

# MSN: Efficient Online Mask Selection Network for Video Instance Segmentation

Vidit Goel<sup>1,2,4</sup>, Jiachen Li<sup>1,3\*</sup>, Shubhika Garg<sup>2\*</sup>, Harsh Maheshwari<sup>2</sup>, Humphrey Shi<sup>1,3,4</sup>

<sup>1</sup>SHI Lab @ University of Oregon, <sup>2</sup>Indian Institute of Technology Kharagpur, <sup>3</sup>University of Illinois at Urbana-Champaign, <sup>4</sup>Picsart AI Research (PAIR)

## Abstract

*In this work we present a novel solution for Video Instance Segmentation (VIS), that is automatically generating instance level segmentation masks along with object class and tracking them in a video. Our method improves the masks from segmentation and propagation branches in an online manner using the **Mask Selection Network (MSN)** hence limiting the noise accumulation during mask tracking. We propose an effective design of MSN by using patch-based convolutional neural network. The network is able to distinguish between very subtle differences between the masks and choose the better masks out of the associated masks accurately. Further we make use of temporal consistency and process the video sequences in both forward and reverse manner as a post processing step to recover lost objects. The proposed method can be used to adapt any video object segmentation method for the task of VIS. Our method achieves a score of 49.1 mAP on 2021 YouTube-VIS Challenge and was ranked third place among more than 30 global teams. Our code will be available at: <https://github.com/SHI-Labs/Mask-Selection-Networks>.*

## 1. Introduction

Video understanding has gained a lot of attention with the abundance of data as well as methods for analysing and processing videos. In this work, we particularly focus on Video Instance Segmentation. The task here is to perform instance segmentation and classification on each frame and track the instances through out the video. A very similar task is that of unsupervised video object segmentation [5]. The major difference is that there is no need to predict object classes. Due to the large similarity between the two tasks we develop our solution based on the best performing method FrameSelect [13].

The extra challenges that are present in VIS compared to unsupervised video object segmentation are majorly due to

classification task. In case of VIS the method is penalized for false positives, hence even if a method produces very good masks if the classification is wrong the score would be low. Due to this only the objects whose classification score is very good should be added as a true detection in a frame. However, this can lead to under detection of instances. Other problems which are common to both the task are also present, such as occlusion, motion blur, re-identification, presence of multiple objects of similar visual appearances and absence of information of objects that needs to be tracked. Keeping the challenges in mind we adapt and improve the FrameSelect [13] method.

Our method consists of three stages 1) instance segmentation 2) efficient online mask selection and temporal propagation 3) post processing. From the discussion above, it is clear that object classification is very important for good performance. Hence, we develop a strong instance segmentation module based on Swin transformer [25] trained on various open source datasets including ImageNet [17] MS COCO [23] and OpenImages [18]. In stage 2, we handle the task related to mask propagation and object addition. In this work we propose Mask selection network (MSN) to choose better masks from the available masks of an instance in an online manner. This limits the noise accumulation at the time of mask propagation. Finally, in the stage 3 we use temporal consistency of a video sequence to handle under detection of objects. Most of the recent approaches [22, 2, 26] divide the task into instance segmentation and mask propagation which is similar to ours. However, our method is significantly different from them as it provides a mechanism to assess the quality of masks at the time of propagation of masks and choose the best mask from the available masks. Further, we provide an efficient yet effective post processing method. Our method proposes a general pipeline to adapt any video object segmentation method for the task of VIS.

---

\*Equal contribution

## 2. Related work

**Video Object Segmentation** In semi-supervised setting the annotations of first frame are given and they need to segment and track those objects throughout the video in a class-agnostic manner. Many approaches in this field deal with online learning and fine tuning [27, 32, 38] using the ground truth information from the first frame. These types of methods show impressive performance, however they are not suitable for real time video object segmentation as they take a lot of time. Some other variations of method involve mask propagation [30, 20, 29] using previous frames and feature matching [31, 28, 42, 41] of the embedding of the current frame with the stored templates. CFBI [42] does local feature matching at multiple scales and treats foreground and background as separate classes, hence, giving better results.

