

# Two-Stream Networks for Object Segmentation in Videos

Hannan Lu<sup>1</sup>, Zhi Tian<sup>2</sup>, Lirong Yang<sup>2</sup>, Haibing Ren<sup>2</sup>, Wangmeng Zuo<sup>1</sup>  
 {hannanlu, wmzuo}@hit.edu.cn, {tianzhi02, yanglirong, renhaibing}@meituan.com

## Abstract

Existing matching-based approaches perform video object segmentation (VOS) via retrieving support features from a pixel-level memory, while some pixels may suffer from lack of correspondence in the memory (i.e., unseen), which inevitably limits their segmentation performance. In this paper, we present a **Two-Stream Network (TSN)**. Our TSN includes (i) a pixel stream with a conventional pixel-level memory, to segment the *seen* pixels based on their pixel-level memory retrieval. (ii) an instance stream for the *unseen* pixels, where a holistic understanding of the instance is obtained with dynamic segmentation heads conditioned on the features of the target instance. (iii) a pixel division module generating a routing map, with which output embeddings of the two streams are fused together. The compact instance stream effectively improves the segmentation accuracy of the *unseen* pixels, while fusing two streams with the adaptive routing map leads to an overall performance boost. Through extensive experiments, we demonstrate the effectiveness of our proposed TSN, and we also report state-of-the-art performance of 86.1% on YouTube-VOS 2018 and 87.5% on the DAVIS-2017 validation split.

## 1 Introduction

Video Object Segmentation (VOS) is a fundamental task in video analysis, and has been widely applied in computer vision applications, such as video content editing and automatic driving. In recent years, semi-supervised VOS has been developed extensively and achieved great progress. Its goal is to separate specific object(s) from the background for each frame in a video sequence, given the mask of the target object(s) in the initial frame.

Top-performing works for semi-supervised VOS primarily exploit the template matching technique, in which labels are propagated from a reference set (i.e., the first frame with a given annotation and the historical segmented frames) to the query frame (i.e., the current frame) through matching. Inspired from the memory network [35], STM [21] constructs a pixel-level memory with multiple historical frames,




Figure 1: Visualization of segmenting with a pixel stream alone. In the forward pass, the pixel stream (top row) fails to segment the leg area correctly since it is unable to find a correspondence in the reference frame. Our two-stream network (bottom row) can address this issue effectively.

where pixel-level features are extracted and stored in the memory. Afterwards, the query pixels are segmented based on their retrieved reference features in the memory with their pixel-wise matching affinity. Based on STM, most follow-up works[27, 6, 28, 36, 2, 3, 40] are devoted to investigating solutions to obtain more accurate retrieval from pixel-level memory. However, these methods still have trouble in handling pixels not appeared in the reference frames (i.e., unseen), which lack a correspondence in the pixel-level memory. As shown in Figure 1, such approaches relying on pixel-level memory fail to segment the emerging leg area due to its absence in the reference.

In this paper, we propose a **Two-Stream Network (TSN)** to address this issue. TSN mainly includes a pixel stream and an instance stream, which are fused with a pixel division module. The pixel stream is responsible for the *seen* pixels as in previous methods, and these pixels can often be accurately segmented because it is able to find similar correspondences in the pixel-level memory. In contrast, the instance stream is mainly in charge of segmenting the *unseen* pixels, which lack a correspondence in the pixel-level memory. For segmentation, the instance stream generates segmentation heads composed of dynamical kernels, which are conditioned on the feature of each instance of the target object. Integrating the instance-level information endows the dynamic kernels in the instance stream with a holistic understanding of the instance and reasoning ability, which makes it more suitable for segmenting the *unseen* pixels. To obtain the final segmentation results, the features of the

two streams are weighted-aggregated with a routing map, which is generated with a proposed pixel division module and decide how much the a pixel should rely on the two streams, respectively. We evaluate our TSN on two popular VOS benchmarks, *i.e.*, DAVIS and YouTube, and report the new state-of-the-art results. We summarize the main contributions as follows,

- We propose a two-stream network for semi-supervised VOS, which mainly includes a pixel stream, an instance stream and a pixel division module to fuse them together.
- The proposed instance stream improves the segmentation accuracy of the unseen pixels, attributing to its dynamic convolution bank generated conditioned on the features of the target instance. While the pixel division module adaptively merges the outputs of the two streams through the routing map, which effectively improves the overall performance.
- TSN achieves the state-of-the-art results on two popular VOS benchmarks, *i.e.*, DAVIS and YouTube-VOS.

## 2 Related Work

**Semi-supervised video object segmentation.** Early semi-supervised video object segmentation (VOS) methods [1, 18] apply segmentation-by-detection scheme, where general segmentation models are fine-tuned on the first annotated frame to focus on specific targets. To tackle with the evolution of object appearance over time, OnAVOS [33] online updates the object-specific model along the video. For efficient inference, subsequent solutions [7, 22, 17] resort to motion cues for temporal propagation, some of which [17, 9] propagates object mask from the first frame to the succeeding frames leveraging optical flow.

