

In Defense of Online Models for Video Instance Segmentation

Junfeng Wu^{1*}, Qihao Liu^{2*}, Yi Jiang³, Song Bai^{3†}, Alan Yuille², Xiang Bai¹

¹ Huazhong University of Science and Technology

² Johns Hopkins University

³ ByteDance

Abstract. In recent years, video instance segmentation (VIS) has been largely advanced by offline models, while online models gradually attracted less attention possibly due to their inferior performance. However, online methods have their inherent advantage in handling long video sequences and ongoing videos while offline models fail due to the limit of computational resources. Therefore, it would be highly desirable if online models can achieve comparable or even better performance than offline models. By dissecting current online models and offline models, we demonstrate that the main cause of the performance gap is the error-prone association between frames caused by the similar appearance among different instances in the feature space. Observing this, we propose an online framework based on contrastive learning that is able to learn more discriminative instance embeddings for association and fully exploit history information for stability. Despite its simplicity, our method outperforms all online and offline methods on three benchmarks. Specifically, we achieve 49.5 AP on YouTube-VIS 2019, a significant improvement of 13.2 AP and 2.1 AP over the prior online and offline art, respectively. Moreover, we achieve 30.2 AP on OVIS, a more challenging dataset with significant crowding and occlusions, surpassing the prior art by 14.8 AP. The proposed method won **first place** in the video instance segmentation track of the 4th Large-scale Video Object Segmentation Challenge (CVPR2022). We hope the simplicity and effectiveness of our method, as well as our insight on current methods, could shed light on the exploration of VIS models. The code is available at <https://github.com/wjf5203/VNext>.

Keywords: Video instance segmentation, Online model, Contrastive learning

1 Introduction

Video instance segmentation aims at detecting, segmenting, and tracking object instances simultaneously in a given video. It attracted considerable attention after first defined [45] in 2019 due to the huge challenge and the wide applications in video understanding, video editing, autonomous driving, augmented reality,

* First two authors contributed equally. Work done during an internship at ByteDance.

† Corresponding author: songbai.site@gmail.com

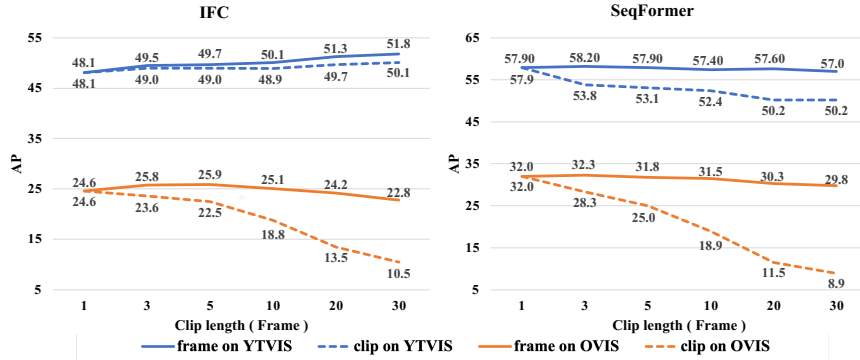


Fig. 1. Oracle experiments on SOTA offline methods: We analyze current VIS methods and visualize some results here. More results can be found in Sec. 4.2. ‘YTVIS’ is short for YouTube-VIS 2019. ‘Frame’ and ‘clip’ stand for frame oracle and clip oracle experiments, respectively. For frame oracles, we provide the ground-truth instance ID both within each clip and between adjacent clips. Thus the performance only depends on the quality of the estimated segmentation masks. For clip oracles, we only provide the ground-truth instance ID between adjacent clips, and the method is required to do association within the clips by itself. Therefore the gaps between frame oracles and clip oracles show the effect of the black-box association done within current offline models. When clip length is 1, the method is doing per-frame segmentation.

etc. Current VIS methods can be categorized as online or offline methods. Online methods [45, 4, 8, 20, 46, 18] take as input a video frame by frame, detecting and segmenting objects per frame while tracking instances and optimizing results across frames. Offline methods [3, 21, 1, 37, 16, 40], in contrast, take the whole video as input and generate the instance sequence of the entire video with a single step.

