

Self-Supervised Video Object Segmentation by Motion-Aware Mask Propagation

Author

Miao, B, Bennamoun, M, Gao, Y, Mian, A

Published

2022

Conference Title

2022 IEEE International Conference on Multimedia and Expo (ICME)

Version

Accepted Manuscript (AM)

DOI

[10.1109/ICME52920.2022.9859966](https://doi.org/10.1109/ICME52920.2022.9859966)

Rights statement

© 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Downloaded from

<http://hdl.handle.net/10072/421556>

Griffith Research Online

<https://research-repository.griffith.edu.au>

Self-Supervised Video Object Segmentation by Motion-Aware Mask Propagation

Bo Miao¹, Mohammed Bennamoun¹, Yongsheng Gao², Ajmal Mian¹

¹ The University of Western Australia

² Griffith University

bo.miao@research.uwa.edu.au, {mohammed.bennamoun, ajmal.mian}@uwa.edu.au, yongsheng.gao@griffith.edu.au

Abstract

We propose a self-supervised spatio-temporal matching method, coined Motion-Aware Mask Propagation (MAMP), for video object segmentation. MAMP leverages the frame reconstruction task for training without the need for annotations. During inference, MAMP extracts high-resolution features from each frame to build a memory bank from the features as well as the predicted masks of selected past frames. MAMP then propagates the masks from the memory bank to subsequent frames according to our proposed motion-aware spatio-temporal matching module to handle fast motion and long-term matching scenarios. Evaluation on DAVIS-2017 and YouTube-VOS datasets show that MAMP achieves state-of-the-art performance with stronger generalization ability compared to existing self-supervised methods, i.e., 4.2% higher mean $\mathcal{J}\&\mathcal{F}$ on DAVIS-2017 and 4.85% higher mean $\mathcal{J}\&\mathcal{F}$ on the unseen categories of YouTube-VOS than the nearest competitor. Moreover, MAMP performs at par with many supervised video object segmentation methods.

Introduction

Video object segmentation (VOS) is a fundamental problem in visual understanding where the aim is to segment objects of interest from the background in unconstrained videos. VOS enables machines to sense the motion pattern, location, and boundaries of the objects of interest in videos, which is useful in a wide range of applications. For example, in video editing, manual frame-wise segmentation is laborious and does not maintain temporal consistency whereas VOS can segment all frames automatically using the mask of one frame as a guide. The problem of segmenting objects of interest in a video using the ground truth object masks provided for only the first frame is referred to as semi-supervised VOS. This is challenging because the appearance of objects in a video change significantly due to fast motion, occlusion, scale variation, etc. Moreover, other similar looking non-target objects may confuse the model to segment incorrect objects.

Semi-supervised VOS techniques fall into two categories: supervised and self-supervised. Supervised approaches (Oh et al. 2019; Yang, Wei, and Yang 2020) use rich annotation information from training data to learn the model achieving




Figure 1: Comparison on DAVIS-2017 validation set with other methods. MAMP outperforms existing self-supervised methods, and is at par with some supervised methods trained with large amounts of annotated data.

great success in VOS. However, these methods are unattractive given their reliance on accurate pixel-level annotations for training (see Fig. 1), which are expensive to generate. Moreover, supervised approaches struggle to maintain the same performance in the wild. In contrast, self-supervised methods (Lu et al. 2020; Lai, Lu, and Xie 2020) learn feature representations based on the intrinsic properties of the video frames, and thus do not require any annotations and can better generalize to unseen objects. Even though the motivations behind existing self-supervised methods are different, they share the same objective of learning to extract general feature representations and construct precise spatio-temporal correspondences to propagate the object masks in video sequences. Exploiting the spatio-temporal coherence in videos, the pretext tasks of self-supervised methods can be designed to either maximize the temporal correspondence consistency (Lu et al. 2020) or minimize the reconstruction or prediction loss (Lai, Lu, and Xie 2020). Once trained, the models are able to extract general feature representations and build spatio-temporal correspondences between the reference and query frames. Therefore, pixels in the query frames can be classified according to the mask labels of their corresponding region of interests (ROI) in

the reference frames. Despite their simplicity, existing self-supervised methods perform poorly in cases of fast motion and long-term matching scenarios.

To overcome the above challenges, we propose a self-supervised method coined motion-aware mask propagation (MAMP). Similar to previous self-supervised methods, MAMP learns image features and builds spatio-temporal correspondences without any annotations during training. During inference, MAMP first leverages the feature representations and the given object masks for the first frame to build a memory bank. The proposed motion-aware spatio-temporal matching module in MAMP then exploits the motion cues to mitigate the issues caused by fast motion and long-term correspondence mismatches, and propagates the mask from the memory bank to subsequent frames. Moreover, the proposed size-aware image feature alignment module fixes the misalignment during mask propagation and the memory bank is constantly updated by the past frames to provide the most appropriate spatio-temporal guidance. We evaluate MAMP on the DAVIS-2017 and YouTube-VOS benchmarks to verify its effectiveness as well as generalization ability. Our contributions are summarized as follows:

- We propose Motion-Aware Mask Propagation (MAMP) for VOS that trains the model end-to-end without any annotations and effectively propagates the masks across frames.
- We propose a motion-aware spatio-temporal matching module to mitigate errors caused by fast motion and long-term correspondence mismatches. This module improves the performance of MAMP on YouTube-VOS by 6.4%.
- Without any bells and whistles (e.g., fancy data augmentations, online adaptation, and external datasets), MAMP significantly outperforms existing self-supervised methods by 4.2% on DAVIS-2017 and 4.85% on the unseen categories of YouTube-VOS.
- Experiment on YouTube-VOS dataset shows that MAMP has the best generalization ability compared to existing self-supervised and supervised methods.