**Video Instance Segmentation** The works in this category can be divided into two categories, i) treats instance segmentation and tracking as two separate tasks [26, 33] and ii) has end to end trainable pipeline [3, 6, 1, 21, 40, 22, 12, 11]. The methods in first category fine-tune each stage individually and are able to achieve good results. For example Ensemble VIS [26] has separate modules for classification and segmentation and finally tracks using a ReID module. The methods in second category train their method in end to end manner. Methods such as [3, 22] are based of propagating masks detected in previous frame whereas methods like [40, 6, 10] solve the task by first detecting and then tracking the objects. Different from others [1] treated the video as 3D volumes and track and segment objects in a single stage. Our method falls in the first category.

**Instance Segmentation** Previous instance segmentation methods can be categorized into two groups: two-stage methods like Mask RCNN [15] and one-stage methods like TensorMask [9, 39]. For two-stage methods, they follow the idea from two-stage object detector that localize proposals then segment instances from proposals. Mask RCNN adds a segmentation branch along with detection head. Following methods HTC [7] and CNet [24] employ multi-stage head for accurate instance segmentation. For one-stage methods, CenterMask [19] and Yolact [4] achieve real-time instance segmentation without generating proposals. SOLO [34] and SOLOv2 [36] further segments each instance from individual feature map. Recently, with the development of vision transformer [37, 14], using transformer-based model as backbone like Swin [25] for instance segmentation achieves new state-of-the-art performance on COCO [23] benchmark.

## 3. Method

### 3.1. Overall framework

Our overall framework leverages an effective multi-stage design for video instance segmentation and consists of 3 stages. Its overall pipeline is similar to Frame Select [13] which was proposed for unsupervised video object segmentation [5]. In our stage 1, instance segmentation is performed using Swin transformer backbone [25] based cascaded mask R-CNN. In stage 2 we adapt a video object segmentation method for VIS task. Specifically, we used CFBI [42] and initialized it using the masks generated in stage 1. The masks are improved in an online manner using a Mask selection network (MSN) which enables to capture very fine details. MSN is able to select better masks even with very subtle differences between the masks. In stage 3, we propose an efficient yet effective post processing pipeline which utilizes temporal consistency in a video in both forward and backward direction recover and track missing objects.

#### Stage 1: Instance segmentation

For instance segmentation, we develop our model based on Swin [25] transformer and cascade mask R-CNN with mmdetection [8] toolbox. We use ImageNet [17] pretrained Swin-base model as our backbone and build a cascade mask R-CNN model with 3 stages. The VIS dataset contains 40 classes and we add images from MS COCO [23] and OpenImages [18] for data augmentation. Since COCO contains 80 classes and OpenImages contains hundreds of classes, we map similar classes from these two datasets to VIS dataset. We also observe that data distribution is imbalanced in VIS dataset and to solve this issue, we add a data balanced sampler to improve sampling rate of rare classes and improve instance segmentation result. Only the instance masks which have object scores above a certain threshold are passed to stage 2.

#### Stage 2: Efficient online mask selection and temporal propagation

In order to propagate the masks we adapt a video object segmentation method CFBI [42] for our task. In case of video object segmentation the annotations of the objects are given when they first appear in a video. In this case, we initialize CFBI by using masks from stage 1 as annotations for first frame. Using CFBI gives 2 major benefits, first one is being able to generate instance masks by incorporating the temporal information and second it helps in tracking. However, there can be cases when the annotations of first frame are noisy which can lead to noise accumulation and improper propagation of masks. There can also be cases when object gets detected in some intermediate frames. In




Figure 1. Block diagram of stage 2 of our algorithm. Given an input frame, two sets of mask are generated using segmentation from swin and propagation from CFBI. It is followed by mask association and identifying new objects detected by Swin. The associated pairs are sent to the MSN and best mask is selected to be propagated further. The new objects are added only if it does not intersect with an object of the same class. The outputs from this stage are sent to post processing step in i.e. stage 3.

order to handle the mentioned challenges we propose the following algorithm.

