This WACV 2020 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

RPM-Net: Robust Pixel-Level Matching Networks for Self-Supervised Video Object Segmentation

Youngeun Kim Seokeon Choi Hankyeol Lee Taekyung Kim Changick Kim KAIST KAIST KAIST KAIST KAIST KAIST {youngeunkim, seokeon, hankyeol, tkkim93, changick}@kaist.ac.kr

Abstract

In this paper, we introduce a self-supervised approach for video object segmentation without human labeled data. Specifically, we present Robust Pixel-level Matching Networks (RPM-Net), a novel deep architecture that matches pixels between adjacent frames, using only color information from unlabeled videos for training. Technically, RPM-Net can be separated in two main modules. The embedding module first projects input images into high dimensional embedding space. Then the matching module with deformable convolution layers matches pixels between reference and target frames based on the embedding features. Unlike previous methods using deformable convolution, our matching module adopts deformable convolution to focus on similar features in spatio-temporally neighboring pixels. Our experiments show that the selective feature sampling improves the robustness to challenging problems in video object segmentation such as camera shake, fast motion, deformation, and occlusion. Also, we carry out comprehensive experiments on three public datasets (i.e., DAVIS-2017, SegTrack-v2, and Youtube-Objects) and achieve state-ofthe-art performance on self-supervised video object segmentation. Moreover, we significantly reduce the performance gap between self-supervised and fully-supervised video object segmentation (41.0% vs. 52.5% on DAVIS-2017 validation set).

1. Introduction

Video object segmentation, segmenting a foreground object along an entire video sequence, is one of the challenging tasks in computer vision. Most of the previous work [9, 36, 8, 15, 39, 1, 25, 4, 14] focus on increasing the performance with human labeled annotations. However, compared to other tasks using videos (e.g., video object tracking and video object detection) and image segmentation [17], video object segmentation suffers from generating pixel-level annotations for every frame [42, 35]. For example, the



Figure 1: Sampling locations (i.e., receptive fields of the matching module, marked in red) for target pixel (marked in green) matching are adaptively adjusted according to the change of frame. We overlay two consecutive frames for visualization. The best view is in color and zoomed in.

DAVIS dataset [28], which is the most widely used dataset in the video object segmentation, contains 4,219 manually pixel-wise annotated frames.

The main motivation of this paper begins with relieving the amount of efforts to generate annotations for video object segmentation. To this end, we introduce Robust Pixel-Level Matching Networks (RPM-Net) for self-supervised video object segmentation. Rather than using segmentation annotations, RPM-Net leverages only color information for training. The proposed network consists of two parts: an embedding module and a matching module. The embedding module extracts high-dimensional features from input RGB images and the matching module matches pixels between two frames according to embedding features.

However, since there is no human supervision for training, the network is vulnerable to challenges in video object segmentation such as camera shake, fast motion, deformation, and occlusion. To address this problem, we use deformable convolution layers [6] in the matching module to focus on similar features in spatio-temporally neighboring locations (see Fig. 1), and show that this leads to robust pixel-level matching. Moreover, we analyze how deformable convolution works in self-supervised learning with extensive experiments and visualizations. In fully-supervised learning, deformable convolution discovers class geometric-invariant features for semantic segmentation [23, 41], object detection [31, 40, 3], and other computer vision tasks [20, 19, 37]. Unlike the conventional approaches, in our self-supervised scheme, deformable convolution in the matching module pays attention to spatiotemporally similar features (see Fig. 2).

There are several advantages of RPM-Net: Firstly, the proposed RPM-Net can be trained in an end-to-end way with a single forward path, where all layers are differentiable. Secondly, although the model is only trained with unlabeled videos, RPM-Net well tracks the objects without online training at inference procedure, which allows the network to operate with a high speed. Moreover, RPM-Net can be used in real-world applications, since it segments target objects without accessing future frames.