Related Work

Semi-supervised VOS aims to leverage the ground truth object mask given (only) in the first frame to segment the objects of interest in subsequent frames. Existing semi-supervised VOS methods can be divided into online-learning and offline-learning methods depending on whether online adaptation is needed during inference. The former usually update the networks dynamically during inference based on the first frame (Caelles et al. 2017; Maninis et al. 2018; Perazzi et al. 2017; Huang et al. 2020), synthetic frames (Khoreva et al. 2019; Luiten, Voigtlaender, and Leibe 2018), or multiple selected past frames (Voigtlaender and Leibe 2017; Meinhardt and Leal-Taixe 2020) of each video making the networks object-specific. Literature has shown the effectiveness of online-learning (Wang et al. 2019c), however, it is time consuming and adversely affects the models' generalization ability. Offline-learning methods usually propagate the given object mask or features of the first

frame either explicitly or implicitly to subsequent frames, making the expensive online-adaptation no longer necessary (Yang et al. 2018; Wang et al. 2019b). Recurrent neural networks (Ventura et al. 2019; Li and Loy 2018) and recurrent training strategies (Oh et al. 2018) can implicitly propagate the spatio-temporal features from past frames to the current frame. Memory networks with attention mechanism (Oh et al. 2019; Liang et al. 2020; Li, Shen, and Shan 2020; Seong, Hyun, and Kim 2020; Xie et al. 2021), correlation mechanism (Hu, Huang, and Schwing 2018; Voigtlaender et al. 2019; Yang, Wei, and Yang 2020; Lai, Lu, and Xie 2020), or clustering methods (Chen et al. 2018) can explicitly propagate the spatio-temporal features from multiple different past frames to the current frame. Existing offline-learning methods usually perform local or non-local spatio-temporal matching for temporal association and mask propagation. However, non-local matching is noisy and has a large memory footprint, while local matching struggles to cope with problems from fast motion and long-term correspondence mismatches.

The proposed MAMP is an offline-learning method and does not require time-consuming online adaptation. MAMP leverages a dynamically updated memory bank to store features and masks from selected past frames, and propagates masks effectively according to our proposed motion-aware spatio-temporal matching module. Unlike previous matching methods, our motion-aware spatio-temporal matching module not only excludes the noisy matching results but also mitigates the problems caused by fast motion and long-term correspondence mismatches.

Memory Networks aim to capture the long-term dependencies by storing temporal features or different categories of features in a memory module. LSTM (Hochreiter and Schmidhuber 1997) and GRU (Cho et al. 2014) implicitly represent spatio-temporal features with local memory cells in a highly compressed way limiting the representation ability. Memory networks (Weston, Chopra, and Bordes 2014) were introduced to explicitly store the important features. A common memory network in VOS is STM (Oh et al. 2019) which incrementally adds the features of past frames to the memory bank, and leverages the non-local spatio-temporal matching to provide spatio-temporal features. However, the incremental memory bank updates are impractical when segmenting long videos due to the growing memory cost. In this work, we divide the memory into long-term and short-term memory. The former is fixed, whereas the latter is updated dynamically using the past few frames making our MAMP memory efficient.

Self-supervised Learning can learn general feature representations and spatio-temporal correspondences based on the intrinsic properties of videos. It has shown promising capacity on various downstream tasks as it does not require annotations and can better generalize (Vondrick et al. 2018; Han, Xie, and Zisserman 2019; Li et al. 2019; Kim, Cho, and Kweon 2019; Wang, Jiao, and Liu 2020; Tao, Wang, and Yamasaki 2020; Pan et al. 2021). Many pretext tasks have been explored for self-supervised learning such as future frame prediction (Liu et al. 2018), query frame reconstruction (Lai and Xie 2019; Kim et al. 2020; Lai, Lu, and Xie 2020), patch




Figure 2: Framework of the proposed MAMP. During training, a random pair of neighboring video frames is sampled. The frames are converted to *Lab* color space and channel dropout is used only on the *ab* channels to generate the reconstruction target for self-supervision. During inference, a parameter sharing encoder is used to encode the frames into feature maps. The proposed motion-aware spatio-temporal matching module is used to compute spatio-temporal correspondences between the past frames in the memory bank and the current frame and propagate the masks across frames.

re-localization (Wang, Jabri, and Efros 2019; Lu et al. 2020), and motion statistics prediction (Wang et al. 2019a).

Method

Overview

Fig. 2 shows an overview of our method for VOS. MAMP is trained with the reconstruction task to learn feature representations and construct robust spatio-temporal correspondences between pairs of frames from the same video. Hence, *zero* annotation is required to train the model. During inference, MAMP segments the frames in a sequential manner. The parameter sharing encoder is used to extract frame-wise features and the motion-aware spatio-temporal matching module propagates the masks from the past frames to the current frame according to the spatio-temporal affinity matrix. The size-aware image feature alignment module also facilitates the mask propagation to prevent misalignment.

Self-supervised Feature Representation Learning

We use the reconstruction task for self-supervised feature representation learning and robust spatio-temporal matching. Since the channel correlation in *Lab* color space is smaller than that of *RGB* (Reinhard et al. 2001), we randomly dropout one of the *ab* channels and make the *Lab* color as the reconstruction target. Dropout preserves enough information for the input while avoiding trivial solutions. Therefore, the model learns the general feature representations and the spatio-temporal correspondences between the reference and query frames instead of learning how to predict the missing channel from the observed channels. To minimize the reconstruction loss, semantically similar pixels between reference and query frames are forced to have highly correlated feature representations, while semantically dissimilar pixels are forced to have weakly correlated feature representations. Finally, the reconstruction target of the

query frames is predicted according to the highly correlated ROIs in the reference frames.

