

SwiftNet: Real-time Video Object Segmentation

Haochen Wang^{1*}, Xiaolong Jiang^{1*}, Haibing Ren¹, Yao Hu¹, Song Bai^{1,2}

¹Alibaba Youku Cognitive and Intelligent Lab

²University of Oxford

{zhinong.whc, xainglu.jxl, haibing.rhb, yaohu}@alibaba-inc.com

songbai.site@gmail.com

Abstract

In this work we present SwiftNet for real-time semisupervised video object segmentation (one-shot VOS), which reports 77.8% $J&F$ and 70 FPS on DAVIS 2017 validation dataset, leading all present solutions in overall accuracy and speed performance. We achieve this by elaborately compressing spatiotemporal redundancy in matching-based VOS via Pixel-Adaptive Memory (PAM). Temporally, PAM adaptively triggers memory updates on frames where objects display noteworthy inter-frame variations. Spatially, PAM selectively performs memory update and match on dynamic pixels while ignoring the static ones, significantly reducing redundant computations wasted on segmentation-irrelevant pixels. To promote efficient reference encoding, light-aggregation encoder is also introduced in SwiftNet deploying reversed sub-pixel. We hope SwiftNet could set a strong and efficient baseline for real-time VOS and facilitate its application in mobile vision. The source code of SwiftNet can be found at <https://github.com/haochenheheda/SwiftNet>.

1. Introduction

Given the first frame annotation, semi-supervised video object segmentation (one-shot VOS) localizes the annotated object(s) on pixel-level throughout the video. One-shot VOS generally adopts a matching-based strategy, where target objects are first modeled from historical reference frames, then precisely matched against the incoming query frame for localization. Being a video-based task, VOS finds vast applications in surveillance, video editing, and mobile visions, most of which ask for real-time processing [41].

Nonetheless, although pursued in fruitful endeavors [2, 22, 35, 6, 17, 29, 11], real-time VOS remains unsolved, as object variation over-time poses heavy demands for sophisticated object modeling and matching computations. As a compromise, most existing methods solely focus on

*These authors contribute equally.




Figure 1. Accuracy and speed performance of state-of-the-art methods on DAVIS2017 validation dataset, methods locate on the right side of the red vertical dotted line meet the real-time requirement. Our solutions (ResNet-18 and 50 versions) are marked in red.

improving segmentation accuracy while at the expense of speed. Amongst, memory-based methods [20, 46, 45] reveal exceptional accuracy with comprehensively modeling object variations using all historical frames and expressive non-local [33] reference-query matching. Unfortunately, deploying more reference frames and complicated matching scheme inevitably slow down segmentation. Accordingly, recent attempts seek to accelerate VOS with reduced reference frames and light-weight matching scheme [9, 26, 13, 3, 34, 29, 4, 35]. For the first aspect, solutions proposed in [26, 13, 3, 34, 29, 4, 35] follow a mask-propagation strategy, where only the first and last historical frames are considered reference for current segmentation. For the second aspect, efficient pixel-wise matching [13, 29, 34], region-wise distance measuring [26, 9, 4], and correlation filtering [3, 32] are deployed to reduce computations. However, as shown in Fig.1, although these accelerated methods enjoy faster segmentation speed, they still barely meet real-time requirement, and more critically, they are far from state-of-the-art segmentation accuracy.



Figure 2. An illustration of pixel-wise memory update in video clips. In each row, pixels updated into memory (red) are the ones previously invisible and remain static afterwards.

We argue that, the accurate solutions are less efficient due to the spatiotemporal redundancy inherently resides in matching-based VOS, and the fast solutions suffer degraded accuracy for reducing the redundancy indiscriminately. Considering its pixel-wise modeling, matching, and estimating nature, matching-based VOS manifests positive correlation between processing time T and multiplied number of pixels \mathcal{N}_r and \mathcal{N}_q in reference and query frames as described in Eq. 1. The spatiotemporal redundancy denotes that \mathcal{N}_r is populated with pixels not beneficial for accurate segmentation. Temporally, existing methods [20, 45] carelessly involve all historical frames (mostly by periodic sampling) for reference modeling, resulting in the fact that static frames showing no object evolution are repeatedly modeled, while dynamic frames containing incremental object information are less attended. Spatially, full-frame modeling and matching are adopted as defaults [20, 40], wherein most static pixels are redundant for segmentation. Illustration in Fig. 2 vividly advocates the above point. From this standpoint, explicitly compressing pixel-wise spatiotemporal redundancy is the best way to yield accurate and fast one-shot VOS.

$$T \propto \mathcal{O}(\mathcal{N}_r \times \mathcal{N}_q), \quad (1)$$

Accordingly, we propose SwiftNet for real-time one-shot video object segmentation. Overall, as depicted in Fig.3, SwiftNet instantiates matching-based segmentation with an encoder-decoder architecture, where spatiotemporal redundancy is compressed within the proposed Pixel-Adaptive Memory (PAM) component. Temporally, instead of involving all historical frames indiscriminately as reference, PAM introduces a variation-aware trigger module, which computes inter-frame difference to adaptively activate memory update on temporally-varied frames while overlook the static ones. Spatially, we abolish full-frame operations and design pixel-wise update and match modules in PAM. For pixel-wise memory update, we explicitly evaluate inter-frame pixel similarity to identify a subset of pixels beneficial for memory, and incrementally add their feature representation into the memory while bypassing the redundant ones. For pixel-wise memory match, we compress the time-consuming non-local computation to accom-

modate the pixel-wise memory as reference, thus achieving efficient matching without degradation of accuracy. To further accelerate segmentation, PAM is equipped with a novel light-aggregation encoder (LAE), which eschews redundant feature extraction and enables multi-scale mask-frame aggregation leveraging reversed sub-pixel operations.

