

LIP: Learning Instance Propagation for Video Object Segmentation

Ye Lyu
University of Twente
y.lyu@utwente.nl

George Vosselman
University of Twente
george.vosselman@utwente.nl

Gui-Song Xia
Wuhan University
guisong.xia@whu.edu.cn

Michael Ying Yang
University of Twente
michael.yang@utwente.nl

Abstract

In recent years, the task of segmenting foreground objects from background in a video, i.e. video object segmentation (VOS), has received considerable attention. In this paper, we propose a single end-to-end trainable deep neural network, convolutional gated recurrent Mask-RCNN, for tackling the semi-supervised VOS task. We take advantage of both the instance segmentation network (Mask-RCNN) and the visual memory module (Conv-GRU) to tackle the VOS task. The instance segmentation network predicts masks for instances, while the visual memory module learns to selectively propagate information for multiple instances simultaneously, which handles the appearance change, the variation of scale and pose and the occlusions between objects. After offline and online training under purely instance segmentation losses, our approach is able to achieve satisfactory results without any post-processing or synthetic video data augmentation. Experimental results on DAVIS 2016 dataset and DAVIS 2017 dataset have demonstrated the effectiveness of our method for video object segmentation task.

1. Introduction

Video object segmentation (VOS) aims at segmenting foreground objects from background in a video with coherent object identities. Such visual object tracking task serves for many applications including video analysis and editing, robotics and autonomous cars. Compared to the video object tracking task in bounding box level, this task is more challenging as pixel level segmentation is more detailed description of an object.

The VOS task is defined as a semi-supervised problem if ground truth annotations are given for the first several frames. It is otherwise an unsupervised problem if no annotation is provided. The ground truth annotations are masks that mark the objects that need to be tracked through the whole video. In our work, we focus on semi-supervised video object segmentation task, where the ground truth annotations are provided only for the first frame.

There are several challenges that make VOS a difficult task. First, both the appearance of the target objects and the background surroundings may change significantly over time. Second, there could be a large pose and scale variation over time. Third, there could be occlusions between different objects, which hinder the performance of tracking. Examples of the above three challenges are shown in Fig. 1. A notable and challenging dataset for the VOS task is the DAVIS 2016 dataset [43], which is designed for single-object video segmentation. Later the DAVIS 2017[44] is brought out focusing on segmentation of multiple video objects. Both of the datasets are provided with mask annotations of extremely high accuracy.

Most of the current methods for the VOS task, such as VPN [26], MSK [42] and RGMP [54], are based on the pixel level mask propagation. However, those methods fail to give a coherent label within an instance. In this paper, we introduce a single end-to-end trainable network to predict masks on instance level, namely the convolutional gated recurrent Mask-RCNN. It integrates instance segmentation network (Mask-RCNN [18]) with visual memory module (Conv-GRU [1]). Instance segmentation network is designed for foreground object segmentation, which is extended with visual memory for foreground object segmentation in a video. The incorporated visual memory helps to propagate information across frames to handle the appearance change, the pose and scale variation and the occlusions between objects. Our network gives a coherent label to a detected instance and assigns one label to only one detected instance. The model structure is shown in Fig. 2.

Our **Contributions** are:

- We propose a novel convolutional gated recurrent Mask-RCNN to learn instance propagation (**LIP**) for video object segmentation (VOS) task. Our model simultaneously segments all the target objects in the images.
- We design a single end-to-end trainable network for VOS task, enabling both mask propagation in the long term and bottom-up path augmentation.
- A strategy to successfully train the model for VOS task




Figure 1. Example predictions by our method from DAVIS [44, 43] dataset. Top row: Parkour sequence. An example of large appearance change over time. One every 20 frames shown of 100 in total. Middle row: Drift-straight sequence. An example of large scale and pose variation over time. One every 10 frames shown of 50 in total. Bottom row: Dogs-jump sequence. An example of occlusions between objects. One every 5 frames shown of first 20.

has been brought out. All the training processes are guided by the instance segmentation losses only.

2. Related work

In this section, we will discuss some relevant work.

Object detectors. Object detection starts with box level prediction and has a great improvement over the years. Single-stage detectors [45, 46, 36, 14, 33] have faster running speed while two-stage networks [15, 47] are more accurate in general. Later, Mask-RCNN [18] merges object detection with semantic segmentation by combining Faster-RCNN [47] and FCN [37], which form a conceptually simple, flexible yet effective network for instance segmentation task. Mask-RCNN network is suitable for instance segmentation on static images, but lacks the ability for temporal inference. Our work is to further extend Mask-RCNN with Conv-GRU module to solve video object segmentation task.

Recurrent neural networks (RNNs). RNNs [22, 48] are widely used for tasks with sequential data, such as image captioning [28], image generation [17] and speech recognition [16]. The key for RNNs is the hidden state, which selectively accumulates information from current input and the previous hidden state over time. However, RNN has its limitation as it fails to propagate information for a long sequence due to the problem of gradient vanishing or explosion in training [20, 40]. Two RNN variants, LSTM [21] and GRU [8] are more effective for the long term prediction by taking advantage of gating mechanism. To further encode spatial information, they are extended to Conv-LSTM [56] and Conv-GRU [1] respectively and have been used for video prediction [13] and action recognition [1].

