

# F2Net: Learning to Focus on the Foreground for Unsupervised Video Object Segmentation

Daizong Liu<sup>1†</sup>, Dongdong Yu<sup>2†</sup>, Changhu Wang<sup>2</sup>, Pan Zhou<sup>1\*</sup>

<sup>1</sup>Huazhong University of Science and Technology, China

<sup>2</sup>ByteDance AI Lab, China

{dzliu, panzhou}@hust.edu.cn, {yudongdong, wangchanghu}@bytedance.com

## Abstract

Although deep learning based methods have achieved great progress in unsupervised video object segmentation, difficult scenarios (e.g., visual similarity, occlusions, and appearance changing) are still not well-handled. To alleviate these issues, we propose a novel *Focus on Foreground Network (F2Net)*, which delves into the intra-inter frame details for the foreground objects and thus effectively improve the segmentation performance. Specifically, our proposed network consists of three main parts: Siamese Encoder Module, Center Guiding Appearance Diffusion Module, and Dynamic Information Fusion Module. Firstly, we take a siamese encoder to extract the feature representations of paired frames (reference frame and current frame). Then, a Center Guiding Appearance Diffusion Module is designed to capture the inter-frame feature (dense correspondences between reference frame and current frame), intra-frame feature (dense correspondences in current frame), and original semantic feature of current frame. Specifically, we establish a Center Prediction Branch to predict the center location of the foreground object in current frame and leverage the center point information as spatial guidance prior to enhance the inter-frame and intra-frame feature extraction, and thus the feature representation considerably focus on the foreground objects. Finally, we propose a Dynamic Information Fusion Module to automatically select relatively important features through three aforementioned different level features. Extensive experiments on DAVIS2016, Youtube-object, and FBMS datasets show that our proposed **F2Net** achieves the state-of-the-art performance with significant improvement.

## Introduction

Unsupervised video object segmentation (UVOS) aims to separate foreground objects from their background in a video sequence without any prior information. Due to the lack of prior knowledge about the foreground objects, this task not only suffers from the common challenges (e.g. object deformation and occlusion) in other video-related tasks, but also faces huge difficulties in accurately discovering the most prominent and distinct objects across video frames from a complex and diverse background.

<sup>†</sup>Equal contributions. \*Corresponding author.  
Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Compared to other methods, our network shows better segmentation performance under the challenging scenarios of visual similarity, occlusion, and appearance changing by considering the center information of the foreground.

Traditional methods tend to address this task by using handcrafted or learnable features, such as objectness (Zhang, Javed, and Shah 2013), motion boundary (Papazoglou and Ferrari 2013), saliency (Wang, Shen, and Porikli 2015), and trajectories (Ochs and Brox 2011). These are typically non-learning methods working in a purely unsupervised manner without using any training data. Recently, more research efforts have been devoted to tackling this task in deep learning frameworks, leading to an unsupervised solution (no annotation used for any testing frame). Many of these (Zhou et al. 2020; Cheng et al. 2017; Li et al. 2018c) employ two-stream networks to combine local motion and appearance information. Due to the usage of optical flow, they may fail to correctly infer the foreground when the object is occluded or nearly static (Tokmakov, Alahari, and Schmid 2017). Latest appearance matching based methods (Wang et al. 2019; Yang et al. 2019; Lu et al. 2019) are proposed to explore higher-order relationships among video frames, which obtain more optimal results from a global view by attention mechanisms. However, they lack robustness when there are visually similar object existed in the background.

For example, as shown in Figure 1, there are three challenging scenarios: foreground-background visual similarity, object occlusion, and appearance changing. For appearance matching method Anchor Diffusion Network (AnDiff) (Yang et al. 2019), it fails to separate the foreground camel from the background one in the first video as the two camels have similar appearance. It also can not handle the appearance changing problem as shown in the third video. For two stream network MAT (Zhou et al. 2020), it only segments

part of the occluded object in the second video because optical flow is not robustness to the object occlusion. The above methods pay less attention to the foreground object, leading to the inaccurate segmentation result. Considering center point can be taken as the spatial prior guidance (Zhou, Wang, and Krähenbühl 2019; Zhou, Koltun, and Krähenbühl 2020; Wang et al. 2020), we tend to firstly predict the center point of the primary object and then segment the mask from such point to its surroundings. Therefore, our network can focus more on the foreground object, and alleviate the visually similarity, occlusion and appearance changing problems.

Towards this end, we propose a novel *Focus on Foreground Network (F2Net)* for unsupervised video object segmentation, which exploits center point information to focus on the foreground object. Different from the common appearance matching based methods, we additionally establish a Center Prediction Branch to estimate the center location of the primary object. Then, we encode the predicted center point into a gauss map as the spatial guidance prior to enhance the intra-frame and inter-frame feature matching in our Center Guiding Appearance Diffusion Module, leading the model to focus on the foreground object. After the appearance matching process, we can get three kinds of information flows: inter-frame features, intra-frame features, and original semantic features of current frame. Instead of fusing these three features by simple concatenation like previous methods, an attention based Dynamic Information Fusion Module is developed to automatically select the most discriminative features across the three features, providing more optimal representations for final segmentation.

To summarize, our main contributions are three-folds:

- To the best of our knowledge, we are the first to take center point information into UVOS task for spatial guidance prior, which helps model focus on the foreground object. Specifically, our proposed Center Guiding Appearance Diffusion Module leverages the center prior in appearance matching procedure to extract foreground attentive representations.
- We develop a Dynamic Information Fusion Module to aggregate information from different level features, which generates more discriminative representations for the final foreground object segmentation.
- Extensive experiments are conducted on three popular UVOS benchmarks, DAVIS2016, FBMS, and Youtube-Objects. Compared to the state-of-the-art methods, we achieve the significant improvement with a large margin.