To sum up, our contributions can be summarized as follows: 1) We propose novel self-supervised video object segmentation framework, which is annotation-free, end-to-end trainable, and no online training at inference process. 2) In order to achieve that, we adopt deformable convolution [6] which is widely used in fully-supervised learning. Our experiments and visualizations suggest that deformable convolution aggregates similar spatio-temporal features, so that the robust pixel matching is available. To the best of our knowledge, this is the first time that deformable convolution is adopted in self-supervised learning and pixel-wise matching. 3) To provide a reference work on self-supervised video object segmentation, and also to show the generality of RPM-Net, we report the performance on three public video object segmentation datasets (DAVIS-2017 [28], SegTrack-v2 [21], and Youtube-Objects [29]). Our RPM-Net outperforms the latest self-supervised method. Most importantly, we significantly reduce the performance gap between self-supervised and fully-supervised video object segmentation.

The paper is organized as follows. Section 2 presents the related work. Section 3 describes the details of the proposed frameworks. Section 4 shows the performance on video object segmentation datasets and also presents the analysis of RPM-Net. Finally, Section 5 concludes the paper.



Figure 2: Illustration of the role of deformable convolution in fully-supervised and our self-supervised learning. (a) With human-given class, the deformable convolution layers extract class geometric-invariant features for pixel-wise classification (we illustrate a segmentation example). (b) In our proposed method, the deformable convolution layers in the matching module focus on spatio-temporally neighboring pixels which have similar features. The best view is in color and zoomed in.

2. Related Work

2.1. Video Object Segmentation

There is a large literature on video object segmentation since it is one of the most important tasks in video analysis. Recently, most research focus on unsupervised and semi-supervised video object segmentation. Unsupervised methods segment primary objects without any information of these objects. Since there is no given target object information in inference time, they use optical flow [16, 10, 32] and saliency maps [34] for obtaining information about the target object. Note that even though they are referred to as "unsupervised" methods, these methods require annotations for network training.

Semi-supervised methods segment foreground objects where object mask is given in the first frame. Recent semisupervised approaches are based on deep neural networks [11, 1, 7, 5, 26, 9] using semantic information from the first frame and achieve reliable performance on given datasets. Our method is included in semi-supervised video object segmentation. However, unlike the above two methods, we do not use annotations for training.

2.2. Self-Supervised Tracking

Our method is comparable to recent work by Vondrick et al. [33], which proposes a self-supervised tracking method using gray-scale features from target and reference frames. They outperform the high-performance optical flow network [12] on the video object segmentation dataset [28]. However, this method is vulnerable to object deformation and scale change during the video sequence since they use



Figure 3: The training and inference scheme of RPM-Net. The RPM-Net architecture consists of two parts: an *embedding* module and a matching module. The embedding module extracts features from input RGB images with up-convolutional layers to keep the resolution of embedding feature maps (128×128). The matching module is composed of deformable convolution layers and offset convolution layers. In our self-supervised training scheme, these layers are trained to focus on similar embedding features. After training, RPM-Net infers the target segmentation mask S_t using RGB images and predicted segmentation mask S_{t-1} of the frame t - 1. Note that a ground-truth segmentation mask is only given in the first frame.

only the several frames in the beginning of the video for inferring the target frame mask. Also, the boundary distortion problem can be occurred since the features of object boundary usually have a large variation during the video sequence. Our RPM-Net deals with this problem by pixelwise matching between adjacent frames so that the more reliable segmentation results are achieved.

2.3. Deformable Convolution (Background)

Deformable convolution [6] is designed to improve the model geometric transformation capability of CNNs [23, 31]. It shifts standard grid sampling locations R by adding 2D offsets. For example, considering the 3×3 kernel, standard convolution operation at location \mathbf{p}_0 on the input feature map \mathbf{x} and the output feature map \mathbf{y} becomes

$$\mathbf{y}(\mathbf{p}_0) = \sum_{\mathbf{p}_n \in R} \mathbf{w}(\mathbf{p}_n) \cdot \mathbf{x}(\mathbf{p}_0 + \mathbf{p}_n), \quad (1)$$

where $R = \{(-1, -1), (-1, 0), ..., (0, 1), (1, 1)\}$. We use the same notation as in [6] for consistency.

