

LVOS: A Benchmark for Long-term Video Object Segmentation

Lingyi Hong¹, Wenchao Chen¹, Zhongying Liu¹, Wei Zhang^{1,*},
 Pinxue Guo², Zhaoyu Chen², Wenqiang Zhang^{1,2,*}

¹Shanghai Key Laboratory of Intelligent Information Processing,
 School of Computer Science, Fudan University

²Academy for Engineering and Technology, Fudan University
 lyhong22@m.fudan.edu.cn {weizh, wqzhang}@fudan.edu.cn




Figure 1: Example sequences of our Long-term Video Object Segmentation (LVOS).

Abstract

Existing video object segmentation (VOS) benchmarks focus on short-term videos which just last about 3-5 seconds and where objects are visible most of the time. These videos are poorly representative of practical applications, and the absence of long-term datasets restricts further investigation of VOS on the application in realistic scenarios. So, in this paper, we present a new benchmark dataset named **LVOS**, which consists of 220 videos with a total duration of 421 minutes. To the best of our knowledge, **LVOS** is the first densely annotated long-term VOS dataset. The videos in our **LVOS** last 1.59 minutes on average, which is 20 times longer than videos in existing VOS datasets. Each video includes various attributes, especially challenges deriving from the wild, such as long-term reappearing and cross-temporal similar objects. Based on **LVOS**, we assess

existing video object segmentation algorithms and propose a **Diverse Dynamic Memory network (DDMemory)** that consists of three complementary memory banks to exploit temporal information adequately. The experimental results demonstrate the strength and weaknesses of prior methods, pointing promising directions for further study. Data and code are available at <https://lingyihongfd.github.io/lvos.github.io/>.

1. Introduction

Given a specific object mask at the first frame, video object segmentation (VOS) aims to highlight target in a video. VOS plays a significant role in video understanding and has many potential downstream applications, such as video editing [45], augmented reality [43], robotics [11, 14], self-driving cars [78, 54, 55]. For most practical applications, objects may experience frequent disappearing and videos

*Corresponding author.

Dataset	Videos	Mean Frames	Total Frames	Mean Duration	Total Duration	Frame Rate	Object Classes	Objects	Annotations	Annotations Type
FBMS [44]	59	235	13,860	0.13	7.7	30	16	139	1,465	M
DAVIS [50]	90	69	6,298	0.04	5.17	24	-	205	13,543	M
YouTube-VOS [72]	3,252	27	107,181	0.06	217.2	6	78	<u>6,048</u>	133,886	M
YouTube-VIS [73]	<u>2,883</u>	28	78,000	0.06	216.7	6	<u>40</u>	<u>4,883</u>	~131,000	M
OVIS [52]	901	90	~68,650	0.21	190.7	6	25	5,223	~296,000	M
UVO [67]	1,200	28	~108,000	0.05	<u>511</u>	30	-	14,748	<u>~1,327,000</u>	M
VOT-ST 2021 [28]	60	324	19,447	0.18	10.8	30	-	60	17,248	M
VOT-LT 2019 [29]	50	4,305	215,298	2.39	119	30	-	50	215,298	B
UAV20L [42]	20	<u>2,934</u>	~59,000	<u>1.63</u>	32.6	30	5	20	~59,000	B
LaSOT [15]	<u>1,400</u>	<u>2,506</u>	~ 3,520,000	1.39	1,950	30	<u>70</u>	1,400	~ 3,520,000	B
YouTube-VIS 2022 Long [72]	121	75	9,014	0.8	100	1.5	-	-	-	N
YouTube-VOS 2022 Long [73]	116	67	7,873	0.74	87	1.5	-	116	-	N
Long-time Video [33]	3	2,470	7,411	1.3	4	30	-	3	60	M
LVOS	220	574	<u>126,280</u>	<u>1.59</u>	<u>421</u>	6	27	282	156,432	M

Table 1: Comparison of LVOS with the most popular video segmentation and tracking benchmarks. The top part is existing short-term video datasets and the bottom part is long-term video datasets. Duration denotes the total duration (in minutes) of the annotated videos. Annotations type means the type of groundtruth annotations. **M** and **B** denote mask and box annotations. **N** means that the groundtruth annotations are unavailable. The largest value is in bold, and the second and third largest values are underlined.

always last more than 1 minute. It is crucial for VOS model to precisely re-detect and segment target objects in videos of arbitrary length.

However, existing VOS models are specifically designed for short-term situation, which struggle to tackle unforeseen challenges in long-term videos. They are vulnerable to long-term disappearance and error accumulation over time [63, 75, 76]. [46, 8, 20, 70, 57] may suffer from the poor efficiency and out-of-memory crash due to the ever-expanding memory bank, especially in a long video. However, the lack of the densely annotated long-term VOS datasets restricts the development of VOS in practice. To date, almost all VOS benchmark datasets, such as DAVIS [50] and YouTube-VOS [72], just focus on short-term videos, which are a poor reflection of practitioners’ demands. The average video length is less than 6 seconds and target objects are always visible, while the average duration is much more longer (*i.e.*, 1-2 minutes) and target objects disappear and reappear frequently in real-world scenarios.

To this end, we propose the first **long-term** video object segmentation benchmark dataset, named **Long-term Video Object Segmentation (LVOS)**. LVOS contains 220 videos with an average duration of 1.59 minutes. The emphasized properties of LVOS are summarised as follows. (1) **Long-term**. Videos in LVOS last 1.59 minutes on average (*vs* 6 seconds in short-term videos), which is much closer to real applications (Table 1). These videos cover multiple challenges, especially attributes specific in long-term videos such as frequent reappearance and long-term similar object confusion. Figure 1 shows some sample videos. (2) **Dense and high-quality annotations**. All frames in LVOS are manually and precisely annotated at 6 FPS. To annotate

the target object accurately and efficiently, we build a semi-automatic annotation pipeline. There are 156K annotated objects in LVOS, about 18% times more annotations than the largest VOS dataset [72]. (3) **Comprehensive labeling**. Videos in LVOS feature 27 categories to represent the daily scenarios. Among the 27 categories, there are 7 unseen categories to better assess the generalization ability of models.

Extensive experiments on LVOS are conducted to assess existing VOS models. To capture the different temporal context in long-term videos adequately, we propose **Diverse Dynamic Memory (DDMemory)**. DDMemory consists of three complementary memory banks: reference memory, global memory, and local memory to encode historical information into fixed-size features. Due to the diverse and dynamic memory mechanism, DDMemory can handle videos of any length with constant memory cost and high efficiency. Oracle experiment demonstrates that error accumulation and complex motion are the primary cause for the unsatisfactory performance.