## Related Work

**Unsupervised Video Object Segmentation.** UVOS aims to automatically separate foreground object(s) from their background in a video without any human intervention. Early methods typically utilize handcrafted features (e.g., color, optical flow) (Papazoglou and Ferrari 2013; Faktor and Irani 2014; Tsai, Yang, and Black 2016; Hu, Huang, and Schwing 2018). Recently, benefiting from the large datasets (Perazzi et al. 2016), more research efforts have been devoted to tackling this task in deep learning frameworks. Tokmakov et al. (Tokmakov, Alahari, and Schmid 2017) proposed a purely

optical flow based network that discards appearance modelling and casts segmentation as foreground motion prediction, thus poorly deals with static foreground objects. To address this problem, two-stream networks are introduced to fuse appearance and motion information (Li et al. 2018a; Jain, Xiong, and Grauman 2017; Cheng et al. 2017; Li et al. 2018c; Zhou et al. 2020). However, above methods utilize optical flow information, and significantly suffer from not only the large computation of the optical flows, but also the deterioration in the quality of their predictions over time. Targeting this issue, several approaches (Wang et al. 2019; Yang et al. 2019; Lu et al. 2019; Chen et al. 2018; Fathi et al. 2017; Li et al. 2018b; Oh et al. 2019) tackle video object segmentation by simply learning similarities between pixel embeddings without motion contexts. AGNN (Wang et al. 2019) provides an unified, end-to-end trainable network to capture the higher-order correlated information with graph attention network. AnDiff (Yang et al. 2019) performs appearance similarity learning, feature propagation and binary segmentation in a single network. COS (Lu et al. 2019) utilizes co-attention to comprehensively use the rich, inherent correlation information within videos.

**Attention Mechanism.** Differentiable attentions have been widely used in recent neural networks for various tasks, such as visual question answering (Lu et al. 2016), human pose estimation (Chu et al. 2017; Su et al. 2019), and image classification (Hu, Shen, and Sun 2018; Li et al. 2019). It allows networks to focus on the most informative parts of the inputs. Latest appearance matching methods in UVOS (Wang et al. 2019; Yang et al. 2019; Lu et al. 2019) generally match the pixel embeddings across the whole frame based on attention mechanism, thus may wrongly segment the visually similar backgrounds. They also fuse the matched features by simple concatenation which may lose fine-grained details. To pay more attention on the foreground, we predict the foreground object center and encode it to a spatial gauss map. Such spatial guidance is injected into the attention mechanism to selectively focus on the foreground pixels. Besides, instead of fusing features by simple concatenation, we devise an attention based Dynamic Information Fusion Module to aggregate information from different level features.

## Method

### Overview

UVOS aims to automatically segment the primary object(s) from videos. To focus on the foreground and neglect the background, we propose the *Focus on Foreground Network F2Net* illustrated in Figure 2, which consists of three main parts: a Siamese Encoder Module, a Center Guiding Appearance Diffusion Module, and a Dynamic Information Fusion Module. Given a pair of frames, we first generate their embeddings by a siamese encoder. After that, the Center Guiding Appearance Diffusion Module first predicts the foreground object center point by a Center Prediction Branch, then encodes it into a gauss map as the spatial-prior condition during the appearance matching process. After appearance matching, we can get three kinds of features. At last, we apply a Dynamic Information Fusion Module to fuse differ-

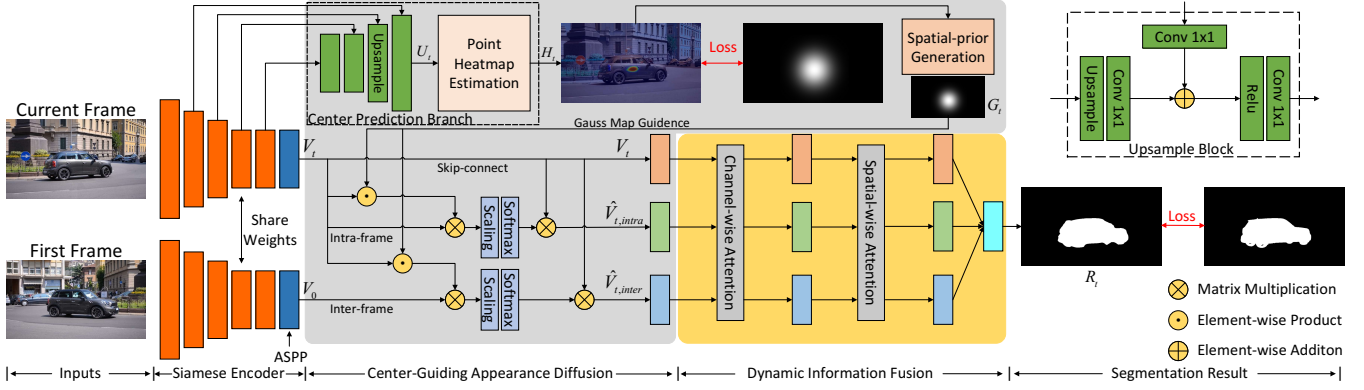


Figure 2: Overall pipeline of the proposed network architecture. A pair of frames is first fed into a siamese encoder to obtain the feature representations. After that, we develop a Center Guiding Appearance Diffusion Module to first predict the center point of the foreground object in current frame, and then generate a gauss map as spatial guidance prior for the following up appearance matching procedure. At last, we devise a Dynamic Information Fusion Module to aggregate different level features for foreground object segmentation.