Let  $S_t$  and  $P_t^i$  be the set of segmented masks produced by Swin and propagated mask by CFBI for frame  $t$ . The two sets of masks are passed to association module, where a bipartite matching is done between the object masks present in both the sets. A 2D matrix is formed whose rows and columns are the objects present in  $S_t$  and  $P_t$  respectively and  $v_{i,j}$  is the IOU between the  $i^{th}$  object mask in  $S_t$  and  $j^{th}$  object mask in  $P_t$ . The assignment is done using Hungarian algorithm. Object masks in  $S_t$  having a IOU higher than 0.1 are associated to corresponding object masks in  $P_t$  and the rest of the objects are kept in consideration for new objects. A new object is added only if it does not intersect with an object of the same class. This ensures to limit the addition of false positives. Now for every associated object we have two mask proposal one from Swin and another from CFBI. We use the proposed Mask selection network (explained in Sec. 3.2) for selecting the better mask of the two. This ensures that only the best mask out of the available masks gets propagated hence limiting the noise accumulation. The complete schematic of the pipeline is shown in the Fig. 1. Note that any other network other than CFBI can also be used. Our method proposes a general pipeline to adapt video object segmentation methods effectively for the task of VIS.

### Stage 3: Post processing

We leverage the temporal consistency in a video and process the video sequence using stage 2 in both forward and backward manner. In the forward implementation, we start with

the instance masks in the first frame of the video sequence and move ahead in time using stage 2 algorithm. In the backward implementation, we start from the instance masks in the last frame and move backward in time using stage 2 algorithm. Hence, at the end of this step, there are 2 outputs for every frame in the video sequence. The next step is to combine these outputs to create a single unified output. To understand the importance of this step, consider a scenario where an object is present in the video from the beginning of the sequence, however that object does not get detected till the 10<sup>th</sup> frame. Hence in the forward implementation, the object will be propagated only from the 10<sup>th</sup> onwards and in the backward implementation, it will only be present in the first 10 frames. So combining both of them, will ensure that the object is present in the entire video. This is achieved by associating all the objects in every frame using Hungarian algorithm in the forward and backward pass. The sequences of the objects are combined if an object is associated in at least one frame. Having this approach gives us freedom of choosing high threshold value for object score because even if the object is present in few frames it will be tracked through out the video. This helps in reducing false positives. Fig 2 shows an example of such association.

Further we also employ human-object association similar to [40]. Specifically, we associate boat, motorbike, skateboard, snowboard, surfboard, and tennis racket with humans. This helps in accurate tracking of objects.

### 3.2. Mask selection network

In this work we propose an improved version of *selector net* [13] named Mask selection network (MSN). The network proposed in [13] consists of a feature extractor




Figure 2. Results for forward and backward association




Figure 3. Mask selection network for selecting better mask of the two input masks.

backbone and the feature vectors are aggregated in spatial dimension before feeding them to fully connected layers. This architecture leads to the loss of spatial information due to feature aggregation in the spatial dimensions. Further, the network was trained with the task of predicting scores of the two masks. However, the task of selecting a better mask is more similar to the task of discriminator in GAN. Given two masks, the network needs to tell which mask is better than the other mask. Hence, for this task a relativistic discriminator [35] based architecture would be more suitable.

We use a patch-based network [16] because, in a candidate mask, some region could be better and other much worse compared to other mask in input pair. For each patch, the network outputs a score indicating which of the masks in the input pair is better. Finally, the information from each patch is aggregated to make the final decision. The schematic representing the MSN is shown in Fig. 3. The data preparation is similar to that in FrameSelect [13].

The input to MSN consists of 2 binary masks and RGB image corresponding to both the masks, hence the input consists of 8 channels. A tight bounding box around the mask is cropped from the complete image and is used as input as shown in Fig. 3. This is different from [13] where the complete image is used as input. The images are resized to  $256 \times 256$  before feeding to the network. MSN is trained using binary cross entropy loss. The new network is able to achieve an accuracy of 74% on the validation split whereas the one proposed in [13] achieved 69%. Note, both the networks were trained on the same training dataset generated using youtube training dataset [40] for fair comparison between them. MSN is more accurate as well as efficient compared to the *selector net* [13] with only 2.22

Network	Acc. %	FLOPs (GMac)	Params(M)
Selector Net [13]	69	19.45	11.7
<b>MSN</b>	<b>74</b>	<b>2.22</b>	<b>11.1</b>

Table 1. Comparison between MSN and selector net [13] in terms of accuracy and compute complexity

GMac operations compared to 19.45 GMac operations. Table 1. summarizes the comparison between MSN and selector network [13].