Specifically, given a reference and query frame $\{I_r, I_q\} \in \mathbb{R}^{H \times W \times 3}$ from a video, a parameter-sharing convolutional encoder $\Phi(g(I); \theta)$ is used to extract their feature representations $\{F_r, F_q\} \in \mathbb{R}^{h \times w \times c}$, where $g(\cdot)$ is the information bottleneck to prevent trivial solutions. The dropped channels of the two frames $\{C_r, C_q\} \in \mathbb{R}^{H \times W \times 1}$ are downsampled to the resolution of the features $\{C_{r,d}, C_{q,d}\} \in \mathbb{R}^{h \times w \times 1}$ based on the size-aware image feature alignment module.

To enable C_r to represent and reconstruct C_q , the spatio-temporal affinity matrix $A_{q,r} \in \mathbb{R}^{hw \times R}$ that represents the strength of the correlation between F_q and F_r is:

$$A_{q,r}^{i,j} = \frac{\exp(\langle F_q^i, F_r^j \rangle / \sqrt{c})}{\sum_{j \in R} \exp(\langle F_q^i, F_r^j \rangle / \sqrt{c})}, \quad (1)$$

where i and j are the locations in F_q and F_r , respectively. $\langle \cdot, \cdot \rangle$ is the dot product between two vectors, and c refers to the number of channels to re-scale the correlation value and R is the ROI of i . Next, a location i in $C_{q,d}$ is represented by the weighted sum of the corresponding ROI in $C_{r,d}$:

$$\widehat{C}_{q,d}^i = \sum_{j \in R} A_{q,r}^{i,j} C_{r,d}^j \quad (2)$$

Finally, $\widehat{C}_{q,d}^i$ is upsampled to \widehat{C}_q^i , and Huber Loss \mathcal{L} is used to force \widehat{C}_q^i to be close to C_q^i :

$$\mathcal{L} = \frac{1}{n} \sum_{i=1}^N \mathcal{Z}_i \quad (3)$$

where

$$\mathcal{Z}_i = \begin{cases} 0.5(\widehat{C}_q^i - C_q^i)^2, & \text{if } |\widehat{C}_q^i - C_q^i| < 1 \\ |\widehat{C}_q^i - C_q^i| - 0.5, & \text{otherwise.} \end{cases} \quad (4)$$




Figure 3: (a) Proposed motion-aware spatio-temporal matching module. F_q , F_r are the features of query, reference frames. M_r is the masks of reference frames, and \hat{M}_q is the predicted mask. (b) Comparison of the matching results between the vanilla local spatio-temporal matching module and our motion-aware spatio-temporal matching module.

Motion-aware Spatio-temporal Matching

To account for moving objects, VOS models should be able to retrieve the corresponding ROIs from the reference frames for mask propagation under motion. To meet this constraint, non-local spatio-temporal matching methods (Oh et al. 2019) consider all locations in the reference frames as potential ROIs. However, non-local matching generates many noisy matches and has a large memory footprint. Local spatio-temporal matching methods (Lai, Lu, and Xie 2020) retrieve the ROIs in the reference frames based on the location coordinates in the query frame and a pre-defined retrieval radius. Although local matching is more efficient, existing local methods have limited receptive fields limiting their ability to localize the most correlated ROIs when they encounter fast motion or after long-term matching.

We propose a motion-aware spatio-temporal matching module that leverages optical flow to enable the query locations to retrieve the most correlated ROIs from the reference frames, see Fig. 3(a). We use RAFT (Teed and Deng 2020) which costs only about 18ms and 13ms to compute the motion offsets between frame pairs in DAVIS-2017 and YouTube-VOS datasets, respectively. As shown in Fig. 3(b), the vanilla local spatio-temporal matching method cannot retrieve the most correlated ROIs for the query point. However, with our motion-aware spatio-temporal matching, the query point can find its most correlated ROIs even if the ROI

pixels are not consecutive in raw space.

Our motion-aware spatio-temporal matching module takes the features of a query frame $F_q \in \mathbb{R}^{h \times w \times c}$, the features of reference frames $F_r \in \mathbb{R}^{h \times w \times c \times n}$, the downsampled masks of reference frames $M_r \in \mathbb{R}^{h \times w \times o \times n}$, and the downsampled motion offsets $MO(\Delta x, \Delta y) \in \mathbb{R}^{h \times w \times n \times 2}$ between the query and reference frames as inputs, where Δx and Δy are the displacement vectors along the horizontal and vertical directions, respectively. F_r and M_r are first warped according to $MO(\Delta x, \Delta y)$ making the locations with the same coordinates in F_q and the warped F_r and M_r to be the most similar pairs. Therefore, for one location $i \in F_q(x, y)$ in the query frame, the initial corresponding ROIs can be retrieved, i.e., $R = \{j \in F_r, |w(j) - i| \leq r\}$, where r is the radius of ROIs and $w(j) \in F_r(x + \Delta x, y + \Delta y)$. Due to different input resolutions, we empirically set r to 6 during training and 12 during inference. Subsequently, the initial spatio-temporal affinity matrix $Corr_{q,r} \in \mathbb{R}^{h \times w \times n(2r+1)^2}$ is computed and the TopK selection block is used to filter out the weakly correlated locations in the ROIs to save the Softmax operation from being affected by noise. k is set to 36 for all experiments (see Supplementary Material). Finally, mask propagation is achieved by multiplying the selected spatio-temporal affinity matrix with the corresponding ROIs in M_r .