Our contributions are summarized as follows: (1) We construct a new long-term, densely and high-quality annotated, and comprehensively labeled VOS dataset named LVOS with 220 videos whose average duration is 1.59 minutes. (2) We propose the DDMemory to handle long-term videos better. (3) We assess existing VOS models and DDMemory on LVOS and analyze the cause of errors to discover cues for the development of robust VOS methods.

2. Related work

Semi-supervised Video Object Segmentation. The key to semi-supervised VOS lies in the construction and utiliza-

tion of feature memory. [3, 39, 64, 69, 53, 41, 47, 2] employ online learning approaches to finetune pretrained networks at test time on the first frame and groundtruth, which require a large amount of time. [23, 5, 60, 1] employ the manually annotated first frame to guide the segmentation of the rest frames, while [49, 77, 9, 24, 4, 45, 71, 21, 62, 22, 31, 66, 27] use the already segmented previous frame as a reference to propagate mask frame-to-frame. [63, 75, 74, 68, 25] combine the first frame and previous frame as feature memory, and the temporal context is limited. To address the limitation, [46, 19, 7, 57, 20, 70, 65, 37, 8, 58, 40, 76, 48, 36] develop a feature memory bank to store all historical frames, while the ever-expanding memory bank may encounter an out-of-memory crash when handling long-term videos.

The adaptive feature bank is developed in [33] to dynamically manage key features of objects by using exponential moving averages. [32] introduces a global context module to summarize target information. [30] proposes a recurrent dynamic embedding (RDE) to construct the memory bank of fixed size in a recurrent manner. Xmem [6] develops three kinds of memory banks and connects them to segment the current frame. By compressing the memory bank, these methods achieve constant memory cost, but they still struggle with losing track after a long period of disappearing in long-term videos. Thus, we propose DDMemory to efficiently exploit the temporal context and maintain the fixed-sized memory cost, which is robust to various challenges in long-term videos.

Short-term Video Object Segmentation Dataset. The existing video object segmentation benchmark datasets are all short-term video datasets. FBMS [44] has 59 sequences with 13,860 frames in total and is divided into 29 and 30 videos as training and evaluation set, respectively. DAVIS 2017 [51] is a popular benchmark dataset with 60 and 30 videos for train and validation sets. There are 6,298 frames in total. DAVIS 2017 provides pixel-level and high-quality annotations for each frame. YouTube-VOS [72], as a large-scale dataset, has 3,252 sequences with precise annotations at 6 FPS. YouTube-VOS includes 78 diverse categories. All these benchmarks are short-term video datasets, where the average duration of videos is about 3-6 seconds. Although some VOS methods [33, 32, 30, 6] claim to scale well to long-term videos, they do not conduct quantitative experiments on a long-term VOS benchmark because of the lack of such a dataset. Videos in LVOS are long-term with an average duration of about 1.59 minutes, which is relevant to scenarios for actual applications.

Long-term Tracking Dataset. There are several benchmark datasets specific to long-term tracking. UAV20L [42] is a small scale dataset with only 20 long videos. OxUvA [61] consists of 366 sequences, but each video is sparse-annotated every 30 frames. LaSOT [15] is the first large-scale and densely annotated long-term tracking dataset,

which provides 1,400 videos totaling 3.52M frames. The average length of sequences in LaSOT is 2,512 frames at 30 FPS. Each frame is manually annotated with a bounding box. These long-term tracking datasets demonstrate the significance of long-term tasks. However, these datasets only provide box-level annotations, and pixel-level annotations are unavailable, which is more crucial for fine-grained study. Long-time Video [33] is a dataset of 3 long videos with 2,470 frames on average per video, where only 20 frames are uniformly annotated for each video. Note that YouTube-VOS 2022 Long and YouTube-VIS 2022 Long proposed at the CVPR 2022 workshop also include long-term videos, while no groundtruth is available. There lacks a comprehensive long-term VOS dataset which is with training data and available all the time. LVOS focuses on long-term video object segmentation with 220 videos in total including both training, valid, and test sets. Each frame in LVOS is manually and precisely annotated. We propose LVOS to promote the development of robust VOS models and provide a more suitable evaluation benchmark in practical application.

3. LVOS: Long-term Video Object Segmentation Benchmark Dataset

3.1. Dataset Construction

Dataset Design. To make up for the lack of a dedicated dataset, LVOS aims to provide the community with a novel and dedicated VOS dataset for training and evaluating robust VOS models. We adhere to the three principles listed below to construct LVOS.

1) **Long-term VOS.** Compared with current VOS datasets [50, 72] where the average length of each video is only 3-6 seconds, we ensure videos in LVOS last about 1.59 minutes (*i.e.*, 574 frames at 6 FPS), about 20 times longer than short-term videos, which is much closer to the real application.

2) **Dense and high-quality annotation.** The time-consuming mask annotation processing severely constrains the duration and scale of current VOS datasets. High-quality and densely annotated masks are essential for training robust VOS models and assessing their performance in practical applications. So, all frames in LVOS are manually and precisely annotated by leveraging the semi-automatic annotation pipeline proposed in Sec 3.2.

3) **Comprehensive labeling.** We design a set of categories that are relevant to daily life and have 5 parent classes and 27 subclasses. It is worth noting that the 27 categories are not limited to COCO dataset [35] and also include some categories not present in the COCO dataset, such as frisbee. Among the 27 categories, there are 7 unseen categories to better assess the generalization ability of models.

Data Collection. To construct LVOS, we carefully se-

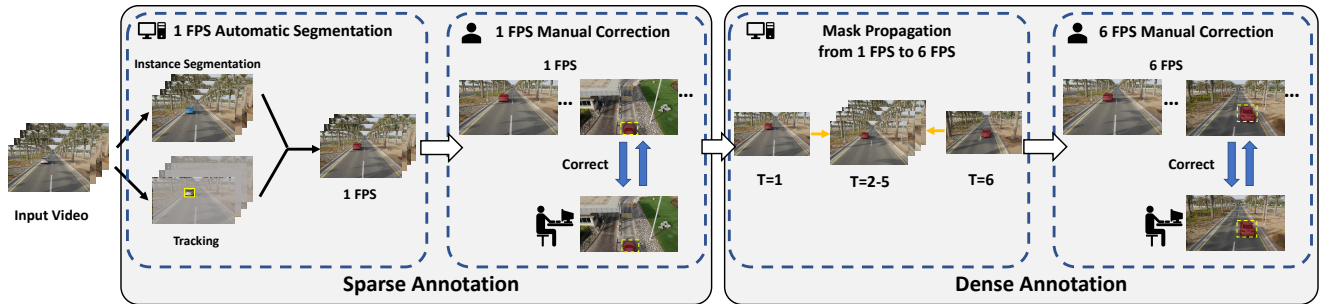


Figure 2: Annotation Pipeline, including four steps. Step 1: 1 FPS Automatic Segmentation. We utilize instance segmentation [26] and tracking [12] models to get the mask of target object at 1 FPS automatically. Step 2: 1 FPS Manual Correction. We refine masks obtained in Step 1 manually. Step 3: Mask Propagation from 1 FPS to 6 FPS. We propagate masks from 1 FPS to 6 FPS by using a VOS model[76]. Step 4: 6 FPS Manual Correction. We correct the masks obtained in Step 3 manually.