In deformable convolution, convolution operation is formulated as follows:

$$\mathbf{y}(\mathbf{p}_0) = \sum_{\mathbf{p}_n \in R} \mathbf{w}(\mathbf{p}_n) \cdot \mathbf{x}(\mathbf{p}_0 + \mathbf{p}_n + \Delta \mathbf{p}_n).$$
(2)

Here, offsets for each sampling point are denoted as $\{\Delta \mathbf{p}_n | n = 1, ..., |R|\}$. Offset $\Delta \mathbf{p}_n$ is assigned from the output of another convolution layer (i.e., offset convolution layer) which takes the feature map x as input.

Here, the shifted sampling location ($\mathbf{p} = \mathbf{p}_0 + \mathbf{p}_n + \Delta \mathbf{p}_n$) is the fractional location. Therefore, $\mathbf{x}(\mathbf{p})$ is calculated by bilinear interpolation as

$$\mathbf{x}(\mathbf{p}) = \sum_{\mathbf{q}\in Z} K(\mathbf{q}, \mathbf{p}) \cdot \mathbf{x}(\mathbf{q}), \tag{3}$$

where $K(\cdot, \cdot)$ is the bilinear interpolation kernel, Z is the set of all integral spatial locations on the input feature map

x. Moreover, the bilinear interpolation kernel consists of two one dimensional kernels as

$$K(\mathbf{q}, \mathbf{p}) = k(q_x, p_x) \cdot k(q_y, p_y), \tag{4}$$

where k(a, b) = max(0, 1 - |a - b|). Refer to [6] for more details. In our method, deformable convolution is adopted in the matching module which helps to achieve robust pixel-level matching without training annotations.

3. Proposed Method

3.1. Overview

Our goal is to segment foreground objects along an entire video without annotations for training. Rather than training with dense annotations, RPM-Net leverages color information from unlabeled videos. Our proposed network architecture is illustrated in Fig. 3.

Also, our intention is not just building the network without interpretation, but specifying how the network behaves in self-supervised learning. Therefore, we divide the network into two parts: an embedding module and a matching module. Moreover, we verify that these modules work according to their purpose by showing experiment results in Section 4. In the remainder of this section, we first describe each module, and then demonstrate training and inference procedures in detail.

3.2. Embedding Module

The purpose of the embedding module is extracting a deeper representation from RGB images for more reliable matching. Our embedding module produces 64dimensional embedding feature maps from two given RGB images. The extracted embedding features are taken into the matching module with concatenation.

A reference frame I_{ref} and a target frame I_{tar} are given as the input for the embedding module. The inputs are resized to a resolution of 512×512 . We use FCN-ResNet101 [24] for our embedding module and modify up-convolution layers to obtain 128×128 feature maps. The process of the embedding module can be formulated as follows:

$$\widehat{f}_{cat} = concat(E(I_{ref}), E(I_{tar})), \tag{5}$$

where $E(\cdot)$ represents the modified FCN-ResNet101 in our embedding module, and \hat{f}_{cat} denotes the concatenated features obtained from two input frames I_{ref} and I_{tar} .

3.3. Matching Module

The matching module matches pixels between reference and target frames using deformable convolution layers. To simplify notations, we use D(a, b) to represent a deformable convolution layer, where a and b denote the input feature map and offsets, respectively. Also, we represent a Table 1: The layer-by-layer definition of the matching module. Batch normalization and ReLU non-linearity are omitted for brevity. D_1 takes I_{M1} as input, and D_4 directly takes I_{M2} as input. Please see Appendix A in supplementary material for detailed layer-by-layer definition including the embedding module.

	Layer Description	Output Tensor Dim.
I_{M1}	Concatenated features	(64, H/4, W/4)
I_{M2}	Reference image (training) or	(3, H, W) or
	t-1 Segmentation mask (inference)	(1, H, W)
D_1	3×3 deformable conv, 32, stride 1	(22 H/4 W/4)
	offset conv: 3×3 conv, 18, stride 1	(32, 11/4, W/4)
D_2	3×3 deformable conv, 16, stride 1	(16 H/A W/A)
	offset conv: 3×3 conv, 18, stride 1	(10, 11/4, 11/4)
D_3	3×3 deformable conv, 2, stride 1	$(2 \mathbf{H}/4 \mathbf{W}/4)$
	offset conv: 3×3 conv, 18, stride 1	(2, 11/4, W/4)
D_4	1×1 deformable conv, 3, stride 1	(3, H/4, W/4) or
	offset: feature maps from layer 108	(1, H/4, W/4)