Methods for VOS. Conv-GRU has already been used for video object segmentation. It serves as visual memory in [51] and has been proved to boost the performance for

VOS task. However, their model performs binary semantic segmentation only, which is not suitable for video object segmentation task with multiple objects.

VPN [26], MSK [42] and RGMP [54] learn to propagate mask for the VOS task. VPN utilizes learnable bilateral filters to achieve video-adaptive information propagation across frames. MSK learns to utilize both current frame and mask estimation from the previous frame for mask prediction. RGMP utilizes the first frame and mask as reference for instant information propagation besides the usage of current frame and previous mask estimation. Both MSK and RGMP achieve good results, but they can only propagate information for instances one by one.

Specially, OSVOS [3], OSVOS+ [39] and OnAVOS [52] tackle video object segmentation from static images, achieving temporal consistency as a by-product. They learn a general object segmentation model from image segmentation datasets and transfer the knowledge for video object segmentation. They all rely on additional post processing for better segmentation result. OnAVOS further applies online adaptation to continuously fine-tune the model, which is very time consuming.

[29] explores the benefits from in-domain training data synthesis with the labelled frames of the test sequences. [54] synthesizes video training data from static image dataset to add to limited video training samples. [25, 50] explore fast prediction without online training through matching based method. CINM [2] achieves good prediction by spatial-temporal post-processing based on results from OSVOS [3]. To handle the problem of long term occlusion, [31, 30] apply re-identification network to retrieve the missing objects, which complements their mask propagation methods. Recently, there are still many researches focusing on single-object video segmentation [55, 23, 9],

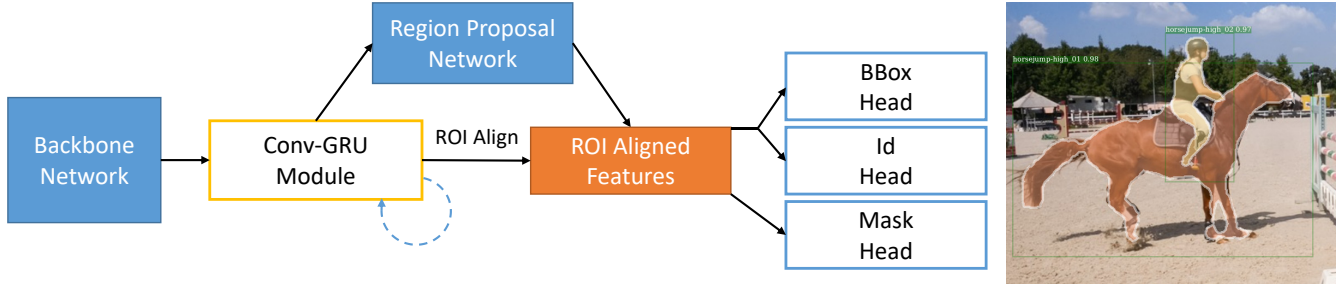


Figure 2. Overall model structure. The backbone network distills useful features from each input image. The features are then sent to Conv-GRU module (visual memory) for feature propagation. The output features from Conv-GRU module are utilized by region proposal network for proposal generation. Multiple heads finally take the ROI aligned features for video object segmentation. An example output is shown on the right, including bounding boxes, id predictions and object segments. The class of an instance is named by video sequence name plus object index.

which are not easily transformed for video segmentation of multiple objects. MaskRNN [24] is another method for instance level segmentation, but it only predicts for one instance at a time. The best results are achieved by ensemble of multiple specialized networks. PRemVOS [38] takes the 1st place of recent DAVIS2018 semi-supervised VOS task by utilizing complex pipeline with multiple specialized networks trained on multiple datasets.

3. Method

In this section, we first introduce the structure of our convolutional gated recurrent Mask-RCNN, which extracts and propagates information for multiple objects in a video. It is comprised of mainly three parts. They are the feature extraction backbone, the visual memory module and the prediction heads. The backbone network extracts features that are forwarded to visual memory module. The visual memory module then selectively remembers the new input features and forgets the old hidden states. On top of Conv-GRU, region proposal network (RPN), bounding box regression head, id classification head and mask segmentation head are constructed to solve the VOS task. The whole network is end-to-end trainable under the guidance of instance segmentation losses.

3.1. Mask-RCNN

Mask-RCNN [18] is one of the most popular framework designed for instance segmentation task. It is used for instance-wise object detection, classification and mask segmentation, which makes it naturally suitable for multiple video objects segmentation. Roles of different components in Mask-RCNN directly shift to fit VOS task as illustrated below.

Backbone: The backbone network still serves to extract features from images, but more focused on generating useful features adaptively for gates of Conv-GRU module. ResNet101-FPN [19, 32] with group normalization [53] is

used as our backbone network. Detailed structure is shown in Fig.3.

RPN: Mask-RCNN is known as a two stage instance segmentation network. Bounding boxes of general objects are proposed in the first stage, while classes and masks are predicted instance-wisely in the second stage. Such two stage framework adopts the same philosophy as the training stages of OSVOS [3]. For OSVOS, the network first learns to segment binary mask for general objects in a class-agnostic manner. Then it learns to segment specific objects during online training. In Mask-RCNN, RPN learns to reject background objects and to propose foreground objects in the first stage, which is also class-agnostic. It is in the second stage that classes and masks of different objects are determined.