lect a set of categories comprising 5 parent classes and 27 subclasses from the videos in existing long-term tracking datasets such as VOT-LT [29] and LaSOT [15]. These datasets, containing more than 1,800 videos in total, have been customized for long-term tracking are similar to VOS in terms of tracking tasks. As such, the videos in VOT-LT and LaSOT are suitable for the long-term VOS task. After selecting a set of object categories, we screen about 600 videos with a resolution of 720P as candidate videos. 220 videos are selected to make up for LVOS after a comprehensive consideration of video quality. Because videos in VOT-LT and LaSOT have been processed specially for tracking task, such as removing irrelevant content, we did not apply any treatment to these videos. VOT-LT and LaSOT are single-object datasets, where only one target is annotated in a video, while our LVOS is multiple-object, where there may be several target objects in a video. For target selection, we may either follow the target object in VOT-LT and LaSOT, or select different objects as targets.

3.2. Semi-Automatic Annotation Pipeline

The exhaustion of the mask annotation process limits the scale of VOS datasets to a large extent. We propose a novel semi-automatic annotation pipeline to annotate frames efficiently. Concretely, the pipeline can be divided into four steps, as shown in Figure 2.

Step 1: 1 FPS Automatic Segmentation. Firstly, transfiner [26] is adopted to generate the pixel-wise segmentation of each object in the frames at 1 FPS. Then we manually mark the bounding box of the target objects when they first appeared and utilize MixFormer [12] to propagate the box from the first frame to all subsequent frames. Based on the pixel-wise segmentation and the bounding box of each frame, we obtain the masks of target objects at 1 FPS.

Step 2: 1 FPS Manual Correction. Tracking errors, segmentation defects, and other prediction mistakes may lead to inaccuracy or the absence of the target object mask


in some frames. Thus, we use EISeg [17] (An Efficient Interactive Segmentation Tool based on PaddlePaddle [38]) to refine masks. On average, about 30% frames need to be corrected.

Step 3: Mask Propagation. By using a VOS model (*i.e.*, AOT [76]) to propagate the annotation masks at the frame rate of 1 FPS obtained in Step 2 to their adjacent unlabeled frames, we extend the masks from 1 FPS to 6 FPS automatically.


Step 4: 6 FPS Manual Correction. Because of flaws in masks segmented by VOS model, we correct every frame artificially until the results are satisfactory. In this step, about 40% of frames require further refinement.

Time and Quality Analysis. To examine the annotation quality, we randomly choose 100 videos from HQYouTubeVIS [26] training set and relabel them using our semi-automatic annotation pipeline. Then we compare the results with the groundtruth, and the average IoU score is 0.93. The score shows that the annotation results obtained by our pipeline are largely consistent with groundtruth and also proves the validity of our pipeline. Moreover, we ask annotators to record the total time overheads. It takes 60 minutes for one annotator to label an entire long-term video (500 frames at 6 FPS) on average by utilizing our pipeline, while a skilled annotator spends 1500 minutes labeling the same video (3 minutes for one frame). The pipeline significantly reduces the labeling cost when ensuring annotation quality.


Discussion. Similar semi-automatic annotation pipelines have also been proposed in UVO [67] and EPIC-VISOR [13], which can be separated into two parts: manual annotation of videos sparsely and propagating masks. The mask propagation part is similar to the Step 3 and Step 4 in our pipeline. For sparsely annotating video, pipelines in UVO and EPIC-VISOR require manual annotation, while our pipeline also adopts models, which is faster. To obtain the pixel-wise segmentation of each object in Step 1, we



(a) Attributes distribution of each sequence in LVOS.



(b) Main Mutual dependencies among attributes.



(c) Distribution comparison with DAVIS 2017 [49].

Figure 3: Attributes distribution in LVOS and comparison with DAVIS2017 [49]. In sub-figure (b), the link indicates the high likelihood that more than one attributes will appear in a sequence. Best viewed in color.

employ Transfiner trained on COCO. Classes in LVOS are not restricted to classes in COCO or LVIS [16]. Transfiner can detect and segment the segment objects whose classes are not among the 80 categories in COCO, although the category classification is incorrect. Because VOS is class-agnostic, we simply ignore the category and utilize the mask segmentaion, where classification errors have no influence on the annotation pipeline.




Figure 4: The histogram of instance masks for five parent classes and sub-classes. Objects are sorted by frequency. The entire category set roughly covers diverse objects and motions that occur in everyday scenarios.

3.3. Dataset Statics

Video-level Statics. The video-level information of LVOS is shown in Table 1. We collect 220 videos in LVOS, whose average duration is 1.59 minutes with 574 frames on average at 6 FPS (vs 3-6 seconds in short-term dataset). There are 126,280 frames and 156,432 annotations in total, which is larger than the sum of the other data sets [50, 72, 33, 44]. Videos are categorized into 5 parent classes and 27 sub-classes. The detail and distribution of instance masks can be seen in Figure 4. Notably, there are 7 categories which are not present in the training set. We sample frames with a frame rate of 6 FPS. By keeping the distribution of subsets and video length, videos are divided into 120 training, 50 validation, and 50 testing. Annotations of the training and validation sets are publicly released for the

Attribute	Definition
BC	<i>Background Clutter.</i> The appearances of background and target object are similar.
DEF	<i>Deformation.</i> Target appearance deform complexly.
MB	<i>Motion Blur.</i> Boundaries of target object is blurred because of camera or object fast motion.
FM	<i>Fast Motion.</i> The per-frame motion of target is larger than 20 pixels, computed as the centroids Euclidean distance.
LR	<i>Low Resolution.</i> The average ratio between target box area and image area is smaller than 0.1 .
OCC	<i>Occlusion.</i> The target is partially or fully occluded in the video.
OV	<i>Out-of-view</i> The target leaves the video frame completely.
SV	<i>Scale Variation</i> The ratio of any pair of bounding-box is outside of range [0.5,2.0].
DB	<i>Dynamic Background</i> Background regions undergoes deformation.
SC	<i>Shape Complexity</i> Boundaries of target object is complex.
AC	<i>Appearance Change</i> Significant appearance change, due to rotations and illumination changes .
LRA	<i>Long-term Reappearance</i> Target object reappears after disappearing for at least 100 frames.
CTC	<i>Cross-temporal Confusion</i> There are multiple different objects that are similar to target object but do not appear at the same time.