## 4. Experiments

### 4.1. Implementation details

*Swin-based Instance Segmentation* For training instance segmentation model, we use Swin-base as our backbone and cascade mask R-CNN as instance segmentation model. For data augmentation, we make class project map and convert Openimage and MS COCO data to classes in VIS dataset and pretrain our model on these open-sourced data. Then, we add a balanced data sampler to improve sampling rate of rare classes and a ViT-based strong classifier to re-classify wrong classifications from the Swin model. The whole network is built based open-sourced implementation of Swin model and trained on 8 RTX 2080Ti GPU. The embedding channel of Swin-base model is 128 and the number of cascade mask R-CNN head is 3. The total batch size is 16 with 2 on each GPU and learning rate is 0.0001. We use Adamw as optimizer with weight decay at 0.05 and warm-up training with 1000 iterations. The total training epoch is 6 with learning rate decay at epoch 3 and 5. For data augmentation in training, we use multi-scale training scales at [(1333, 800), (1333, 768), (1333, 736), (1333, 704), (1333, 672), (1333, 640)] and random flip with probability at 0.5. During inference, we tried different testing test augmentations including multi-scale inference and image flip but they didn't bring further improvements based on our experiments. More experimental results are introduced and shared in section 4.2

*Mask selection network* was trained using the data generated from youtube training dataset [40]. The resulting dataset consisted of 80k training images and 8k validation images. Network was trained using a batch size of 512 and learning rate of 0.01 using Adam optimizer. The network was trained for 40 epochs and learning rate was scaled by a factor 0.1 after every 10 epochs. The input pair of masks were randomly shuffled at the time of training to prevent the network getting biased towards the ordering of masks. The training was conducted on 8 RTX 2080 GPU.

We used the publicly available implementation of CFBI [42]. At the time of inference we used multi-scale and flip strategy in CFBI. Specifically, we apply the scales of {1.0, 1.15, 1.3, 1.5}. Further, we limited the maximum






Figure 4. Qualitative results on Youtube Video Instance Segmentation 2021 test dataset.

Team	mAP	AP		AR	
		50	75	1	10
tuantng	54.1	74.2	61.6	43.3	58.9
eastonssy	52.3	76.7	57.7	43.9	57.0
<b>Ours</b>	<b>49.1</b>	<b>68.1</b>	<b>54.5</b>	<b>41.0</b>	<b>55.0</b>
linhj	47.8	69.3	52.7	42.2	59.1
hongsong.wang	47.6	68.4	52.9	41.4	54.6
gb7	47.3	66.5	51.1	40.5	51.6
zfonemore	46.1	64.4	51.0	38.3	50.6
DeepBlueAI	46.0	64.6	52.0	38.7	54.2

Table 2. Results on Youtube VIS 2021 test dataset compared to other participants in the challenge.

number of objects in a video to be 15.

## 4.2. Results and analysis

The proposed method is tested on the Youtube VIS test dataset which contains 453 video sequences. Further, extensive ablation studies were conducted on the validation split of the dataset which contains 421 videos. Table 2 shows our results in comparison to top 8 teams. It can be observed that although our AP50 ranks fifth, still we are able to achieve third position. This implies that our tracking and mask quality is more consistent across various IOU thresholds compared to other method which can be attributed to our online mask improvement using MSN.

In order to understand the importance of various tricks applied in instance segmentation we did an ablation study for each component. The results are shown in Table 2. Other parameters for CFBI such as multiscale inferencing

Model description	mAP
COCO Pretrained + VIS finetune	41.3
COCO + VIS trained	47.1
COCO + VIS + OpenImages trained	48.3
+ Balanced sampling	48.8

Table 3. Ablation study on instance segmentation (stage 1). Experiments are done on Youtube VIS 2021 validation dataset.