convolution layer for deformable offsets as $O(\cdot)$. Since the matching module consists of 4 cascaded deformable convolution layers, we can write the output feature map \hat{f}_{k+1} from the deformable convolution layer as follows:

$$\widehat{f}_{k+1} = D_k(\widehat{f}_k, O_k(\widehat{f}_k)), \tag{6}$$

where \hat{f}_k denotes the input feature map of the layer $k \in \{1, 2, 3\}$ in the matching module. Note that \hat{f}_1 obtained from the concatenated embedding feature map \hat{f}_{cat} in eq. 5.

As shown in Table 1, the first three layers are 3×3 deformable convolution layers, which sample useful features from the concatenated embeddings for pixel-level matching. And the last layer is a 1×1 deformable convolution layer with fixed weights of value 1. Also, we do not use the offset convolution layer O_4 for the last deformable convolution layer D_4 , and the layer directly use the features \hat{f}_4 for offsets. Thus, the last layer matches each target pixel with a fractional location in the reference frame based on offsets \hat{f}_4 . The input feature map for the last deformable convolution layer is varied according to the training and inference. In the training procedure, the layer directly takes the RGB image from the reference frame as input and predicts the target frame image:

$$\widehat{I}_{tar} = D_4(I_{ref}, \widehat{f}_4). \tag{7}$$

Otherwise, in the inference procedure, the reference segmentation mask S_{t-1} is used for computing target segmentation mask:

$$S_t = D_4(S_{t-1}, \hat{f_4}).$$
 (8)

3.4. Training and Inference

In training and inference procedures, we assume that the image intensity does not change significantly between adjacent frames. Therefore, we can match pixels with similar features and predict that these pixels belong to the same object.

Training. Since our goal is to train RPM-Net without annotations, we use unlabeled videos for training. Therefore, we do not represent network outputs as binary values (e.g., background: 0, target: 1), which requires annotations for training. Instead, we use the L2 loss to penalize the color difference between the target frame I_{tar} and the predicted target frame \hat{I}_{tar} from eq. 7. Our loss function L has the form:

$$L = \frac{1}{N} \sum_{x} \sum_{y} \left\| I_{tar}(x, y) - \widehat{I}_{tar}(x, y) \right\|_{2}^{2}, \quad (9)$$

where (x, y) is the pixel location and N denotes the total number of pixels.

Inference. In video object segmentation, the segmentation mask contains the integer class value (e.g., background: 0, dog: 1, and cow: 2) at each pixel. Therefore, a fractional pixel value obtained from bilinear interpolation kernel in the deformable convolution layer (eq. 3) reduces the segmentation performance of RPM-Net.

To address this problem, we first separate the segmentation mask S_{t-1} into the binary maps B_c of each class c including background class. Then the binary maps are given as a input to the matching module. After that, we select the object class with the highest value at each pixel position **p**. Therefore, in the last deformable convolution layer in the matching module, the target segmentation mask S_t is calculated as follows:

$$S_t(\mathbf{p}) = \operatorname*{argmax}_{c \in \mathcal{C}} D_4(B_c(\mathbf{p}), \widehat{f}_4), \tag{10}$$

where C is the set of class labels in the video sequence, and $S_t(\mathbf{p})$ and $B_c(\mathbf{p})$ denote the pixel value of the segmentation mask S_t and the class binary map B_c at the location **p**, respectively. Moreover, we add a refinement stage using dense CRFs [18].

Note that RPM-Net can be directly used at inference procedure without online training. Also, our RPM-Net segments primary objects without accessing future frames (e.g., I_{t+1}), which allows RPM-Net can be applied in real-world applications.

4. Experimental Results

In this section, we show the experimental results of RPM-Net. We first describe the implementation details and the datasets used for video object segmentation before we present evaluation metrics, quantitative and qualitative results, and the analysis of our proposed network.

Table 2: Quantitative evaluation on the DAVIS-2017 validation set [28]. We report the both \mathcal{J} -score (mIOU) and \mathcal{F} -score (contour accuracy).