During inference, with all the ROIs being dynamically sampled from different reference frames in the memory bank based on the motion-aware spatio-temporal matching module, the mask propagation becomes more accurate and the problems caused by fast motion and long-term correspondence mismatches are alleviated.

Size-aware Image Feature Alignment

To reduce memory consumption, previous methods perform bilinear downsampling on the supervision signals (masks) of the reference frames and propagate these signals at the feature resolution. However, this operation introduces misalignment (see Fig. 4(b)) between the strided convolution layers and the supervision signals from naïve bilinear downsampling and upsampling. MAST (Lai, Lu, and Xie 2020) has an image feature alignment module to deal with this problem but it does not cater for the misalignment caused at the upsampling stage. To solve this problem, we propose a size-aware image feature alignment module where the supervision signals are automatically padded (see Fig. 4(a)) so that the input size can be divisible by the size after downsampling, and the supervision signals are sampled at the convolution kernel centers. Hence, the misalignment is removed at both downsampling and upsampling stages.

Implementation Details

Training: We modify ResNet-18 and use it as the encoder to extract image features with a spatial resolution of 1/4 of the input images (see Supplementary Material for details). The encoder parameters are randomly initialized *without* pre-training. A pair of nearby video frames are randomly sampled as the reference and query frame, and the reconstruction task with Huber Loss is used to train the model. For

Method	Year	Backbone	Param.	T. Data	Vid. Length	Supervised	\mathcal{J} & \mathcal{F} (Mean)	\mathcal{J} (Mean)	\mathcal{F} (Mean)
Vid. Color.	2018	ResNet-18	5M	K	800 hrs	✗	34.0	34.6	32.7
CycleTime	2019	ResNet-50	9M	V	344 hrs	✗	48.7	46.4	50.0
CorrFlow	2019	ResNet-18	5M	O	14 hrs	✗	50.3	48.4	52.2
UVC	2019	ResNet-18	3M	K	800 hrs	✗	59.5	57.7	61.3
RPM-Net	2020	ResNet-101	43M	DY	5.67 hrs	✗	41.6	41.0	42.2
Mug	2020	ResNet-50	9M	O	14 hrs	✗	56.1	54.0	58.2
MAST	2020	ResNet-18	5M	Y	5.58 hrs	✗	65.5	63.3	67.6
Ours	2021	ResNet-18	5M	Y	5.58 hrs	✗	69.7	68.3	71.2
OSVOS	2017	VGG-16	15M	ID		✓	60.3	56.6	63.9
OSMN	2018	VGG-16	15M	ICD		✓	54.8	52.5	57.1
OSVOS-S	2018	VGG-16	15M	IPD		✓	68.0	64.7	71.3
SiamMask	2019	ResNet-50	9M	ICY		✓	56.4	54.3	58.5
FEELVOS	2019	Xception-65	38M	ICDY		✓	71.5	69.1	74.0
STM	2019	ResNet-50	9M	ICPSEDY		✓	81.8	79.2	84.3
GC	2020	ResNet-50	9M	ISEHD		✓	71.4	69.3	73.5
AFB-URR	2020	ResNet-50	9M	ICPSED		✓	74.6	73.0	76.1
CFBI	2020	ResNet-101	43M	ICDY		✓	81.9	79.1	84.6
RMNet	2021	ResNet-50	9M	ICPSEDY		✓	83.5	81.0	86.0

Table 1: Evaluation on DAVIS-2017 validation set. Note that each method modifies vanilla backbone models to suit their framework. Training Dataset (T. Data) notations: C=COCO, D=DAVIS, E=ECSSD, H=HKU-IS, I=ImageNet, K=Kinetics, M=Mapillary, O=OxUvA, P=PASCAL-VOC, S=MSRA10K, V=VLOG, Y=YouTube-VOS.




Figure 4: Size-aware image feature alignment module in comparison to bilinear sampling and image feature alignment. n is the ratio of the input to downsampled size. The proposed size-aware image feature alignment fixes the misalignment in both downsampling and upsampling stages.

pre-processing, the frames are resized to 256×256 , and no data augmentation is used. We train our model with pairwise frames for 33 epochs on YouTube-VOS using a batch-size of 24 for all experiments. We adopt Adam optimizer with the base learning rate of $1e-3$, and the learning rate is divided by 2 after 0.4M, 0.6M, 0.8M, and 1.0M iterations, respectively. Our model is trained end-to-end *without* any multi-stage training strategies. The training takes about 11 hours on one NVIDIA GeForce 3090 GPU.

Testing: The proposed MAMP does *not* require time-

consuming online adaption to fine-tune the model during testing. To be consistent with benchmarks, MAMP is evaluated on the YouTube-VOS validation set at half resolution and DAVIS-2017 validation set at full resolution. Results on DAVIS-2017 and YouTube-VOS are obtained using the official evaluation code and server. During testing, MAMP leverages the size-aware image feature alignment module to fix the misalignment, and uses the trained encoder to extract features from each frame. After that, MAMP uses the proposed motion-aware spatio-temporal matching to propagate the masks from the memory bank to subsequent frames. The memory bank of MAMP is updated dynamically to include I_0 and I_5 as long-term memory and I_{t-5} , I_{t-3} , and I_{t-1} as short-term memory.

Experiments

Datasets

DAVIS-2017 (Pont-Tuset et al. 2017) is commonly for VOS in *short* video clips and complex scenes. It contains 150 videos with over 200 objects. The validation set of DAVIS-2017 contains 30 videos.