In summary, we highlight three main contributions:

- We propose SwiftNet to set the new record in overall segmentation accuracy and speed, thus providing a strong baseline for real-time VOS with publicized source code.
- We pinpoint spatiotemporal redundancy as the Achilles heel of real-time VOS, and resolve it with Pixel-Adaptive Memory (PAM) composing variation-aware trigger and pixel-wise update & matching. Light-Aggregation Encoder (LAE) is also introduced for efficient and thorough reference encoding.
- We conduct extensive experiments deploying various backbones on DAVIS 2016 & 2017 and YouTube-VOS datasets, reaching the best overall segmentation accuracy and speed performance at 77.8% $\mathcal{J}\&\mathcal{F}$ and 70 FPS on DAVIS2017 validation set.

2. Related Work

2.1. One-shot VOS

One-shot VOS establishes a spatiotemporal matching problem, such that objects annotated in the first frame are localized in upcoming query frames by searching pixels best-matched to object template modeled in the reference frames. From this perspective, we categorize one-shot VOS methods *w.r.t* different reference modeling and reference-query matching strategies. Reference modeling builds object template by exploiting object evolution in historical frames, and methods either follows the last-frame or all-frame approaches. For the former one, [35, 34, 42, 43, 14, 29, 10] utilize only the first and/or last frame as reference, demonstrate favorable segmentation speed but suffer uncompetitive accuracy due to inadequate modeling over object variation. For the latter one, methods proposed in [20, 36, 28, 8, 27, 44, 23] leverage all previous frames and reveal improved accuracy, but they suffer slower speed for heavy computation overhead even with periodic sampling.

Considering reference-query matching, we classify methods as two-stage [43, 38, 42] and one-stage [14, 35, 34, 36, 20, 10, 29, 28, 7] basing on whether matching with proposed region-of-interest. Similar to object detection, two-stage method is more accurate while one-stage leads in speed. Besides, the key of matching is similarity measuring, where convolutional networks [14, 35, 36, 10], cross correlation [34, 32], and non-local computation [20, 46, 45, 29] are widely adopted. Amongst, non-local [33] reveals best




Figure 3. An illustration of SwiftNet. Solid black lines represent operation executed first to generate segmentation mask, dotted lines perform afterwards for memory update. Colored dots in memory update and match modules denote pixel-wise key feature vectors. Corresponding value features are omitted in memory update and match for brevity.

accuracy for capturing all-pairs pixel-wise dependency but are computationally heavy. In addition to matching-based VOS, propagation-based methods [11, 22, 39, 44, 25] leverage temporal motion consistency to reinforce segmentation, which is highly effective when appearance matching fails due to severe variations. Additionally, time-consuming on-line fine-tuning are exploited in [2, 18, 31] to improve segmentation accuracy, which however is impractical for real-time application.

2.2. Fast VOS

For efficiency, most fast VOS solutions deploy the single-frame reference strategy [29, 11, 34, 35, 39]. Besides, methods proposed in [32, 3, 4, 26, 9] employ segmentation-by-tracking where pixel-wise estimation is gated within tracked bounding-boxes to avoid full-frame estimation. To expedite time-consuming pixel-wise matching, RGMP [35] computes similarity responses with convolutions; AGAME [11] discriminates object from background with a probabilistic generative appearance model; RANet [34] adopts cross-correlation on ranked pixel-wise features to match query with reference. In addition, OSNM [39] proposes to spur VOS with network modulation.

2.3. Memory-based VOS

Memory-based VOS exploits all historical frames in an external memory for object modeling, an alternative approach for modeling all-frame evolution is via the implementation of recurrent neural networks [8, 37, 28]. First proposed in [20], STM is the seminal memory-based method which boosts segmentation accuracy by a large margin. As follows, [46, 45] modify STM by introducing Siamese-based semantic similarity and motion-guided attention. To induce heavy computations, GCNet [13] designs a global context module using attentions to reduce temporal

complexity executed in the memory.

3. SwiftNet

In this section, we present SwiftNet by first briefly formulating the problem of matching-based one-shot VOS. As follows, PAM is discussed in details, including variation-aware trigger as well as pixel-wise memory update and match modules. LAE is explained afterwards.

3.1. Problem Formulation

Given a video sequence $V = [x_1, x_2, \dots, x_T]$, its first frame x_1 is annotated with mask y_1 . The goal of one-shot VOS is to delineate objects from the background by generating mask y_t for each frame t . Particularly, matching-based VOS computes mask via object modeling and matching.