Methods	\mathcal{J} -score	\mathcal{F} -score
Copy from the Initial Frame	22.1	23.6
Optical Flow (Coarse-to-Fine [22])	13.0	15.1
Optical Flow (FlowNet 2.0 [12])	26.7	25.2
Colorization [33]	34.6	32.7
RPM-Net w/ Conv (Baseline)	31.7	33.5
RPM-Net w/ Diliated Conv	30.8	33.1
RPM-Net	35.7	38.8
\mathbf{RPM} -Net _R	41.0	42.2
Supervised [39]	52.5	57.1

Table 3: Quantitative evaluation on SegTrack-v2 [21] and Youtube-Objects [29] dataset. \mathcal{J} -score is only used for evaluation.

Methods	SegTrack-v2	Youtube
Copy from the Initial Frame	31.9	54.2
RPM-Net w/ Conv. (Baseline)	41.6	55.2
RPM-Net w/ Dilated Conv.	41.8	54.4
RPM-Net	45.2	56.2
$RPM-Net_R$	48.4	57.4

4.1. Implementation Details

We trained RPM-Net on the training set from Youtube-VOS [38] and DAVIS-2016 [27] (total 3,501 videos). In the training procedure, we used frame t-5 as a reference frame and frame t as a target frame. Moreover, we did not leverage pre-trained weights since our ultimate goal is to train the network without any human supervision. Therefore, the weights of all network layers were initialized with Gaussian distribution. We trained RPM-Net with 100 epochs on the training dataset and use the SGD optimizer with momentum 0.9, weight decay 0.0005, and initial learning rate 0.001. Here, we used "ploy" learning rate policy [2] for every epoch as $lr_{init} \cdot (1 - \frac{epoch_{curr}}{epoch_{max}})^{0.9}$.

4.2. Datasets

To show the generality of our method, we report the performance on three public datasets (DAVIS-2017, SegTrackv2, and Youtube-Objects).

DAVIS-2017. The DAVIS-2017 dataset [28] is the well-known video object segmentation dataset which includes multiple objects for each video sequence. The dataset consists of 60 training videos and 30 validation videos. Also, the dataset contains several videos from its previous version (i.e., DAVIS-2016 dataset). However, it is more difficult to segment since multiple and similar objects appear in each video sequence.

SegTrack-v2. To validate the capability of RPM-Net,



Figure 4: Qualitative results of RPM-Net (without the refinement stage) on three public video object segmentation datasets (i.e., DAVIS-2017, Youtube-Objects, and SegTrack-v2).



Figure 5: Visualization of attributes-based analysis on the DAVIS-2017 validation dataset. The matching module with deformable convolution achieves the best performance in all categories. In this experiments, we fix the embedding module to FCN-ResNet101.

we further conduct further experiments on the SegTrack-v2 dataset [21]. The SegTrack-v2 dataset consists of 14 test video sequences with 24 objects. The instance-level masks are available for multiple objects tracking along the video sequence.

Youtube-Objects. The Youtube-Objects dataset [29] provides videos with 10 object categories (e.g., bird, boat, and train). Note that pixel-level segmentation masks are provided by [13].

4.3. Evaluation Metrics

We use the intersection over union metric \mathcal{J} [27] for evaluating region similarity. The \mathcal{J} score of DAVIS-2017, SegTrack-v2, and Youtube-Objects datasets are used for evaluation. Also, like the previous methods, we further compute contour accuracy \mathcal{F} [27] for the DAVIS-2017 validation dataset.

4.4. Video Object Segmentation Results

We present the performance of self-supervised and recent fully-supervised methods [1, 39] on the DAVIS-2017, SegTrack-v2, and Youtube-Objects datasets. Firstly, we compared RPM-Net with the very recent self-supervised method [33]. As shown in Table 2, RPM-Net shows better performance in both region similarity \mathcal{J} and contour accuracy \mathcal{F} on the DAVIS-2017 validation dataset. This is because [33] uses only the several frames in the beginning of the video for inference, and also the features of object boundary usually have a large variation during the video sequence. Otherwise, RPM-Net deals with this problem by robust pixel-level matching between adjacent frames. Moreover, with the refinement step using CRFs,



Figure 6: The importance of feature representation capability in the embedding module. We change the backbone network of the FCN model in the embedding module.