YouTube-VOS (Xu et al. 2018) is the largest dataset for VOS with *long* video clips. It contains over 4000 high-resolution videos with 7000+ objects. The validation set of YouTube-VOS 2018 contains 474 videos. Unlike previous methods (Oh et al. 2019; Xie et al. 2021) that leverage several external datasets to train the model, we only train our model on YouTube-VOS and test our model on both DAVIS-2017 and YouTube-VOS. Unless specified otherwise, the YouTube-VOS dataset in this paper refers to the 2018 version which is consistent with previous benchmarks.

Method	Vid. Color.	CorrFlow	MAST	Ours	OSMN	RGMP	OnAVOS	S2S	A-GAME	STM	GC	RMNet
Sup.	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓	✓
Overall	38.9	46.6	64.2	68.2	51.2	53.8	55.2	64.6	66.1	79.4	73.2	81.5
\mathcal{J} Seen	43.1	50.6	63.9	67.0	60.0	59.5	60.1	71.0	67.8	79.7	72.6	82.1
\mathcal{F} Seen	38.6	46.6	64.9	68.4	60.1	-	62.7	70.0	-	84.2	75.6	85.7
\mathcal{J} Unseen	36.6	43.8	60.3	64.5	40.6	45.2	46.6	55.5	60.8	72.8	68.9	75.7
\mathcal{F} Unseen	37.4	45.6	67.7	73.2	44.0	-	51.4	61.2	-	80.9	75.7	82.4
Gen. Gap.	3.9	3.9	0.4	-1.2	17.8	14.3	12.4	12.2	7.0	5.1	1.8	4.9

Table 2: Evaluation on YouTube-VOS 2018 validation set for “seen” and “unseen” categories (“unseen” object category does not appear in training). For *overall*, \mathcal{J} , and \mathcal{F} , higher values are better and for Gen. (Generalization) Gap, lower values are better.

Evaluation Metrics

We use Region Similarity \mathcal{J} and Countour Accuracy \mathcal{F} for evaluation. We also report the Generalization Gap as in (Lai, Lu, and Xie 2020) to evaluate the generalization ability of MAMP on YouTube-VOS. Generalization Gap computes the model’s performance gap between inference on seen and unseen categories. Its value is inversely proportional to the generalization ability of a model.

Quantitative Results

We compare MAMP with existing methods on DAVIS-2017 and YouTube-VOS 2018. Results are obtained using the official evaluation code and server. We do our best to compare the results as fairly as possible. For example, multi-stage training strategies, external datasets, data augmentations, and online adaptation are *not* used in this work.

Table 1 and Table 2 summarize the performance of the state-of-the-art methods and MAMP on DAVIS-2017 and YouTube-VOS 2018. MAMP significantly outperforms benchmark self-supervised methods by over 4.2% on DAVIS-2017, 4% on YouTube-VOS 2018, and by 4.85% on the unseen categories of YouTube-VOS 2018. Moreover, MAMP is also comparable to some supervised methods. These results demonstrate the effectiveness of MAMP.

To evaluate the generalization ability of MAMP, we evaluate it on both “seen” and “unseen” categories of YouTube-VOS 2018. Objects in “unseen” categories do not appear in the training set. Table 2 shows that MAMP performs well on “unseen” categories and has the best generalization ability. Surprisingly, it performs better on “unseen” categories than on “seen” categories because of the better boundary segmentation performance on “unseen” objects. These results indicate that MAMP can learn general feature representations that are not restricted by the specific object categories in the training set. The most comparable supervised method in generalization ability is GC (Li, Shen, and Shan 2020) (1.8 vs -1.2). However, GC is trained with several external datasets with precise ground truth annotations.

In addition to YouTube-VOS 2018, we also evaluate MAMP on YouTube-VOS 2019 which has more videos and object instances. The results are included in supplementary material and show that MAMP achieves the best performance and generalization ability compared to other self-supervised methods on YouTube-VOS 2019 as well.

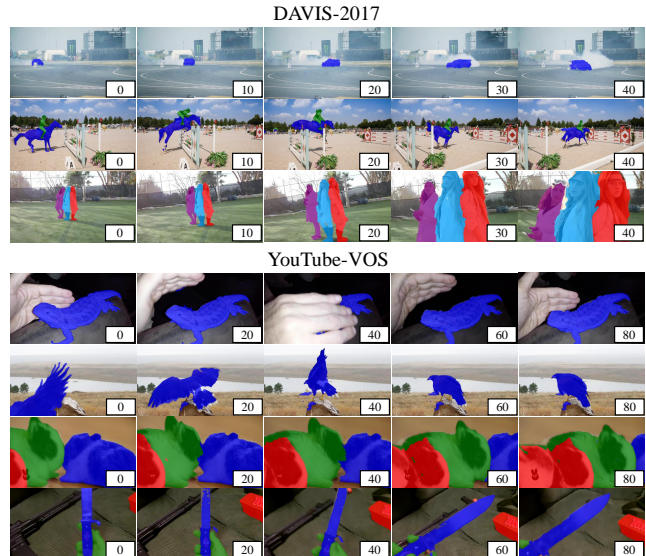


Figure 5: Qualitative results of MAMP. The frames are sampled at fixed intervals and the first frame in each video is assigned index 0. MAMP performs well in challenging scenarios of occlusion/dis-occlusion, fast motion, large deformations, and scale variations.

Qualitative Results

Figure 5 shows qualitative results of MAMP under various challenging scenarios, e.g., occlusion, fast motion, large deformations, and scale variations. MAMP is able to handle these challenging scenarios effectively.

