, *e.g.*, clicks or scribbles. Existing interactive VOS mainly follows a paradigm of interaction-propagation way. Besides the feature encoder that extracts discriminative pixel features, there are other two modules placed on the feature encoder to achieve interactive video object segmentation, *i.e.*, interactive segmentation module that corrects prediction based on user’s interaction and mask propagation module that propagates user-corrected masks to other frames.

- **Referring VOS.** Referring video object segmentation [19, 79, 105] is an emerging setting that involves multi-modal information. It gives a natural language expression to indicate the target object and aims at segmenting the target object throughout the video clips. Existing referring video object segmentation methods can be categorized into two ways: bottom-up methods and top-down methods. Bottom-up methods [20, 56, 79] directly segment the target object at the first frame and propagate the mask to the remaining frames, or conduct image segmentation on each frame independently and then associate these masks. Top-down methods [2, 49, 93] exhaustively propose all potential tracklets and select the one that is best matched with language expression as output.

- **Unsupervised VOS.** It is also known as automatic VOS or zero-shot VOS [24, 44, 54, 58, 61, 78, 83, 88, 100, 101]. Different from the above VOS settings, unsupervised VOS does not require any manual clues to indicate the objects but aims to automatically find the primary objects in a video. Unsupervised VOS can only deal with objects of pre-defined categories. Early methods usually use post-processing steps [24]. Then end-to-end training methods become the mainstream in unsupervised VOS, which can be categorized into local content encoding and contextual content encoding. The local content encoding methods [14, 33, 43, 44, 82, 83, 112] typically employ two parallel networks to extract features from optical flow and RGB image. The contextual content encoding methods [59, 60, 87] focus on capturing long-term global information.

2.2. Related Video Segmentation Tasks

There are some other video segmentation tasks that are related to video object segmentation, for example, video

instance segmentation, video semantic segmentation, and video panoptic segmentation.

- **Video Instance Segmentation (VIS).** Video instance segmentation is extended from image instance segmentation by Yang *et al.* [99], it simultaneously conducts detection, segmentation, and tracking of instances of predefined categories in videos. Thanks to the large-scale VIS dataset YouTube-VIS [99], a series of learning methods have been developed and greatly advanced the performance of VIS [36, 48]. Then, occluded video instance segmentation is proposed by [76] to study the video instance segmentation under occluded scenes. Similar to [76], we study video understanding under complex scenarios like occlusions, but different from [76], we focus on video object segmentation (VOS) and our proposed MOSE dataset contains more videos and more complex scenes than [76], especially in terms of object’s disappearance-reappearance.

- **Video Semantic Segmentation (VSS).** Driven by the great success in image semantic segmentation [7, 57] and large-scale video semantic segmentation datasets [3, 17, 65], video semantic segmentation has drawn lots of attention and achieved significant achievement. Compared to image domain, temporal consistency and model efficiency are the new efforts in the video domain. For example, Sun *et al.* [80] propose Coarse-to-Fine Feature Mining (CFFM) to capture both static context and motional context.

- **Video Panoptic Segmentation (VPS).** Kim *et al.* [38] introduce panoptic segmentation to the video domain to simultaneously segment and track both the foreground instance objects and background stuff. They build a VPS dataset with 124 videos. Then, Miao *et al.* [64] build a larger VPS dataset called VIPSeg with 3,536 videos. Existing methods [77, 92] mainly add temporal refinement or cross-frame association modules upon image panoptic segmentation models to enhance temporal conformity and instance tracking performance.

2.3. Complex Scene Understanding

Complex scene understanding has become a research focus in the image understanding domain [15, 18, 39, 41, 47, 63, 90, 108, 110]. For example, Ke *et al.* [37] propose Bilayer Convolutional Network (BCNet) to decouple overlapping objects into occluder and occludee layers. Zhang *et al.* [108] propose a self-supervised approach to conduct de-occlusion by ordering recovery, amodal completion, and content completion. On the video domain, however, occlusion understanding is still underexplored with only several multi-object tracking works [16, 55, 96, 113]. For example, Chu *et al.* [16] propose a spatial temporal attention mechanism (STAM) to capture the visible parts of targets and deal with the drift brought by occlusion. Zhu *et al.* [113] propose dual matching attention networks (DMAN) to deal with the noisy occlusions in multi-object tracking. Li *et al.* [46] propose

Table 1. A complete list of object categories and their #instances in the MOSE dataset. The object categories are sorted in descending order of their frequency of occurrence.

Category	No.	Category	No.	Category	No.	Category	No.
Person-other	519	Monkey	181	Bicycle	130	Boat	73
Fish	343	Dog	175	Cyclist	116	Lizard	68
Horse	272	Boat	166	Tiger	108	Duck	64
Sheep	264	Sedan	160	Giraffe	108	Goose	54
Zebra	254	Motorcyclist	153	Panda	106	Horse-rider	46
Rabbit	237	Turtle	142	Driver	90	Bus	29
Bird	226	Cat	139	Airplane	89	Truck	15
Elephant	207	Parrot	139	Chicken	86	Vehicle-other	14
Motorcycle	192	Cow	139	Bear	84	Poultry-other	12

to track every thing in the open world by performing class-agnostic association. In this work, we build a Complex Video Object Segmentation dataset to support future work on complex scene understanding in VOS.

3. MOSE Dataset

In this section, we introduce the newly built MOSE dataset. We first present the video collection and annotation process in Section 3.1 and then give the dataset statistics and analysis in Section 3.2. Finally, we give the evaluation metrics in Section 3.3.

3.1. Video Collection and Annotation

Video Collection. The videos in MOSE are obtained from two sources. The first source is inherited from the OVIS dataset [76], which was designed for video instance segmentation. The second source consists of newly captured videos from real-world scenarios. We chose to include videos from the OVIS dataset because they contain numerous heavy occlusions and meet the requirements of complex scenes. However, for videos inherited from OVIS, since many objects appear for the first time after the first frame, they cannot be directly used for video object segmentation, which requires a first-frame reference. One solution would be to discard all objects that do not appear in the first frame, but this would result in a significant waste of mask annotations. To make full use of OVIS videos and adapt them for video object segmentation tasks, we cut each of these videos into several videos according to the frame where each object first appears in the videos. We then discard videos that do not meet our requirements. For the newly captured videos, we first select a set of semantic categories that are common in the real world, including vehicles (*e.g.*, bicycle, bus, airplane, boat), animals (*e.g.*, bird, panda, dog, giraffe), and human beings in different activities (*e.g.*, motorcycling, driving, riding, running), and

then collect/film videos containing these categories from the campus, zoo, indoor, city street, etc. A complete list of object categories in the MOSE dataset can be seen in Table 1. These categories are common in our daily lives, where crowds and occlusions frequently occur. Moreover, they are included in popular large-scale image segmentation datasets such as MS-COCO [52], which makes it easy to use image-pretrained models.