Table 2: Definitions of video attributes in LVOS. We extend and modify the short-term video challenges defined in [50] (top), which is exanded with a complementary set of long-term video attributes (bottom).

development of VOS methods, while annotations of the testing set are kept private for competition use.

Attributes. For a further and comprehensive analysis of VOS approaches, it is of great significance to identify video attributes. We label each sequence with 13 challenges, which are defined in Table 2. These attributes include short-term video challenges, which are extended from DAVIS [50], and are expanded with a complementary set of challenges specific to long-term videos. It is important to note that these attributes are not exclusive, and a video can contain multiple challenges. The distribution of each video and the main mutual dependencies are shown in Figure 3a and 3b. Scale variation (SV), occlusion (OCC), low resolution (LR), and fast motion (FM) are the most common challenges in LVOS. The comparison of attributes distribution between LVOS and DAVIS is demonstrated in Figure

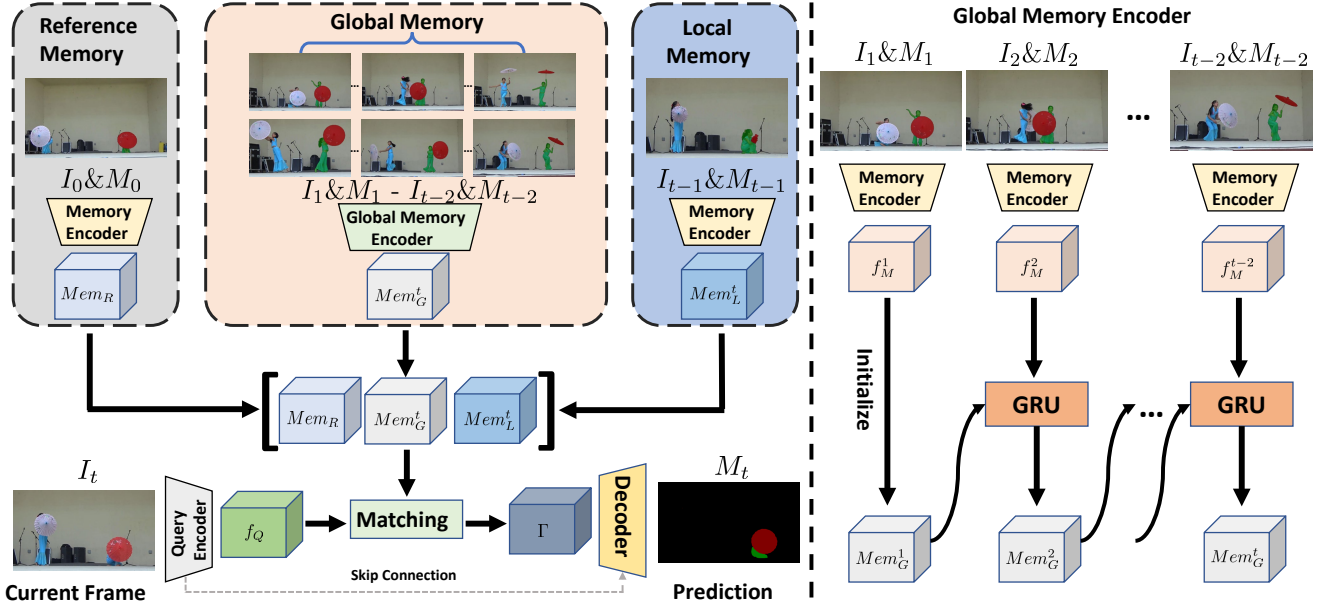


Figure 5: Model Overview. DDMemory consists of three memory banks: reference memory, global memory, and local memory. The global memory encoder is responsible for encoding the historical information into a fixed-size feature.

3c. We observe the difference in challenges between short-term and long-term videos. Because of the longer length of videos, the object motion and background changes are much more complex and varied, which is not obvious in short-term videos. The variation in the distribution of attributes places differences and higher demanding requirements on the design of VOS models.

4. Method

Although many VOS methods [33, 32, 30, 6, 36, 65] attempt to compress memory bank to achieve the trade-off of efficiency and accuracy, these models still struggle with the loss of critical temporal information in long videos. We propose a novel VOS method, **Diverse Dynamic Memory (DDMemory)**, specifically designed for long-term VOS task, which includes diverse memory banks with constant size: *reference memory*, *global memory*, and *local memory*. Diverse memory banks can compress the global temporal memory into three memory features with rich temporal information and maintain low GPU memory usage while achieving high performance.

4.1. Method Overview

Let I_1 and M_1 denote the first frame and its groundtruth mask, and I_t denotes the current frame whose segmentation mask M_t is to be predicted. $\{I_n, M_n\}_{n=2}^{t-1}$ denote the intermediate frames from the second frame to the previous frame, along with their estimated masks. Following [76], for the query frame I_t of size $H \times W$, query encoder

takes the image as input to extract visual features $f_Q \in \mathbb{R}^{C \times \frac{H}{16} \times \frac{W}{16}}$, where C is the channel dimension. Each intermediate frame and corresponding mask are fed into memory encoder to obtain memory feature $f_M^i \in \mathbb{R}^{C \times \frac{H}{16} \times \frac{W}{16}}$ ($i = 0, 1, \dots, n-1$). The memory feature of first and previous frame, f_M^0 and f_M^{t-1} , act as the reference memory Mem_R and local memory Mem_L^t , respectively. For $\{f_M^i\}_{i=2}^{t-2}$, we feed them into global memory encoder (Sec 4.2) to generate global memory feature $Mem_G^t \in \mathbb{R}^{C \times \frac{H}{16} \times \frac{W}{16}}$. The three memory features (Mem_R , Mem_G , Mem_L) are fed into the matching module together with the query feature f_Q to get the matching output Γ . Finally, Γ and the low-level features from the decoder are used to generate the segmentation mask M_t .

4.2. Global Memory Encoder

To use fixed-size features to construct global historical information with as little loss as possible, we adopt a recurrent manner to build a global feature memory bank. Specifically, at time t , we utilize a Gated Recurrent Unit (GRU) [10, 59] as global memory encoder to compress $\{f_M^i\}_{i=1}^{t-2}$ into global memory Mem_G^t .