Recently, more researchers tend to formulate the semi-supervised VOS task as a template matching problem, where labels are propagated from a reference set to the query frame through pixel-wise matching and retrieval. Prior works [29, 8] employs siamese structure network to extract features from both template frames and the query frame. VideoMatch [8] updates the template feature set with segmented frames. Under the assumption of temporal continuity, FEELVOS [32] narrows the matching area with the latest frame to a local range. CFBI [39] further promotes FEELVOS [32] by establishing template sets for foreground and background respectively. Utilizing memory network [35, 4, 26], STM [21] constructs spatio-temporal memory with multiple segmented frames. The success of STM greatly encourages the development of matching based solutions [16, 11, 12, 34, 27, 36, 6, 28]. EGMN [16] develops a graph structure memory network, where memory

reading and writing are performed on each node sequentially. STCN [2] improves matching accuracy applying the negative squared Euclidean distance as similarity metric. Later works [3, 40, 19] introduce transformer [31] to VOS, where matching and retrieval are carried out multiple times with multiple attention layers. Joint [19] complements pixel-level retrieval with an online updated discriminative branch additionally.

The aforementioned solutions classify pixels to foreground/background based on their pixel-level retrieval. However, pixels not appeared in the past frames emerge along the video, which may be unable to find correspondence in the reference. We propose to disentangle the query pixels with a routing map calculated based on the residual of the feature of query and their corresponding retrieval. The value of residual indicates whether pixels can find similar correspondence in the reference.

**Instance-level segmentation.** Through global pooling the feature of all foreground pixels, CFBI and RPCM [39, 38] generate object-level embedding for channel re-weighting on query feature. For instance segmentation, CondInst [30] generates dynamic segmentation heads for each instance. In this work, we construct an instance-level memory with segmentation heads dynamically generated integrating each instance of the target object.

## 3 Our Approach

### 3.1 Overview of the Two-Stream Network

Given a video sequence of length  $T$  and the mask annotation of its first frame, semi-supervised VOS aims to segment the subsequent frames sequentially. For the current frame  $I^t$ , *i.e.*, query frame, a reference set  $\mathcal{S} = \{(I^i, M^i)\}_{i=1}^N$  is usually constructed with  $N$  historical frames and their corresponding masks. Under the semi-supervised VOS setting, we propose a Two-Stream Network (TSN). As illustrated in Figure 2, the proposed TSN includes an instance stream, a pixel stream, a pixel division module and a fusion module. The pixel stream is deployed to segment the query frame with a conventional pixel-level memory, in which the reference features are retrieved through affinity calculation and feature aggregation with features of query frame. The instance stream is specially designed to tackle with those *unseen* pixels which lack correspondance in the memory. Toward this end, we construct the segmentation heads composed of dynamically generated kernels. In order to determine the *unseen* pixels, we present the pixel division module. It generates a routing map based on the difference between the aggregated feature and the feature of the query frame, which is then leveraged to select pixels without correspondance in the memory. Finally, the fusion module merges the output embeddings of the two streams with the routing map, and

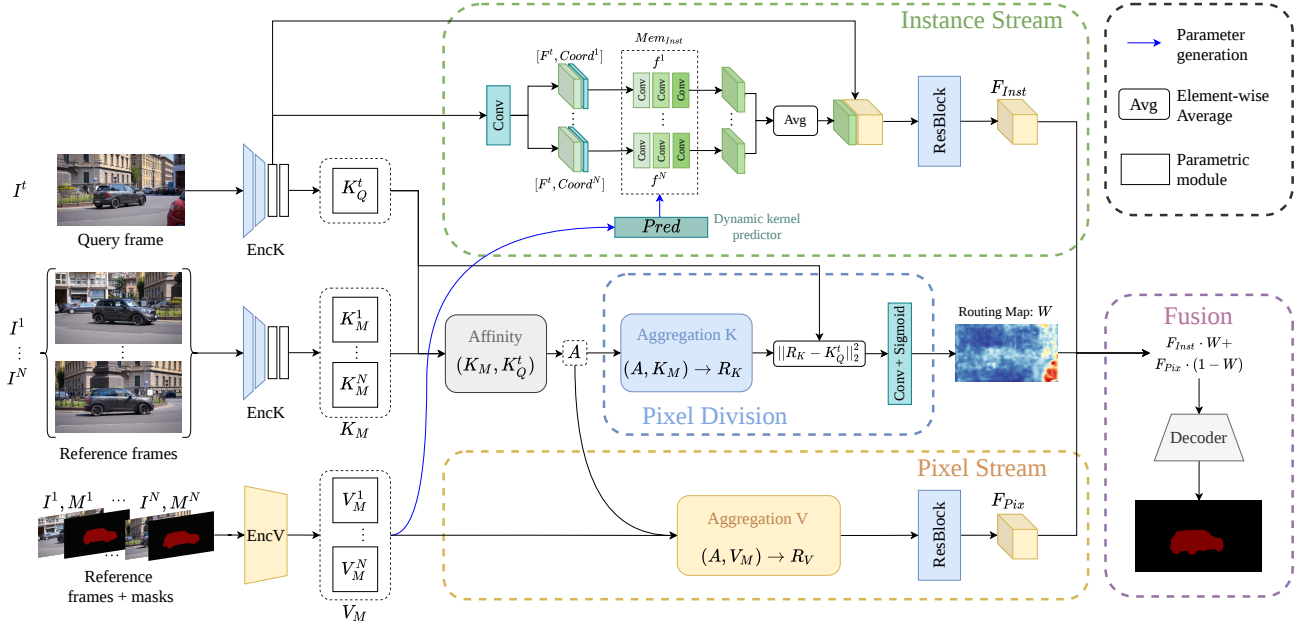


Figure 2: The overview of Two-Stream Network (TSN), which consists of an instance stream, a pixel stream, a pixel division module and a fusion module. The instance stream and the pixel stream process the query frame respectively. The pixel division module generates a routing map to separate the seen and the unseen pixels, and the fusion module merges the output embeddings of the two streams weighed on the routing map.

decodes to a segmentation result. In the following sections, we describe each module in details.