For object modeling at frame t , historical information embedded in reference frames $[x_1, \dots, x_{t-1}]$ and $[y_1, \dots, y_{t-1}]$ is exploited to establish object model M_{t-1} for up till frame $t - 1$:

$$M_{t-1} = \phi(I(1) \cdot EnR(x_1, m_1), I(2) \cdot EnR(x_2, m_2), \dots, I(t-1) \cdot EnR(x_{t-1}, m_{t-1})), \quad (2)$$

here $I(t)$ is an indicator function denoting whether frame t involves in modeling, $EnR(\cdot)$ indicates reference encoder for feature extraction, and $\phi(\cdot)$ generalizes the object modeling process.

For reference-query matching, the task is to search M_{t-1} within x_t on a pixel-level and generate object affinity map A_t :

$$A_t = \gamma(M_{t-1}, EnQ(x_t)), \quad (3)$$

Here $\gamma(\cdot)$ denotes pixel-wise matching and $EnQ(\cdot)$ refers to the query encoder. The final segmentation mask is produced by a decoder integrating encoded features and A_t .

At test time with SwiftNet, upon the arrival of query frame x_t , it is first processed by the query encoder and then passed into the pixel-wise memory match module. The matching output and encoded query features are aggregated in the decoder to generate mask y_t . Subsequently, x_t , y_t , x_{t-1} and y_{t-1} are jointly fed into the variation-aware trigger module, and if triggered, they are then handled by LAE for later pixel-wise memory update. This overall workflow is illustrated in Fig.3.

3.2. Pixel-Adaptive Memory

As the core component of SwiftNet, PAM models object evolution and performs object matching with explicitly compressed spatiotemporal redundancy. PAM mainly composes the variation-aware trigger as well as the pixel-wise memory update and match modules.

3.2.1 Variation-Aware Trigger

Instead of utilizing merely the first and last frames for object modeling, incorporating all historical frames as reference help establish temporally-coherent object evolution [20, 37]. Nonetheless, this approach is rather impractical considering its prohibitive temporal redundancy and computation overhead. As a straightforward solution, previous methods sample historical frames at a predefined pace [20, 23], which indiscriminately reduces temporal redundancy and leads to accuracy degradation.

To explicitly compress temporal redundancy, variation-aware trigger module evaluates inter-frame variation frame-by-frame, and activates memory update once the accumulated variation surpass threshold P_{th} . Specifically, given x_t , y_t and x_{t-1} and y_{t-1} , we separately compute image difference D_f and mask difference D_m as:

$$D_f^i = \sum_{c \in \{R, G, B\}} |x_t^{i,c} - x_{t-1}^{i,c}| / 255, \quad (4)$$

$$D_m^i = |y_t^i - y_{t-1}^i|, \quad (5)$$

at each pixel i we update the overall running variation degree P as:

$$P = \begin{cases} P + 1, & \text{if } D_f^i > th_f \text{ or } D_m^i > th_m \\ P, & \text{otherwise} \end{cases} \quad (6)$$

Once P exceeds P_{th} , PAM triggers a new round of memory update as described in 3.2.2. Empirically, P_{th} , th_f and th_m equal 200, 1, 0 respectively yields best performance.

3.2.2 Pixel-wise Memory Update

In terms of matching-based VOS, memory infers a temporally-maintained template which characterizes object evolution over time. In the existing literature, memory update and matching typically adopt full-frame operations,

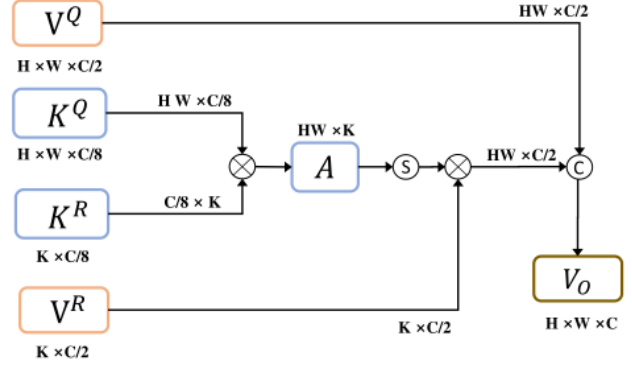


Figure 4. A diagram of the pixel-wise memory match computation, sub-script t is omitted for brevity. The dotted s and c stands for Softmax and concatenation operations.

where reference frames are concatenated into memory and matched with query frame intactly [20, 45]. This strategy induces heavy storage and computation overhead, as redundant pixels showing no benefit for object modeling are incorporated without discrimination.

To compress redundant pixels from full frames, PAM introduces pixel-wise memory update and match modules. For memory update, if frame x_t is triggered, PAM first discovers pixels in x_t that demonstrates significant variations from itself in memory M_{t-1} , then incrementally updates newly discovered features (as displayed in x_t) into the memory. Through EnR , x_t is encoded into key $K_{Q,t} \in \mathbb{R}^{H \times W \times C/8}$ and value $V_{Q,t} \in \mathbb{R}^{H \times W \times C/2}$ features, key features are with shallower depth to facilitate efficient matching. In the experiment C is set to 256. Similarly, memory M_{t-1} containing k_{t-1} pixels in forms of $K_{R,t-1} \in \mathbb{R}^{k_{t-1} \times C/8}$ and $V_{R,t-1} \in \mathbb{R}^{k_{t-1} \times C/2}$. To discover varied pixels, we flatten $K_{Q,t}$ and compute cosine similarity matrix $S \in \mathbb{R}^{HW \times k_{t-1}}$ as:

$$S^{i,j} = \frac{K_{Q,t}^i \cdot K_{R,t-1}^j}{\|K_{Q,t}^i\| \|K_{R,t-1}^j\|}, \quad (7)$$

for each row i in S , we find the largest score as the feature similarity between pixel i in query frame t and in pixel j in the memory. In formulation, we compute pixel similarity vector s as:

$$s_i = \arg \max_j S[i, :], \quad (8)$$

we sort s in increasing order of similarity (original index is kept), then the select top β percents pixels for memory update. These set of pixels exhibit most severe feature variations. Here β is a hyper-parameter controlling the balance between method efficiency and update comprehensiveness, and is experimentally set to 10% for the best performance. To execute the memory update, we find feature vectors of the selected set of pixels from $K_{Q,t}$ and $V_{Q,t}$ according to indexes as in s , then directly add them into memory M which is instantiated as an array of feature vectors.

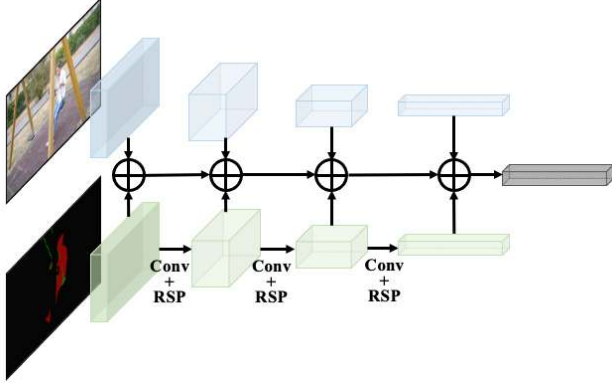


Figure 5. An illustration of LAE. Image features are pre-computed by EnQ , mask features are parallelly generated via convolutions and reversed sub-pixel (RSP) modules.

3.2.3 Pixel-wise Memory Match

As illustrated in Fig. 3, segmentation mask is decoded utilizing query value V_Q and output of reference-query matching, which provides strong spatial prior *w.r.t.* the foreground object. In essence, this matching process computes similarity between pixels from the reference and query frames, and can be instantiated with cross-correlation [3, 32], neural networks [35, 11], distance measuring [29], and non-local computation [20], etc. Comparatively, non-local leads to excellent accuracy performance but suffers heavy computation expenses in the context of full-frame operations. In PAM, we implement pixel-wise matching to achieve efficient and accurate segmentation.

In Fig. 4 we illustrate the pipeline of pixel-wise matching computation. At first, query frame key $K_{Q,t}$ and the memory key $K_{R,t-1}$ are reshaped into vectors of size $HW \times C/8$ and $C/8 \times K$. We then calculate dot-product similarity between corresponding vectors to produce affinity map $A \in \mathbb{R}^{HW \times K}$ as:

$$A^{i,j} = \exp(K_{Q,t}^i \odot K_{R,t-1}^j), \quad (9)$$

I_t is passed through a Softmax layer and further multiplied with memory value $V_{R,t-1}$. The resulted tensor ($\in \mathbb{R}^{HW \times C/2}$) is concatenated with $V_{Q,t}$ to form the activated feature as:

$$V_D = [A \times V_{R,t-1}, V_{Q,t}], \quad (10)$$

which is then input into the decoder. We emphasize that, our approach is different to normal non-local computation implemented in [20] as we eliminate redundant pixels from full-frame memory. Affinity map A , being the computation bottleneck with size $\mathbb{R}^{HW \times HW \times K}$, is significantly reduced to $\mathbb{R}^{HW \times K}$. K is the size of pixel-wise memory and is strongly controlled by the update ratio β . By explicitly compressing redundancy during memory update and matching, storage requirement and computation speed are both optimized in SwiftNet without considerable accuracy loss.

3.3. Light-Aggregation Encoder

In existing memory-based VOS solutions [20, 45], the image encoding process is time-consuming as both query and reference encoders adopt heavy backbone networks. In SwiftNet we expedite encoding by sparing the backbone resides in EnR . Particularly, we instantiate EnQ with ResNet-based [5] networks, and after x_t is encoded by EnQ , we buffer the generated feature maps. If frame t is triggered for update, these buffered features are directly utilized in EnR for reference encoding. Efficiency comparison *w.r.t.* encoding strategies are listed in Table. 1.

To facilitate the described encoding process, we design the novel light-aggregation encoder as shown in Fig. 5. The upper blue entities represent buffered feature maps encoded by EnQ , the bottom green ones show feature transformation hierarchy of the input mask. Features aligned vertically in the same column are with the same size and concatenated together to facilitate multi-scale aggregation. In particular, to instantiate feature transformation of the input mask, we implement reversed sub-pixel for down-samplings and 1×1 convolutions for channel manipulation. Reversed sub-pixel technique is motivated by the popular up-sampling method in super-resolution [24], which shrinks spatial dimension of features without information loss.