Ablation Studies

Motion-aware Spatio-temporal Matching. In Table 3, we replaced the vanilla local spatio-temporal matching module of (Lai, Lu, and Xie 2020) with the proposed motion-aware spatio-temporal matching module and the performance increased by 3.1% on DAVIS-2017 and 6.4% on YouTube-VOS 2018. Vanilla local spatio-temporal matching module retrieves the corresponding ROIs according to the pre-defined radius and ROI localization method which is not accurate enough as the most correlated ROIs are prone to be outside the search radius under fast motion and long-term

M	A	DAVIS-2017			YouTube-VOS
		$\mathcal{J} \& \mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J} \& \mathcal{F}$
		65.8	63.6	68.0	61.8
✓		68.9 +3.1	66.9 +3.3	70.9 +2.9	68.2 +6.4
	✓	67.4 +1.6	65.9 +2.3	68.9 +0.9	61.8 +0.0
✓	✓	69.7 +3.9	68.3 +4.7	71.2 +3.2	68.2 +6.4

Table 3: Ablation experiment for motion-aware spatio-temporal matching module (M) and size-aware image feature alignment module (A).

matching scenarios. Without these most correlated ROIs, the label of one location in the query frame will be determined by the labels of several uncorrelated or weakly-correlated locations in the reference frames. Therefore, the segmentation results could be further improved if we can retrieve the most correlated ROIs for each location in the query frame.

Our motion-aware spatio-temporal matching module leverages the motion cues to register the reference frames in the memory bank to the query frame before computing the local spatio-temporal correspondences and filters out weakly correlated noise before mask propagation. Therefore, the above issues are alleviated even if the reference frames are temporally far from the query frame. As shown in Table 3, the motion-aware spatio-temporal matching module brings more performance gains in YouTube-VOS 2018, as this dataset has longer video clips and lower frame rates compared to DAVIS-2017. To further demonstrate the effectiveness of our motion-aware spatio-temporal matching module, we choose different k values to select the TopK correlated locations in the ROIs for mask propagation, the results show that MAMP still achieves the best performance compared to previous benchmarks even if only one of the 3125 locations in the ROIs is used for mask propagation (details in Supplementary Material).

Size-aware Image Feature Alignment. We removed the image feature alignment module of (Lai, Lu, and Xie 2020) to replace it with our size-aware image feature alignment module. The performance increased by 1.6% on DAVIS-2017 and remained unchanged on YouTube-VOS 2018 (see Table 3). This is because our size-aware image feature alignment module fixes the misalignment at the upsampling stage caused by the improper input size. After computing the statics, we found that the number of videos that have an improper input size is 96.7% for DAVIS-2017 validation set and only 1.9% for the YouTube-VOS 2018 validation set.

Long-term Memory and Short-term Memory. In Table 4, we compare results for different memory settings. We can see that all memory settings have reasonable performance. Long-term memory provides accurate ground truth information for query frames, while short-term memory offers up-to-date information from past neighboring frames. The results show that MAMP with short-term memory performs better than using long-term memory. This is because the appearance and scale of objects usually change significantly

Memory	DAVIS-2017			YouTube-VOS
	$\mathcal{J} \& \mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J} \& \mathcal{F}$
Long & Short	69.7	68.3	71.2	68.2
Long	52.3	49.7	54.9	58.7
Short	64.1	62.5	65.7	66.4

Table 4: Ablation study for long and short-term memory.

over time and long-term memory alone is unable to adapt to these changes. Furthermore, it can be seen that MAMP using both memory types has the best performance as both memories are complementary.

Comparison with MAST

Our nearest competitor is MAST (Lai, Lu, and Xie 2020) which leverages local spatio-temporal matching module with ROI localization for long-term mask propagation. The framework of MAMP is inspired from MAST, however, MAMP is different in various aspects: (1) The local spatio-temporal matching in MAST is sub-optimal for handling fast motions and long-term matching scenarios whereas the proposed motion-aware spatio-temporal matching in MAMP can better handle (see Table 3) such situations by exploiting motion cues and noise filters. (2) The local spatio-temporal matching module of MAST has a larger memory footprint compared to the proposed motion-aware spatio-temporal matching in MAMP. One 3090 GPU supports only 5 reference frames in the memory bank for MAST, but 18 reference frames for MAMP. (3) The image feature alignment module of MAST does not fix the misalignment at the upsampling stage, whereas the size-aware image feature alignment module of MAMP alleviates this problem as shown in Fig. 4(b) and Table 3. (4) MAST leverages multi-stage training strategies for training, i.e., finetunes the model using multiple reference frames. However, MAMP is trained end-to-end with pairwise frames once only. (5) MAST does not converge well using new releases of PyTorch and only obtains about 50 mean $\mathcal{J} \& \mathcal{F}$ on the validation set of DAVIS-2017 when training with PyTorch 1.9. MAMP solves this issue by normalizing the spatio-temporal affinity matrix with the channel number and achieves 69.7 mean $\mathcal{J} \& \mathcal{F}$ on the validation set of DAVIS-2017. Finally, qualitative comparisons of MAST and MAMP further demonstrate the superiority of MAMP (see videos in Supplementary Material).

Conclusion

In this paper, we proposed MAMP that enables general feature representation and motion-guided mask propagation. MAMP model can be trained without the need for annotations, and outperforms existing self-supervised methods by a large margin. Moreover, MAMP demonstrates the best generalization ability compared to previous methods. We believe that MAMP has the potential to propagate spatio-temporal features and masks in practical video segmentation tasks. In the future, we will develop more effective pretext tasks and adaptive memory selection methods to further improve the performance of MAMP.

Acknowledgements

This research was supported by the ARC Industrial Transformation Research Hub IH180100002.