Bounding box regression head: This branch is used to refine the bounding box proposals. Each predicted box contains one object. The boxes serve to separate different objects in an image.

Classification head: This branch is used to assign the object a correct class label. However, class type is unknown for VOS task. Instead, different objects are associated with different ids, which need to be predicted coherently in a video sequence. Classification branch is naturally transformed into an id classification branch.

Mask segmentation head: This branch is used to extract a mask for each foreground object in the image, which is the main target of VOS task.

Clearly, for the components in Mask-RCNN, there is a direct responsibility mapping from instance segmentation task to VOS task.

3.2. Convolutional gated recurrent unit

One difficulty for video object segmentation is the problem of long term dependency. The ground truth is provided only for the first frame, but the objects still need to be predicted after tens or hundreds of frames based on the

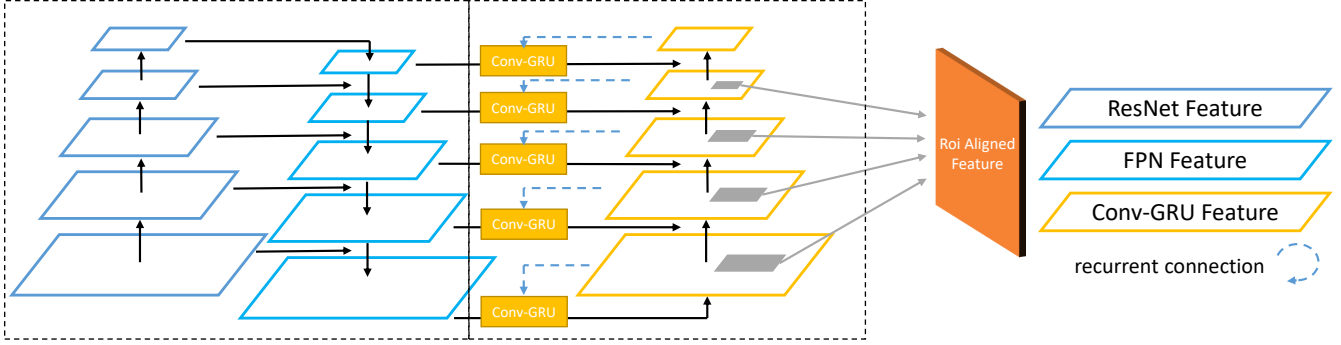


Figure 3. Model structure details. The left black dashed box shows the ResNet101-FPN backbone structure. The right black dashed box shows the Conv-GRU module. Our network brings bottom-up path augmentation for output features in Conv-GRU module. The augmented output features are used for both RPN and the prediction heads. All 5 layers are utilized for multi-level RPN, but only 4 bottom layers are used for multi-level ROIs.

ground truth from the first frame. The appearance of different objects in the videos may vary greatly and the objects sometimes get partially or even completely occluded, which makes coherent prediction more difficult.

In order to handle the above problem, we utilize the convolutional gated recurrent unit, serving as a visual memory to handle appearance morphing and occlusion. The memory module learns to selectively propagate the memorized features and to merge them with the newly observed ones. The key role for Conv-GRU module is to maintain a good feature over time for prediction of region proposal, bounding box regression, id classification and mask segmentation. Compared to the instance segmentation task, where each training batch is comprised of multiple randomly sampled images, the batch in temporal training has less variation as consecutive images from one sequence are highly correlated. This is similar to the problem of small batch size. To relieve such effect, we further replace the bias term in Conv-GRU with the group normalization (GN) layer, which are proved to give consistent performance across different batch sizes [53]:

$$z_t = \sigma(GN(W_{hz} * h_{t-1} + W_{xz} * x_t)) \quad (1a)$$

$$r_t = \sigma(GN(W_{hr} * h_{t-1} + W_{xr} * x_t)) \quad (1b)$$

$$\hat{h}_t = \Phi(GN(W_h * (r_t \odot h_{t-1}) + W_x * x_t)) \quad (1c)$$

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \hat{h}_t, \quad (1d)$$

where x_t is the input feature of time t , h_t is the hidden state of time t . z_t, r_t are update gate and reset gate respectively. W are convolutional filter parameters. σ and Φ are sigmoid function and tanh function respectively. $*$ and \odot denote the convolution operation and element-wise multiplication respectively.

For each level of the feature pyramid network [32] (FPN), we create a corresponding Conv-GRU layer. The layers at different levels learn different transition functions

for the hidden states. As bottom up path augmentation has been proved to be useful for instance segmentation [35], we easily achieve it by down-sampling and addition operation with output features from multi-level Conv-GRU module. The structure is shown in Fig. 3. The output features after path augmentation are used for RPN and prediction heads. Conv-GRU module is deliberately directly inserted after backbone network. In this way, information for both region proposal and instance prediction can be propagated through time.

3.3. Online inference

As our model predicts roi mask for each unique instance, there naturally exist constraints for prediction.

One maximum constraint. For each instance, there should be at most one object detected. This constraint is achieved by selecting highest id prediction score.

Location continuity constraint. If an instance is detected in the previous frame with high enough id prediction score, the location of the current detection should not be far from its previous location. To achieve this constraint, we suppress the prediction for the instance, whose boxes iou between consecutive frames is low.