4. Experiments

In this section we first discuss implementation details of the experiments, then elaborate the ablation study specifying contributions of different components proposed in SwiftNet. Comparisons with other state-of-the-art methods on DAVIS 2016 & 2017 and YouTube-VOS datasets are provided as follows, where SwiftNet demonstrates the best overall segmentation accuracy and inference speed. All experiments are implemented in PyTorch [21] on 1 NVIDIA P100 GPU. Particularly, SwiftNet adopting both ResNet-18 and ResNet-50 [5] backbones are experimented to show the favorable compatibility and efficacy of our method. We employ a decoder similarly constructed as in STM [20], which adopts three refinement modules to gradually restore the spatial scale of the segmentation mask.

4.1. Datasets and Evaluation Metrics

DAVIS 2016 & 2017. DAVIS 2016 dataset contains in total 50 single-object videos with 3455 annotated frames. Considering its confined size and generalizability, it is soon supplemented into DAVIS 2017 dataset comprising 150 sequences with 10459 annotated frames, a subset of which exhibit multiple objects. Following the DAVIS standard, we utilize mean Jaccard \mathcal{J} index and mean boundary \mathcal{F} score, along with mean $\mathcal{J}\&\mathcal{F}$ to evaluate segmentation accuracy. We adopt the Frames-Per-Second (FPS) metric to measure segmentation speed.

Method	w/o pixel-wise		w pixel-wise	
	J&F	FPS	J&F	FPS
low-level	78.0	22	77.5	51
high-level	75.4	37	73.6	71
LAE	78.2	35	77.8	70

Table 1. Ablation study of LAE on Davis 2017 validation set.

YouTube-VOS. Being the largest dataset at the present, YouTube-VOS encompasses totally 4453 videos annotated with multiple objects. In particular, its validation set possesses 474 sequences covering 91 object classes, 26 of which are not visible in the training set, and thus facilitating evaluations *w.r.t.* seen and unseen object classes to reflect method generalizability. On YouTube-VOS we report $\mathcal{J}\&\mathcal{F}$ for accuracy assessment, the overall score \mathcal{G} is generated by averaging $\mathcal{J}\&\mathcal{F}$ on seen and unseen classes.

4.2. Training and Inference

4.2.1 Training

SwiftNet is first pre-trained on simulated data generated upon MS-COCO dataset [15], then finetuned on DAVIS 2017 and YouTube-VOS Dataset respectively. In both training stages, input image size is set to 384×384 , and we adopt Adam optimizer with learning rate starting at $1e-5$. The learning rate is adjusted with polynomial scheduling using the power of 0.9. All batch normalization layers in the backbone are fixed at its ImageNet pre-trained value during training. We use batch size of 4, which is realized on 1 GPU via manual accumulation.

MS-COCO Pre-train. Considering the scarcity of video data and to ensure the generalizability of SwiftNet, we perform pre-training on simulated video clips generated upon MS-COCO dataset [16]. Specifically, we randomly crop foreground objects from a static image, which are then pasted onto a randomly sampled background image to form a simulated image. Affine transformations such as rotation, resizing, sheering, and translation are applied to foreground and background separately to generate deformation and occlusion, and we maintain an implicit motion model to generate clips with length of 5. SwiftNet is trained with simulated clips for 150000 iterations and the $\mathcal{J}\&\mathcal{F}$ reaches 65.6 on DAVIS 2017 validation set, which demonstrates the efficacy of our simulated pre-training.

DAVIS 2017 & YouTube-VOS Finetune. After pre-training, we finetune SwiftNet on DAVIS 2017 and YouTube-VOS training set for 200000 iterations. At each iteration, we randomly sampled 5 images consecutively (with random skipping step smaller than 5 frames) and estimate corresponding segmentation masks one after another. Pixel-wise memory update and match are executed on every frame within the 5-frame clip.

	Metric	periodical	variation-aware
		sampling (5)	trigger
full frame	J&F	78.2	78.1
	FPS	35	52
pixel-wise update & match	J&F	77.8	77.8
	FPS	65	70

Table 2. Ablation study of PAM on Davis 2017 validation set.

4.2.2 Inference

Given a test video accompanied by its first frame annotation mask, at inference time we frame-by-frame segment the video using SwiftNet. Particularly, memory at the first frame, M_0 , is initialized with feature maps output by the encoder given first frame image and mask, then it is updated online throughout the inference. At frame t , we utilize memory M_{t-1} and frame image I_t to compute segmentation mask m_t with SwiftNet. If frame t is triggered, m_t is feed into the LAE and to update the memory for further computations.

4.3. Ablation Study

Ablation study is conducted on DAVIS 2017 validation set to show contributions of different SwiftNet modules.

4.3.1 Light-Aggregation Encoder

To demonstrate the efficacy of the proposed LAE, we additionally develop two baseline reference encoders for comparison. The first baseline instantiates low-level aggregation as adopted in STM [20], where mask produced by the last frame is directly concatenated with raw image. This encoder maintains high-resolution mask but requires two separate encoders for reference and query frames, hence heavier model size. The second baseline implements high-level aggregation following CFBI [40], where segmentation mask is first down-sampled to the minimal feature resolution and then fused with query feature for foreground discovery. This baseline enables encoder reuse between reference query frames, but spatial details of the mask are lost during pooling-based down-samplings. As shown in Table 1, low-level baseline reveals better accuracy while high-level baseline runs faster, conforming to the fact that the low-level one involves more sophisticated feature aggregations between image and mask. Notably, LAE surpasses the low-level baseline in both $\mathcal{J}\&\mathcal{F}$ and FPS (by 19), and outperforms the high-level baseline by 4.2% in $\mathcal{J}\&\mathcal{F}$ while keeping comparable FPS. This results strongly suggest that LAE promotes thorough mask-frame aggregation and elevates segmentation speed.