Figure 7: The relationship between performance and the receptive field size of the matching module. We overlay the theoretical size of receptive fields.

RPM-Net_{*R*} further reduces the performance gap between self-supervised and fully-supervised video object segmentation. Our qualitative results in Fig. 4 show that although the model is only trained with unlabeled videos, RPM-Net achieves reasonable tracking results.

Obviously, there is a performance gap between fullysupervised and self-supervised way. Nevertheless, through the self-supervised approach, the cost for human supervision is heavily decreased for training. Moreover, it is possible that further improvement in self-supervised learning might be achieved by more accurate pixel-level matching.

4.5. Experimental Analysis

In this section, we analyze the importance of each module with extensive experiments. Before we present our experimental setup, a short note on terminology is required for clarity. We use the "sample locations" for referring the receptive fields of deformable convolutions, following the original paper [6].

Experimental Setup. We experimented on the network with different combinations of two modules. In order to show the importance of feature representation capability in the embedding module, we replaced the backbone network of FCN with ResNet-18, ResNet-34, and ResNet-50. We also constructed matching modules using the standard convolution and dilated convolution [2], to present the effect of receptive fields in the matching module. We varied the number of standard convolution layers (L = 1, 3, 5), and the dilation rate of dilated convolution layers (D = 3, 6, 9). In the experiments, we fixed the number of dilated convolution layers to 3, which shows the best performance with dilated convolution. Note that we maintain the last 1×1 deformable convolution layer for pixel-level matching.

All experiments were performed on the DAVIS-2017 validation set. From now on, standard convolution with k layers denotes as conv-k, and also dilated convolution with dilation rate r represents as dilation-r for brevity. Refer to Appendix B in supplementary material for quantitative performances of network configurations using in our experiments.

Embedding Module for Better Feature Representation. The embedding module generates high-dimensional features that help the robust matching. As shown in Fig. 6, the network performance is increased with deeper network architectures. The results imply that rich feature information is essential for obtaining good performance in our selfsupervised training scheme. Moreover, the figure shows that deformable convolution achieves higher performance gain compared to others, which means that deformable convolution efficiently uses features from the embedding module.

Matching Module for Robust Pixel-Level Matching. The performance of pixel-level matching is closely related with the receptive fields in the matching module. Figure 7 shows the performance with varying the size of receptive fields. In the figure, deformable convolution is excluded from the comparison, since it has adaptive receptive field size that can not be represented by fixed value. In the figure, convolution with smallest receptive fields (i.e., conv-1) shows the lowest performance. This result demonstrates that the receptive fields size of 1-layer convolution is not enough to cover the matching region. However, the figure also shows that, except for conv-3, as the size of the receptive fields increases, the average performance is decreased. We further discuss this observations in the following subsection (i.e., feature-level analysis of receptive fields). In Fig. 8, we compare the variants of our RPM-Net.

To show the effects of the matching module on challenging situations, we illustrate the performance of each video attribute [27] in Fig. 5. Deformable convolution achieves the best performance on the problems from the change of scene (i.e., fast motion, motion blur, camera shake, deformation, and scale-variation). Also, it is beneficial for the problems from the objects configurations (i.e., heterogeneus object and interacting object). For detailed description of



Figure 8: The effect of the embedding and matching modules. We visualize the last frame of the "car-shadow" sequence in the DAVIS-2017 validation set.



Figure 9: Feature similarities between embedding features in the receptive fields (or sampling locations). We use the DAVIS-2017 validation set for the experiments.

the attributes, see [27].

Feature-Level Analysis of Receptive Fields. Moreover, we further analyze the embedding features in the receptive fields of the matching module. To do this, in Fig. 9, we computed the average of cosine similarity between the feature vector at target pixel and other feature vectors located in the receptive fields. The results demonstrate that larger receptive fields are likely to contain redundant features for pixel matching, which causes performance degradation as already shown in Fig. 7. Otherwise, deformable convolution shows significantly higher similarity value, which means that the selective feature sampling is helpful to reduce the redundant features in the receptive fields.