Figure 3: Illustration of the heatmap estimation in current frame, which considers 1) spatial clues  $G_{t-1}$  from previous frame, and 2) semantic information  $U_t$  from current frame.

ent level matched features for final segmentation.

### Siamese Encoder

The siamese encoder takes a pair of RGB images as inputs, including a current frame  $I_t \in \mathbb{R}^{W \times H \times 3}$  and a reference frame  $I_0 \in \mathbb{R}^{W \times H \times 3}$ . Here, we take the first frame  $I_0$  as the reference frame because it is guaranteed to contain the foreground objects. The backbone network of the siamese encoder is DeepLabv3 (Chen et al. 2017), which consists of five convolution blocks (*res1*, *res2*, *res3*, *res4*, *res5*) from ResNet (He et al. 2016) and an atrous spatial pyramid pooling (ASPP) module. We denote the extracted embeddings of  $I_t, I_0$  as  $V_t, V_0 \in \mathbb{R}^{\frac{W}{8} \times \frac{H}{8} \times C}$ , where  $W, H, C$  are the width, height, and the number of channels, respectively.

### Center Guiding Appearance Diffusion

To focus more on the foreground object, we delve into the intra-inter features with center point information. In detail, we build a Center Prediction Branch to detect the object center point, and then utilize it as spatial guidance prior for further Spatial-Prior Guided Appearance Matching.

**Center Prediction Branch.** This branch tends to predict the center point of the foreground object. Specially, following

(Tompson et al. 2014), we transform the point prediction into a heatmap  $H_t$  estimation task, which consists of two main steps: *Feature Upsampling* and *Heatmap Generation*. In *Feature Upsampling*, we adopt an upsample module to efficiently merge features of different res-blocks (*res2*, *res3*, *res4*, *res5*) in different scales to enhance the information for high-resolution features. As shown in Figure 2, we illustrate the detailed structure of upsample block. The input of previous layer is first upsampled to the same size to the skip-connected features, and then both inputs are element-wise added. The final output of the whole upsampling blocks can be denoted as  $U_t \in \mathbb{R}^{\frac{W}{4} \times \frac{H}{4} \times D}$ , where  $D$  is the channel number. In *Heatmap Generation*, to predict the final point-wise heatmap  $H_t \in \mathbb{R}^{\frac{W}{4} \times \frac{H}{4} \times 1}$ , we consider two aspects as shown in Figure 3: 1) Spatial clues from previous frame: we propagate previous gauss map  $G_{t-1}$  to current frame for better locating the foreground (if the current frame is the first frame, we utilize a zero map as the previous gauss map). We concatenate the current embedding  $U_t$  and the spatial clues  $G_{t-1}$  to learn  $S, b \in \mathbb{R}^{\frac{W}{4} \times \frac{H}{4} \times 1}$ , which are scale and bias parameters (Yang et al. 2018) to control the weight to adjust the  $G_{t-1}$  under the guidance of the appearance information of current frame. We can formulate such process as:

$$S = \text{Conv2d}(\text{Concat}[U_t, G_{t-1}]), \quad (1)$$

$$b = \text{Conv2d}(\text{Concat}[U_t, G_{t-1}]), \quad (2)$$

$$\hat{G}_{t-1} = \text{Sigmoid}(S) \odot G_{t-1} + \text{Sigmoid}(b), \quad (3)$$

where  $\odot$  denotes the element-wise product. 2) Semantic information from current feature: Since  $U_t$  also contains enough information to predict the target center point, we can directly estimate the heatmap  $F_t$  on  $U_t$  by applying a  $3 \times 3$  convolutional layer with ReLU, followed by a  $1 \times 1$  convolutional layer. At last, we add this two estimation branch into one with a sigmoid function by:

$$H_t = \text{Sigmoid}(\hat{G}_{t-1} + F_t). \quad (4)$$

To choose the best center point  $o_t = (x_t, y_t) \in \mathbb{R}^2$  from  $H_t$ , instead of directly choosing the point with the maximum score, we additionally consider the motion clues across the sequential frames for better accuracy. We first rank top  $K$  points using NMS (Lin et al. 2017) from  $H_t$  as candidates, and then utilize a motion mechanism (Xu et al. 2019) to predict a coarse center point  $p_t = p_{t-1} + \frac{1}{n} \sum_{m=t-n}^{t-1} (p_m - p_{m-1})$  using  $n$  history object centers in previous frames. At last, we compute the distance from each candidate to  $p_t$ , and choose the closest one as the final center  $o_t$  of the foreground object in current frame. In our experiments, we find it can achieve better performance than the maximum strategy.

**Spatial-Prior Guided Appearance Matching.** To determine the foreground object, there are two essential properties: 1) distinguishable in an individual frame (locally saliency), and 2) frequently appearing throughout the video sequence (globally consistent). To achieve the first goal, we apply a non-local operation (Wang et al. 2018) on the current feature  $V_t$  to locate the salient object in an intra-wise matching way. To achieve our second goal, we utilize another non-local operation on both current and the reference features  $V_t, V_0$  to capture the inter-frame correlated information for alleviating appearance drift. We also employ a skip-connection on current feature  $V_t$  to preserve the semantic information. All three features contain different level information as shown in Figure 2. To focus on the foreground, we encode the predicted center point  $o_t$  into a gauss map  $G_t$  as spatial guidance prior for the intra-frame and inter-frame appearance matching. Compared to other appearance matching methods, the key difference of our method is that the encoded feature representation is weighted by the gauss-guided spatial prior.