### 3.2 Pixel Stream

The pixel stream is responsible for segmenting the *seen* pixels, for they are able to retrieve reliable correspondences from a pixel-level memory to support segmentation. Following STCN [2], a conventional pixel-level memory  $\{K_M, V_M\}$  is constructed with reference frames and their masks in the pixel stream. For memory retrieval, a pixel-wise affinity is calculated and then used to aggregate reference features. Finally, A mask embedding carrying label information is generated.

Specifically, key encoder  $EncK$  is applied to encode the query frame  $I^t$  and the reference frames into key features  $K_Q = \{K_Q(p)\} \in \mathbb{R}^{HW \times C_k}$  and  $K_M = \{K_M(q)\} \in \mathbb{R}^{NHW \times C_k}$ , where  $p$  and  $q$  are the spatial locations,  $N$  is the number of frames in the reference set. Spatial dimensions  $H = sH_{im}$  and  $W = sW_{im}$ , where  $H_{im}$  and  $W_{im}$  corresponds to the spatial dimension of the image and  $s$  to the stride of the backbone in TSN. The affinity matrix  $A \in \mathbb{R}^{HW \times NHW}$  is then calculated through non-local matching between the query key and the reference key. For each  $p$  and  $q$ , the affinity is calculated as,

$$A(p, q) = \frac{\exp(\mathcal{C}(K_Q(p), K_M(q)))}{\sum_q \exp(\mathcal{C}(K_Q(p), K_M(q)))} \quad (1)$$

$\mathcal{C}$  is a similarity function, where we adopt negative squared Euclidean distance as STCN [2]. Support feature  $R_V \in \mathbb{R}^{HW \times C_v}$  is aggregated through weighted sum of the value feature  $V_M$  of reference frames with the matching affinity  $A$ ,

$$R_V(p) = \sum_q (A(p, q) \cdot V_M^M(q)) \quad (2)$$

The aggregated feature  $R_V$  is then concatenated with  $F^t$  and processed with a residual block, obtaining the mask embedding  $F_{Pix} \in \mathbb{R}^{HW \times C_v}$  of the pixel stream.

### 3.3 Instance Stream

To address the issue of the *unseen* pixels lacking reliable correspondence in the pixel-level memory, we design the instance stream. Segmentation heads  $\{f^1, \dots, f^N\}$  composed of dynamic convolutions are generated by a kernel predictor  $Pred$ , conditioned on the value feature  $\{V_M^1, \dots, V_M^N\}$  of each instance of the target object. The deep feature of the query frame is processed by all heads, and the results are averaged element-wisely to generate the final mask embedding.

Particularly, an instance-level memory  $Mem_{inst} = \{f^i\}, i = 1, 2, \dots, N$  composed of  $N$  segmentation heads is constructed. Each segmentation head is composed of three  $1 \times 1$  convolution layers with channels of 8, and each layer is activated by a ReLU function except for the last one. The

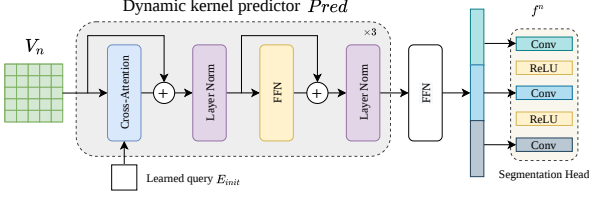


Figure 3: Details of the dynamic kernel predictor  $Pred$ , which generates parameters for each segmentation head dynamically.

parameters of each segmentation head are dynamically generated from an instance of the target object, characterizing the instance-level cues of the target object, such as the overall appearance and the spatial size. In practice, we do not generate segmentation heads for the reference frames where objects do not appear. A dynamic kernel predictor  $Pred$  is deployed to generate parameters, taking as input a learnable embedding  $E_{init} \in \mathbb{R}^{1 \times C_v}$  as well as the value feature of the  $n$ -th reference frame  $V_M^n \in \mathbb{R}^{HW \times C_v}$ , and outputs a vector  $\theta \in \mathbb{R}^{1 \times K}$  for each segmentation head. The value feature  $V_M^n \in \mathbb{R}^{HW \times C_v}$  is encoded by value encoder  $EncV$ , taking as input the reference frame  $I^n$  and its corresponding mask  $M_n$ . As shown in Figure 3,  $Pred$  is composed of three transformer layers as [31], where  $E_{init}$  learns to gather the object information from  $V_M^n$  adaptively. The output vector  $\theta \in \mathbb{R}^{1 \times K}$  can be seen as the parameters of all three convolution layers in a segmentation head concatenated together.