Our primary concern is video object segmentation under complex scenes that contain crowds, occlusion, multi-scale objects, and other challenging scenarios. Therefore, we established several rules to ensure that complex environments are included when collecting/shooting videos:

- R1. Each video must contain several objects, while videos with only a single object are excluded. Specifically, videos with crowded objects of similar appearance are highly valued.
- R2. Occlusions should be present in the video. Videos that do not have any occlusions throughout the entirety of the frames are discarded. We encourage occlusions caused by other moving objects.
- R3. Great emphasis should be placed on scenarios where objects disappear and then reappear due to occlusions or crowding.
- R4. The target objects should be of a variety of scales and types, such as small-scale or large-scale, conspicuous or inconspicuous.
- R5. Objects in the video must show sufficient motion. Videos with completely still objects or very little motion are discarded.

Besides the rules mentioned above, we also require the video quality to be of high standard, which means the video resolution should be 1920×1080 and the video length should be between 5 to 60 seconds in general.

Video Annotation. Having all videos for MOSE been collected, our research team looks through them and figure out a set of targets-of-interest for each of the videos. Then we slightly clip the start and the end of videos, to reduce the number of less motional or simple frames in the video. Next we annotate the first-frame mask of the target objects, as VOS input. Following this, the videos are sent to the annotation team along with the first-frame masks for annotation of the subsequent video frames.

Using the given first-frame mask as a reference, the annotation team is required to identify the target object in the given first-frame mask, then track and annotate the segmentation mask of the target object in all frames following the first frame. The process of annotating videos has been made easier with the help of an interactive annotation tool

Table 2. Scale comparison between MOSE and existing VOS datasets. “Annotations” denotes the number of annotated object masks. “Duration” denotes the total duration (in minutes) of the annotated videos. “mBOR” denotes the mean of the Bounding-box-Occlusion Rate. “Disapp. Rate” represents the frequency of objects that disappear in at least one frame. The newly built MOSE has the longest video duration and the largest number of annotations. More importantly, the most notable feature of MOSE is that it contains lots of crowds, occlusions, and disappearance-reappearance objects, which provide much more complex scenarios for video object segmentation.

Dataset	Year	Videos	Categories	Objects	Annotations	Duration	mBOR	Disapp. Rate
YouTube-Objects [75]	2012	96	10	96	1,692	9.01	-	-
SegTrack-v2 [42]	2013	14	11	24	1,475	0.59	0.12	8.3%
FBMS [67]	2014	59	16	139	1,465	7.70	0.01	11.2%
JumpCut [23]	2015	22	14	22	6,331	3.52	0	0%
DAVIS ₁₆ [73]	2016	50	-	50	3,440	2.88	-	-
DAVIS ₁₇ [74]	2017	90	-	205	13,543	5.17	0.03	16.1%
YouTube-VOS [97]	2018	4,453	94	7,755	197,272	334.81	0.05	13.0%
MOSE (ours)	2023	2,149	36	5,200	431,725	443.62	0.23	41.5%

that we developed. The annotation tool automatically loads videos and all target objects. Annotators use the tool to load and preview videos and first-frame masks, annotate and visualize the segmentation masks in the subsequent frames, and save them. The annotation tool also has a built-in interactive object segmentation network XMem [10], to assist annotations in producing high-quality masks. To ensure the annotation quality under complex scenes, the annotators are required to clearly track the object that disappears and reappears due to heavy occlusions and crowd. For frames in which the target object is disappeared or is fully occluded, annotators also need to confirm that the output masks of such frames shall be blank. It is a requirement that all of our videos be annotated every five frames at the very least. For the purpose of testing the frame-rate robustness of the models, some videos are annotated every frame.

After annotation, the videos are reviewed by our verification team to ensure high-quality annotations, particularly for videos containing occlusions, crowds, or instances of objects disappearing and reappearing.

3.2. Dataset Statistics

In Table 2, we analyze the data statistics of the new MOSE dataset using previous datasets for the video object segmentation as a reference for the analysis, including JumpCut [23], SegTrack-v2 [42], YouTube-Objects [75], FBMS [67], DAVIS [73, 74], and YouTube-VOS [97]. As shown in Table 2, there are 2,149 videos and 431,725 mask annotations for 5,200 objects contained in the proposed MOSE dataset. Comparing with previous largest VOS dataset YouTube-VOS [97], MOSE has much more mask annotations, 431k vs. 197k. Among the 8 video object segmentation datasets that we analyzed in Table 2, MOSE has the longest video duration (433.62 minutes), which is about 100 minutes longer than the second longest dataset YouTube-VOS (334.81 minutes) and is hundreds of times

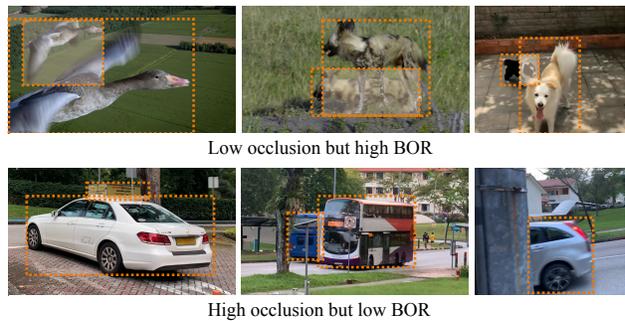


Figure 2. Failure cases of the BOR indicator. It can be seen from the first row of samples that they have high BOR values, but there is less or no occlusion present. Samples in the second row have very small BOR values, but there are severe occlusions in the samples.

longer than the other seven datasets. We include more long videos in the MOSE compared to previous VOS datasets to ensure that our dataset contains a diverse range of occlusion scenarios, motion scenarios, and instances of objects disappearing and reappearing.