As shown in Figure 2, given the intra-frame and inter-frame feature pairs  $(V_t, V_t)$  and  $(V_0, V_t)$ , we first flatten  $V_0, V_t$  into shape  $\frac{WH}{64} \times C$ , then compute the gauss-guided correlation matrices  $M_{intra}, M_{inter} \in \mathbb{R}^{\frac{WH}{64} \times \frac{WH}{64}}$  as:

$$M_{intra} = \text{Softmax}\left(\frac{1}{\sqrt{C}} V_t (V_t \odot G_t)^T\right), \quad (5)$$

$$M_{inter} = \text{Softmax}\left(\frac{1}{\sqrt{C}} V_0 (V_t \odot G_t)^T\right), \quad (6)$$

where  $C$  is the channel number of  $V_t$ . After that, we reconstruct feature  $\hat{V}_t$  in which the pixel embeddings are weighted according to their similarity with the foreground:

$$\hat{V}_{t,intra} = M_{intra} V_t, \quad (7)$$

$$\hat{V}_{t,inter} = M_{inter} V_t. \quad (8)$$

At last, we can get three kinds of information flows and reshape them back to a 3D tensor with the size of  $\frac{W}{8} \times \frac{H}{8} \times C$  for further segmentation: 1) initial current feature  $V_t$  by a skip connection (He et al. 2016), 2) intra-frame feature  $\hat{V}_{t,intra}$  for intra-frame discriminability and 3) inter-frame feature  $\hat{V}_{t,inter}$  for inter-frame consistency.

## Dynamic Information Fusion

Original current feature  $V_t$  only contains a coarse clue for inferring the target foreground object without a global view.




Figure 4: Scheme of the Dynamic Information Fusion Module, in which we apply attentions on channel and spatial dimensions for different level features aggregation.

Intra-frame feature  $\hat{V}_{t,intra}$  contains more accurate salient object information in current frame, but fails to address the appearance changes in a video sequence. Inter-frame feature  $\hat{V}_{t,inter}$  contains more contexts to adapt to the appearance changes of the foreground objects. Instead of directly fusing these three features by simple concatenation (Lu et al. 2019; Yang et al. 2019), we need to selectively aggregate them to generate more discriminative features. According to the above analysis, we develop a Dynamic Information Fusion Module to emphasize meaningful features along both channel and spatial dimensions with attention mechanism.

**Channel-wise attention.** The basic idea of channel-wise attention is to use gates to control the information flows from three-level features in channel dimension. To achieve this goal, the gates need to integrate information from all features. As shown in Figure 4, we first fuse three kinds of features by an element-wise addition as:

$$\tilde{V}_t = V_t + \hat{V}_{t,intra} + \hat{V}_{t,inter}. \quad (9)$$

Then we utilize global average pooling on  $\tilde{V}_t$  to conduct spatial squeeze, where each element is calculated by shrinking  $\tilde{V}_t$  through spatial dimensions. We further apply a simple fully connected (FC) on the squeezed feature to enable the guidance for adaptive selection. After that, we employ another three FC layers to generate the gate vectors  $\{Z_1, Z_2, Z_3\} \in \mathbb{R}^{1 \times 1 \times C}$  for  $\{V_t, \hat{V}_{t,intra}, \hat{V}_{t,inter}\}$ , respectively, and utilize channel-wise softmax function to generate adaptive weights  $\{W_1, W_2, W_3\} \in \mathbb{R}^{1 \times 1 \times C}$  corresponding to different level features  $\{Z_1, Z_2, Z_3\}$  as:

$$W_i^c = \frac{\exp(Z_i^c)}{\sum_{j=1}^3 \exp(Z_j^c)}, \quad i \in \{1, 2, 3\}, \quad (10)$$

in which,  $W_i^c$  represents relative importance of feature  $Z_i$  at channel  $c \in C$ , and  $\sum_{i=1}^3 W_i^c = 1$ . The gated feature maps for three information flows can be formulated as:

$$V_t = V_t \odot W_1, \quad \hat{V}_{t,intra} = \hat{V}_{t,intra} \odot W_2, \quad (11)$$

$$\hat{V}_{t,inter} = \hat{V}_{t,inter} \odot W_3.$$

**Spatial-wise attention.** Similar to the channel-wise attention, we also first fuse three features by element-wise addition. To strengthen the information on the spatial dimension, we then employ  $1 \times 1$  convolutional filters to conduct

the channel squeeze and compress the channels to 3, where the feature map of each channel corresponding to the spatial weights for each level feature. After that, a softmax operation is conducted across channels to rescale activations and the rescaled activations are sliced along the channel dimension to generate pixel-wise adaptive weights. At last, we sum the three features to one according to the generated weights.

After the channel-wise and spatial-wise attention modules, the information flows from three-level input features can be adaptively aggregated into one feature map. Like most of previous works, we employ a decoder which consists of two convolutional layers to generate the final binary mask result  $\mathbf{R}_t \in \mathbb{R}^{W \times H}$ .