Visualizations from Two Modules. One may wonder whether our self-supervised training scheme is correctly implemented as we intended. To alleviate this concern, we present the results from two modules in Fig. 10. For visualization, we project the 64-dimensional concatenated features (i.e., embeddings from frame t - 1 and t) onto 3dimensional space using PCA [30]. From the results, we observe that the nearby pixels in the same object are represented in similar embedding features. Also, we overlay sample locations of the matching module on the projected embedding features. It shows that the matching module enables selective feature sampling by focusing on similar em-



Figure 10: Visualization of the contributions of embedding and matching modules. For given frame t - 1 and frame t(two images are overlaid in the first row), we show examples of embedding features (second row) and sampling locations (third row) for target pixel along the video sequence.

Table 4: Computation time per frame on the DAVIS-2017validation dataset.

J-score 31.7 35.7 41.0 F-score 33.5 38.8 42.2		RPM_{conv-3}	RPM-Net	$RPM-Net_R$
<i>F</i> -score 33.5 38.8 42.2	\mathcal{J} -score	31.7	35.7	41.0
	$\mathcal{F} ext{-score}$	33.5	38.8	42.2
time $0.02s$ $0.09s$ $0.55s$	time	0.02s	0.09s	0.55s

bedding features.

Runtime of RPM-Net We evaluated the runtime of RPM_{conv-3}, RPM-Net, and RPM-Net_R. As shown in Table 4, deformable convolution does not require large computation time. Also, we use two CRF iterations in RPM-Net_R for obtaining reasonable speed. Although the network is trained without annotations, RPM-Net well tracks the objects apace.

5. Conclusion

In this paper, we have proposed a novel self-supervised RPM-Net for video object segmentation. Particularly, we adopted deformable convolution to improve the robustness to challenging situations in video. Our experiments showed that RPM-Net provides reasonable tracking results without annotations. Moreover, our analysis on two modules showed that the proposed RPM-Net works well with deformable convolution in our self-supervised scheme. Our future work will be focused on the improvement of the robustness of the pixel matching in video object segmentation.

6. Acknowledgement

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (NRF-2018R1A5A7025409).

References

- S. Caelles, K. K. Maninis, J. Pont-Tuset, L. Leal-Taixé, D. Cremers, and L. V. Gool. One-shot video object segmentation. In *CVPR*. 2017.
- [2] L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2018.
- [3] X. Chen, Z. Wu, and J. Yu. Dual refinement network for single-shot object detection. *arXiv:1807.08638*, 2018.
- [4] 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*. 2018.
- [5] J. Cheng, Y. H. Tsai, S. Wang, and M. H. Yang. Segflow: Joint learning for video object segmentation and optical flow. In *ICCV*. 2017.
- [6] J. Dai, H. Qi, Y. Xiong, Y. Li, G. Zhang, H. Hu, and Y. Wei. Deformable convolutional networks. In *ICCV*. 2017.
- [7] J. Han, L. Yang, D. Zhang, X. Chang, and X. Liang. Reinforcement cutting-agent learning for video object segmentation. In *CVPR*. 2018.
- [8] P. Hu, G. Wang, X. Kong, J. Kuen, and Y. P. Tan. Motionguided cascaded refinement network for video object segmentation. In *CVPR*. 2018.
- [9] Y.-T. Hu, J.-B. Huang, and A. Schwing. Maskrnn: Instance level video object segmentation. In *NIPS*. 2017.
- [10] Y. T. Hu, J. B. Huang, and A. Schwing. Unsupervised video object segmentation using motion saliency-guided spatiotemporal propagation. In *CVPR*. 2018.
- [11] Y. T. Hu, J. B. Huang, and A. G. Schwing. Videomatch: Matching based video object segmentation. In ECCV. 2018.
- [12] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox. Flownet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*. 2017.
- [13] S. D. Jain and K. Grauman. Supervoxel-consistent foreground propagation in video. In ECCV. 2014.
- [14] V. Jampani, R.Gadde, and P. V. Gehler. Video propagation networks. In CVPR. 2017.
- [15] W. D. Jang and C. S. Kim. Online video object segmentation via convolutional trident network. In *CVPR*. 2017.
- [16] W. D. Jang, C. Lee, and C. S. Kim. Primary object segmentation in videos via alternate convex optimization of foreground and background distributions. In *CVPR*. 2016.
- [17] Y. Kim, T. Kim, S.and Kim, and C. Kim. Cnn-based semantic segmentation using level set loss. In WACV. 2019.
- [18] P. Krähenbühl and V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In NIPS. 2011.
- [19] B. Kwolek. Person re-identification using multi-region triplet convolutional network. In *ACM*. 2017.
- [20] P. Lei and S. Todorovic. Temporal deformable residual networks for action segmentation in videos. In CVPR. 2018.
- [21] F. Li, T. Kim, A. Humayun, D. Tsai, and J. M. Rehg. Video segmentation by tracking many figure-ground segments. In *ICCV*. 2013.