As probability for id prediction decays over time, we further apply a very light weighted fine-tuning process for the last linear layer of the id head during online prediction. If there exists a target object detected with a high enough id prediction score, its predicted bounding box is set as ground truth for fine-tuning the id head only. By saving and reusing intermediate tensors, the speed for fine-tuning is fast.

4. Training the network

In this section, we will describe our training strategy in detail. The training modality for video object segmentation can be divided into offline training and online training [29, 42]. During offline training, the model is trained with the




Figure 4. Shortcut in prediction head. In order to let output from Conv-GRU module have more direct influence towards final prediction, we add a shortcut connection between ROI aligned feature and head logits by simple addition operation.

training set only. During online training, the model is fine-tuned with the first frame from the test set. As the class types of the test set are not known and objects may never be seen during offline training, online fine-tuning is necessary to help the model to generalize better for test set.

Our network needs both offline training and online training. During offline training stage, our network learns the features to differentiate all the object instances and learns to predict class-agnostic boxes and masks. During online training stage, our network is fine-tuned to differentiate objects for each test sequence and trained with boxes and masks in a class-specific manner.

4.1. Class-agnostic offline training

To provide our model with as much generality as possible, we apply class-agnostic training for bounding box and mask through the whole offline training process. Offline training for our model can be divided into two steps. First, our model is trained with instance segmentation dataset. This step is to provide our model with general object detection ability. Then, we train the model with video dataset to learn to propagate information over time for video object segmentation.

4.1.1 Pre-train on instance segmentation dataset

Pre-training on additional dataset is a common practice [52, 3, 39, 30]. We initialize our model by pre-training on Microsoft COCO dataset [34]. Ms-COCO dataset has been widely used for object detection task. It targets common objects in context with annotations including boxes, classes and masks. By first training on Ms-COCO dataset, our model learns to extract useful features for general object detection. As the training is on static images, we set hidden states to be zeros without update for Conv-GRU module.

After this step, our model gains general region proposal ability, general bounding box prediction ability and general object segmentation ability. Our model also learns to differentiate general objects by classes defined on Ms-COCO dataset.




Figure 5. Transforming class-agnostic weights to class-specific weights. During online fine-tuning, the class-agnostic bounding box and mask predictions are altered to class-specific. The rectangles are weights in the last linear layer of bbox head or the last convolutional layer of mask head. The grey color marks weights for background and the blue for foreground. Foreground weights are copied for each foreground instance to be fine-tuned uniquely.

4.1.2 Fine-tune on VOS dataset

In this stage, we train all the modules except the backbone network. By fine-tuning our model on video object segmentation dataset, the Conv-GRU module learns to tune its gates to best propagate information. It should be noted that the class number has changed as the video object segmentation dataset does not share the class definition with instance segmentation dataset. Instead, we replace the last linear layer right before softmax layer in the class prediction head with a new one, which now predicts the ids in the dataset. The class prediction head turns into an id prediction head.

The network is trained purely with instance segmentation losses. The different losses guide our model to have different abilities. The mask loss helps our model to propagate mask segmentation. The losses from id head and bbox head help our model to propagate information differently for each instance. Although the mask head and bbox head are trained in a class-agnostic manner, the id head and bbox head provide a chance to learn to propagate class-specific information.

To facilitate the information propagation, we further add a shortcut connection between the ROI aligned feature and the head logits as shown in Fig. 4.

4.2. Class-specific online fine-tuning

As the instances in test sequences are not the same as in training sequences, the last linear layer in id head needs to be re-initialized and trained to differentiate instances in the current sequence. We replace the last linear layer in the same way as in section 4.1.2. We also adopt focal loss [33] for id head to balance the training for multiple instances.

During online fine-tuning, the parameters in backbone network and Conv-GRU module are frozen to keep the learned propagation property. All other parts are fine-tuned for the new objects in test image. The class-agnostic prediction in mask head and bbox head are altered to be class-specific in order to have less competition for different instances. The process is illustrated in Fig. 5.