## Implementation Details

**Training loss.** Given a pair of input images, our **F2Net** predicts a center point heatmap  $\mathbf{H}$  and a mask result  $\mathbf{R}$  of the current frame. For point heatmap estimation, we apply an element-wise focal loss (Lin et al. 2017) on the heatmap  $\mathbf{H}$  in Eq.(4) and the corresponding ground-truth gauss map  $\hat{\mathbf{H}}$ :

$$\mathcal{L}_f = \sum_{(x,y)} \begin{cases} (1 - H_{(x,y)})^\alpha \log(H_{(x,y)}), & \text{if } \hat{H}_{(x,y)} = 1 \\ (1 - \hat{H}_{(x,y)})^\beta (H_{(x,y)})^\alpha \log(1 - H_{(x,y)}), & \text{else} \end{cases} \quad (12)$$

where  $H_{(x,y)}$  is the score at location  $(x, y)$  in the predicted heatmap  $\mathbf{H}$ , and we set  $\alpha$  as 2 and  $\beta$  as 4 following the default setting in (Law and Deng 2018). For segmentation mask result, we employ the binary cross entropy loss on the predict mask  $\mathbf{R}$  and the ground truth  $\hat{\mathbf{R}}$  as:

$$\mathcal{L}_b = - \sum_{(x,y)} \hat{R}_{(x,y)} \log(R_{(x,y)}) + (1 - \hat{R}_{(x,y)}) \log(1 - R_{(x,y)}). \quad (13)$$

The overall loss function can be obtained by:

$$\mathcal{L} = \mathcal{L}_f + \mathcal{L}_b. \quad (14)$$

**Training settings.** Following (Wang et al. 2019; Lu et al. 2019), we adopt two alternated steps to train our model. In the static-image iteration, we utilize image saliency dataset MSRA10K (Cheng et al. 2014) to fine-tune the DeepLabV3 based feature embedding module and the Center Prediction Branch. This allows the backbone to extract more discriminative foreground features and the center prediction module to locate the center of salient object more accurately. Meanwhile, in the dynamic-video iteration, we train the whole model with the training set in DAVIS16 (Perazzi et al. 2016). During training the Center Guiding Appearance Diffusion module, we first separately train the Center Prediction Branch and Spatial-Prior Guided Appearance Matching in parallel in the first 20 epochs, where we feed the appearance matching module with the center point of ground truth. Then, in latter epochs, we feed the appearance matching module with the predicted point from the Center Prediction Branch to jointly train them. For the above alternated training process, the size of input RGB frame is  $473 \times 473 \times 3$ . The entire network is trained using the SGD optimizer with an initial learning rate of  $2.5 \times 10^{-4}$ . We set the batchsize as 16. All the experiments are conducted using 4 V100 GPUs on a server. The overall training time is about 9 hours, and it takes about 0.1s with one image in a forward pass.

Network Variant	Mean $\mathcal{J} \uparrow$	$\Delta \mathcal{J}$	Mean $\mathcal{F} \uparrow$	$\Delta \mathcal{F}$
Baseline	77.6	-5.5	76.8	-7.6
+ AD	79.6	-3.5	79.0	-5.4
+ CGAD	82.5	-0.6	82.9	-1.5
+ CGAD&SA	82.7	-0.4	83.4	-1.0
+ CGAD&CA	82.8	-0.3	83.6	-0.8
+ CGAD&SCA	83.0	-0.1	84.2	-0.2
+ CGAD&CSA	<b>83.1</b>	-	<b>84.4</b>	-

Table 1: Ablation study of the proposed model on DAVIS-2016 dataset, measured by Mean  $\mathcal{J}$  and Mean  $\mathcal{F}$ .

Mechanism	Mean $\mathcal{J} \uparrow$	Mean $\mathcal{F} \uparrow$
maximum	82.4	84.2
motion based (Xu et al. 2019)	<b>83.1</b>	<b>84.4</b>

Table 2: Ablation study of different strategies to choose the center point from the predicted heatmap.

## Experiments

### Experimental Setup

We conduct experiments on three well-known datasets: DAVIS2016 (Perazzi et al. 2016), Youtube-Objects (Prest et al. 2012), and FBMS (Ochs, Malik, and Brox 2013).

**DAVIS2016** is a challenging video object segmentation dataset which consists of 50 videos in total (30 for training and 20 for validation) with pixel-wise annotations for every frame. Three evaluation criteria are used following the standard evaluation protocol (Perazzi et al. 2016): region similarity  $\mathcal{J}$ , boundary accuracy  $\mathcal{F}$ , and time stability  $\mathcal{T}$ .

**Youtube-Objects** contains 126 video sequences which belong to 10 objects categories with more than 20,000 frames in total. Following its protocol, we use the region similarity  $\mathcal{J}$  to measure the segmentation performance.

**FBMS** is comprised of 59 video sequences (29 training videos and 30 test videos). The ground-truth of FBMS is sparsely labeled. Following (Yang et al. 2019; Lu et al. 2019; Zhou et al. 2020), we use region similarity  $\mathcal{J}$  to evaluate our method on test set without training.

### Ablation Study

We conduct the ablation study on DAVIS2016 dataset as shown in Table 1 and 2, where the baseline model is a the DeepLabV3 network.

**How does the gauss based spatial guidance help?** We first investigate the effectiveness of the proposed Center Guiding Appearance Diffusion (CGAD) Module, which relies on the center point estimated by the Center Prediction Branch. We denote AD as normally appearance diffusion (AnDiff without multi-scale and pruning strategies), and CGAD as the center-prior guided one. Compare to the baseline model, as shown in Table 1, AD brings the improvement of 2.0 on  $\mathcal{J}$  and 2.2 on  $\mathcal{F}$  which indicates the effectiveness of the general appearance matching based method. Compare to it, our CGAD outperforms AD by 2.9 on  $\mathcal{J}$  and 3.9 on  $\mathcal{F}$  with a large margin. It demonstrates the effectiveness of the gauss based spatial prior, which helps to focus on the foreground