Despite high performance, the requirement of full videos limits the application scenarios of offline models, especially for scenarios involving long video sequences (videos exceed 50 frames on GPU of 32G RAM [40]) and ongoing videos (*e.g.*, videos in autonomous driving and augmented reality). However, online models are usually inferior to the contemporaneous offline models by over 10 AP, which is a huge drawback. Previously, little work tries to explain the performance gap between these two paradigms, or gives insight into the high performance of the offline paradigm. A common attempt of the latter is made from the inherent advantages of offline models being able to skip the error-accumulating tracking steps [21] and utilize the richer information provided by multiple frames to improve segmentation [16, 1, 40, 38]. However, does that really explain the high performance of current offline methods? What’s the main problem causing the poor performance of online models? Can online models achieve performance comparable to, or even better than, SOTA offline ones?

To deeply understand the performance of both online and offline models, we analyze in detail three SOTA methods (offline: IFC [16] [40],

online: CrossVIS [46]) on two datasets (YouTube-VIS [45] and OVIS [30]) that have different difficulty levels. ‘Simple video’ refers to the video in YouTube-VIS. These videos are much shorter and only contain very slight occlusions, simple motions, and smooth changes in illumination and object shapes. ‘Complex video’ refers to the video in OVIS. Please see Sec. 4.1 for details. The results (Fig. 1) of our oracle experiments give us a deep understanding of current SOTA methods:

From the perspective of instance segmentation, per-clip segmentation doesn’t outperform per-frame segmentation a lot in mask quality, and mask quality is also not the reason for the poor performance of online methods: CrossVIS even outperforms its contemporaneous work (*i.e.* IFC) in frame oracle experiments on both datasets (results in Table 1). What’s more, per-clip segmentation of current SOTA methods is not effective and robust. Multiple frames do provide more information and improve the mask quality by 3.7 AP for IFC on YouTube-VIS (Fig. 1). But it only works for some cases: per-clip segmentation doesn’t improve the performance of SeqFormer a lot. In addition, when testing on more challenging datasets like OVIS, segmentation on multiple frames even degrades the performance by 1.8 and 2.2 AP on IFC and SeqFormer respectively when clip size becomes longer. Although in theory, per-clip segmentation has its inherent advantage to using multiple frames, it still requires further exploration, especially in how to utilize information in multiple frames and how to handle complex motion patterns, occlusion, and object deformation. Currently, we don’t see an obvious gap between the mask qualities of per-clip and per-frame segmentation.

From the perspective of association, a huge advantage of the offline methods is their ability to avoid the use of hand-designed association modules. It works well on simple cases of the YouTube-VIS dataset. We demonstrate that it is the main reason causing the performance gap between the current online and offline paradigms. However, this black-box association process done within offline models also gets worse rapidly when the video becomes complex (degrades the performance of IFC by 12.3 AP and SeqFormer by 20.9 AP on OVIS). In addition, when handling longer videos, *e.g.* videos in the real world or from OVIS dataset, offline methods require splitting the input video into clips to avoid exceeding computational limits, and thus hand-designed clip matching is still inevitable, which will further decrease the overall performance. To sum up, matching/association is the main reasoning for the performance gap, and it is still inevitable and of great importance for offline models.

To improve the matching performance and thus bridge the performance gap, we propose a framework **In Defense of OnLine** models for video instance segmentation, termed IDOL. The key idea is to ensure, in the embedding space, the similarity of the same instance across frames and the difference of different instances in all frames, even for instances that belong to the same category and have similar appearances. It provides more discriminative instance features with better temporal consistency, which guarantees more accurate association results. However, previous method [29] selects positive and negative samples by a hand-craft setting, introducing false positives in occlusions and crowded scenes, thus impairing contrastive learning. To address it, we formulate the problem

of sample selection as an Optimal Transport problem in Optimization Theory, which reduces false positives and further improves the quality of the embedding. During inference, by using one-to-many temporally weighted softmax, we utilize the learned prior of the embedding to re-identify missing instances caused by occlusions and to enforce the consistency and integrality of associations.