Disappearance and Occlusion Analysis. Previous occluded video datasets OVIS [76] defines a Bounding-box Occlusion Rate (BOR) to reflect the occlusion degree. The BOR is calculated by the Intersection-over-Union of bounding boxes in the videos. We provide the mean of BOR of all frames in Table 2. The MOSE has the largest value of mBOR, indicating the more frequency of occlusions in MOSE than previous datasets. However, we find that BOR can only roughly reflect the occlusion. It does not well reveal the degree of occlusion in the MOSE dataset and may be prone to making some mistakes in this regard. As shown in the first row of Figure 2, high BOR is observed in the three images with low occlusion or even no occlusion. In

contrast, in the images of the second row, where there are heavily occluded objects, the BOR is low and even 0.

Therefore, besides counting BOR, we further calculate the number of disappeared objects that disappear in at least one frame of the video, followed by the disappearance rate that reflects the frequency of disappearances, which illustrates how often the disappearance of objects occurs due to complex scenarios in a dataset. As shown in Table 2, the disappearance rate (Disapp. Rate) of MOSE stands at the highest, reaching 41.5%. Meanwhile, a total of 27.0% of objects in MOSE showcase the disappearance-reappearance phenomenon, meaning they completely vanish and subsequently reappear in future frames. These observations underscore the frequent and substantial occurrence of disappearance and occlusions in the proposed MOSE dataset.

3.3. Evaluation Metrics

In accordance with the previous method [73, 74], we compute the region similarity \mathcal{J} and the contour accuracy \mathcal{F} as evaluation metrics. Given segmentation mask predictions $\hat{M} \in \{0, 1\}^{H \times W}$ and the ground-truth masks $M \in \{0, 1\}^{H \times W}$, region similarity \mathcal{J} is obtained by calculating the Intersection-over-Union (IoU) of \hat{M} and M ,

$$\mathcal{J} = \frac{\hat{M} \cap M}{\hat{M} \cup M}. \quad (1)$$

Then, the average region similarity \mathcal{J}_{mean} over all objects is calculated as the final region similarity result. We use \mathcal{J} to represent \mathcal{J}_{mean} for the sake of brevity. To measure the contour quality of \hat{M} , contour recall R_c and precision P_c are calculated via bipartite graph matching [62]. Then, the contour accuracy \mathcal{F} is the harmonic mean of the contour recall R_c and precision P_c , *i.e.*,

$$\mathcal{F} = \frac{2P_c R_c}{P_c + R_c}, \quad (2)$$

which represents how closely the contours of predicted masks resemble the contours of ground-truth masks. The average contour accuracy \mathcal{F}_{mean} over all objects is calculated as the final contour accuracy result. We use \mathcal{F} to represent \mathcal{F}_{mean} for the sake of brevity. Then, $\mathcal{J} \& \mathcal{F} = (\mathcal{J} + \mathcal{F})/2$ is used to measure the overall performance.

4. Experiments

Herein we conduct experiments and benchmarks on four different Video Object Segmentation (VOS) settings, including semi-supervised (or one-shot) VOS with mask-initialization, semi-supervised VOS with box-initialization, unsupervised (or zero-shot) VOS, and interactive VOS, to comprehensively analyze the newly built MOSE dataset.

Implementation Details. The proposed MOSE dataset is consistent with the YouTube-VOS [97] format. We replace

Table 3. Comparisons of state-of-the-art semi-supervised methods on the validation set. “ \mathcal{J} ” and “ \mathcal{F} ” denote the mean of region similarity and the mean of contour accuracy. $\mathcal{J} \& \mathcal{F}$ denotes the mean of \mathcal{J} and \mathcal{F} . BL30K [11] is not added during the training stage to make a fair comparison.

Method	Pub.	MOSE (ours)			DAVIS ₁₇	YT-VOS ₁₈
		\mathcal{J}	\mathcal{F}	$\mathcal{J} \& \mathcal{F}$	$\mathcal{J} \& \mathcal{F}$	\mathcal{G}
AOT [103]	NeurIPS’21	53.1	61.3	57.2	84.9	84.1
STCN [12]	NeurIPS’21	46.6	55.0	50.8	85.4	83.0
RDE [45]	CVPR’22	44.6	52.9	48.8	84.2	-
SWEM [53]	CVPR’22	46.8	54.9	50.9	84.3	82.8
XMem [10]	ECCV’22	53.3	62.0	57.6	86.2	85.7
DeAOT [104]	NeurIPS’22	55.1	63.8	59.4	85.2	86.0

the training dataset of previous methods from YouTube-VOS with our MOSE and strictly follow their training settings on YouTube-VOS [97]. We follow DAVIS [73, 74] to evaluate the performance and report the results of \mathcal{J}_{mean} , \mathcal{F}_{mean} , and $\mathcal{J} \& \mathcal{F}$ on the validation set of MOSE. During the training of these VOS methods, we follow their way of using image pre-trained models but do not use any additional video datasets for pretraining.

Settings. There are 2,149 videos in the whole MOSE dataset. These videos are split into 1,507 training videos, 311 validation videos, and 331 testing videos, for model training, daily evaluation, and competition period evaluation, respectively. Each of the videos gives a first-frame mask or bounding box as the reference of the target object for the semi-supervised (one-shot) VOS setting, and first-frame scribbles for the interactive VOS setting.

4.1. Semi-supervised Video Object Segmentation

Semi-supervised video object segmentation, or one-shot video object segmentation, gives either the first frame mask or bounding box of the target object as a clue and reference. We train and evaluate the very recent 6 mask-initialization semi-supervised VOS methods and 2 box-initialization semi-supervised VOS methods built upon ResNet-50 [26] on the MOSE dataset, as shown in Table 3 and Table 4, respectively. It is our hope that the experiments will provide baselines for future semi-supervised VOS algorithms to be developed.