- [22] C. Liu. Beyond pixels: exploring new representations and applications for motion analysis, 2009. PhD Thesis. MIT.
- [23] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia. Path aggregation network for instance segmentation. In CVPR. 2018.
- [24] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In CVPR. 2015.
- [25] S. W. Oh, J. Y. Lee, K. Sunkavalli, and S. J. Kim. Fast video object segmentation by reference-guided mask propagation. In *CVPR*. 2018.
- [26] F. Perazzi, A. Khoreva, R. Benenson, B. Schiele, and A. Sorkine-Hornung. Learning video object segmentation from static images. In *CVPR*. 2017.
- [27] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. V. Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*. 2016.
- [28] J. Pont-Tuset, F. Perazzi, S. Caelles, P. Arbeláez, A. Sorkine-Hornung, and L. V. Gool. The 2017 davis challenge on video object segmentation. arXiv:1704.00675, 2017.
- [29] A. Prest, C. Leistner, J. Civera, C. Schmid, and V.Ferrari. Learning object class detectors from weakly annotated video. In *CVPR*. 2012.
- [30] J. Shlens. A tutorial on principal component analysis. CoRR, abs/1404.1100, 2014.
- [31] B. Singh and L. S. Davis. An analysis of scale invariance in object detection–snip. In CVPR. 2018.
- [32] P. Tokmakov, K. Alahari, and C. Schmid. Learning video object segmentation with visual memory. In CVPR. 2017.
- [33] C. Vondrick, A. Shrivastava, A. Fathi, S. Guadarrama, and K. Murphy. Tracking emerges by colorizing videos. In *ECCV*. 2018.
- [34] W. Wang, J. Shen, and F. Porikli. Saliency-aware geodesic video object segmentation. In CVPR. 2015.
- [35] M. Wigness and J. G. R. III. Unsupervised semantic scene labeling for streaming data. In CVPR. 2017.
- [36] H. Xiao, J. Feng, G. Lin, Y. Liu, and M. Zhang. Monet: Deep motion exploitation for video object segmentation. In *CVPR*. 2018.
- [37] C. Xin, Y. Li, X. Luo, T. Shao, J. Yu, K. Zhou, and Y. Zheng. Autosweep: Recovering 3d editable objects from a single photograph. *IEEE transactions on visualization and computer graphics*, 2018.
- [38] N. Xu, L. Yang, Y. Fan, D. Yue, Y. Liang, J. Yang, and T. Huang. Youtube-vos: A large-scale video object segmentation benchmark. arXiv:1809.03327, 2018.
- [39] L. Yang, Y. Wang, X. Xiong, J. Yang, and A. K. Katsaggelos. Efficient video object segmentation via network modulation. In *CVPR*. 2018.
- [40] Q. Yang, M. Cheng, W. Zhou, Y. Chen, M. Qiu, and W. Lin. Inceptext: A new inception-text module with deformable psroi pooling for multi-oriented scene text detection. arXiv:1805.01167, 2018.
- [41] M. Zhang, X. Zhang, M. Xu, and Q. Li. Image segmentation and classification for sickle cell disease using deformable unet. arXiv:1710.08149, 2017.
- [42] Y. Zhang, Z. Qiu, T. Yao, D. Liu, and T. Mei. Fully convolutional adaptation networks for semantic segmentation. In *CVPR*. 2018.