Our thorough analysis gives us a deep understanding of current online and offline VIS methods. Based on our observation, we bridge the performance gap from the perspective of feature embeddings and propose IDOL. We conduct extensive experiments on YouTube-VIS 2019, YouTube-VIS 2021, and OVIS datasets. Despite its simplicity, our method sets a new state-of-the-art achieving 62.2 AP, 56.1 AP, and 42.6 AP on the validation set of these three benchmarks, respectively. More importantly, compared with previous online methods, we achieve a consistent improvement ranging from 13.2 AP to 14.7 AP on these datasets. We even surpass the previous SOTA offline method by up to 2.1 AP. We believe the simplicity and effectiveness of our method shall benefit further research. In addition, our thorough analysis provides insights for current methods and suggests useful directions for future work in both online and offline VIS.

2 Related Work

Online Video Instance Segmentation. Most online VIS methods are built upon image-level instance segmentation with an additional tracking head to associate instances across the video. The baseline method MaskTrack R-CNN [45] is built upon Mask R-CNN and proposes to leverage multi cues such as appearance similarity, semantic consistency, spatial correlation, and detection confidence to determine the instance labels. Most online methods [4,8,20,46,17] follow this pipeline. CrossVIS [46] proposes a new learning scheme that uses the instance feature in the current frame to segment the same instance in another frame. Multi-Object Tracking and Segmentation (MOTS) [7,36] aims to simultaneously segment and track all object instances of a given video sequence in real-time, which is similar to online VIS. MOTS methods are usually built upon multiple object trackers [2,29,48,49,32,47,44,18]. Track R-CNN [36] firstly extends the popular task of multi-object tracking to MOTS based on Mask R-CNN [13]. PointTrack [43] proposes a new online tracking-by-points paradigm with learning discriminative instance embeddings. Trackformer [25] adopts the transformer architecture for MOTS and introduces track query embeddings that follow objects through a video sequence. Online models have a wider range of application scenarios, however, they are usually inferior to offline art by over 10 AP. We find that the current SOTA online models fail to achieve accurate associations, causing the performance gap. We aim to tackle this problem in this work.

Offline Video Instance Segmentation. Offline methods for VIS take the whole video as input and predict the instance sequence of the entire video (or video clip) with a single step. MaskProp [3] and Propose-Reduce [21] perform mask propagation in a video clip to improve mask and association. However, the propagation process is time-consuming, which limits its application. Recently, VisTR [37] adopts the transformer [35] to VIS and models the instance queries

for the whole video. However, it learns an embedding for each instance of each frame, which makes it hard to apply to longer videos and more complex scenes. IFC [16] proposes inter-frame communication transformers and significantly reduces computation and memory usage. SeqFormer [40] dynamically allocates spatial attention on each frame and learns a video-level instance embedding, which greatly improves the performance. We deeply analyze the current SOTA offline models, IFC and SeqFormer, find that their improvement mainly comes from the black-box association between frames, but this advantage is gradually lost in complex scenarios. In contrast, our online method can be applied to both ongoing and long videos and complex scenarios, with more stable association quality and higher performance.

Contrastive learning has made significant progress in representation learning [6,12,29,11,15,28,41,33]. MOCO [12] and SimCLR [6] use contrastive learning for image-level self-supervised training and learn strong feature representations for downstream tasks. Some methods [11,19,26] extend the contrastive learning into multiple positive samples format to obtain better feature representations. We absorb ideas from contrastive learning and propose to learn contrastive embeddings between frames for each instance.

3 Method

Given a video clip that consists of multiple image frames, online VIS models [45,46] utilize additional association head upon on instance segmentation models [13,34]. We have already discussed that achieving more stable and discriminative instance embeddings between frames is the key to improve the performance of online models. To achieve this, we propose a contrastive learning framework to extract more discriminative features for instance association. We first introduce the instance segmentation pipeline in Sec. 3.1. Then the details of our contrastive learning framework and the cross-frame instance association strategy are introduced in Sec. 3.2 and Sec. 3.3 respectively. More implementation details can be found in the supplementary.

3.1 Instance Segmentation.

For fair comparisons with the state-of-the-art offline method [40], we take DeformableDETR [50] with dynamic mask head [34] as our instance segmentation pipeline in this paper. Our method can be coupled with other instance segmentation methods with minor modifications.