For the generation and updating of Mem_G^t , we first initialize Mem_G^1 as f_M^1 , and then utilize GRU to propagate it as illustrated in Figure 5. The process is defined as:

$$Mem_G^t = GRU(Mem_G^{t-1}, f_M^{t-2}) \quad (1)$$

where GRU denotes a GRU module. By recurrent refreshment, we achieve the ability to encode global information

Method	FB	FPS	Mem	Before			Finetune		
				$\mathcal{J}\&\mathcal{F}\uparrow$	$\mathcal{J}\uparrow$	$\mathcal{F}\uparrow$	$\mathcal{J}\&\mathcal{F}\uparrow$	$\mathcal{J}\uparrow$	$\mathcal{F}\uparrow$
LWL [2]	OD	14.1	1.2	54.1	49.6	58.6	56.4	51.8	60.9
CFBI [75]	F+P	5.2	3.82	50.0	45.0	55.1	51.5	46.2	56.7
AOT-B [76]	F+P	31.9	0.96	56.9	51.8	61.9	58.9	53.5	64.2
AOT-L [76]	A	20.8	1.32	59.4	53.6	65.2	60.9	55.1	66.8
STCN [8]	A	22.1	0.92	45.8	41.1	50.5	48.9	43.9	54.0
AFB-URR [33]	C	4.8	2.89	34.8	31.3	38.2	36.2	33.1	39.3
RDE [30]	C	22.2	1.0	52.9	47.7	58.1	53.7	48.3	59.2
XMem [6]	C	28.6	1.4	50.0	45.5	54.4	52.9	48.1	57.7
DDMemory	C	30.3	0.88	60.7	55.0	66.3	61.9	56.3	67.4

(a) Results on validation set.

Method	FB	FPS	Mem	Before			Finetune		
				$\mathcal{J}\&\mathcal{F}\uparrow$	$\mathcal{J}\uparrow$	$\mathcal{F}\uparrow$	$\mathcal{J}\&\mathcal{F}\uparrow$	$\mathcal{J}\uparrow$	$\mathcal{F}\uparrow$
LWL [2]	OD	14.1	1.2	50.7	46.5	54.8	50.8	46.4	55.2
CFBI [75]	F+P	5.2	3.82	44.8	40.5	49.0	44.8	40.2	49.4
AOT-B [76]	F+P	31.9	0.96	54.4	49.3	59.4	54.5	49.2	59.8
AOT-L [76]	A	20.8	1.32	54.1	48.7	59.5	54.7	49.2	60.2
STCN [8]	A	22.1	0.92	45.8	41.6	50.0	48.3	44.0	52.5
AFB-URR [33]	C	4.8	2.89	39.9	36.2	43.6	40.8	37.5	44.1
RDE [30]	C	22.2	1.0	49.0	44.4	53.5	50.2	45.7	54.6
XMem [6]	C	28.6	1.4	49.5	45.2	53.7	50.9	46.5	55.3
DDMemory	C	30.3	0.88	55.0	49.9	60.2	55.7	50.3	61.2

(b) Results on test set.

Table 3: Comparisons with state-of-the art models on LVOS validation and test sets. FB denotes the kind of feature bank. **OD**, **F+P**, **A**, and **C** denote online adaption, first and previous frame, all frames, and compressed memory bank respectively. We re-time these models on our hardware (a V100 GPU) for a fair comparison. Mem denotes the maximum GPU memory usage (in GB). Before and Finetune denotes the results just trained on short-term video datasets and finetuned on LVOS training set.

into a fixed size features and discard redundant and noisy information.

4.3. Diverse Dynamic Memory

Our diverse dynamic memory consists of three types of different temporal scale memory banks: *reference memory*, *global memory*, and *local memory*. Due to the fixed-size memory features, the memory cost remains constant no matter how long the video is. The first frame and its groundtruth mask are stored in reference memory, which is responsible for the recovery after disappearance or occlusion. Global memory leverages a recurrent manner to store historical information effectively, which is crucial for the segmentation of long-term videos. Local memory is updated every frame and provides the location and shape cues. The complementary memory banks achieve the storage of rich temporal information and the removal of noisy features. Thanks to the diverse and dynamic memory banks,

DDMemory achieves promising performance with a constant memory cost and high speed in long videos.

4.4. Model Details

We employ MobileNet-V2 [56] as backbone. Memory encoder and query encoder share the same weight, following [76]. LSTT module proposed in [76] is adopted as matching module, and the LSTT layer number is 3. We use FPN [34] as decoder. For Global Memory Encoder, when segmenting current frame I_t , we just save the latest global memory Mem_G^t for the sake of efficiency. After the segmentation of I_t , we update Mem_G^t by utilizing the GRU module to obtain the global memory Mem_G^{t+1} for next frame I_{t+1} . We don't need to restore all intermediate frames and repeat the calculation of global memory. We only need to store a fixed-size global memory and conduct a simple updating per frame.

5. Experiments

5.1. Experiment Setup

Experiment Settings. We evaluate our DDMemory and other existing VOS methods, including CFBI [75], LWL [2], STCN [8], AOT [76], RDE [30], XMem [6] on LVOS validation and test set. We restrict the memory length to 6 when evaluating approaches with memory bank, such as STCN [8], AOT [76], and XMem [6]. For a fair comparison, all the videos are down-sampled to 480p resolution. We also finetune these models on the training set of LVOS for two epochs, with a learning rate of 5×10^{-4} , and reevaluate their performance.

Evaluation Metrics. We adopt the two commonly used evaluation metrics, region similarity \mathcal{J} and contour accuracy \mathcal{F} as metrics, following DAVIS [50, 49] and YouTube-VOS [72]. And we calculate the average their mean value as the final score.

5.2. Benchmark Results

Quantitative Results. As shown in Table 3, DDMemory outperforms other models with different feature banks on both validation (60.7 $\mathcal{J}\&\mathcal{F}$) and test sets (55.0 $\mathcal{J}\&\mathcal{F}$) when maintaining a real-time speed (30.3 FPS) at the lowest GPU memory cost (0.88G). Thanks to the diverse and dynamic memory banks, DDMemory exploits richer temporal contexts and are robust to complex challenges in long-term videos. We finetune these models and DDMemory on the LVOS training sets and assess their performances again. Although each model’s performance has essentially improved to some degree, DDMemory continues to outperform all the competitors on both validation (61.9 $\mathcal{J}\&\mathcal{F}$) and test (55.7 $\mathcal{J}\&\mathcal{F}$) sets.

Qualitative Results. We present the segmentation results in comparison with AOT-L and XMem in Figure 6. As shown in the top section, AOT and XMem lose track or confuse similar objects, while DDMemory can re-detect the target after out-of-view and handle multiple similar objects successfully. The bottom section displays some failure cases. Although DDMemory achieves promising performance, DDMemory is still not robust enough for the complex motion. More details are in supplementary material.

R	G	L	FPS	GPU	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
✓			57.4	0.52	44.2	39.0	49.4
	✓		55.2	0.62	42.7	37.4	48.0
		✓	43.5	0.68	18.3	17.1	19.6
✓	✓		46.7	0.78	47.8	42.4	53.3
✓		✓	35.6	0.82	57.9	53.0	62.8
	✓	✓	35.1	0.76	54.9	51.1	58.7
✓	✓	✓	30.3	0.88	61.9	56.3	67.4

Table 4: Ablation study on LVOS validation set. R, G, and L denote Mem_R , Mem_G , and Mem_L , respectively.