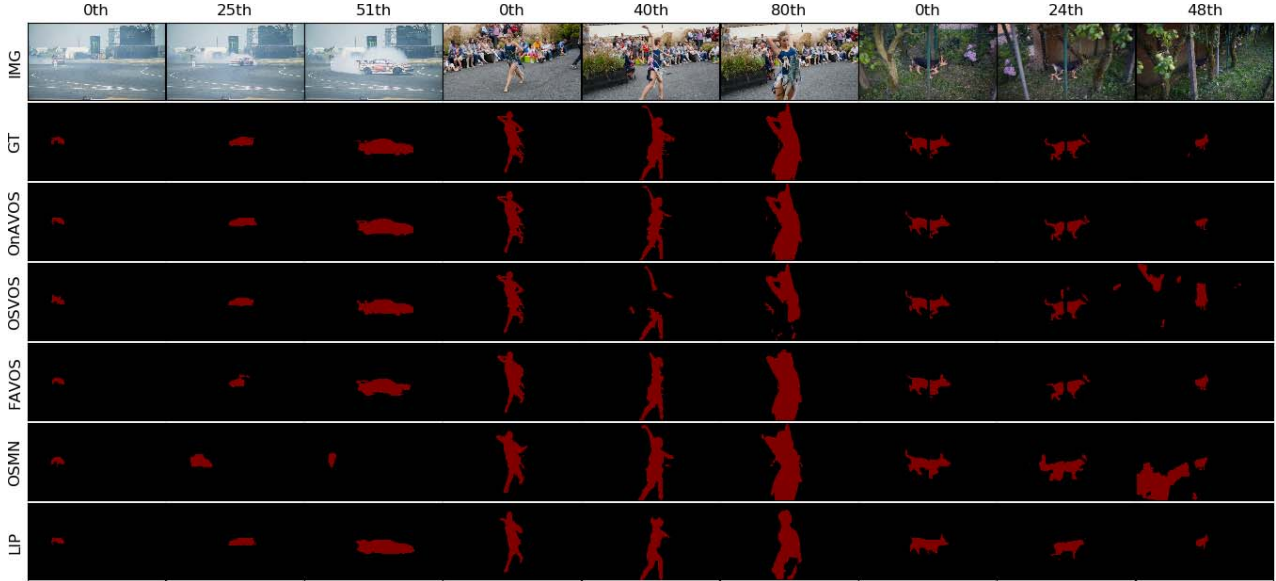


Figure 6. Qualitative results comparison of OnAVOS [52], OSVOS [3], FAVOS [5], OSMN [57] and LIP on DAVIS 2016 dataset [43]. The index of each image in a sequence is shown on the top.

5. Experiments

To test how our model learns to propagate instance information in a long term sequence, we evaluate our model on both DAVIS 2016 [43] and DAVIS 2017 [44] datasets, which contain video sequences of high quality and accurate mask annotations of objects. DAVIS 2016 dataset focuses on single-object video segmentation. It has 30 training and 20 validation videos. As an extension to DAVIS 2016 dataset, DAVIS 2017 dataset brings 30 more video sequences for training set and 10 more for validation set. It also provides another 30 sequences for testing. As DAVIS 2017 dataset focuses on multiple object segmentation, it has been re-annotated for each individual target object.

5.1. Implementation Details

Our model is implemented with PyTorch [41] library. A Nvidia Titan X (Pascal) GPU with 12GB memory is used for experiments. Details of convolutional gated recurrent Mask-RCNN are shown below.

Model structure. Our backbone network is a ResNet101-FPN [19, 32] with group normalization [53]. ResNet101 is initialized with weights pre-trained on Imagenet [10]. In Conv-GRU module, the channel number of each hidden state is 256. Kernels of all convolutions in Conv-GRU are of size 3×3 with 256 filters. We apply multi-level RPN and multi-level ROIs for the network¹. The ROI aligned feature resolution is 28×28 for mask head, and 7×7 for bbox head and id head. In all cases, we adopt image centric training [15].

¹See supplementary material for more details.

Pre-train on Ms-COCO dataset. For each image, we randomly scale it to have its shorter side equal to 1 of 11 different lengths: 640, 608, 576, 544, 512, 480, 448, 416, 384, 352, 320 and its longer size to be maximumly 1333. We sample 512 ROIs with foreground-to-background ratio 1:3. RPN adopts 5 aspect ratios (0.2, 0.5, 1, 2, 5) and 5 scales (32^2 , 64^2 , 128^2 , 256^2 , 512^2). The model is trained with stochastic gradient descent (SGD) for 270K iterations. We fix input hidden states to be zeros for Conv-GRU module, weight decay 0.0001, momentum 0.9. The initial learning rate is 0.02 and dropped by a factor of 10 at 210K and 250K. In the following cases, the configuration is kept the same unless otherwise stated.

Fine-tune on DAVIS dataset. We generate ground truth (GT) bounding boxes from GT masks of DAVIS dataset. The width and height of the boxes are expanded by 10% to prevent incomplete mask prediction caused by inaccurate box prediction. The sequences are randomly shuffled and scaled as in pre-training stage. As there is no causal reasoning in the task, we reverse each sequence for more training data. The backbone network is not trained to prevent overfitting for DAVIS dataset. 128 ROIs are sampled from each image. The model is trained for 12K iterations with an initial learning rate of 0.002 and dropped by a factor of 10 at 8K and 10K. Due to the GPU memory limitation, it only allows to train with maximum recurrence of 4. We extend the length to 8 by stopping gradient back propagation between 4th and 5th frames.

Online fine-tuning. The network is fine-tuned with the GT of the first image for maximally 1000 iterations with early stopping. If the loss for a prediction head is smaller than an

Method	OnAVOS	FAVOS	OSVOS	LIP(Ours)	MSK	PML	SFL	OSMN	CTN	VPN
J&F Mean \uparrow	85.5	81.0	80.2	78.5	77.6	77.4	76.1	73.5	71.4	67.9
J Mean \uparrow	86.1	82.4	79.8	78.0	79.7	75.5	76.1	74.0	73.5	70.2
J Recall \uparrow	96.1	96.5	93.6	88.6	93.1	89.6	90.6	87.6	87.4	82.3
J Decay \downarrow	5.2	4.5	14.9	0.05	8.9	8.5	12.1	9.0	15.6	12.4
F Mean \uparrow	84.9	79.5	80.6	79.0	75.4	79.3	76.0	72.9	69.3	65.5
F Recall \uparrow	89.7	89.4	92.6	86.8	87.1	93.4	85.5	84.0	79.6	69.0
F Decay \downarrow	5.8	5.5	15.0	0.06	9.0	7.8	10.4	10.6	12.9	14.4

Table 1. Results on DAVIS 2016 [43]. Left column shows different metrics. Up-arrow \uparrow means the higher the better. Down-arrow \downarrow means the lower the better. Methods are in descent order according to J&F mean from left to right.