References

- Caelles, S.; Maninis, K.-K.; Pont-Tuset, J.; Leal-Taixé, L.; Cremers, D.; and Van Gool, L. 2017. One-shot video object segmentation. In *CVPR*, 221–230.
- Chen, Y.; Pont-Tuset, J.; Montes, A.; and Van Gool, L. 2018. Blazingly fast video object segmentation with pixel-wise metric learning. In *CVPR*, 1189–1198.
- Cho, K.; van Merriënboer, B.; Gülçehre, Ç.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*.
- Han, T.; Xie, W.; and Zisserman, A. 2019. Video representation learning by dense predictive coding. In *ICCV Workshops*.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Hu, Y.-T.; Huang, J.-B.; and Schwing, A. G. 2018. Video-match: Matching based video object segmentation. In *ECCV*, 54–70.
- Huang, X.; Xu, J.; Tai, Y.-W.; and Tang, C.-K. 2020. Fast Video Object Segmentation With Temporal Aggregation Network and Dynamic Template Matching. In *CVPR*, 8879–8889.
- Khoreva, A.; Benenson, R.; Ilg, E.; Brox, T.; and Schiele, B. 2019. Lucid data dreaming for video object segmentation. *International Journal of Computer Vision*, 127(9): 1175–1197.
- Kim, D.; Cho, D.; and Kweon, I. S. 2019. Self-supervised video representation learning with space-time cubic puzzles. In *AAAI*, 8545–8552.
- Kim, Y.; Choi, S.; Lee, H.; Kim, T.; and Kim, C. 2020. Rpm-net: Robust pixel-level matching networks for self-supervised video object segmentation. In *WACV*, 2057–2065.
- Lai, Z.; Lu, E.; and Xie, W. 2020. MAST: A memory-augmented self-supervised tracker. In *CVPR*, 6479–6488.
- Lai, Z.; and Xie, W. 2019. Self-supervised learning for video correspondence flow. In *BMVC*.
- Li, X.; Liu, S.; De Mello, S.; Wang, X.; Kautz, J.; and Yang, M.-H. 2019. Joint-task self-supervised learning for temporal correspondence. In *NeurIPS*, 318–328.
- Li, X.; and Loy, C. C. 2018. Video object segmentation with joint re-identification and attention-aware mask propagation. In *ECCV*, 90–105.
- Li, Y.; Shen, Z.; and Shan, Y. 2020. Fast Video Object Segmentation using the Global Context Module. In *ECCV*, 735–750.
- Liang, Y.; Li, X.; Jafari, N.; and Chen, J. 2020. Video Object Segmentation with Adaptive Feature Bank and Uncertain-Region Refinement. In *NeurIPS*, 3430–3441.
- Liu, W.; Luo, W.; Lian, D.; and Gao, S. 2018. Future frame prediction for anomaly detection—a new baseline. In *CVPR*, 6536–6545.
- Lu, X.; Wang, W.; Shen, J.; Tai, Y.-W.; Crandall, D. J.; and Hoi, S. C. 2020. Learning video object segmentation from unlabeled videos. In *CVPR*, 8960–8970.
- Luiten, J.; Voigtlaender, P.; and Leibe, B. 2018. Premvos: Proposal-generation, refinement and merging for video object segmentation. In *ACCV*, 565–580.
- Maninis, K.-K.; Caelles, S.; Chen, Y.; Pont-Tuset, J.; Leal-Taixé, L.; Cremers, D.; and Van Gool, L. 2018. Video object segmentation without temporal information. *IEEE transactions on pattern analysis and machine intelligence*, 41(6): 1515–1530.
- Meinhardt, T.; and Leal-Taixé, L. 2020. Make One-Shot Video Object Segmentation Efficient Again. In *NeurIPS*.
- Oh, S. W.; Lee, J.-Y.; Sunkavalli, K.; and Kim, S. J. 2018. Fast video object segmentation by reference-guided mask propagation. In *CVPR*, 7376–7385.
- Oh, S. W.; Lee, J.-Y.; Xu, N.; and Kim, S. J. 2019. Video object segmentation using space-time memory networks. In *ICCV*, 9226–9235.
- Pan, T.; Song, Y.; Yang, T.; Jiang, W.; and Liu, W. 2021. Videomoco: Contrastive video representation learning with temporally adversarial examples. In *CVPR*, 11205–11214.
- Perazzi, F.; Khoreva, A.; Benenson, R.; Schiele, B.; and Sorkine-Hornung, A. 2017. Learning video object segmentation from static images. In *CVPR*, 2663–2672.
- Pont-Tuset, J.; Perazzi, F.; Caelles, S.; Arbeláez, P.; Sorkine-Hornung, A.; and Van Gool, L. 2017. The 2017 DAVIS Challenge on Video Object Segmentation. *arXiv:1704.00675*.
- Reinhard, E.; Adhikhmin, M.; Gooch, B.; and Shirley, P. 2001. Color transfer between images. *IEEE Computer graphics and applications*, 21(5): 34–41.
- Seong, H.; Hyun, J.; and Kim, E. 2020. Kernelized Memory Network for Video Object Segmentation. In *ECCV*, 629–645.
- Tao, L.; Wang, X.; and Yamasaki, T. 2020. Self-supervised video representation learning using inter-intra contrastive framework. In *ACM MM*, 2193–2201.
- Teed, Z.; and Deng, J. 2020. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 402–419.
- Ventura, C.; Bellver, M.; Girbau, A.; Salvador, A.; Marques, F.; and Giro-i Nieto, X. 2019. RVOS: End-to-End Recurrent Network for Video Object Segmentation. In *CVPR*, 5277–5286.
- Voigtlaender, P.; Chai, Y.; Schroff, F.; Adam, H.; Leibe, B.; and Chen, L.-C. 2019. Feelvos: Fast end-to-end embedding learning for video object segmentation. In *CVPR*, 9481–9490.
- Voigtlaender, P.; and Leibe, B. 2017. Online adaptation of convolutional neural networks for video object segmentation. In *BMVC*.