The instance stream segments the query frame  $I^t$  with the constructed instance-level memory, taking as input its deep feature  $F^t \in \mathbb{R}^{HW \times C_v}$ , and outputs mask embedding  $F_{Inst} \in \mathbb{R}^{HW \times C_v}$ . The  $N$  segmentation heads in  $Mem_{Inst}$  segment  $F^t$  respectively and average their results together as follows,

$$O_{Inst} = \frac{1}{N} \cdot \sum_i f^i([w(F^t), Coord_i]) \quad (3)$$

where  $w$  is a linear layer used to reducing the channel of  $F^t$  from 512 to 8, and  $[, ]$  represents concatenation along the channel dimension.  $Coord_i$  is a relative coordinate map taking the centroid of the object in reference frame  $I^n$  as origin. Through concatenating the  $Coord_i$ , the segmentation heads are capable of memorizing the position and size information of an instance, increasing their discriminative ability. The concatenation of  $O_{Inst}$  and  $F^t$  then goes through a residual block, obtaining the mask embedding  $F_{Inst} \in \mathbb{R}^{HW \times C_v}$ .

The previous solutions [39, 38] generate a single head averaging all the reference frames. We find it may degrade the discrimination ability of the segmentation head, since that it is extremely difficult to memorize varying appearances and spatial information of an object in one head. Additionally, the computational costs are practically the same with TSN,

since the generation of a single head in [38] requires to attend to a reference feature of size  $NHW \times C$  every time. We provide the quantitative comparison of these two solutions in Table 3.

### 3.4 Pixel Division and Fusion

The pixel division module is deployed to classify the query pixels into *seen* and *unseen*. A routing map is generated by the pixel division module based on the difference between the aggregates key features and the feature of the query pixels.

Specifically, feature  $R_K = \{R_K(p)\} \in \mathbb{R}^{HW \times C_k}$  is aggregated through weighted sum of the key feature  $K_M$  with the affinity matrix  $A$  as follows,

$$R_K(p) = \sum_q (A(p, q) \cdot V^K(q)) \quad (4)$$

Then, the routing map  $W = \{W(p)\} \in \mathbb{R}^{HW \times 1}$  can be generated as,

$$W = \mathcal{S}(\mathcal{F}(D)) \quad (5)$$

where  $\mathcal{F}$  is a linear layer,  $\mathcal{S}$  is the sigmoid function, and  $D$  is the  $L2$  function calculated as  $\|R_K - K_Q\|^2$ . Through the value of  $W$ , the *unseen* pixels can be distinguished from the *seen* pixels, since they tend to obtain larger residual for lacking of similar correspondence in reference.

The fusion module merges the output embeddings of the two streams weighted on the routing map, and decodes to a final segmentation mask. Fusion with the routing map can ensure that the segmentation results of the *seen* and the *unseen* pixels come from the pixel stream and the instance stream accordingly. The fusion process follows,

$$F = W \cdot F_{Inst} + (1 - W) \cdot F_{Pix} \quad (6)$$

At the end, a decoder of FPN structure is adopted to decode  $F$  to the final mask  $M^t$ .

## 4 Experiments

To verify the effectiveness of the proposed Two-Stream Network (TSN), we train a TSN-R50 which uses ResNet-50 [5] as backbone following STCN [2], and a TSN-SwinB using the strong swin transformer [14] as backbone. We evaluate our two models on two VOS benchmarks, DAVIS [25, 23] and YouTube-VOS [37]. Ablation studies are performed on the challenging YouTube-VOS [37] for in-depth analysis. In the following sections, we first describe the train and evaluation settings and then compare our TSNs with the state-of-the-art semi-supervised VOS methods. Finally, the ablation studies are presented.



Validation 2018 Split					
Method	$\mathcal{G}$	$\mathcal{J}_s$	$\mathcal{J}_u$	$\mathcal{F}_s$	$\mathcal{F}_u$
PreMVOS	66.9	71.4	56.5	75.9	63.7
A-GAME	66.1	67.8	60.8	-	-
STM	79.4	79.7	72.8	84.2	80.9
CFBI	81.4	81.1	75.3	85.8	83.4
RMNet	81.5	82.1	75.7	85.7	82.4
LCM	82.0	82.2	75.7	86.7	83.4
SST	81.7	81.2	76.0	-	-
HMMN	82.6	82.1	76.8	87.0	84.6
JOINT	83.1	81.5	78.7	85.9	86.5
STCN	83.0	81.9	77.9	86.5	85.7
AOT	83.8	82.9	77.7	<u>87.9</u>	86.5
RPCM	84.0	83.1	78.5	87.7	86.7
TSN-R50(Ours)	<u>84.8</u>	<u>83.6</u>	<u>79.8</u>	<u>87.9</u>	<u>87.8</u>
TSN-SwinB(Ours)	<b>86.1</b>	<b>85.1</b>	<b>80.6</b>	<b>89.7</b>	<b>89.1</b>