Method	FSEG	UOVOS	LVO	ARP	PDB	LSMO	MoA	EpO	AGS	COS	AnDiff	MAT	Ours
$\mathcal{J}\&\mathcal{F}$ Mean $\uparrow$	68.0	70.9	74.0	73.4	75.9	77.1	77.3	78.1	78.6	80.0	81.1	81.5	<b>83.7</b>
Mean $\uparrow$	70.7	73.9	75.9	76.2	77.2	78.2	77.2	80.6	79.7	80.5	81.7	82.4	<b>83.1</b>
$\mathcal{J}$ Recall $\uparrow$	83.5	88.5	89.1	91.1	90.1	89.1	87.8	95.2	91.1	93.1	90.9	94.5	<b>95.7</b>
Decay $\downarrow$	1.5	0.6	<b>0.0</b>	7.0	0.9	4.1	5.0	2.2	1.9	4.4	2.2	5.5	<b>0.0</b>
Mean $\uparrow$	65.3	68.0	72.1	70.6	74.5	75.9	77.4	75.5	77.4	79.5	80.5	80.7	<b>84.4</b>
$\mathcal{F}$ Recall $\uparrow$	73.8	80.6	83.4	83.5	84.4	84.7	84.4	87.9	85.8	89.5	85.1	90.2	<b>92.3</b>
Decay $\downarrow$	1.8	0.7	1.3	7.9	<b>-0.2</b>	3.5	3.3	2.4	1.6	5.0	0.6	4.5	0.8
$\mathcal{T}$ Mean $\downarrow$	32.8	39.0	26.5	39.3	29.1	21.2	27.9	19.3	26.7	<b>18.4</b>	21.4	21.6	20.9

Table 3: Quantitative results of UVOS methods on the DAVIS2016 validation set. All the results are borrowed from the public leaderboard maintained by the DAVIS challenge. The best scores are marked in **bold**.

Method	FSEG	LVO	PDB	AGS	COS	MAT	Ours
Airplane	81.7	<b>86.2</b>	78.0	87.7	81.1	72.9	85.8
Bird	63.8	81.0	80.0	76.7	75.7	77.5	<b>82.8</b>
Boat	72.3	68.5	58.9	72.2	71.3	66.9	<b>81.9</b>
Car	74.9	69.3	76.5	78.6	77.6	79.0	<b>81.4</b>
Cat	68.4	58.8	63.0	69.2	66.5	<b>73.7</b>	70.2
Cow	68.0	68.5	64.1	64.6	69.8	67.4	<b>71.0</b>
Dog	69.4	61.7	70.1	73.3	<b>76.8</b>	75.9	75.8
Horse	60.4	53.9	67.6	64.4	67.4	63.2	<b>75.4</b>
Motorbike	62.7	60.8	58.4	62.1	67.7	62.6	<b>71.8</b>
Train	62.2	<b>66.3</b>	35.3	48.2	46.8	51.0	59.6
Mean $\mathcal{J}\uparrow$	68.4	67.5	65.5	69.7	70.5	69.0	<b>75.6</b>

Table 4: Quantitative performance of each category on Youtube-Objects with the Mean  $\mathcal{J}$ .

and thus alleviates the challenges in common appearance matching based methods.

**How does the dynamic information fusion help?** The Dynamic Information Fusion Module consists of channel-wise and spatial-wise attention mechanisms. Here, we do the comparison on four variants: spatial attention only (SA), channel attention only (CA), first spatial attention then channel attention (SCA), and first channel attention then spatial attention (CSA). Detailed experiments on different variant attentions can be found in Table 1. Compared to CGAD, SA can bring improvements of 0.2 on  $\mathcal{J}$ , 0.5 on  $\mathcal{F}$ , and CA can bring improvements of 0.3 on  $\mathcal{J}$ , 0.7 on  $\mathcal{F}$ . To investigate the performance on different combinations of the two attentions, we conduct the experiments on SCA and CSA. The results show that the channel-spatial attention (CSA) achieves the best performance of  $\mathcal{J} = 83.1$  and  $\mathcal{F} = 84.4$ .

**How to select the center point from a heatmap?** Besides, we also do the ablation study on different strategies to choose the center point from the heatmap  $H$  in Eq. (4). Specially, we compare two methods in Table 2: 1) We directly choose the center point in heatmap with the maximum score. 2) We exploit a motion mechanism (Xu et al. 2019) to choose the center point with motion history information under the parameter settings  $K = 5, n = 10$ . We can find that the motion mechanism performs relatively better. It reveals that the motion mechanism can help locate more accurate position of object based on the motion history, whether objects move fast or slow in the video.

## Comparison with State-of-the-arts

**Evaluation on DAVIS2016.** We compare our **F2Net** with the top performing UVOS methods in the public leaderboard on DAVIS2016 dataset. As shown in Table 3, our method outperforms all the reported methods across most metrics. Compared with the second best method MAT (Zhou et al. 2020), our model achieves gains of 2.2 in terms of  $\mathcal{J}\&\mathcal{F}$  Mean. In detail, we obtains improvements of 0.7 and 3.7 on  $\mathcal{J}$  Mean and  $\mathcal{F}$  Mean, respectively. Compared to appearance matching methods COS (Lu et al. 2019) and AnDiff (Yang et al. 2019), we outperform them both on  $\mathcal{J}$  Mean and  $\mathcal{F}$  Mean by a large margin. Compared to methods like MoA (Siam et al. 2019) and EPO (Faisal et al. 2019) which utilize both appearance information and motion cues, our model outperforms them by only utilizing appearance information. In our experiments, we find that the above methods may fail to distinguish the visually similar pixels from both foreground and background. Due to the center based gauss map, our model consider spatial prior during the appearance matching procedure, which helps filter out the background noise and focus more on the foreground object boundary construction (improving  $\mathcal{F}$ ). Besides, the Dynamic Information Fusion Module also can aggregate discriminative features across different level for the final segmentation.