Mask-initialization. This setting is a classic and currently the most popular topic for video object segmentation. A lot of excellent deep-learning-based works have been developed for this setting in the past decade and greatly improve the video object segmentation performance to a saturation level. For example, the most recent state-of-the-art method DeAOT [104] achieves **85.2%** $\mathcal{J} \& \mathcal{F}$ on the DAVIS 2017 [74], **86.0%** \mathcal{G} on YouTube-VOS [97], and **92.3%** $\mathcal{J} \& \mathcal{F}$ on DAVIS 2016 [73], which are excellent

Table 4. Comparisons of state-of-the-art box-initialization semi-supervised methods on the validation set. “ \mathcal{J} ” and “ \mathcal{F} ” denote the mean of region similarity and the mean of contour accuracy. $\mathcal{J}\&\mathcal{F}$ denotes the mean of \mathcal{J} and \mathcal{F} .

Method	Pub.	MOSE (ours)			DAVIS ₁₇	YT-VOS ₁₈
		\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$	\mathcal{G}
SiamMask [86]	CVPR’19	17.3	26.7	22.0	56.4	52.8
FTMU [81]	CVPR’20	19.1	28.5	23.8	70.6	-

performance and almost reach the ground truth. To analyze the newly built MOSE dataset and test the performance of existing methods in complex scenes, we train and evaluate these methods on MOSE. We report the results of 6 recent state-of-the-art methods on the validation set of MOSE, including AOT [103], STCN [12], RDE [45], SWEM [53], XMem [10], and DeAOT [104], as shown in Table 3. Current state-of-the-art methods only achieve the performance from **48.8%** $\mathcal{J}\&\mathcal{F}$ to **59.4%** $\mathcal{J}\&\mathcal{F}$ on the validation set of the newly proposed MOSE dataset, while their results on DAVIS 2017 [74] and YouTube-VOS [97] are usually above **80%** $\mathcal{J}\&\mathcal{F}$ or \mathcal{G} , some of them almost reach **90%** $\mathcal{J}\&\mathcal{F}$ or \mathcal{G} on DAVIS 2017 and YouTube-VOS. The results reveal that although we have achieved excellent VOS performance on previous benchmarks, there are unresolved challenges under complex scenes and more efforts are desired to explore these challenges in the future. Discussion about the MOSE dataset and some potential future directions are provided in Section 5.

Box-initialization. As shown in Table 4, two semi-supervised (one-shot) VOS methods with a bounding box as the first-frame reference are trained on MOSE and evaluated on the validation set of MOSE, including SimMask [86] and FTMU [81]. Current state-of-the-art methods achieve only **22.0%** and **23.8%** $\mathcal{J}\&\mathcal{F}$ on the validation set of the newly proposed MOSE dataset, while their results on DAVIS 2017 [74] are already **70+%** $\mathcal{J}\&\mathcal{F}$.

After careful analysis, we believe that one of the reasons for such a significant drop in performance is that there are many heavy occlusions, small-scale objects, and crowded scenarios in the videos of MOSE dataset, which make the weakly supervised segmentation of the videos much more difficult. There is a possibility that, even within the target bounding box region, the target object may not be the most salient one due to heavy occlusions and crowd. The occlusion of an object can cause it to be broken up into several pieces that are not adjacent to each other, which greatly increases the difficulty of segmenting an object using a box-driven approach.

4.2. Unsupervised Video Object Segmentation

Unsupervised video object segmentation, or zero-shot video object segmentation, does not require any manual clues (e.g., mask or bounding box) as a reference to indicate

Table 5. Comparisons of state-of-the-art multi-object zero-shot VOS methods on the validation set. “ \mathcal{J} ” and “ \mathcal{F} ” denote the mean of region similarity and the mean of contour accuracy. $\mathcal{J}\&\mathcal{F}$ denotes the mean of \mathcal{J} and \mathcal{F} .

Method	Pub.	MOSE (ours)			DAVIS ₁₇
		\mathcal{J}	\mathcal{F}	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}\&\mathcal{F}$
RVOS [84]	CVPR’19	24.1	36.9	30.5	43.7
AGNN [87]	ICCV’19	38.6	48.8	43.7	61.1
STEm-Seg [1]	ECCV’20	43.3	50.5	46.9	64.7

Table 6. Comparisons of state-of-the-art interactive VOS methods on the validation set. $\mathcal{J}\&\mathcal{F}@60s$ denotes the $\mathcal{J}\&\mathcal{F}$ performance reached by methods within 60 seconds interactions.

Method	Pub.	MOSE (ours)	DAVIS ₁₇
		$\mathcal{J}\&\mathcal{F}@60s$	$\mathcal{J}\&\mathcal{F}@60s$
IPNet [68]	CVPR’19	41.2	78.7
STM [70]	ICCV’19	45.3	84.8
MANet [66]	CVPR’20	43.6	79.5
ATNet [27]	ECCV’20	44.5	82.7
GIS [28]	CVPR’21	50.2	86.6
MiVOS [11]	CVPR’21	51.4	88.5
STCN [12]	NeurIPS’21	56.8	88.8

the objects but aims to find the primary objects in a video automatically. The mainstream of zero-shot VOS methods focuses on addressing single-object VOS. However, MOSE is a multi-object VOS dataset like DAVIS 2017 [74], thus we only benchmark multi-object zero-shot VOS methods on MOSE, e.g., STEm-Seg [1], AGNN [87], and RVOS [84]. The results are shown in Table 5. In this setting, only videos with exhaustive first-frame annotations are used. Existing methods rely on off-the-shelf image-trained instance segmentation methods, which can well detect/segment objects in the static image thanks to complex scene learning in the image domain. The performance drop is mainly due to temporal challenges.

4.3. Interactive Video Object Segmentation

Following the interactive track of DAVIS 2019 Challenge [5], we provide initial scribbles for the target object in a specific video sequence, which are used as the first interaction. The interactive video object segmentation methods are required to predict the segmentation mask for the whole video based on the first interaction. Then, by comparing predicted masks and ground truth masks across the video, corrective scribbles on the worst frame are further provided for the methods to refine the video segmentation prediction. The above step is allowed to be repeated up to 8 times with a time limitation of 30s for each object. We report the metric of $\mathcal{J}\&\mathcal{F}@60s$ to encourage the methods to have a good balance between speed and accuracy. As shown in Table 6, seven recent interactive video object segmentation methods are evaluated on the validation set of MOSE.

5. Discussion and Future Directions

Herein we discuss the challenges of the proposed MOSE dataset brought by complex scenes and provide some potential future directions based on the experimental results analysis of the existing methods on MOSE.