Validation 2019 Split					
Method	$\mathcal{G}$	$\mathcal{J}_s$	$\mathcal{J}_u$	$\mathcal{F}_s$	$\mathcal{F}_u$
CFBI	81.0	80.6	75.2	85.1	83.0
SST	81.8	80.9	76.6	-	-
HMMN	82.5	81.7	77.3	86.1	85.0
JOINT	82.8	80.8	79.0	84.8	86.6
STCN	84.2	82.6	79.4	87.0	87.7
AOT	83.7	82.8	78.0	<u>87.5</u>	86.7
RPCM	83.9	82.6	79.1	86.9	87.1
TSN-R50(Ours)	<u>84.6</u>	<u>83.1</u>	<u>80.2</u>	87.2	<u>87.8</u>
TSN-SwinB(Ours)	<b>86.0</b>	<b>84.8</b>	<b>80.9</b>	<b>89.1</b>	<b>89.2</b>

Table 1: Qualitative comparison with different methods on YouTube-VOS. Subscripts  $S$  and  $U$  represents the seen and the unseen category respectively.

#### 4.1 Datasets and Evaluation Metrics

DAVIS has two versions, DAVIS 2016 [23] contains 20 videos for validation, where each video has one annotated object instance. DAVIS 2017 [25] is a multi-object extension of DAVIS 2016, which provides 60 videos for training and 30 videos for validation. YouTube-VOS 2018 [37] is a large-scale and challenging dataset for video object segmentation, where the training set and validation set contains 3471 and 474 videos respectively. YouTube-VOS 2019 further adds additional videos to validation split.

For DAVIS 2016 and 2017, we report the mean of region similarity  $\mathcal{J}$ , contour accuracy  $\mathcal{F}$  and their average  $\mathcal{J\&F}$  for comparison, which are calculated with the standard DAVIS-2017 evaluation toolkit. For YouTube-VOS 2018 and 2019, we report  $\mathcal{J}$  and  $\mathcal{F}$  for both *seen* and *unseen* categories, and the averaged overall score  $\mathcal{G}$ , which are all obtained from the Codalab server.

Method	2017			2016		
	$\mathcal{J\&F}$	$\mathcal{J}$	$\mathcal{F}$	$\mathcal{J\&F}$	$\mathcal{J}$	$\mathcal{F}$
OnAVOS	67.9	64.5	71.3	85.5	86.1	84.9
OSVOS	59.2	56.6	61.8	80.2	79.8	80.6
RGMP	63.2	64.8	68.6	81.7	81.8	81.5
FAVOS	58.2	54.6	61.8	81.7	81.0	82.4
CINN	70.7	67.2	74.2	84.2	83.4	85.0
VideoMatch	62.4	56.5	68.2	81.9	81.0	80.8
PreMVOS	77.8	73.9	81.7	86.8	84.9	88.6
A-GAME	70.0	67.2	72.7	82.1	82.2	82.0
FEELVOS	71.6	69.1	74.0	82.2	81.7	88.1
STM	81.8	79.2	84.3	89.3	88.7	89.9
KMN	82.8	80.0	85.6	90.5	89.5	91.5
CFBI	81.9	79.1	84.6	89.4	88.3	90.5
RMNet	83.5	81.0	86.0	88.8	88.9	88.7
SST	82.5	79.9	85.1	-	-	-
HMMN	84.7	81.9	87.5	90.8	89.6	92
STCN	85.3	82.0	88.6	<u>91.7</u>	<u>90.4</u>	<u>93.0</u>
AOT	83.8	81.1	86.4	90.4	89.6	91.1
RPCM	83.7	81.3	86.0	90.6	87.1	<b>94.0</b>
TSN-R50(Ours)	<u>86.2</u>	<u>82.8</u>	<u>89.6</u>	91.0	90.1	91.8
TSN-SwinB(Ours)	<b>87.5</b>	<b>84.0</b>	<b>91.0</b>	<b>92.2</b>	<b>90.8</b>	<u>93.5</u>

Table 2: Qualitative comparison with different methods on DAVIS. Methods with  $FS$  represents tested with full resolution videos instead of 480p. **Bold** and underline indicate the best and the second-best performance respectively. R50 and SwinB represent adopting ResNet-50 [5] and SwinB[14] as backbone.

#### 4.2 Training and Inference

Follow STM [21], we first pre-train TSN on the synthetic dataset, then conduct main training on video dataset. Image datasets [13] are used to generate synthetic video clips for pre-train, where random affine transformations are adopted for sequence synthesis. Cut&Paste strategy is also adopted for data augmentation. Pre-train stage takes a total of  $2 \times 10^5$  iterations with a constant learning rate of  $1 \times 10^{-5}$ . The training split of DAVIS [25] and YouTube-VOS [37] are used for main training. For each training sample, three frames are randomly collected from a video sequence, augmented by random affine transformation with difference parameters. The temporal interval range of frame collecting increases from 1 to 25 in the first  $1 \times 10^4$  iterations, and decreases from 25 to 5 in the last  $5 \times 10^4$  epochs. The main training takes a total of  $2 \times 10^5$  iterations, where poly learning rate