Ablation Study. We perform an ablation study on LVOS

validation set and analyze the contribution of each memory bank to the segmentation result in Table 4. We experiment with various combinations of three memory banks. Results show the role of each component. The reference memory Mem_R is responsible for the re-detection after occlusion or out-of-view and is sensitive to large appearance changes. The global memory Mem_G encodes the long-term temporal information. The local memory Mem_L provides location cues and appearance prior. For long-term VOS, all three memory banks are crucial and complementary. Please see supplementary material for more details.

Oracle Box	Oracle Mask	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}
		61.9	56.3	67.4
✓		70.2	64.6	75.7
	✓	82.7	76.5	89.0
✓	✓	84.4	77.8	91.1

Table 5: Oracle analysis on LVOS validation set.

Oracle Analysis. To conduct further analysis of object localization and association, we carry out oracle experiments. Results are shown in Table 5. Average performance is improved by 8.3 % when segmentation is provided with an oracle bounding box, proving that segmentation errors result from poor tracking between similar objects. While resolving the segmentation errors, the model achieves a higher score (20.8 % boost). This shows that error accumulation is the primary cause of errors. But even if the correct masks and locations are provided, there is still a large gap between the result (84.4 $\mathcal{J}\&\mathcal{F}$) and groundtruth. The gap demonstrates that complex movements are still very challenging for VOS models. In short, error accumulation is the main cause of unsatisfactory performances, and a robust VOS model must be able to handle the much more complex motion in long-term videos.

5.3. Results on DAVIS Short-term Validation Set

Experiments in Table 3 show the effectiveness of our algorithm on long-term videos. To show the efficacy of DDMemory on short-term videos, we evaluate DDMemory and other models on DAVIS 2017 validation set [49]. The result is shown in Table 6. DDMemory exceeds the majority of models and maintains an efficient speed (28.1 FPS). Despite having higher performance than DDMemory, XMem and STCN employ a stronger backbone ResNet50 [18], while DDMemory only uses MobileNet-V2 [56]. The reason why the improvement resulting from the global temporal information is not very obvious may be that the length of videos is relatively short. More experiment results are in supplementary material.

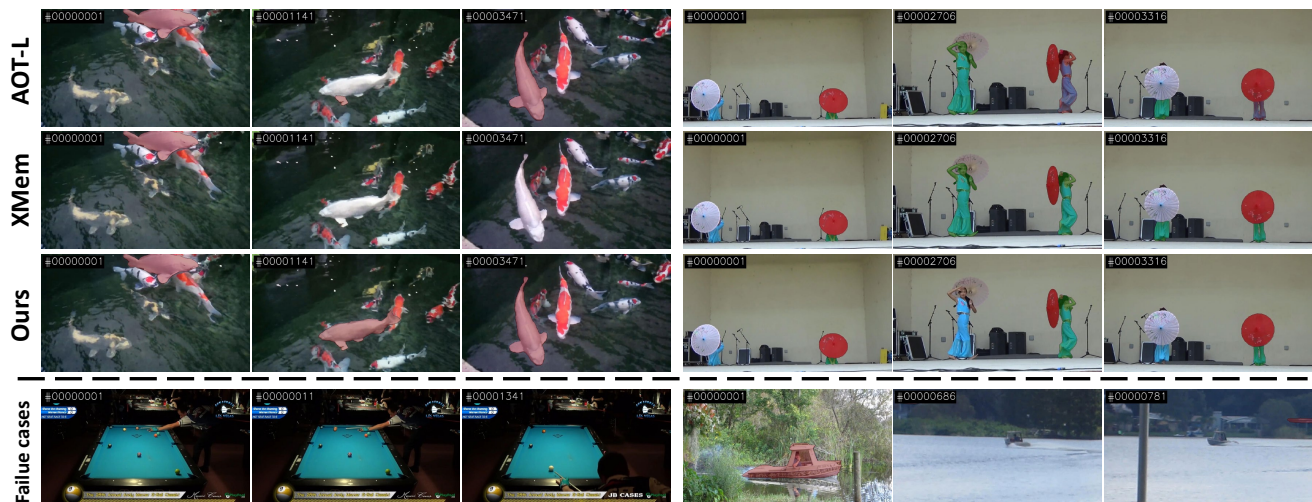


Figure 6: Qualitative results on LVOS validation and test set. We compare the result of DDMemory with XMem[6] and AOT-L[76] in top section, and DDMemory performs better. The bottom row shows failure cases.

Methods	Backbone	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	FPS
CFBI [75]	ResNet101 [18]	81.9	79.1	84.6	5.9
LWL [2]	ResNet50 [18]	81.6	79.1	84.1	13.2
STCN [8]	ResNet50 [18]	<u>85.4</u>	<u>82.2</u>	<u>88.6</u>	20.2
RDE [30]	ResNet50 [18]	84.2	80.8	87.5	27.0
XMem [6]	ResNet50 [18]	86.2	82.9	89.5	22.6
AOT-B [76]	MobileNet-V2 [56]	82.5	79.7	85.2	29.6
AOT-L [76]	MobileNet-V2 [56]	83.8	81.1	86.4	18.7
DDMemory	MobileNet-V2 [56]	84.2	81.3	87.1	<u>28.1</u>

Table 6: Comparisons with state-of-the art models on DAVIS 2017 validation set. Bold and underline denote the best and second-best respectively in each column.

5.4. Attribute-based Evaluation

We report the performance in Table 7 on the validation set characterized by with the most informative attributes. FM (Fast Motion), OCC (Occlusion), OV (Out-of-view), SV (Scale Variation), and AC (Appearance change) are all well known challenges in short-term video segmentation, which also have great influence on performance of long-term task. Furthermore, long-term videos present specific challenges, such as Long-term Reappearance (LRA) and Cross-temporal Confusion (CTC), which have a worse impact. Although these methods achieve promising results on short-term VOS datasets (over 80% $\mathcal{J}\&\mathcal{F}$), they still struggle with complex scenes and frequent reappearance in long-term videos, highlighting the unique value of our LVOS. The ability to recover disappeared object, distinguish target from similar background, detect small object, and model long-term historical information is crucial for robust LVOS.

Attr	AOT-B	AOT-L	XMem	LWL	Ora B	Ora M	Ora B+M
FM	54.3	55.3	46.7	48.2	73.8	82.6	85.5
OCC	50.6	52.1	47.6	50.3	72.5	79.1	83.6
OV	54.4	55.2	53.8	48.6	74.2	79.9	82.8
SV	48.3	50.4	44.7	47.2	66.8	76.6	80.5
AC	53.1	55.7	48.9	52.4	75.7	82.2	84.2
LRA	44.3	45.3	40.7	45.2	63.8	74.6	78.5
CTC	44.5	45.7	45.1	46.1	64.4	75.5	77.7

Table 7: Attribute-based aggregate performance. For each method, we just show \mathcal{J} . Ora B, Ora M, Ora B+M denote oracle box, oracle mask and oracle box + mask in oracle experiments (Table 5), respectively.