- **Stronger Association to Track Reappearing Objects.** It is necessary to develop stronger association/re-identification algorithms for VOS methods in order to be able to track objects that disappear and then reappear. Especially, the most interesting thing is that we have noticed that a number of disappeared-then-reappeared objects have been reappearing with a different appearance from the time they disappeared, *i.e.*, appearance-changing objects. For example, the target player in Figure 1 shows his back before disappearing while showing the front at reappearing. There is a great deal of difficulty in tracking with such an appearance change.

- **Video Object Segmentation of Occluded Objects.** The frequent occlusions in the videos of MOSE provide data support to the research of occlusion video understanding. Tracking and segmenting an object with heavy occlusion have rarely been studied since the target objects in existing datasets are usually salient and dominant. In real-world scenarios, isolated objects rarely appear while occluded scenes occur frequently. Our human beings can well capture those occluded objects, thanks to our contextual association and reasoning ability. With the introduction of an understanding of occlusion into the process of video object segmentation, the VOS methods will become more practical for use in real-world applications. Especially, the occlusions make the box-initialization semi-supervised VOS setting more challenging in terms of segmenting an occluded object by bounding box.

- **Attention on Small & Inconspicuous Objects.** Although the detection and segmentation of small objects is a hot topic in the image domain, tracking and segmenting small and inconspicuous objects in the video object segmentation domain is still to be developed. As a matter of fact, most of the existing video object segmentation methods mainly focus on tracking large and salient objects. The lack of sufficient attention to small-scale objects makes the video object segmentation methods less pronounced in practical applications that may involve target objects of varying sizes and types. There are many objects in the proposed MOSE dataset that are on a small scale and are inconspicuous in the videos, which provides a greater opportunity for the research of tracking and segmenting objects that are small and inconspicuous in realistic scenarios.

- **Tracking Objects in Crowd.** One of the most notable features of the MOSE is crowded scenarios, which are common in real-world applications. There are many videos in the MOSE dataset, which contain crowd objects, such as flocks of sheep moving together, groups of cyclists racing,

pedestrians moving on a crowded street, and *etc.* A scenario like this presents challenges when it comes to segmenting and tracking one object among a crowd of objects that share a similar appearance and motion to the target object with the crowd as a whole. In the image/frame domain, it is desired for video object segmentation algorithms to enhance the identification ability to distinguish between different objects with a similar appearance. What's more, in the temporal domain, it is challenging to perform the object association between the two frames with the crowd consisting of similar-looking objects.

- **Long-Term Video Segmentation:** In terms of practical applications like movie editing and surveillance monitoring, long-term video understanding is much more practical than short-term video understanding. For example, the average length of videos on YouTube is around twelve minutes, which is much longer than the average length of the existing video object segmentation dataset. Although existing VOS methods provide excellent performance, they typically require a lot of computing resources, *e.g.*, GPU memory, for storing the object features of previous frames. Due to the large computation cost, most existing VOS methods cannot well handle long videos. The average video length of the MOSE dataset is longer than existing VOS datasets, bringing more challenges and opportunities in dealing with long-term videos. It is a good and practical direction to design VOS algorithms that can deal with long videos with low computation costs while achieving high-quality segmentation results.

6. Conclusion

We build a large-scale video object segmentation dataset named MOSE to revisit and support the research of video object segmentation under complex scenes. There are 2,149 high-resolution videos in MOSE with 431,725 high-quality object masks for 5,200 objects from 36 categories. The videos in MOSE are commonly long enough to ensure diverse and sufficient occlusion, motion, and disappearance-reappearance scenarios. Based on the proposed MOSE dataset, we benchmark existing VOS methods and conduct a comprehensive comparison. We train and evaluate six methods under the semi-supervised (one-shot) setting with a mask as the first-frame reference, two methods with a bounding box as the first-frame reference, three methods under the unsupervised (zero-shot) setting, and seven methods under the interactive setting. After evaluating existing VOS methods on MOSE dataset and comprehensively analyzing the results, some challenges and potential directions for future VOS research are concluded. We find that we are still at a nascent stage of segmenting and tracking objects in complex scenes where crowds, disappearing, occlusions, and inconspicuous/small objects occur frequently.