decay with initial value of  $4 \times 10^{-5}$  is adopted. We use a batch size of 4 and a patch size of 432 in both per-train and main training stage. It takes about 32 hours to finish the two stage training with 4 Tesla A100 GPUs. The parameters of BatchNorm layers in both key and value encoders are frozen for the whole training. We adopt Adam [10] with standard momentum for the optimization of TSN-R50, and AdamW [15] with weight decay of  $1 \times 10^{-4}$  for TSN-SwinB. To stabilize training, Exponential Moving Average (EMA) [24] is used for TSN-SwinB. The overall loss function of TSN is a combination of bootstrapped cross-entropy loss and mask IoU loss [20]. Bootstrapped cross-entropy loss is calculated following STCN [2], and mask IoU loss is defined as,

$$\mathcal{L}_{mIoU}(P_i, G_i) = 1 - \frac{\sum_{p \in \Omega} \min(P_i^p, G_i^p)}{\sum_{p \in \Omega} \max(P_i^p, G_i^p)} \quad (7)$$

where  $P$  and  $G$  is the predicted mask and ground-truth mask of object  $i$ ,  $\Omega$  represents all pixels in mask  $P$  and  $G$ . Top- $k$  strategy with  $k = 20$  is adopted during inference as STCN [2].

### 4.3 Comparison with State-of-the-arts

The quantitative results of DAVIS 2016 and 2017 [23, 25] are shown in Table 2. On DAVIS 2017 validation set, which is the multi-object extension of DAVIS 2016, the TSN achieves a  $J\&F$  score of 87.5% when testing with 480p resolution videos, which is new state-of-the-art.

The quantitative results of YouTube-VOS 2018 and 2019 validation sets are presented in Table 1. Due to the large amount of test videos, YouTube-VOS poses a huge challenge to VOS approaches. Our TSR achieves state-of-the-art  $J\&F$  performance on both 2018 and 2019 split, where the averaging scores  $\mathcal{G}$  are 86.1% and 86.0%, respectively. Visual comparisons are provided in Figure 4. As demonstrated in video *6ae0cac484*, STCN [2] and RPCM [38] have errors segmenting the emerging person in the box. In contrast, our TSR segments the emerging area more precisely owing to the two-stream network. The segmentation ability of our TSR is also demonstrated in video *62bf7630b3* for successfully avoid the disturbance brought by the emerging object in the background. More visual results will be provided in the supplementary material for comparison.

### 4.4 Ablation study

To evaluate the contribution of each component in the proposed TSN, we conduct ablation study on the Youtube-2018 validation set. In the following, all comparison experiments are conducted based on the TSN-R50 for time efficiency, and models are trained with the video datasets only.

$N_{Head}$	$N_{Depth}$	Predictor	$J\&F$	$J_{seen}$	$J_{unseen}$
N	3	<i>Pred</i>	<b>83.8</b>	<b>82.8</b>	<b>78.4</b>
1	3	<i>Pred</i>	83.0	82.2	77.0
N	3	GAP	83.5	82.3	78.2
N	1	<i>Pred</i>	83.2	82.3	77.3

Table 3: Ablation on several design choices of the instance stream.  $N_{Head}$  represents the number of heads in the instance-level memory.  $N_{Depth}$  represents the depth of each head. GAP represents global average pooling.

Position map	$J\&F$	$J_{seen}$	$J_{unseen}$
-	83.0	82.1	77.5
Sine	83.5	82.3	77.6
Rel coord	<b>83.8</b>	<b>82.8</b>	<b>78.4</b>

Table 4: Ablation on the design choices for position map. Sine: the sine position encoding in [31]. Rel coord: relative coordinate map in 3.

**Instance stream.** We further investigate several design choices for the instance stream. Existing solutions [39, 38] utilize object-level information for VOS. In [39, 38], an object embedding is generated through global pooling on all the reference frames together, which is then used to re-weight the query feature. For comparison, we conduct ablation experiments on the constructed instance-level memory. Model1 utilizes the proposed instance-level memory, whose segmentation heads are generated from each instance of the target object separately with the proposed predictor. Model2 aggregates feature from all reference frames together generating one segmentation head. In model3, parameter predictor *Pred* is replaced with a global average pooling layer. In model4, the output of *Pred* is multiplied directly onto the query feature, rather than forming a segmentation head with multi-layer convolutions. As shown in Table 3, the  $J\&F$  of model2 drops from 83.7 to 83.0. We speculate that mixing all reference together may lead to information confusion, which damages the discrimination capacity of the segmentation heads. The  $J\&F$  also drops when replacing *Pred* with global average pooling, indicating the predictor gathering information with attention learns a better object representation. Model1 achieves better performance than model4, proving the superiority of the dynamically generated convolutions.

**Position map** When segmenting with the instance-level memory, we take as input the concatenation of the query feature and a relative position map, which provides a strong cue for segmentation. We also conduct ablation study to investigate the design choices of the position map. As shown in Table 4, removing position map leads to a significant performance drop. We conjecture that through the relative




Figure 4: Visual comparison of segmentation results on *6ae0cac484* and *62bf7630b3* of the challenging YouTube-VOS. We present the results of STCN [2] and RPCM [38] for comparison.