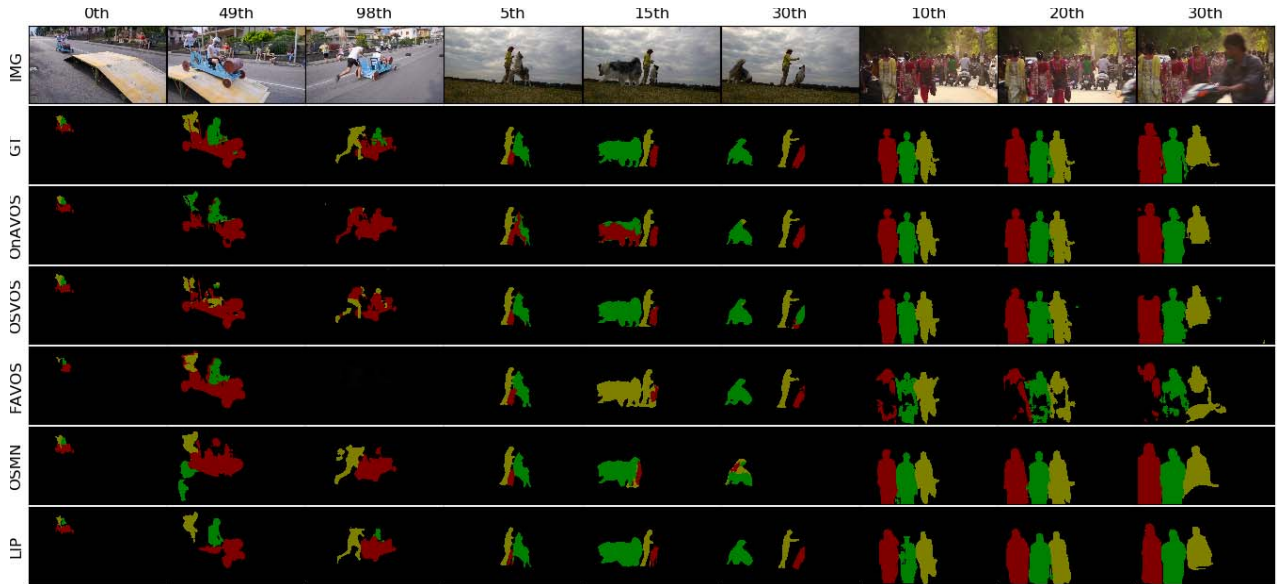


Figure 7. Qualitative results comparison of OnAVOS [52], OSVOS [3], FAVOS [5], OSMN [57] and LIP on DAVIS 2017 dataset [44]. The index of each image in a sequence is shown on the top.

empirically chosen threshold, the loss is ignored. If all the losses are ignored, we stop the training¹. We also stop the loss back-propagation in id head at its last fully connected layer, so the features to distinguish ids will not be affected by the newly initialized head. Focal loss [33] is used to balance id training¹.

Online inference. For each id, we select 10 detected objects that have id score above 0.2 and apply one maximum constraint to select the best candidate. For the location continuity constraint, we suppress the object instance that has IOU lower than 0.3 with the detection from previous frame if the previous id score is higher than 0.4. To relieve the id score from decaying over time, we apply fine-tuning for id head for maximally 500 iterations with early stopping¹.

5.2. Compare with other methods

We compare our method with some state-of-the-art methods on both the DAVIS 2016 benchmark and the

DAVIS 2017 benchmark² by using standard evaluation metrics J and F [43, 44]. The evaluation on DAVIS 2016 benchmark shows the performance for single-object video segmentation, while the evaluation on DAVIS 2017 benchmark shows the performance for video segmentation of multiple objects. It should be noted that our method does not apply any post-processing, but has been pre-trained on Ms-COCO dataset [34]. Among the several top methods, we remove CINM [2] and RGMP [54] to avoid unfair comparison. CINM [2] is built upon OSVOS [3] and further adopts a refinement CNN and MRF for post-processing. The better initial prediction, the better its result. RGMP [54] cannot be successfully trained with static image dataset and DAVIS dataset alone for mask propagation. It has created a large number of synthetic video training data from Pascal VOC [11, 12], ECSSD [49] and MSRA10K [7] datasets. It is not fair to compare with RGMP as the quality of video training data are not the same and cannot be controlled. For

²<https://davischallenge.org/>

Method	OnAVOS	LIP(Ours)	OSVOS	FAVOS	OSMN
J&F Mean \uparrow	65.4	61.1	60.3	58.2	54.8
J Mean \uparrow	61.6	59.0	56.6	54.6	52.5
J Recall \uparrow	67.4	69.0	63.8	61.1	60.9
J Decay \downarrow	27.9	16.0	26.1	14.1	21.5
F Mean \uparrow	69.1	63.2	63.9	61.8	57.1
F Recall \uparrow	75.4	72.6	73.8	72.3	66.1
F Decay \downarrow	26.6	20.1	27.0	18.0	24.3

Table 2. Results on DAVIS 2017 [44]. Left column shows different metrics. Up-arrow \uparrow means the higher the better. Down-arrow \downarrow means the lower the better. Methods are in descent order according to J&F mean from left to right.