References

- [1] Ali Athar, Sabarinath Mahadevan, Aljoša Ošep, Laura Leal-Taixé, and Bastian Leibe. Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In *ECCV*, 2020. 7
- [2] Adam Botach, Evgenii Zheltonozhskii, and Chaim Baskin. End-to-end referring video object segmentation with multimodal transformers. In *CVPR*, 2022. 3
- [3] Gabriel J Brostow, Julien Fauqueur, and Roberto Cipolla. Semantic object classes in video: A high-definition ground truth database. *Pattern Recognition Letters*, 30(2), 2009. 3
- [4] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *CVPR*, 2017. 1, 2
- [5] Sergi Caelles, Jordi Pont-Tuset, Federico Perazzi, Alberto Montes, Kevis-Kokitsi Maninis, and Luc Van Gool. The 2019 davis challenge on vos: Unsupervised multi-object segmentation. *arXiv:1905.00737*, 2019. 7
- [6] Bowen Chen, Huan Ling, Xiaohui Zeng, Gao Jun, Ziyue Xu, and Sanja Fidler. Scribblebox: Interactive annotation framework for video object segmentation. In *ECCV*, 2020. 3
- [7] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE TPAMI*, 40(4), 2017. 3
- [8] Xi Chen, Zuoxin Li, Ye Yuan, Gang Yu, Jianxin Shen, and Donglian Qi. State-aware tracker for real-time video object segmentation. In *CVPR*, 2020. 2
- [9] Yuhua Chen, Jordi Pont-Tuset, Alberto Montes, and Luc Van Gool. Blazingly fast video object segmentation with pixel-wise metric learning. In *CVPR*, 2018. 2, 3
- [10] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *ECCV*. Springer, 2022. 2, 5, 6, 7
- [11] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion. In *CVPR*, 2021. 3, 6, 7
- [12] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. In *NeurIPS*, 2021. 6, 7
- [13] Jingchun Cheng, Yi-Hsuan Tsai, Wei-Chih Hung, Shengjin Wang, and Ming-Hsuan Yang. Fast and accurate online video object segmentation via tracking parts. In *CVPR*, 2018. 2, 3
- [14] Jingchun Cheng, Yi-Hsuan Tsai, Shengjin Wang, and Ming-Hsuan Yang. Segflow: Joint learning for video object segmentation and optical flow. In *ICCV*, 2017. 1, 3
- [15] Meng-Jiun Chiou, Henghui Ding, Hanshu Yan, Changhu Wang, Roger Zimmermann, and Jiashi Feng. Recovering the unbiased scene graphs from the biased ones. In *ACM MM*, 2021. 3
- [16] Qi Chu, Wanli Ouyang, Hongsheng Li, Xiaogang Wang, Bin Liu, and Nenghai Yu. Online multi-object tracking using cnn-based single object tracker with spatial-temporal attention mechanism. In *ICCV*, 2017. 3
- [17] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 3
- [18] Henghui Ding, Xudong Jiang, Bing Shuai, Ai Qun Liu, and Gang Wang. Semantic segmentation with context encoding and multi-path decoding. *IEEE TIP*, 29, 2020. 3
- [19] Henghui Ding, Chang Liu, Shuting He, Xudong Jiang, and Chen Change Loy. MeViS: A large-scale benchmark for video segmentation with motion expressions. In *ICCV*, 2023. 3
- [20] Henghui Ding, Chang Liu, Suchen Wang, and Xudong Jiang. VLT: Vision-language transformer and query generation for referring segmentation. *IEEE TPAMI*, 2023. 3
- [21] Kevin Duarte, Yogesh S. Rawat, and Mubarak Shah. Capsulevos: Semi-supervised video object segmentation using capsule routing. In *ICCV*, 2019. 2
- [22] Brendan Duke, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W Taylor. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *CVPR*, 2021. 2
- [23] Qingnan Fan, Fan Zhong, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. JumpCut: non-successive mask transfer and interpolation for video cutout. *ACM Tran. Graphics*, 34(6), 2015. 5
- [24] Katerina Fragkiadaki, Pablo Arbelaez, Panna Felsen, and Jitendra Malik. Learning to segment moving objects in videos. In *CVPR*, 2015. 3
- [25] Junwei Han, Le Yang, Dingwen Zhang, Xiaojun Chang, and Xiaodan Liang. Reinforcement cutting-agent learning for video object segmentation. In *CVPR*, 2018. 2
- [26] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [27] Yuk Heo, Yeong Jun Koh, and Chang-Su Kim. Interactive video object segmentation using global and local transfer modules. In *ECCV*, 2020. 7
- [28] Yuk Heo, Yeong Jun Koh, and Chang-Su Kim. Guided interactive video object segmentation using reliability-based attention maps. In *CVPR*, 2021. 3, 7
- [29] Li Hu, Peng Zhang, Bang Zhang, Pan Pan, Yinghui Xu, and Rong Jin. Learning position and target consistency for memory-based video object segmentation. In *CVPR*, 2021. 2
- [30] Ping Hu, Gang Wang, Xiangfei Kong, Jason Kuen, and Yap-Peng Tan. Motion-guided cascaded refinement network for video object segmentation. In *CVPR*, 2018. 2
- [31] Xuhua Huang, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Fast video object segmentation with temporal aggregation network and dynamic template matching. In *CVPR*, 2020. 2
- [32] Allan Jabri, Andrew Owens, and Alexei Efros. Space-time correspondence as a contrastive random walk. In *NeurIPS*, 2020. 2
- [33] Suyog Dutt Jain, Bo Xiong, and Kristen Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *CVPR*, 2017. 1, 3
- [34] Varun Jampani, Raghudeep Gadde, and Peter V. Gehler. Video propagation networks. In *CVPR*, 2017. 2
- [35] Won-Dong Jang and Chang-Su Kim. Online video object segmentation via convolutional trident network. In *CVPR*, 2017. 2
- [36] Lei Ke, Henghui Ding, Martin Danelljan, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Video mask transfinder for high-quality video instance segmentation. In *ECCV*. Springer, 2022. 3
- [37] Lei Ke, Yu-Wing Tai, and Chi-Keung Tang. Deep occlusion-aware instance segmentation with overlapping bilayers. In *CVPR*, 2021. 3
- [38] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Video panoptic segmentation. In *CVPR*, 2020. 3
- [39] Adam Kortylewski, Qing Liu, Angtian Wang, Yihong Sun, and Alan Yuille. Compositional convolutional neural networks: A robust and interpretable model for object recognition under occlusion. *IJCV*, 129(3), 2021. 3
- [40] Zihang Lai, Erika Lu, and Weidi Xie. MAST: A memory-augmented self-supervised tracker. In *CVPR*, 2020. 2
- [41] Justin Lazarow, Kwonjoon Lee, Kunyu Shi, and Zhuowen Tu. Learning instance occlusion for panoptic segmentation. In *CVPR*, 2020. 3
- [42] Fuxin Li, Taeyoung Kim, Ahmad Humayun, David Tsai, and James M Rehg. Video segmentation by tracking many figure-ground segments. In *ICCV*, 2013. 5
- [43] Guanbin Li, Yuan Xie, Tianhao Wei, Keze Wang, and Liang Lin. Flow guided recurrent neural encoder for video salient object detection. In *CVPR*, 2018. 3