6. Conclusion

In this paper, we propose a new long-term video object segmentation dataset, LVOS. Different from existing short-term VOS datasets, the average length of videos in LVOS is 1.59 minutes. More complex motion and longer duration place greater demands on VOS models. We assess existing VOS approaches and propose a novel baseline method DDMemory designed for long-term VOS. Based on the baseline model, we analyze the weakness of prior methods and point promising directions for further study. We hope that LVOS can provide a platform to encourage a comprehensive study on long-term VOS.

7. Acknowledge

This work was supported by National Natural Science Foundation of China (No.62072112), Scientific and Technological Innovation Action Plan of Shanghai Science and Technology Committee (No.22511102202), National Key R&D Program of China (2020AAA0108301).

References

- [1] Linchao Bao, Baoyuan Wu, and Wei Liu. Cnn in mrf: Video object segmentation via inference in a cnn-based higher-order spatio-temporal mrf. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5977–5986, 2018.
- [2] Goutam Bhat, Felix Järemo Lawin, Martin Danelljan, Andreas Robinson, Michael Felsberg, Luc Van Gool, and Radu Timofte. Learning what to learn for video object segmentation. In *European Conference on Computer Vision*, pages 777–794. Springer, 2020.
- [3] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 221–230, 2017.
- [4] Xi Chen, Zuoxin Li, Ye Yuan, Gang Yu, Jianxin Shen, and Donglian Qi. State-aware tracker for real-time video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9384–9393, 2020.
- [5] Yuhua Chen, Jordi Pont-Tuset, Alberto Montes, and Luc Van Gool. Blazingly fast video object segmentation with pixel-wise metric learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1189–1198, 2018.
- [6] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. *arXiv preprint arXiv:2207.07115*, 2022.
- [7] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Modular interactive video object segmentation: Interaction-to-mask, propagation and difference-aware fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5559–5568, 2021.
- [8] Ho Kei Cheng, Yu-Wing Tai, and Chi-Keung Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *Advances in Neural Information Processing Systems*, 34:11781–11794, 2021.
- [9] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7415–7424, 2018.
- [10] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [11] Isaac Cohen and Gerard Medioni. Detecting and tracking moving objects for video surveillance. In *Proceedings. 1999 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Cat. No PR00149)*, volume 2, pages 319–325. IEEE, 1999.
- [12] Yutao Cui, Cheng Jiang, Limin Wang, and Gangshan Wu. Mixformer: End-to-end tracking with iterative mixed attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13608–13618, 2022.
- [13] Ahmad Darkhalil, Dandan Shan, Bin Zhu, Jian Ma, Amlan Kar, Richard Ely Locke Higgins, Sanja Fidler, David Fouhey, and Dima Damen. Epic-kitchens visor benchmark: Video segmentations and object relations. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.
- [14] Adám Erdélyi, Tibor Barát, Patrick Valet, Thomas Winkler, and Bernhard Rinner. Adaptive cartooning for privacy protection in camera networks. In *2014 11th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 44–49. IEEE, 2014.
- [15] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. Lasot: A high-quality benchmark for large-scale single object tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5374–5383, 2019.
- [16] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019.
- [17] Yuying Hao, Yi Liu, Zewu Wu, Lin Han, Yizhou Chen, Guowei Chen, Lutaoyang Chu, Shiyu Tang, Zhiliang Yu, Zeyu Chen, et al. Edgeflow: Achieving practical interactive segmentation with edge-guided flow. *arXiv preprint arXiv:2109.09406*, 2021.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Lingyi Hong, Wei Zhang, Liangyu Chen, Wenqiang Zhang, and Jianping Fan. Adaptive selection of reference frames for video object segmentation. *IEEE Transactions on Image Processing*, 31:1057–1071, 2022.
- [20] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4144–4154, 2021.
- [21] Ping Hu, Gang Wang, Xiangfei Kong, Jason Kuen, and Yap-Peng Tan. Motion-guided cascaded refinement network for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1400–1409, 2018.
- [22] Yuan-Ting Hu, Jia-Bin Huang, and Alexander Schwing. Maskrnn: Instance level video object segmentation. *Advances in neural information processing systems*, 30, 2017.
- [23] Yuan-Ting Hu, Jia-Bin Huang, and Alexander G Schwing. Videomatch: Matching based video object segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 54–70, 2018.
- [24] Won-Dong Jang and Chang-Su Kim. Online video object segmentation via convolutional trident network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5849–5858, 2017.
- [25] Joakim Johnander, Martin Danelljan, Emil Brissman, Fahad Shahbaz Khan, and Michael Felsberg. A generative appearance model for end-to-end video object segmentation.