Method	$\Delta$ mAP
Mask selection network	+ 2.6
Multi scale testing(CFBI)	+ 1.0
Forward/Backward	+ 0.8
Object addition criteria	+ 0.2
Object Association	+ 0.2

Table 4. Ablation study on various tricks in stage 2 and 3. Here we state the average improvement in mAP across various experiments on Youtube VIS 2021 validation dataset.

are kept same. From the results it is evident that a strong instance segmentation method is crucial for video instance segmentation. Similar conclusion can be drawn from the work [10], in which they simply improve upon the instance segmentation pipeline to get better results.

Further, we also studied the improvements brought by various tricks namely MSN, multiscale inferencing in CFBI [42], post processing in stage 3 and object intersection criteria while adding new objects in stage 2 and results are shown in the Table 4. It can be observed that MSN contributes highest in improving results.

## 5. Conclusion

In this work we propose a novel approach for adapting any video object segmentation methods for video instance segmentation. We improve upon the mask selection based method by proposing a mask selection network which can accurately choose a better mask of the given two masks. The mask selector network is highly computationally efficient compared to selector net and is more accurate. The proposed method has a modular structure which will improve as each modules become more efficient. We also propose a post processing pipeline which employs forward and backward processing of video to recover lost objects. We believe that instance segmentation is a crucial task in VIS and in order to make significant progress in the task a strong instance segmentation pipeline which incorporates temporal information is necessary.

## References

- [1] Ali Athar, Sabarinath Mahadevan, Aljosa Osep, Laura Leal-Taixé, and Bastian Leibe. Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In *European Conference on Computer Vision*, pages 158–177. Springer, 2020. 2
- [2] Gedas Bertasius and Lorenzo Torresani. Classifying, segmenting, and tracking object instances in video with mask propagation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 1
- [3] Gedas Bertasius and Lorenzo Torresani. Classifying, segmenting, and tracking object instances in video with mask propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9739–9748, 2020. 2
- [4] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact: Real-time instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9157–9166, 2019. 2
- [5] Sergi Caelles, Jordi Pont-Tuset, Federico Perazzi, Alberto Montes, Kevis-Kokitsi Maninis, and Luc Van Gool. The 2019 davis challenge on vos: Unsupervised multi-object segmentation. *arXiv:1905.00737*, 2019. 1, 2
- [6] Jiale Cao, Rao Muhammad Anwer, Hisham Cholakkal, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. Sipmask: Spatial information preservation for fast image and video instance segmentation. *arXiv preprint arXiv:2007.14772*, 2020. 2
- [7] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4974–4983, 2019. 2
- [8] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, et al. Mmdetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 2
- [9] Xinlei Chen, Ross Girshick, Kaiming He, and Piotr Dollár. Tensormask: A foundation for dense object segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2061–2069, 2019. 2
- [10] Yuxin Fang, Shusheng Yang, Xinggang Wang, Yu Li, Chen Fang, Ying Shan, Bin Feng, and Wenyu Liu. Instances as queries. *arXiv preprint arXiv:2105.01928*, 2021. 2, 5
- [11] Yang Fu, Sifei Liu, Umar Iqbal, Shalini De Mello, Humphrey Shi, and Jan Kautz. Learning to track instances without video annotations. In *CVPR*, 2021. 2
- [12] Yang Fu, Linjie Yang, Ding Liu, Thomas S Huang, and Humphrey Shi. Compfeat: Comprehensive feature aggregation for video instance segmentation. In *AAAI*, 2021. 2
- [13] Shubhika Garg and Vidit Goel. Mask selection and propagation for unsupervised video object segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1680–1690, 2021. 1, 2, 3, 4
- [14] Ali Hassani, Steven Walton, Nikhil Shah, Abulikemu Abuduweili, Jiachen Li, and Humphrey Shi. Escaping the big data paradigm with compact transformers. *arXiv preprint arXiv:2104.05704*, 2021. 2
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2
- [16] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 4
- [17] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 1, 2
- [18] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4. *International Journal of Computer Vision*, pages 1–26, 2020. 1, 2
- [19] Youngwan Lee and Jongyoul Park. Centermask: Real-time anchor-free instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13906–13915, 2020. 2
- [20] Siyang Li, Bryan Seybold, Alexey Vorobyov, Alireza Fathi, Qin Huang, and C-C Jay Kuo. Instance embedding transfer to unsupervised video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6526–6535, 2018. 2
- [21] Chung-Ching Lin, Ying Hung, Rogerio Feris, and Linglin He. Video instance segmentation tracking with a modified vae architecture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13147–13157, 2020. 2
- [22] Huajia Lin, Ruizheng Wu, Shu Liu, Jiangbo Lu, and Jia-ya Jia. Video instance segmentation with a propose-reduce paradigm. *arXiv preprint arXiv:2103.13746*, 2021. 1, 2