- [44] Haofeng Li, Guanqi Chen, Guanbin Li, and Yizhou Yu. Motion guided attention for video salient object detection. In *ICCV*, 2019. 3
- [45] Mingxing Li, Li Hu, Zhiwei Xiong, Bang Zhang, Pan Pan, and Dong Liu. Recurrent dynamic embedding for video object segmentation. In *CVPR*, 2022. 6, 7
- [46] Siyuan Li, Martin Danelljan, Henghui Ding, Thomas E Huang, and Fisher Yu. Tracking every thing in the wild. In *ECCV*. Springer, 2022. 3
- [47] Xiangtai Li, Henghui Ding, Wenwei Zhang, Haobo Yuan, Jiangmiao Pang, Guangliang Cheng, Kai Chen, Ziwei Liu, and Chen Change Loy. Transformer-based visual segmentation: A survey. *arXiv preprint arXiv:2304.09854*, 2023. 3
- [48] Xiangtai Li, Hao He, Yibo Yang, Henghui Ding, Kuiyuan Yang, Guangliang Cheng, Yunhai Tong, and Dacheng Tao. Improving video instance segmentation via temporal pyramid routing. *IEEE TPAMI*, 2022. 3
- [49] Chen Liang, Yu Wu, Tianfei Zhou, Wenguan Wang, Zongxin Yang, Yunchao Wei, and Yi Yang. Rethinking cross-modal interaction from a top-down perspective for referring video object segmentation. *arXiv preprint arXiv:2106.01061*, 2021. 3
- [50] Fanchao Lin, Hongtao Xie, Yan Li, and Yongdong Zhang. Query-memory re-aggregation for weakly-supervised video object segmentation. In *AAAI*, 2021. 3
- [51] Huaijia Lin, Xiaojuan Qi, and Jiaya Jia. Agss-vos: Attention guided single-shot video object segmentation. In *CVPR*, 2019. 2
- [52] Tsung-Yi Lin, Michael Maire, Serge J Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 4
- [53] Zhihui Lin, Tianyu Yang, Maomao Li, Ziyu Wang, Chun Yuan, Wenhao Jiang, and Wei Liu. Swen: Towards real-time video object segmentation with sequential weighted expectation-maximization. In *CVPR*, 2022. 6, 7
- [54] Daizong Liu, Dongdong Yu, Changhu Wang, and Pan Zhou. F2net: Learning to focus on the foreground for unsupervised video object segmentation. In *AAAI*, 2021. 3
- [55] Qiankun Liu, Qi Chu, Bin Liu, and Nenghai Yu. Gsm: Graph similarity model for multi-object tracking. In *IJCAI*, 2020. 3
- [56] Si Liu, Tianrui Hui, Shaofei Huang, Yunchao Wei, Bo Li, and Guanbin Li. Cross-modal progressive comprehension for referring segmentation. *IEEE TPAMI*, 2021. 3
- [57] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. 3
- [58] Xinkai Lu, Wenguan Wang, Martin Danelljan, Tianfei Zhou, Jianbing Shen, and Luc Van Gool. Video object segmentation with episodic graph memory networks. In *ECCV*, 2020. 3
- [59] Xiankai Lu, Wenguan Wang, Chao Ma, Jianbing Shen, Ling Shao, and Fatih Porikli. See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In *CVPR*, 2019. 3
- [60] Xiankai Lu, Wenguan Wang, Jianbing Shen, David Crandall, and Luc Van Gool. Segmenting objects from relational visual data. *IEEE TPAMI*, 2021. 3
- [61] Xiankai Lu, Wenguan Wang, Jianbing Shen, Yu-Wing Tai, David J Crandall, and Steven CH Hoi. Learning video object segmentation from unlabeled videos. In *CVPR*, 2020. 3
- [62] David R Martin, Charless C Fowlkes, and Jitendra Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE TPAMI*, 26(5), 2004. 6
- [63] Bo Miao, Mohammed Bennamoun, Yongsheng Gao, and Ajmal Mian. Region aware video object segmentation with deep motion modeling. *arXiv preprint arXiv:2207.10258*, 2022. 3
- [64] Jiaxu Miao, Xiaohan Wang, Yu Wu, Wei Li, Xu Zhang, Yunchao Wei, and Yi Yang. Large-scale video panoptic segmentation in the wild: A benchmark. In *CVPR*, 2022. 3
- [65] Jiaxu Miao, Yunchao Wei, Yu Wu, Chen Liang, Guangrui Li, and Yi Yang. Vspw: A large-scale dataset for video scene parsing in the wild. In *CVPR*, 2021. 3
- [66] Jiaxu Miao, Yunchao Wei, and Yi Yang. Memory aggregation networks for efficient interactive video object segmentation. In *CVPR*, 2020. 3, 7
- [67] Peter Ochs, Jitendra Malik, and Thomas Brox. Segmentation of moving objects by long term video analysis. *IEEE TPAMI*, 36(6), 2014. 5
- [68] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Fast user-guided video object segmentation by interaction-and-propagation networks. In *CVPR*, 2019. 2, 3, 7
- [69] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *ICCV*, 2019. 2
- [70] Seoung Wug Oh, Joon Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *ICCV*, 2019. 7
- [71] Hyojin Park, Jayeon Yoo, Seohyeong Jeong, Ganesh Venkatesh, and Nojun Kwak. Learning dynamic network using a reuse gate function in semi-supervised video object segmentation. In *CVPR*, 2021. 1
- [72] Federico Perazzi, Anna Khoreva, Rodrigo Benenson, Bernt Schiele, and Alexander Sorkine-Hornung. Learning video object segmentation from static images. In *CVPR*, 2017. 2
- [73] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 1, 2, 5, 6
- [74] Jordi Pont-Tuset, Federico Perazzi, Sergi Caelles, Pablo Arbeláez, Alex Sorkine-Hornung, and Luc Van Gool. The 2017 davis challenge on video object segmentation. *arXiv preprint arXiv:1704.00675*, 2017. 1, 2, 5, 6, 7
- [75] Alessandro Prest, Christian Leistner, Javier Civera, Cordelia Schmid, and Vittorio Ferrari. Learning object class detectors from weakly annotated video. In *CVPR*, 2012. 5
- [76] Jiyang Qi, Yan Gao, Yao Hu, Xinggang Wang, Xiaoyu Liu, Xiang Bai, Serge Belongie, Alan Yuille, Philip HS Torr, and Song Bai. Occluded video instance segmentation: A benchmark. *IJCV*, 130(8), 2022. 3, 4, 5
- [77] Siyuan Qiao, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Vip-deeplab: Learning visual perception with depth-aware video panoptic segmentation. In *CVPR*, 2021. 3
- [78] Sucheng Ren, Wenxi Liu, Yongtuo Liu, Haoxin Chen, Guoqiang Han, and Shengfeng He. Reciprocal transformations for unsupervised video object segmentation. In *CVPR*, 2021. 3
- [79] Seonguk Seo, Joon-Young Lee, and Bohyung Han. Urvos: Unified referring video object segmentation network with a large-scale benchmark. In *ECCV*, 2020. 3
- [80] Guolei Sun, Yun Liu, Henghui Ding, Thomas Probst, and Luc Van Gool. Coarse-to-fine feature mining for video semantic segmentation. In *CVPR*, 2022. 3
- [81] Mingjie Sun, Jimin Xiao, Eng Gee Lim, Bingfeng Zhang, and Yao Zhao. Fast template matching and update for video object tracking and segmentation. In *CVPR*, 2020. 3, 7
- [82] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning video object segmentation with visual memory. In *ICCV*, 2017. 3
- [83] Pavel Tokmakov, Cordelia Schmid, and Karteek Alahari. Learning to segment moving objects. *IJCV*, 127(3), 2019. 3
- [84] Carles Ventura, Miriam Bellver, Andreu Girbau, Amaia Salvador, Ferran Marques, and Xavier Giro-i Nieto. Rvos: End-to-end recurrent network for video object segmentation. In *CVPR*, 2019. 7
- [85] Paul Voigtlaender, Yuning Chai, Florian Schroff, Hartwig Adam, Bastian Leibe, and Liang-Chieh Chen. Feelvos: Fast end-to-end embedding learning for video object segmentation. In *CVPR*, 2019. 2