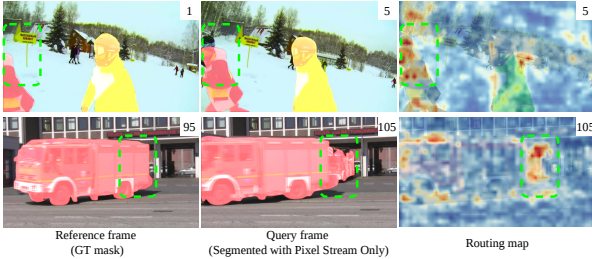


Figure 5: Visualization of the routing map. The area in the dashed box on the query frame is the unseen area in the reference frame. They have relatively large values on the routing map (heat area).

position, the segmentation head can better realize the spatial information of the object, such as spatial position and size. Replacing the relative position map with the absolute sine position encoding also decrease the performance slightly. We speculate that the normalized relative coordinates are easier to recognize than the sine position encoding, for each head has only 8 channel.

**Pixel division module.** Being a critical module in the two-stream network, pixel division module splits the query pixels into *seen* and *unseen* through a routing map. To investigate the dividing ability of the pixel division module, a visualization of the routing map is presented in Figure 5. Pixels that have been unseen in the reference frame appear in the query frame, while the pixel stream is unable to classify them correctly. On the routing map, the *unseen* pixels have relatively large values, which makes their segmentation results mainly determined by the instance branch in our TSN. Figure 6 plots the mean values of routing maps calculated

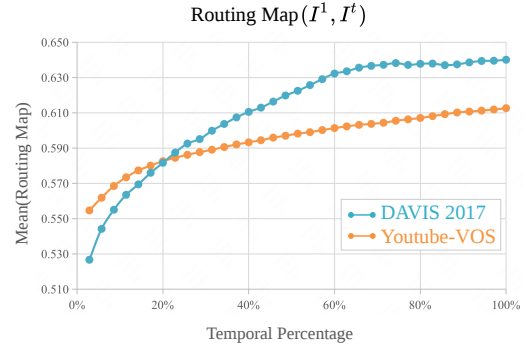


Figure 6: Mean value of routing maps, calculated with the difference between the  $t$ -th frame  $I^t$  and the first annotated frame  $I^1$ . The value of routing map increase over time, indicating that unseen pixels are gradually emerging along the video.

by the first annotated frame ( $I^1$ ) and each subsequent frame ( $I^t$ ). The mean values which characterize the differences between each two frames raise along the video.

**Fusion with Routing Map.** The output embeddings are fused weighted on the generated routing map. We provide quantitative results to verify the effectiveness of the proposed fusion strategy. Models with only instance stream, only pixel stream and combining the two streams with an equal weight are trained for comparison. As shown in Table 5, model fusing the two streams with the routing map achieves the best results. Performance of the model with only pixel stream drops since its lack the ability to deal with the *unseen* pixels. Combining the two streams with an equal weight also leads to performance drop, because they may interfere with each

Inst	Pix	$\mathcal{J}\&\mathcal{F}$	$\mathcal{J}_{seen}$	$\mathcal{J}_{unseen}$	FPS
$W$	$1-W$	<b>83.8</b>	<b>82.8</b>	<b>78.4</b>	17.2
0	1	82.6	82.1	76.5	18.6
0.5	0.5	82.8	81.5	77.2	17.2
1	0	55.4	62.8	45.5	<b>25.0</b>

Table 5: Ablation on the two stream of TSN. Inst and Pix represent the instance stream and the pixel stream, while their value indicates the way of combining the two streams. 1st: combining the two streams with the generated routing map. 2nd: only the pixel stream. 3rd: combining the two streams with equal weights, which are spatial-invariant. 4th: only the instance stream.

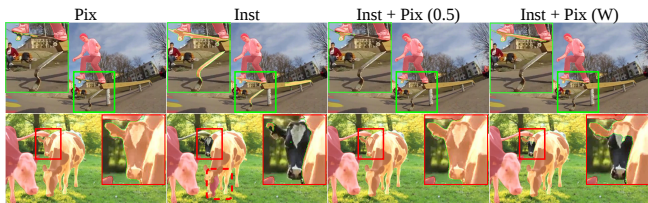


Figure 7: Visual comparison of the proposed two streams. Each stream owns its inherent characteristics. Simply fusing them with equal weights is not enough to combine their strengths.

other in areas where they are not skilled. The performance drops severely when leaving only instance stream, indicating TSR still heavily relies on the pixel-level cue. Visual comparison of the two streams are provided in Figure 7, presenting that the pixel stream and the instance stream owns their inherent characteristics. Simply fusing them with equal weights is not enough to combine their strengths.

## 5 Conclusion

We have presented a Two-Stream Network for video object segmentation task. Our work shows that it is greatly beneficial in VOS to separate the pixels of the query image into unseen and seen pixels with a routing map, and process them with the object-wise and pixel-wise segmentation module respectively. We also propose a novel dynamic kernel-based module to obtain the instance-level segmentation, which can encode the target object in its weights and efficiently and accurately segment it in the query image. Our framework achieves the new state-of-the-art performance on both DAVIS and the large-scale YouTube-VOS and we believe that the proposed framework is a simple and strong baselines for further research.