Figure 6. Qualitative results of SwiftNet (ResNet-18) generated on DAVIS17 validation set.




Figure 7. \mathcal{J} & \mathcal{F} and FPS evolution *w.r.t.* update ratio β .

4.3.2 Pixel-Adaptive Memory

In this section we showcase the efficacy of PAM in elevating accuracy and speed. Table 2 row-wise illustrates the contribution of pixel-wise memory update and match in eliminating spatial redundancy, where it significantly boosts processing speed by 30 and 18 PFS in both temporal strategies, and only experience around 0.4% drop in \mathcal{J} & \mathcal{F} . Column-wise reveals the contribution of variation-aware trigger in compressing temporal redundancy, where it raises segmentation speed by 17 and 5 FPS in both spatial strategies, and at most 0.1% \mathcal{J} & \mathcal{F} is reported. Notably here we experiment with periodic sampling at a pace of 5 frame, which is tested to be the optimal parameter as in [20]. To provide a up-closer view of PAM, in Fig. 7 we illustrate the variation of \mathcal{J} & \mathcal{F} and FPS *w.r.t.* different spatial update ratio β and temporal trigger strategy. As shown in green, \mathcal{J} & \mathcal{F} increases in accordance with enlarged β , *i.e.* segmentation accuracy will grow if more percentage of pixels are updated. It is worth noting that, $\beta = 10\%$ yields the best accuracy while larger value shows no significant improvement. Besides, temporal trigger brings minute effect in accuracy. The blue color draws variations *w.r.t.* FPS, where

larger β steadily decreases FPS, and variation-aware trigger constantly increases FPS in under different β . Notably, the gap between blue curves are enlarged with larger β , echos that more spatiotemporal redundancy are compressed by the trigger during heavy spatial update.

4.4. State-of-the-art Comparison

4.4.1 DAVIS 2017

Comparison results on Davis 2017 validation set are listed in Table 3. As shown, both SwiftNet versions demonstrate better \mathcal{J} & \mathcal{F} , \mathcal{J} , and \mathcal{F} scores than all other real-time methods by a large margin. In particular, SwiftNet with ResNet-18 runs the fastest at 70 FPS, outperforming the second fastest SAT-fast [3] in \mathcal{J} & \mathcal{F} by 8.3%. This considerable lead is because that SAT updates global feature with cropped regions containing heavy background noise, while SwiftNet updates memory with useful and discriminative pixels and filters out redundant and noise regions. SwiftNet with ResNet-50 not only meets real-time requirement, but also reaches 81.1 in \mathcal{J} & \mathcal{F} score, which ranks the second best in both real-time and slow methods. STM [20] reports the best \mathcal{J} & \mathcal{F} at 81.8, which is 0.7% better than ours, while we run almost 4 times faster than STM. This significant improvement of SwiftNet is achieved by explicitly compressing spatiotemporal redundancy resides in STM, which adopts heavy periodical sampling and full-frame matching. In addition, GCNet [13] also strives to accelerate memory-based VOS by designing light-weight memory reading and writing strategies. As shown, it runs at comparable speed with our ResNet-50 version, while we exceeds GCNet in term of \mathcal{J} & \mathcal{F} by 9.7%. Fig 6 shows qualitative results on DAVIS17 validation set produced by SwiftNet with ResNet-18. The first row demonstrates that SwiftNet is robust against deformation, the second to the fourth row reveal that SwiftNet is highly capable of handling fast motion, similar distractor, and tremendous occlusion, respectively.

Method	OL	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	FPS
PReMVOS [17]	✓	77.8	73.9	81.7	0.01
CINM [1]	✓	67.5	64.5	70.5	0.01
OnAVOS [31]	✓	67.9	64.5	70.5	0.08
OSVOS [2]	✓	60.3	56.7	63.9	0.22
OSVOS-s[19]	✓	68.0	64.7	71.3	0.22
STCNN [36]	×	61.7	58.7	64.6	0.25
FAVOS [4]	×	58.2	54.6	61.8	0.56
FEELVOS [29]	×	71.5	69.1	74.0	2.2
Dyenet [12]	✓	69.1	67.3	71.0	2.4
STM [20]	×	81.8	79.2	84.3	6.3
Fasttan [9]	×	75.9	72.3	79.4	7
RGMP [35]	×	66.7	64.8	68.6	7.7
Fasttmu [26]	×	70.6	69.1	72.1	11
AGAME [11]	×	70.0	67.2	72.7	14
FRTM-VOS [23]	✓	76.7	-	-	22
<hr/>					
GCNet [13]	×	71.4	69.3	73.5	25
RANet [34]	×	65.7	63.2	68.2	30
SiamMask [32]	×	56.4	64.3	58.5	35
TVOS [44]	×	72.3	69.9	74.7	37
SAT [3]	×	72.3	68.6	76.0	39
FRTM-VOS-fast [23]	✓	70.2	-	-	41
SAT-fast [3]	×	69.5	65.4	73.6	60
<hr/>					
SwiftNet(ResNet-50)	×	81.1	78.3	83.9	25
SwiftNet(ResNet-18)	×	77.8	75.7	79.9	70

Table 3. Quantitative results on DAVIS 2017 validation set. In all following tables, OL denotes online learning and real-time methods reside below the horizontal line.