- [86] Qiang Wang, Li Zhang, Luca Bertinetto, Weiming Hu, and Philip HS Torr. Fast online object tracking and segmentation: A unifying approach. In *CVPR*, 2019. 3, 7
- [87] Wenguan Wang, Xiankai Lu, Jianbing Shen, David J Crandall, and Ling Shao. Zero-shot video object segmentation via attentive graph neural networks. In *ICCV*, 2019. 3, 7
- [88] Wenguan Wang, Hongmei Song, Shuyang Zhao, Jianbing Shen, Sanyuan Zhao, Steven CH Hoi, and Haibin Ling. Learning unsupervised video object segmentation through visual attention. In *CVPR*, 2019. 3
- [89] Wenguan Wang, Tianfei Zhou, Fatih Porikli, David Crandall, and Luc Van Gool. A survey on deep learning technique for video segmentation. *arXiv preprint arXiv:2107.01153*, 2021. 2
- [90] Xinlong Wang, Tete Xiao, Yuning Jiang, Shuai Shao, Jian Sun, and Chunhua Shen. Repulsion loss: Detecting pedestrians in a crowd. In *CVPR*, 2018. 3
- [91] Ziqin Wang, Jun Xu, Li Liu, Fan Zhu, and Ling Shao. Ranet: Ranking attention network for fast video object segmentation. In *ICCV*, 2019. 2
- [92] Sanghyun Woo, Dahun Kim, Joon-Young Lee, and In So Kweon. Learning to associate every segment for video panoptic segmentation. In *CVPR*, 2021. 3
- [93] Dongming Wu, Xingping Dong, Ling Shao, and Jianbing Shen. Multi-level representation learning with semantic alignment for referring video object segmentation. In *CVPR*, 2022. 3
- [94] Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, and Seon Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *CVPR*, 2018. 2
- [95] Huaxin Xiao, Jiashi Feng, Guosheng Lin, Yu Liu, and Maojun Zhang. Monet: Deep motion exploitation for video object segmentation. In *CVPR*, 2018. 2
- [96] Jiarui Xu, Yue Cao, Zheng Zhang, and Han Hu. Spatial-temporal relation networks for multi-object tracking. In *ICCV*, 2019. 3
- [97] Ning Xu, Linjie Yang, Yuchen Fan, Jianchao Yang, Dingcheng Yue, Yuchen Liang, Brian Price, Scott Cohen, and Thomas Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *ECCV*, 2018. 1, 2, 5, 6, 7
- [98] Shuangjie Xu, Daizong Liu, Linchao Bao, Wei Liu, and Pan Zhou. Mhp-vos: Multiple hypotheses propagation for video object segmentation. In *CVPR*, 2019. 2
- [99] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *ICCV*, 2019. 3
- [100] Yanchao Yang, Brian Lai, and Stefano Soatto. Dystab: Unsupervised object segmentation via dynamic-static bootstrapping. In *CVPR*, 2021. 3
- [101] Zhao Yang, Qiang Wang, Luca Bertinetto, Weiming Hu, Song Bai, and Philip HS Torr. Anchor diffusion for unsupervised video object segmentation. In *ICCV*, 2019. 3
- [102] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *ECCV*, 2020. 2
- [103] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *NeurIPS*, volume 34. Curran Associates, Inc., 2021. 6, 7
- [104] Zongxin Yang and Yi Yang. Decoupling features in hierarchical propagation for video object segmentation. In *NeurIPS*, 2022. 2, 6, 7
- [105] Linwei Ye, Mrigank Rochan, Zhi Liu, Xiaoqin Zhang, and Yang Wang. Referring segmentation in images and videos with cross-modal self-attention network. *IEEE TPAMI*, 2021. 3
- [106] Zhaoyuan Yin, Jia Zheng, Weixin Luo, Shenhan Qian, Hanling Zhang, and Shenghua Gao. Learning to recommend frame for interactive video object segmentation in the wild. In *CVPR*, 2021. 3
- [107] Jae Shin Yoon, François Rameau, Jun-Sik Kim, Seokju Lee, Seunghak Shin, and In So Kweon. Pixel-level matching for video object segmentation using convolutional neural networks. In *ICCV*, 2017. 2
- [108] Xiaohang Zhan, Xingang Pan, Bo Dai, Ziwei Liu, Dahua Lin, and Chen Change Loy. Self-supervised scene de-occlusion. In *CVPR*, 2020. 3
- [109] Lu Zhang, Zhe Lin, Jianming Zhang, Huchuan Lu, and You He. Fast video object segmentation via dynamic targeting network. In *ICCV*, 2019. 2
- [110] Shifeng Zhang, Longyin Wen, Xiao Bian, Zhen Lei, and Stan Z Li. Occlusion-aware r-cnn: Detecting pedestrians in a crowd. In *ECCV*, 2018. 3
- [111] Yizhuo Zhang, Zhirong Wu, Houwen Peng, and Stephen Lin. A transductive approach for video object segmentation. In *CVPR*, 2020. 2
- [112] Tianfei Zhou, Shunzhou Wang, Yi Zhou, Yazhou Yao, Jianwu Li, and Ling Shao. Motion-attentive transition for zero-shot video object segmentation. In *AAAI*, 2020. 3
- [113] Ji Zhu, Hua Yang, Nian Liu, Minyoung Kim, Wenjun Zhang, and Ming-Hsuan Yang. Online multi-object tracking with dual matching attention networks. In *ECCV*, 2018. 3