## References

- [1] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixe, D. Cremers, and L. Van Gool. One-shot video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [2] H. K. Cheng, Y.-W. Tai, and C.-K. Tang. Rethinking space-time networks with improved memory coverage for efficient video object segmentation. *Advances in Neural Information Processing Systems*, 34, 2021.
- [3] B. Duke, A. Ahmed, C. Wolf, P. Aarabi, and G. W. Taylor. Sstvos: Sparse spatiotemporal transformers for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5912–5921, 2021.
- [4] A. Graves, G. Wayne, and I. Danihelka. Neural turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- [5] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [6] L. Hu, P. Zhang, B. Zhang, P. Pan, Y. Xu, and R. 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.
- [7] Y.-T. Hu, J.-B. Huang, and A. Schwing. Maskrnn: Instance level video object segmentation. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [8] Y.-T. Hu, J.-B. Huang, and A. G. Schwing. Video-match: Matching based video object segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 54–70, 2018.
- [9] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele. Lucid data dreaming for video object segmentation. *International Journal of Computer Vision*, 127(9):1175–1197, 2019.
- [10] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [11] Y. Li, Z. Shen, and Y. Shan. Fast video object segmentation using the global context module. In *European Conference on Computer Vision*, pages 735–750. Springer, 2020.

- [12] Y. Liang, X. Li, N. Jafari, and J. Chen. Video object segmentation with adaptive feature bank and uncertain-region refinement. In *Advances in Neural Information Processing Systems*, volume 33, pages 3430–3441. Curran Associates, Inc., 2020.
- [13] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European conference on computer vision (ECCV)*, pages 740–755. Springer, 2014.
- [14] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.
- [15] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- [16] X. Lu, W. Wang, M. Danelljan, T. Zhou, J. Shen, and L. Van Gool. Video object segmentation with episodic graph memory networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, volume 12348, pages 661–679. Springer International Publishing, 2020.
- [17] J. Luiten, P. Voigtlaender, and B. Leibe. Premvos: Proposal-generation, refinement and merging for video object segmentation. In *Asian Conference on Computer Vision (ACCV)*, 2018.
- [18] K.-K. Maninis, S. Caelles, Y. Chen, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. Video object segmentation without temporal information. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(6):1515–1530, 2019.
- [19] Y. Mao, N. Wang, W. Zhou, and H. 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.
- [20] J. Mei, M. Wang, Y. Lin, Y. Yuan, and Y. Liu. Transvos: Video object segmentation with transformers. *arXiv preprint arXiv:2106.00588*, 2021.
- [21] S. W. Oh, J.-Y. Lee, N. Xu, and S. J. Kim. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 9226–9235, 2019.
- [22] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3491–3500, 2017.
- [23] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 724–732, 2016.
- [24] B. T. Polyak and A. B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization*, 30(4):838–855, 1992.
- [25] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbelaez, A. Sorkine-Hornung, and L. V. Gool. The 2017 DAVIS challenge on video object segmentation. *CoRR*, abs/1704.00675, 2017.
- [26] A. Santoro, S. Bartunov, M. Botvinick, D. Wierstra, and T. Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning (ICML)*, pages 1842–1850. PMLR, 2016.
- [27] H. Seong, J. Hyun, and E. Kim. Kernelized memory network for video object segmentation. In *European Conference on Computer Vision*, pages 629–645. Springer, 2020.
- [28] H. Seong, S. W. Oh, J.-Y. Lee, S. Lee, S. Lee, and E. Kim. Hierarchical memory matching network for video object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12889–12898, 2021.
- [29] J. Shin Yoon, F. Rameau, J. Kim, S. Lee, S. Shin, and I. So Kweon. Pixel-level matching for video object segmentation using convolutional neural networks. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, pages 2167–2176, 2017.
- [30] Z. Tian, C. Shen, and H. Chen. Conditional convolutions for instance segmentation. In *European Conference on Computer Vision*, pages 282–298. Springer, 2020.
- [31] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [32] P. Voigtlaender, Y. Chai, F. Schroff, H. Adam, B. Leibe, and L.-C. Chen. Feelvos: Fast end-to-end embedding



learning for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.

- [33] P. Voigtlaender and B. 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.
- [34] H. Wang, X. Jiang, H. Ren, Y. Hu, and S. Bai. Swiftnet: Real-time video object segmentation. *arXiv preprint arXiv:2102.04604*, 2021.
- [35] J. Weston, S. Chopra, and A. Bordes. Memory networks. *arXiv preprint arXiv:1410.3916*, 2014.
- [36] H. Xie, H. Yao, S. Zhou, S. Zhang, and W. Sun. Efficient regional memory network for video object segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [37] N. Xu, L. Yang, Y. Fan, J. Yang, D. Yue, Y. Liang, B. Price, S. Cohen, and T. Huang. Youtube-vos: Sequence-to-sequence video object segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 585–601, 2018.
- [38] X. Xu, J. Wang, X. Li, and Y. Lu. Reliable propagation-correction modulation for video object segmentation. *arXiv preprint arXiv:2112.02853*, 2021.
- [39] Z. Yang, Y. Wei, and Y. Yang. Collaborative video object segmentation by foreground-background integration. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [40] Z. Yang, Y. Wei, and Y. Yang. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34, 2021.