4.4.2 DAVIS 2016

Results on DAVIS 2016 dataset is shown in Table. 4. Since DAVIS 2016 only contains single-object sequences, most methods experience considerable performance gains when transferred from DAVIS 2017, and the accuracy gap between ResNet-18 and ResNet-50 SwiftNet is reduced because the demands for highly semantical features are alleviated. It is worth noting that, SwiftNet with both ResNet-18 and ResNet-50 outperform all other methods in segmentation accuracy, where the ResNet-50 version leads the second best STM by 1.1% and 18.7 in terms of $\mathcal{J}\&\mathcal{F}$ and FPS.

4.4.3 Youtube-VOS

As testing on the large YouTube-VOS validation set is time-consuming, here we show comparison results with most representative methods. Besides, considering the varied formulation of test sequences, we omit FPS readings on YouTube-VOS by default. As shown in Table 5, SwiftNet with ResNet-50 considerably outperform all other real-time methods in accuracy, leading the second best GCNet by 4.6% in term of overall score G , while running at a comparable speed with GCNet. SwiftNet with ResNet-18 performs comparably with GCNet, but runs significantly faster. Moreover, SwiftNet performs stably across seen and unseen classes, demonstrating its favorable generalizability.

Method	OL	$\mathcal{J}\&\mathcal{F}$	\mathcal{J}	\mathcal{F}	FPS
PReMVOS [17]	✓	86.8	84.9	88.6	0.01
OnAVOS [30]	✓	85.5	86.1	84.9	0.08
OSVOS [2]	✓	80.2	79.8	80.6	0.22
RANet+ [34]	✓	87.1	86.6	87.6	0.25
FAVOS [4]	×	80.8	82.4	79.5	0.56
FEELVOS [29]	×	81.7	81.1	82.2	2.2
Dyenet [12]	✓	-	86.2	-	2.4
STM [20]	×	89.3	88.7	89.9	6.3
Fasttan [9]	×	75.9	72.3	79.4	7
RGMP [35]	×	81.8	81.5	82.0	7.7
Fasttmu [26]	×	78.9	77.5	80.3	11
AGAME [11]	×	-	82.0	-	14
FRTM-VOS [23]	✓	83.5	-	-	22
<hr/>					
GCNet [13]	×	86.6	87.6	85.7	25
SiamMask [32]	×	70.0	71.7	67.8	35
SAT [3]	×	83.1	82.6	83.6	39
FRTM-VOS-fast [23]	✓	78.5	-	-	41
<hr/>					
SwiftNet(ResNet-50)	×	90.4	90.5	90.3	25
SwiftNet(ResNet-18)	×	90.1	90.3	89.9	70

Table 4. Quantitative results on DAVIS 2016 validation set.

Method	OL	G	\mathcal{J}_s	\mathcal{J}_u	\mathcal{F}_s	\mathcal{F}_u
RGMP [35]	×	53.8	59.5	45.2	-	-
OnAVOS [30]	✓	55.2	60.1	46.1	62.7	51.4
PReMVOS [17]	✓	66.9	71.4	56.5	75.9	63.7
OSVOS [2]	✓	58.8	59.8	54.2	60.5	60.7
FRTM-VOS [23]	✓	72.1	72.3	65.9	76.2	74.1
STM [20]	×	79.4	79.7	84.2	72.8	80.9
<hr/>						
SiamMask [32]	×	52.8	60.2	45.1	58.2	47.7
SAT [3]	×	63.6	67.1	55.3	70.2	61.7
FRTM-VOS-fast [23]	✓	65.7	68.6	58.4	71.3	64.5
TVOS [44]	×	67.8	67.1	63.0	69.4	71.6
GCNet [13]	×	73.2	72.6	68.9	75.6	75.7
<hr/>						
SwiftNet(ResNet-50)	×	77.8	77.8	72.3	81.8	79.5
SwiftNet(ResNet-18)	×	73.2	73.3	68.1	76.3	75.0

Table 5. Quantitative results on Youtube-VOS validation set. G denotes overall score, subscript s and u denote scores in seen and unseen categories.

5. Conclusion

We propose a real-time semi-supervised video object segmentation (VOS) solution, named SwiftNet, which delivers the best overall accuracy and speed performance. SwiftNet achieves real-time segmentation by explicitly compressing spatiotemporal redundancy via Pixel-Adaptive Memory (PAM). In PAM, temporal redundancy is reduced using variation-aware trigger, which adaptively selects incremental frames for memory update and ignores static ones. Spatial redundancy is eliminated with pixel-wise memory update and match modules, which abandon full-frame operations and incrementally process with temporally-varied pixels. Light-aggregation encoder is also introduced to promote thorough and expedite reference encoding. Overall, SwiftNet is effective and compatible, we hope it could set a strong baseline for real-time VOS.