DAVIS 2017 benchmark, we exclude PReMVOS [38] and OSVOS+ [39] as they both use multiple specialized networks in multiple processes to refine their results.

For DAVIS 2016, we compare with OnAVOS [52], FAVOS [5], OSVOS [3], MSK [42], PML [4], SFL [6], OSMN [57], CTN [27] and VPN [26]. We detect multiple objects and evaluate in the way for single-object. Our method ranks the 4th among the compared methods as shown in Table 1. It should be noted that our results are better than FAVOS and OSVOS if they are without post-processing. FAVOS achieves J mean and F mean of 77.9% and 76% respectively without tracker and CRF [5]. OSVOS achieves J mean and F mean of 77.4% and 78.1% respectively without boundary snapping post-processing [3]. OnAVOS achieves J mean of 82.8% without CRF post-processing [52]. In addition, we compare our method with another visual memory (Conv-GRU) based VOS method [51]. Both of the methods are trained with additional image dataset, but we achieve 4.5% gain in J&F mean without optical flow and CRF post-processing.

For DAVIS 2017, we compare LIP with OnAVOS [52], OSVOS [3], FAVOS [5] and OSMN [57] as shown in Table 2. LIP has relatively better performance as it is better at separating different instances and keeping coherent label within an instance.

Qualitative results on DAVIS 2016 and DAVIS 2017 are shown in Fig. 6 and Fig. 7, respectively³. Fig. 6 shows that our LIP can track single object well on instance level and preserve good mask extent for an instance. OSMN [57] and OSVOS [3] fail to keep the mask within an instance. In Fig. 7, it is obvious that the information of an instance in our LIP helps segment multiple objects. All the other methods either assign one label to multiple objects or assign multiple labels to one object, while LIP handles those issues better.

5.3. Ablation study

We perform ablation study on DAVIS 2017 dataset by comparing the model with and without dynamic visual memory as shown in Table 3. We first evaluate the static

³More quantitative results and qualitative examples on DAVIS 2016 and DAVIS 2017 are shown in the supplementary material.

model by fixing input hidden states (h_{t-1}) to zeros for Conv-GRU module. This is Mask-RCNN with static Conv-GRU module and bottom up path augmentation. Fine-tuning on video dataset is done by training with static images only. The J&F mean score is 59.2%, which is 1 percent lower than the performance of OSVOS [3] with post-processing. The full version of our model is trained with dynamic video images. It reaches the best J&F mean score of 61.1%. The dynamic visual memory contributes as it learns to propagate masks. The static model lacks such property to handle large appearance change, as shown in Fig. 8.




Figure 8. A qualitative example of prediction with (top row) and without (bottom row) dynamic visual memory.

Mask-RCNN	Conv-GRU	J Mean	F Mean	J&F Mean
\checkmark	input zero hidden states	56.9	61.5	59.2
\checkmark	\checkmark	59.0	63.2	61.1

Table 3. Ablation study results on DAVIS 2017 dataset.

6. Conclusions

We have presented a single end-to-end trainable neural network for video segmentation of multiple objects. We extend the powerful instance segmentation network with visual memory for inference ability across time. Such design serves as an instance segmentation based baseline for VOS task. The newly designed convolutional gated recurrent Mask-RCNN learns to extract and propagate information for multiple instances simultaneously and achieves the state of the art results.