- [23] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 1, 2
- [24] Yudong Liu, Yongtao Wang, Siwei Wang, TingTing Liang, Qijie Zhao, Zhi Tang, and Haibin Ling. Cbnet: A novel composite backbone network architecture for object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11653–11660, 2020. 2
- [25] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021. 1, 2
- [26] Jonathon Luiten, Philip Torr, and Bastian Leibe. Video instance segmentation 2019: A winning approach for combined detection, segmentation, classification and tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 1, 2
- [27] Jonathon Luiten, Paul Voigtlaender, and Bastian Leibe. Premvos: Proposal-generation, refinement and merging for video object segmentation. In *Asian Conference on Computer Vision*, pages 565–580. Springer, 2018. 2
- [28] Seoung Wug Oh, Joon-Young Lee, Kalyan Sunkavalli, and Seon Joo Kim. Fast video object segmentation by reference-guided mask propagation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7376–7385, 2018. 2
- [29] Seoung Wug Oh, Joon-Young Lee, Ning Xu, and Seon Joo Kim. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9226–9235, 2019. 2
- [30] Federico Perazzi, Anna Khoreva, Rodrigo Benenson, Bernt Schiele, and Alexander Sorkine-Hornung. Learning video object segmentation from static images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2663–2672, 2017. 2
- [31] Paul Voigtlaender, Yuning Chai, Florian Schroff, Hartwig Adam, Bastian Leibe, and Liang-Chieh Chen. Feelvos: Fast end-to-end embedding learning for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9481–9490, 2019. 2
- [32] Bofei Wang, Chengjian Zheng, Ning Wang, Shunfei Wang, Xiaofeng Zhang, Shaoli Liu, Si Gao, Kaidi Lu, Diankai Zhang, Lin Shen, et al. Object-based spatial similarity for semi-supervised video object segmentation. In *Conference on Computer Vision and Pattern Recognition Workshops*, 2019. 2
- [33] Qiang Wang, Yi He, Xiaoyun Yang, Zhao Yang, and Philip Torr. An empirical study of detection-based video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 2
- [34] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. Solo: Segmenting objects by locations. In *European Conference on Computer Vision*, pages 649–665. Springer, 2020. 2
- [35] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018. 4
- [36] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 2
- [37] A Waswani, N Shazeer, N Parmar, J Uszkoreit, L Jones, AN Gomez, L Kaiser, and I Polosukhin. Attention is all you need. In *NIPS*, 2017. 2
- [38] Shuangjie Xu, Daizong Liu, Linchao Bao, Wei Liu, and Pan Zhou. Mhp-vos: Multiple hypotheses propagation for video object segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 314–323, 2019. 2
- [39] Xingqian Xu, Mang Tik Chiu, Thomas S Huang, and Honghui Shi. Deep affinity net: instance segmentation via affinity. *arXiv preprint arXiv:2003.06849*, 2020. 2
- [40] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5188–5197, 2019. 2, 3, 4
- [41] Zongxin Yang, Yuhang Ding, Yunchao Wei, and Yi Yang. Cfbi+: Collaborative video object segmentation by multi-scale foreground-background integration. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, volume 1, 2020. 2
- [42] Zongxin Yang, Yunchao Wei, and Yi Yang. Collaborative video object segmentation by foreground-background integration. In *European Conference on Computer Vision*, pages 332–348. Springer, 2020. 2, 4, 5