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. **8**
- [2] 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. **1, 3, 8**
- [3] 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. **1, 3, 5, 7, 8**
- [4] 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. **1, 3, 8**
- [5] 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. **5**
- [6] 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. **1**
- [7] 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. **2**
- [8] Yuan-Ting Hu, Jia-Bin Huang, and Alexander Schwing. Maskrcnn: Instance level video object segmentation. In *Advances in neural information processing systems*, pages 325–334, 2017. **2, 3**
- [9] Xuhua Huang, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Fast video object segmentation with temporal aggregation network and dynamic template matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8879–8889, 2020. **1, 3, 8**
- [10] Suyog Dutt Jain, Bo Xiong, and Kristen Grauman. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, pages 2117–2126. IEEE, 2017. **2**
- [11] 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 Conference on Computer Vision and Pattern Recognition*, pages 8953–8962, 2019. **1, 3, 5, 8**
- [12] 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. **8**
- [13] Yu Li, Zhuoran Shen, and Ying Shan. Fast video object segmentation using the global context module. *arXiv preprint arXiv:2001.11243*, 2020. **1, 3, 7, 8**
- [14] Huaijia Lin, Xiaojuan Qi, and Jiaya Jia. Agss-vos: Attention guided single-shot video object segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3949–3957, 2019. **2**
- [15] 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. **6**
- [16] 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. **6**
- [17] Jonathon Luiten, Paul Voigtlaender, and Bastian Leibe. Premvos: Proposal-generation, refinement and merging for video object segmentation. In *Asian Conference on Computer Vision*, pages 565–580. Springer, 2018. **1, 8**
- [18] 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. **3**
- [19] 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. **8**
- [20] 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 International Conference on Computer Vision*, pages 9226–9235, 2019. **1, 2, 3, 4, 5, 6, 7, 8**
- [21] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. **5**
- [22] 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. **1, 3**
- [23] 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. **2, 4, 8**
- [24] W. Shi, J. Caballero, F. Huszár, J. Totz, A. P. Aitken, R. Bishop, D. Rueckert, and Z. Wang. Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network. In *2016 IEEE Conference*

- on *Computer Vision and Pattern Recognition (CVPR)*, pages 1874–1883, 2016. 5
- [25] 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. 3
- [26] Mingjie Sun, Jimin Xiao, Eng Gee Lim, Bingfeng Zhang, and Yao Zhao. Fast template matching and update for video object tracking and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10791–10799, 2020. 1, 3, 8
- [27] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning video object segmentation with visual memory. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4481–4490, 2017. 2
- [28] 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 Conference on Computer Vision and Pattern Recognition*, pages 5277–5286, 2019. 2, 3
- [29] 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 Conference on Computer Vision and Pattern Recognition*, pages 9481–9490, 2019. 1, 2, 3, 5, 8
- [30] Paul Voigtlaender and Bastian Leibe. Online adaptation of convolutional neural networks for the 2017 davis challenge on video object segmentation. In *The 2017 DAVIS Challenge on Video Object Segmentation-CVPR Workshops*, volume 5, 2017. 8
- [31] Paul Voigtlaender and Bastian Leibe. Online adaptation of convolutional neural networks for video object segmentation. *arXiv preprint arXiv:1706.09364*, 2017. 3, 8
- [32] 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 conference on computer vision and pattern recognition*, pages 1328–1338, 2019. 1, 2, 3, 5, 8
- [33] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. 1, 2
- [34] 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 international conference on computer vision*, pages 3978–3987, 2019. 1, 2, 3, 8
- [35] 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. 1, 2, 3, 5, 8
- [36] Kai Xu, Longyin Wen, Guorong Li, Liefeng Bo, and Qingming Huang. Spatiotemporal cnn for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1379–1388, 2019. 2, 8
- [37] 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. 3, 4
- [38] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5188–5197, 2019. 2
- [39] 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. 3
- [40] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. *arXiv preprint arXiv:2003.08333*, 2020. 2, 6
- [41] Rui Yao, Guosheng Lin, Shixiong Xia, Jiaqi Zhao, and Yong Zhou. Video object segmentation and tracking: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(4):1–47, 2020. 1
- [42] Xiaohui Zeng, Renjie Liao, Li Gu, Yuwen Xiong, Sanja Fidler, and Raquel Urtasun. Dmm-net: Differentiable mask-matching network for video object segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3929–3938, 2019. 2
- [43] Lu Zhang, Zhe Lin, Jianming Zhang, Huchuan Lu, and You He. Fast video object segmentation via dynamic targeting network. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5582–5591, 2019. 2
- [44] Yizhuo Zhang, Zhirong Wu, Houwen Peng, and Stephen Lin. A transductive approach for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6949–6958, 2020. 2, 3, 8
- [45] Qiang Zhou, Zilong Huang, Lichao Huang, Yongchao Gong, Han Shen, Wenyu Liu, and Xinggang Wang. Motion-guided spatial time attention for video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. 1, 2, 3, 4, 5
- [46] Zhishan Zhou, Lejian Ren, Pengfei Xiong, Yifei Ji, Peisen Wang, Haoqiang Fan, and Si Liu. Enhanced memory network for video segmentation. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 0–0, 2019. 1, 2, 3