Vondrick, C.; Shrivastava, A.; Fathi, A.; Guadarrama, S.; and Murphy, K. 2018. Tracking emerges by colorizing videos. In *ECCV*, 391–408.

Wang, J.; Jiao, J.; Bao, L.; He, S.; Liu, Y.; and Liu, W. 2019a. Self-supervised spatio-temporal representation learning for videos by predicting motion and appearance statistics. In *CVPR*, 4006–4015.

Wang, J.; Jiao, J.; and Liu, Y.-H. 2020. Self-supervised video representation learning by pace prediction. In *ECCV*, 504–521.

Wang, Q.; Zhang, L.; Bertinetto, L.; Hu, W.; and Torr, P. H. 2019b. Fast online object tracking and segmentation: A unifying approach. In *CVPR*, 1328–1338.

Wang, X.; Jabri, A.; and Efros, A. A. 2019. Learning correspondence from the cycle-consistency of time. In *CVPR*, 2566–2576.

Wang, Z.; Xu, J.; Liu, L.; Zhu, F.; and Shao, L. 2019c. Ranet: Ranking attention network for fast video object segmentation. In *ICCV*, 3978–3987.

Weston, J.; Chopra, S.; and Bordes, A. 2014. Memory networks. *arXiv preprint arXiv:1410.3916*.

Xie, H.; Yao, H.; Zhou, S.; Zhang, S.; and Sun, W. 2021. Efficient Regional Memory Network for Video Object Segmentation. In *CVPR*, 1286–1295.

Xu, N.; Yang, L.; Fan, Y.; Yue, D.; Liang, Y.; Yang, J.; and Huang, T. 2018. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*.

Yang, L.; Wang, Y.; Xiong, X.; Yang, J.; and Katsaggelos, A. K. 2018. Efficient video object segmentation via network modulation. In *CVPR*, 6499–6507.

Yang, Z.; Wei, Y.; and Yang, Y. 2020. Collaborative video object segmentation by foreground-background integration. In *ECCV*, 332–348.

Backbone Architecture

We modify ResNet-18 and use it as the encoder to extract image features with a spatial resolution of 1/4 of the input images, the specific architecture is shown in Table 5.

Layer Name	Output Size	Configuration
Conv1	$H/2 \times W/2$	$7 \times 7, 64, \text{stride } 2$
Conv2	$H/2 \times W/2$	$\begin{bmatrix} 3 \times 3, 64 \\ 3 \times 3, 64 \end{bmatrix} \times 2$
Conv3	$H/4 \times W/4$	$\begin{bmatrix} 3 \times 3, 128 \\ 3 \times 3, 128 \end{bmatrix} \times 2$
Conv4	$H/4 \times W/4$	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 2$
Conv5	$H/4 \times W/4$	$\begin{bmatrix} 3 \times 3, 256 \\ 3 \times 3, 256 \end{bmatrix} \times 2$

Table 5: Architecture of the modified ResNet-18.

TopK Correlated Locations for Mask Propagation

If we retrieve ROIs with a radius of 12 on 5 reference frames, the corresponding ROIs for one location in the query frame will include 3125 locations. However, noisy matches to 3125 locations may adversely affect the model’s performance. Hence, the proposed motion-aware spatio-temporal matching filters out redundant and noisy ROIs by selecting TopK correlated locations only for mask propagation. As shown in Table 6, leveraging the top 36 or top 9 correlated locations in the ROIs for mask propagation improves the performance of MAMP compared to using all 3125 locations. Using the top 36 correlated locations obtains the best performance. Moreover, compared to other state-of-the-art self-supervised methods, MAMP still maintains the best performance even if only one of the 3125 locations in the ROIs is used for mask propagation. These results further demonstrate the effectiveness of the proposed motion-aware spatio-temporal matching module.

TopK	DAVIS-2017			YouTube-VOS
	$\mathcal{J} \ \& \ \mathcal{F}$	\mathcal{J}	\mathcal{F}	$\mathcal{J} \ \& \ \mathcal{F}$
ALL	69.0	67.6	70.3	67.9
Top 1	66.5 -2.5	64.4 -3.2	68.6 -1.7	66.3 -1.6
Top 9	69.4 +0.4	67.7 +0.1	71.2 +0.9	68.4 +0.5
Top 36	69.7 +0.7	68.3 +0.7	71.2 +0.9	68.2 +0.3

Table 6: Ablation of TopK correlated locations in the ROIs for mask propagation.

Evaluation on YouTube-VOS 2019

We also evaluate the proposed MAMP on YouTube-VOS 2019. YouTube-VOS 2019 includes more videos and object instances compared to YouTube-VOS 2018. As shown

in Table 7, MAMP still significantly outperforms other self-supervised methods and has the best generalization ability.

Method	Sup.	Overall	Seen		Unseen		Gen. Gap
			\mathcal{J}	\mathcal{F}	\mathcal{J}	\mathcal{F}	
Vid. Color.	\times	39.0	43.3	38.2	36.6	37.5	3.7
CorrFlow	\times	47.0	51.2	46.6	44.5	45.9	3.7
MAST	\times	64.9	64.3	65.3	61.5	68.4	0.15
Ours	\times	68.2	66.3	67.5	65.4	73.7	-2.6

Table 7: Evaluation on YouTube-VOS 2019 validation set for “seen” and “unseen” categories.