- In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8953–8962, 2019.
- [26] Lei Ke, Martin Danelljan, Xia Li, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Mask transfiner for high-quality instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4412–4421, 2022.
- [27] Anna Khoreva, Rodrigo Benenson, Eddy Ilg, Thomas Brox, and Bernt Schiele. Lucid data dreaming for object tracking. In *The DAVIS challenge on video object segmentation*, 2017.
- [28] Matej Kristan, Jiří Matas, Aleš Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kämäräinen, Hyung Jin Chang, Martin Danelljan, Luka Cehovin, Alan Lukežič, et al. The ninth visual object tracking vot2021 challenge results. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2711–2738, 2021.
- [29] Matej Kristan, Jiri Matas, Ales Leonardis, Michael Felsberg, Roman Pflugfelder, Joni-Kristian Kamarainen, Luka Cehovin Zajc, Ondrej Drbohlav, Alan Lukežic, Amanda Berg, et al. The seventh visual object tracking vot2019 challenge results. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [30] Mingxing Li, Li Hu, Zhiwei Xiong, Bang Zhang, Pan Pan, and Dong Liu. Recurrent dynamic embedding for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1332–1341, 2022.
- [31] Xiaoxiao Li and Chen Change Loy. Video object segmentation with joint re-identification and attention-aware mask propagation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 90–105, 2018.
- [32] Yu Li, Zhuoran Shen, and Ying Shan. Fast video object segmentation using the global context module. In *European Conference on Computer Vision*, pages 735–750. Springer, 2020.
- [33] Yongqing Liang, Xin Li, Navid Jafari, and Jim Chen. Video object segmentation with adaptive feature bank and uncertain-region refinement. *Advances in Neural Information Processing Systems*, 33:3430–3441, 2020.
- [34] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017.
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [36] Zhihui Lin, Tianyu Yang, Maomao Li, Ziyu Wang, Chun Yuan, Wenhao Jiang, and Wei Liu. Swem: Towards real-time video object segmentation with sequential weighted expectation-maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1362–1372, 2022.
- [37] Xiankai Lu, Wenguan Wang, Martin Danelljan, Tianfei Zhou, Jianbing Shen, and Luc Van Gool. Video object segmentation with episodic graph memory networks. In *European Conference on Computer Vision*, pages 661–679. Springer, 2020.
- [38] Yanjun Ma, Dianhai Yu, Tian Wu, and Haifeng Wang. PaddlePaddle: An open-source deep learning platform from industrial practice. *Frontiers of Data and Computing*, 1(1):105–115, 2019.
- [39] K-K Maninis, Sergi Caelles, Yuhua Chen, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. Video object segmentation without temporal information. *IEEE transactions on pattern analysis and machine intelligence*, 41(6):1515–1530, 2018.
- [40] Yunyao Mao, Ning Wang, Wengang Zhou, and Houqiang Li. Joint inductive and transductive learning for video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9670–9679, 2021.
- [41] Tim Meinhardt and Laura Leal-Taixé. Make one-shot video object segmentation efficient again. *Advances in Neural Information Processing Systems*, 33:10607–10619, 2020.
- [42] Matthias Mueller, Neil Smith, and Bernard Ghanem. A benchmark and simulator for uav tracking. In *European conference on computer vision*, pages 445–461. Springer, 2016.
- [43] King Ngi Ngan and Hongliang Li. *Video segmentation and its applications*. Springer Science & Business Media, 2011.
- [44] Peter Ochs, Jitendra Malik, and Thomas Brox. Segmentation of moving objects by long term video analysis. *IEEE transactions on pattern analysis and machine intelligence*, 36(6):1187–1200, 2013.
- [45] Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, and Seon Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7376–7385, 2018.
- [46] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9226–9235, 2019.
- [47] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8405–8414, 2021.
- [48] Kwanyong Park, Sanghyun Woo, Seoung Wug Oh, In So Kweon, and Joon-Young Lee. Per-clip video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1352–1361, 2022.
- [49] Federico Perazzi, Anna Khoreva, Rodrigo Benenson, Bernt Schiele, and Alexander Sorkine-Hornung. Learning video object segmentation from static images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2663–2672, 2017.
- [50] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016.
- [51] 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.
- [52] 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. *International Journal of Computer Vision*, 130(8):2022–2039, 2022.
- [53] Andreas Robinson, Felix Jaremo Lawin, Martin Danelljan, Fahad Shahbaz Khan, and Michael Felsberg. Learning fast and robust target models for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7406–7415, 2020.
- [54] German Ros, Sebastian Ramos, Manuel Granados, Amir Bakhtiary, David Vazquez, and Antonio M Lopez. Vision-based offline-online perception paradigm for autonomous driving. In *2015 IEEE Winter Conference on Applications of Computer Vision*, pages 231–238. IEEE, 2015.
- [55] Khaled Saleh, Mohammed Hossny, and Saeid Nahavandi. Kangaroo vehicle collision detection using deep semantic segmentation convolutional neural network. In *2016 International Conference on Digital Image Computing: Techniques and Applications (DICTA)*, pages 1–7. IEEE, 2016.
- [56] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [57] Hongje Seong, Junhyuk Hyun, and Euntai Kim. Kernelized memory network for video object segmentation. In *European Conference on Computer Vision*, pages 629–645. Springer, 2020.
- [58] Hongje Seong, Seoung Wug Oh, Joon-Young Lee, Seongwon Lee, Suhyeon Lee, and Euntai Kim. Hierarchical memory matching network for video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12889–12898, 2021.
- [59] Xingjian Shi, Zhourong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. *Advances in neural information processing systems*, 28, 2015.
- [60] Jae Shin Yoon, Francois Rameau, Junsik Kim, Seokju Lee, Seunghak Shin, and In So Kweon. Pixel-level matching for video object segmentation using convolutional neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2167–2176, 2017.
- [61] Jack Valmadre, Luca Bertinetto, Joao F Henriques, Ran Tao, Andrea Vedaldi, Arnold WM Smeulders, Philip HS Torr, and Efstratios Gavves. Long-term tracking in the wild: A benchmark. In *Proceedings of the European conference on computer vision (ECCV)*, pages 670–685, 2018.
- [62] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5277–5286, 2019.
- [63] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9481–9490, 2019.
- [64] Paul Voigtlaender and Bastian Leibe. Online adaptation of convolutional neural networks for video object segmentation. *arXiv preprint arXiv:1706.09364*, 2017.
- [65] Haochen Wang, Xiaolong Jiang, Haibing Ren, Yao Hu, and Song Bai. Swiftnet: Real-time video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1296–1305, 2021.
- [66] Qiang Wang, Li Zhang, Luca Bertinetto, Weiming Hu, and Philip HS Torr. Fast online object tracking and segmentation: A unifying approach. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 1328–1338, 2019.
- [67] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10776–10785, 2021.
- [68] Ziqin Wang, Jun Xu, Li Liu, Fan Zhu, and Ling Shao. Ranet: Ranking attention network for fast video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3978–3987, 2019.
- [69] Huaxin Xiao, Jiashi Feng, Guosheng Lin, Yu Liu, and Maojun Zhang. Monet: Deep motion exploitation for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1140–1148, 2018.
- [70] Haozhe Xie, Hongxun Yao, Shangchen Zhou, Shengping Zhang, and Wenxiu Sun. Efficient regional memory network for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1286–1295, 2021.
- [71] Kai Xu, Longyin Wen, Guorong Li, Liefeng Bo, and Qingming Huang. Spatiotemporal cnn for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1379–1388, 2019.
- [72] 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 *Proceedings of the European conference on computer vision (ECCV)*, pages 585–601, 2018.
- [73] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5188–5197, 2019.
- [74] Linjie Yang, Yanran Wang, Xuehan Xiong, Jianchao Yang, and Aggelos K Katsaggelos. Efficient video object segmentation via network modulation. In *Proceedings of the IEEE*

- Conference on Computer Vision and Pattern Recognition*, pages 6499–6507, 2018.
- [75] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *European Conference on Computer Vision*, pages 332–348. Springer, 2020.
- [76] Zongxin Yang, Yunchao Wei, and Yi Yang. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34:2491–2502, 2021.
- [77] Lu Zhang, Zhe Lin, Jianming Zhang, Huchuan Lu, and You He. Fast video object segmentation via dynamic targeting network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5582–5591, 2019.
- [78] Ziyu Zhang, Sanja Fidler, and Raquel Urtasun. Instance-level segmentation for autonomous driving with deep densely connected mrfs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 669–677, 2016.