**Evaluation on Youtube-Objects.** Table 4 illustrates the results of all compared methods for different categories on Youtube-Objects dataset. Our approach brings improvement of 5.1 on  $\mathcal{J}$  than the second best method COS (Lu et al. 2019) by a large margin. It is also worth to note that we outperform all compared methods on almost all categories. The main reason lies in two folds: First, for optical guided methods MAT, FSEG (Jain, Xiong, and Grauman 2017) and LVO (Tokmakov, Alahari, and Schmid 2017), sequences in the Airplane and Boat categories contain objects that have quick appearance variation or move slowly. Both factors result in inaccurate estimation of optical flow. Compared to them, our matching based framework can handle these scenarios well. Second, compared to other matching based methods COS, our estimated center point provides the spatial prior for better segmentation focusing more on the actual foreground.

**Evaluation on FBMS.** As shown in Table 5, we also conduct experiments on FBMS for completeness. Compared to others, our method performs the best result with 77.5 over the Mean  $\mathcal{J}$ , outperforming the second best one by 1.4. Con-




Figure 5: Qualitative results on three datasets. From top to bottom: *breakdance* from the DAVIS2016 dataset, *cats01* from the FBMS dataset, and *dog0006* from the Youtube-Objects dataset.

Method	NLC	FST	SFL	APR	MSTP	FSEG
Mean $\mathcal{J}$ $\uparrow$	44.5	55.5	56.0	59.8	60.8	68.4
Method	IET	OBN	PDB	COS	MAT	Ours
Mean $\mathcal{J}$ $\uparrow$	71.9	73.9	74.0	75.6	76.1	<b>77.5</b>

Table 5: Quantitative results on FBMS test set over Mean  $\mathcal{J}$ .

sidering lots of foreground objects in FBMS share similar appearance with the background, our Center Guiding Appearance Diffusion Module exploits center information to focus on the foreground objects and filter out the visually similar background ones for better segmentation.

## Qualitative Results

**Does center prediction branch estimate the center point well?** To investigate the performance of Center Prediction Branch, we give some visualization results on the generated center point heatmap, especially for the first frame of each video sequence. As shown in Figure 6, there are four challenging sequences in which the surroundings has a similar appearance to the foreground object. Without any previous motion history, our Center Prediction Branch still can estimate relatively accurate object center point of the first frame mainly based on the semantic features. It demonstrates that our Center Prediction Branch can effectively capture the contextual information across the frame to locate the target point. In the latter frames, besides the current semantic features, the center point based gauss map from previous frame is additionally fed to refine the point position, which provides more precisely center point estimation.

**Segmentation visualization.** Figure 5 shows qualitative results sampled from the three datasets. The *breakdance* sequence from DAVIS2016 contains many challenging factors, such as fast motion, deformation and multiple instances of the same category. We can find that our method is robust to such complex scenarios and can accurately segment out primary objects from the cluttered background. The effectiveness is further proved in the *cat01* sequence of FBMS dataset. In addition, our method performs well in the *dog0006* sequence of Youtube-Objects, in which the target




Figure 6: Visualization of center point heatmaps on each frame of the DAVIS2016 dataset. From top to down: *camel*, *dog*, *goat*, and *parkour*. Although the first frame has no history motion knowledge of the foreground object, our Center Prediction Branch can still estimate accurate point result among noisy neighbors. The other frames receive additional point information from previous frame and thus predict more accurate center point.

suffers from large scale variations.

## Conclusion

In this paper, we propose a novel *Focus on Foreground Network (F2Net)* for unsupervised video object segmentation. Compared to recent appearance matching based methods, we additionally estimate the center point of the foreground object and encode it into a gauss map as spatial guidance for the appearance matching procedure. This Center Guiding Appearance Diffusion Module is flexible and can be easily adapted to other segmentation frameworks. We also develop a Dynamic Information Fusion Module to aggregate multi-level features for generating more discriminative features for final segmentation. Extensive results on three public datasets demonstrate the effectiveness of our proposed method.