References

- [1] N. Ballas, L. Yao, C. Pal, and A. C. Courville. Delving deeper into convolutional networks for learning video representations. In *ICLR*, 2016.
- [2] L. Bao, B. Wu, and W. Liu. Cnn in mrf: Video object segmentation via inference in a cnn-based higher-order spatiotemporal mrf. In *CVPR*, June 2018.
- [3] S. Caelles, K.-K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. Van Gool. One-shot video object segmentation. In *CVPR*, 2017.
- [4] Y. Chen, J. Pont-Tuset, A. Montes, and L. Van Gool. Blazingly fast video object segmentation with pixel-wise metric learning. In *CVPR*, 2018.
- [5] J. Cheng, Y.-H. Tsai, W.-C. Hung, S. Wang, and M.-H. Yang. Fast and accurate online video object segmentation via tracking parts. In *CVPR*, pages 7415–7424, 2018.
- [6] J. Cheng, Y.-H. Tsai, S. Wang, and M.-H. Yang. Segflow: Joint learning for video object segmentation and optical flow. In *ICCV*, pages 686–695, 2017.
- [7] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. Torr, and S.-M. Hu. Global contrast based salient region detection. *TPAMI*, 37(3):569–582, 2015.
- [8] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *EMNLP*, pages 1724–1734. Association for Computational Linguistics, 2014.
- [9] H. Ci, C. Wang, and Y. Wang. Video object segmentation by learning location-sensitive embeddings. In *ECCV*, September 2018.
- [10] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.
- [11] M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes challenge: A retrospective. *IJCV*, 111(1):98–136, Jan. 2015.
- [12] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88(2):303–338, June 2010.
- [13] C. Finn, I. Goodfellow, and S. Levine. Unsupervised learning for physical interaction through video prediction. In *NeurIPS*, pages 64–72, 2016.
- [14] C. Fu, W. Liu, A. Ranga, A. Tyagi, and A. C. Berg. DSSD : Deconvolutional single shot detector. *CoRR*, abs/1701.06659, 2017.
- [15] R. Girshick. Fast r-cnn. In *ICCV*, pages 1440–1448, 2015.
- [16] A. Graves and N. Jaitly. Towards end-to-end speech recognition with recurrent neural networks. In *ICML*, pages 1764–1772, 2014.
- [17] K. Gregor, I. Danihelka, A. Graves, D. Rezende, and D. Wierstra. Draw: A recurrent neural network for image generation. In F. Bach and D. Blei, editors, *ICML*, volume 37 of *Proceedings of Machine Learning Research*, pages 1462–1471, Lille, France, 07–09 Jul 2015. PMLR.
- [18] K. He, G. Gkioxari, P. Dollar, and R. Girshick. Mask r-cnn. In *ICCV*, Oct 2017.
- [19] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, June 2016.
- [20] S. Hochreiter. The vanishing gradient problem during learning recurrent neural nets and problem solutions. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 6(02):107–116, 1998.
- [21] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [22] J. J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982.
- [23] P. Hu, G. Wang, X. Kong, J. Kuen, and Y.-P. Tan. Motion-guided cascaded refinement network for video object segmentation. In *CVPR*, June 2018.
- [24] Y.-T. Hu, J.-B. Huang, and A. Schwing. Maskrnn: Instance level video object segmentation. In *NeurIPS*, pages 325–334, 2017.
- [25] Y.-T. Hu, J.-B. Huang, and A. G. Schwing. Videomatch: Matching based video object segmentation. In *ECCV*, September 2018.
- [26] V. Jampani, R. Gadde, and P. V. Gehler. Video propagation networks. In *CVPR*, July 2017.
- [27] W.-D. Jang and C.-S. Kim. Online video object segmentation via convolutional trident network. In *CVPR*, pages 5849–5858, 2017.
- [28] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, pages 3128–3137, 2015.
- [29] A. Khoreva, R. Benenson, E. Ilg, T. Brox, and B. Schiele. Lucid data dreaming for object tracking. In *The 2017 DAVIS Challenge on Video Object Segmentation - CVPR Workshops*, 2017.
- [30] X. Li and C. Change Loy. Video object segmentation with joint re-identification and attention-aware mask propagation. In *ECCV*, September 2018.
- [31] X. Li, Y. Qi, Z. Wang, K. Chen, Z. Liu, J. Shi, P. Luo, X. Tang, and C. C. Loy. Video object segmentation with re-identification. In *The 2017 DAVIS Challenge on Video Object Segmentation - CVPR Workshops*, 2017.
- [32] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie. Feature pyramid networks for object detection. In *CVPR*, July 2017.
- [33] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *ICCV*, 2017.
- [34] 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 D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *ECCV*, 2014.
- [35] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia. Path aggregation network for instance segmentation. In *CVPR*, June 2018.
- [36] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg. Ssd: Single shot multibox detector. In *ECCV*, pages 21–37. Springer, 2016.
- [37] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, pages 3431–3440, 2015.

- [38] J. Luiten, P. Voigtlaender, and B. Leibe. Premvos: Proposal generation, refinement and merging for video object segmentation. In *ACCV*, 2018.
- [39] 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. *TPAMI*, 2018.
- [40] R. Pascanu, T. Mikolov, and Y. Bengio. On the difficulty of training recurrent neural networks. In *ICML*, pages 1310–1318, 2013.
- [41] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS-W*, 2017.
- [42] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *CVPR*, 2017.
- [43] 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 *CVPR*, 2016.
- [44] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. Van Gool. The 2017 davis challenge on video object segmentation. *arXiv:1704.00675*, 2017.
- [45] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016.
- [46] J. Redmon and A. Farhadi. YOLO9000: better, faster, stronger. *CoRR*, abs/1612.08242, 2016.
- [47] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, pages 91–99, 2015.
- [48] D. E. Rumelhart, G. E. Hinton, R. J. Williams, et al. Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1, 1988.
- [49] J. Shi, Q. Yan, L. Xu, and J. Jia. Hierarchical image saliency detection on extended cssd. *TPAMI*, 38(4):717–729, 2016.
- [50] 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 *ICCV*, pages 2167–2176, 2017.
- [51] P. Tokmakov, K. Alahari, and C. Schmid. Learning video object segmentation with visual memory. In *ICCV*, Oct 2017.
- [52] P. Voigtlaender and B. Leibe. Online adaptation of convolutional neural networks for video object segmentation. In *BMVC*, 2017.
- [53] Y. Wu and K. He. Group normalization. In *ECCV*, September 2018.
- [54] S. Wug Oh, J.-Y. Lee, K. Sunkavalli, and S. Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *CVPR*, June 2018.
- [55] H. Xiao, J. Feng, G. Lin, Y. Liu, and M. Zhang. Monet: Deep motion exploitation for video object segmentation. In *CVPR*, June 2018.
- [56] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *NeurIPS*, pages 802–810, 2015.
- [57] L. Yang, Y. Wang, X. Xiong, J. Yang, and A. K. Katsaggelos. Efficient video object segmentation via network modulation. In *CVPR*, pages 6499–6507, 2018.