## References

- Chen, L.-C.; Papandreou, G.; Schroff, F.; and Adam, H. 2017. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*.
- Chen, Y.; Pont-Tuset, J.; Montes, A.; and Van Gool, L. 2018. Blazingly fast video object segmentation with pixel-wise metric learning. In *CVPR*, 1189–1198.
- Cheng, M.-M.; Mitra, N. J.; Huang, X.; Torr, P. H.; and Hu, S.-M. 2014. Global contrast based salient region detection. *TPAMI* 37(3):569–582.
- Cheng, J.; Tsai, Y.-H.; Wang, S.; and Yang, M.-H. 2017. Segflow: Joint learning for video object segmentation and optical flow. In *ICCV*, 686–695.
- Chu, X.; Yang, W.; Ouyang, W.; Ma, C.; Yuille, A. L.; and Wang, X. 2017. Multi-context attention for human pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1831–1840.
- Faisal, M.; Akhter, I.; Ali, M.; and Hartley, R. 2019. Exploiting geometric constraints on dense trajectories for motion saliency. In *WACV*.
- Faktor, A., and Irani, M. 2014. Video segmentation by non-local consensus voting. In *BMVC*, volume 2, 8.
- Fathi, A.; Wojna, Z.; Rathod, V.; Wang, P.; Song, H. O.; Guadarrama, S.; and Murphy, K. P. 2017. Semantic instance segmentation via deep metric learning. *arXiv preprint arXiv:1703.10277*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *CVPR*, 770–778.
- Hu, Y.-T.; Huang, J.-B.; and Schwing, A. G. 2018. Unsupervised video object segmentation using motion saliency-guided spatio-temporal propagation. In *ECCV*, 786–802.
- Hu, J.; Shen, L.; and Sun, G. 2018. Squeeze-and-excitation networks. In *CVPR*, 7132–7141.
- Jain, S. D.; Xiong, B.; and Grauman, K. 2017. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *CVPR*, 2117–2126.
- Law, H., and Deng, J. 2018. Cornernet: Detecting objects as paired keypoints. In *ECCV*, 734–750.
- Li, G.; Xie, Y.; Wei, T.; Wang, K.; and Lin, L. 2018a. Flow guided recurrent neural encoder for video salient object detection. In *CVPR*, 3243–3252.
- Li, S.; Seybold, B.; Vorobyov, A.; Fathi, A.; Huang, Q.; and Jay Kuo, C.-C. 2018b. Instance embedding transfer to unsupervised video object segmentation. In *CVPR*, 6526–6535.
- Li, S.; Seybold, B.; Vorobyov, A.; Lei, X.; and Jay Kuo, C.-C. 2018c. Unsupervised video object segmentation with motion-based bilateral networks. In *ECCV*, 207–223.
- Li, X.; Wang, W.; Hu, X.; and Yang, J. 2019. Selective kernel networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 510–519.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *ICCV*, 2980–2988.
- Lu, J.; Yang, J.; Batra, D.; and Parikh, D. 2016. Hierarchical question-image co-attention for visual question answering. In *Advances in Neural Information Processing Systems (NIPS)*, 289–297.
- Lu, X.; Wang, W.; Ma, C.; Shen, J.; Shao, L.; and Porikli, F. 2019. See more, know more: Unsupervised video object segmentation with co-attention siamese networks. In *CVPR*, 3623–3632.
- Ochs, P., and Brox, T. 2011. Object segmentation in video: a hierarchical variational approach for turning point trajectories into dense regions. In *ICCV*, 1583–1590.
- Ochs, P.; Malik, J.; and Brox, T. 2013. Segmentation of moving objects by long term video analysis. *TPAMI* 36(6):1187–1200.
- Oh, S. W.; Lee, J.-Y.; Xu, N.; and Kim, S. J. 2019. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 9226–9235.
- Papazoglou, A., and Ferrari, V. 2013. Fast object segmentation in unconstrained video. In *ICCV*, 1777–1784.
- Perazzi, F.; Pont-Tuset, J.; McWilliams, B.; Van Gool, L.; Gross, M.; and Sorkine-Hornung, A. 2016. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 724–732.
- Prest, A.; Leistner, C.; Civera, J.; Schmid, C.; and Ferrari, V. 2012. Learning object class detectors from weakly annotated video. In *CVPR*, 3282–3289.
- Siam, M.; Jiang, C.; Lu, S.; Petrich, L.; Gamal, M.; Elhoseiny, M.; and Jagersand, M. 2019. Video object segmentation using teacher-student adaptation in a human robot interaction (hri) setting. In *ICRA*, 50–56.
- Su, K.; Yu, D.; Xu, Z.; Geng, X.; and Wang, C. 2019. Multi-person pose estimation with enhanced channel-wise and spatial information. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 5674–5682.
- Tokmakov, P.; Alahari, K.; and Schmid, C. 2017. Learning motion patterns in videos. In *CVPR*, 3386–3394.
- Tompson, J. J.; Jain, A.; LeCun, Y.; and Bregler, C. 2014. Joint training of a convolutional network and a graphical model for human pose estimation. In *NIPS*, 1799–1807.
- Tsai, Y.-H.; Yang, M.-H.; and Black, M. J. 2016. Video segmentation via object flow. In *CVPR*, 3899–3908.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *CVPR*, 7794–7803.
- Wang, W.; Lu, X.; Shen, J.; Crandall, D. J.; and Shao, L. 2019. Zero-shot video object segmentation via attentive graph neural networks. In *ICCV*, 9236–9245.
- Wang, Y.; Xu, Z.; Shen, H.; Cheng, B.; and Yang, L. 2020. Centermask: single shot instance segmentation with point representation. In *CVPR*, 9313–9321.
- Wang, W.; Shen, J.; and Porikli, F. 2015. Saliency-aware geodesic video object segmentation. In *CVPR*, 3395–3402.
- Xu, S.; Liu, D.; Bao, L.; Liu, W.; and Zhou, P. 2019. Mhp-



- vos: Multiple hypotheses propagation for video object segmentation. In *CVPR*, 314–323.
- Yang, L.; Wang, Y.; Xiong, X.; Yang, J.; and Katsaggelos, A. K. 2018. Efficient video object segmentation via network modulation. In *CVPR*, 6499–6507.
- Yang, Z.; Wang, Q.; Bertinetto, L.; Hu, W.; Bai, S.; and Torr, P. H. 2019. Anchor diffusion for unsupervised video object segmentation. In *ICCV*, 931–940.
- Zhang, D.; Javed, O.; and Shah, M. 2013. Video object segmentation through spatially accurate and temporally dense extraction of primary object regions. In *CVPR*, 628–635.
- Zhou, T.; Wang, S.; Zhou, Y.; Yao, Y.; Li, J.; and Shao, L. 2020. Motion-attentive transition for zero-shot video object segmentation. In *AAAI*, volume 2, 3.
- Zhou, X.; Koltun, V.; and Krähenbühl, P. 2020. Tracking objects as points. *arXiv:2004.01177*.
- Zhou, X.; Wang, D.; and Krähenbühl, P. 2019. Objects as points. *arXiv preprint arXiv:1904.07850*.