Given an input frame $x \in R^{3 \times H \times W}$ of a video, a CNN backbone extracts multi-scale feature maps. The Deformable DETR module takes the feature maps with additional fixed positional encodings [5] and N learnable object queries as input. The object queries are first transformed into output embeddings $E \in R^{N \times C}$ by the transformer decoder. After that, they are decoded into box coordinates and class labels by 3-layer feed-forward network (FFN) separately. For per-frame mask generation, we employ an FPN-like [22] mask branch to make the use of multi-scale feature maps from transformer encoder and generate feature map F_{mask} that are 1/8 resolution of the input frame. Another FFN encode

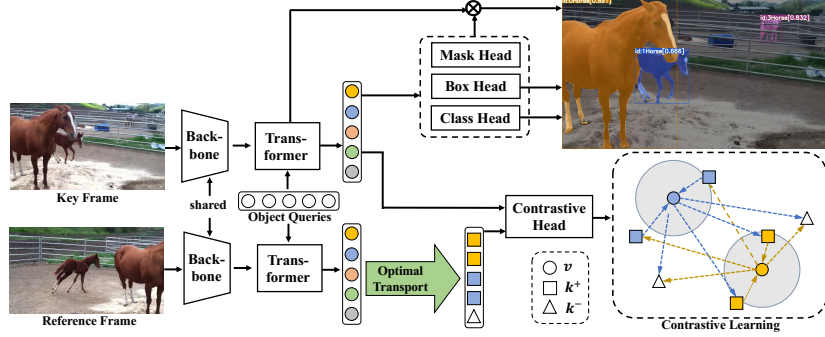


Fig. 2. The training pipeline of IDOL. Given a key frame and a reference frame, the shared-weight backbone and transformer predict the instance embeddings on them respectively. The embeddings on the key frame are used to predict masks, boxes, and categories, while the embeddings on the reference frame are selected as positive and negative embeddings dynamically by our optimal transport progress. Embeddings of the same color belong to the same video instance. Best viewed in color.

outputs embeddings into parameters ω of mask head, which performs three-layer 1×1 convolution on the given feature map F_{mask} :

$$\mathbf{m}_i = \text{MaskHead}(\mathbf{F}_{mask}, \omega_i). \quad (1)$$

Then we calculate pair-wise matching cost which takes into account both the class prediction and the similarity of predicted and ground truth boxes. For each ground truth, we assign multi predictions to it by selecting the top k predictions with the least cost by an optimal transport method [9,10]. Finally, the whole model is optimized with a multi-task loss function

$$\mathcal{L} = \mathcal{L}_{cls} + \lambda_1 \mathcal{L}_{box} + \lambda_1 \mathcal{L}_{mask} + \lambda_2 \mathcal{L}_{embed}, \quad (2)$$

where loss weights λ_1 and λ_2 are set to 2.0 and 1.0 by default. For \mathcal{L}_{box} , we use a combination of \mathcal{L}_1 loss and the generalized IoU loss [31]. The \mathcal{L}_{mask} is defined as a combination of the Dice loss [27] and Focal loss [23]. \mathcal{L}_{embed} is the contrastive loss described in the next section.

3.2 Contrastive Learning between Frames.

More discriminative feature embeddings can help distinguish instances on different frames, thereby improving the quality of cross-frame association. To this end, we introduce contrastive learning between frames to make the embedding of the same object instance closer in the embedding space, and the embedding of different object instances farther away. Object queries are used to query the features of instances from each frame in our instance segmentation pipeline. Therefore, the output embeddings can be regarded as features of different instances. We employ an extra light-weighted FFN as a contrastive head to decode the contrastive embeddings from the instance features.

Given a key frame for instance segmentation training, we select a reference frame from the temporal neighborhood. The instances appearing on the key frame may have different positions and appearances on the reference frame, but their contrastive embeddings should be as close as possible in embedding space. For each instance in the key frame, we send the output embedding with the lowest cost to the contrastive head and get the contrastive embedding \mathbf{v} . Different from previous method [29], which selects positive and negative samples by a hand-craft setting, if the same instance appears on the reference frame, we take top $m1$ predictions with the least cost as positive and top $m2$ predictions with the highest cost as negatives. $m1$ and $m2$ are calculated dynamically by the optimal transport method [9,10]. Please refer to the supplementary for more details. The contrastive loss function for a positive pair of examples is defined as follows:

$$\begin{aligned}\mathcal{L}_{embed} &= -\log \frac{\exp(\mathbf{v} \cdot \mathbf{k}^+)}{\exp(\mathbf{v} \cdot \mathbf{k}^+) + \sum_{\mathbf{k}^-} \exp(\mathbf{v} \cdot \mathbf{k}^-)} \\ &= \log[1 + \sum_{\mathbf{k}^-} \exp(\mathbf{v} \cdot \mathbf{k}^- - \mathbf{v} \cdot \mathbf{k}^+)],\end{aligned}\tag{3}$$

where \mathbf{k}^+ and \mathbf{k}^- are positive and negative feature embeddings from the reference frame, respectively. We extend Eq. 3 to multiple positive scenarios:

$$\mathcal{L}_{embed} = \log[1 + \sum_{\mathbf{k}^+} \sum_{\mathbf{k}^-} \exp(\mathbf{v} \cdot \mathbf{k}^- - \mathbf{v} \cdot \mathbf{k}^+)].\tag{4}$$

3.3 Instance Association.

Previous online methods [45,46] take semantic consistency, spatial correlation, and detection confidence as cues. They are then leveraged to determine the instance labels. Other clip-based nearly online methods [3,1,16] match instances using the predicted masks of overlapping frames by masking soft IoU metric between clips. However, the online models perform instance segmentation on each frame independently, and therefore, the prediction quality on each frame is unstable. What makes it worse is the complex motion patterns, severe occlusions, false positives, duplicate predictions, error accumulation in long videos, and the frequently disappear and reappear objects, which makes instance association very challenging. Therefore, a strong instance association method should be robust to these cases. To this end, we propose a temporally weighted softmax score for instance matching and a memory bank-based association strategy to address these problems and improve the association quality of the online model.

Temporally Weighted Softmax. Considering scenarios with fast motion, occlusion, and crowded objects, box-based matching introduces wrong association due to ambiguous location priors. However, the contrastive embedding learned by our method is able to maintain discriminative in embedding space when the position and shape change. Therefore, the contrastive embeddings are used to calculate the embedding similarity between the instances on the current and the previous frames. Assume there are N predicted instances with N contrastive

embeddings $\mathbf{d}_i \in \mathbb{R}^C$, and M instances in the memory bank, each of which has multiple temporal contrastive embeddings $\{\mathbf{e}_j^t\}_{t=1}^T, \mathbf{e}_j^t \in \mathbb{R}^C$ from previous T frames. These embeddings are combined by a temporally weighted sum:

$$\hat{\mathbf{e}}_j = \frac{\sum_{t=1}^T \mathbf{e}_j^t \times (\tau + T/t)}{\sum_{t=1}^T \tau + T/t}. \quad (5)$$

Then we compute bi-directional similarity f between predicted instance i and memory instance j by:

$$\mathbf{f}(i, j) = [\frac{\exp(\hat{\mathbf{e}}_j \cdot \mathbf{d}_i) + \sigma_j}{\sum_{k=1}^M \exp(\hat{\mathbf{e}}_k \cdot \mathbf{d}_i) + \sigma_k} + \frac{\exp(\hat{\mathbf{e}}_j \cdot \mathbf{d}_i)}{\sum_{k=1}^N \exp(\hat{\mathbf{e}}_j \cdot \mathbf{d}_k)}]/2, \quad (6)$$

where σ_j is the existing time of instance j in the memory, it serves as the confidence scores of each instance in the memory. By introducing the temporal contrastive embeddings and the confidence scores determined by the duration of existence, the learned prior information is able to reidentify missing instances caused by occlusions, enforcing the consistency and integrality of associations.

Association Strategy. To take full advantage of the learned contrastive embedding, we propose a new association strategy during inference. Given a test video, we initialize an empty memory bank for it and perform instance segmentation on each frame sequentially in an online scheme. For the prediction of each frame, we first perform inter-class duplicate removal by NMS with a threshold of 0.5. Then we compute matching scores $\mathbf{f}(i, j)$ between predictions and memory bank by Eq. 6, and search for the best assignment for instance i by:

$$\hat{j} = \arg \max \mathbf{f}(i, j), \forall j \in \{1, 2, \dots, M\}. \quad (7)$$

If $\mathbf{f}(i, \hat{j}) > 0.5$, we assign the instance i on current frame to the memory instance \hat{j} . For the prediction without an assignment but has a high class score, we start a new instance ID in the memory bank.

4 Experiments

4.1 Dataset and Metrics

We report our results on YouTube-VIS 2019 [45], YouTube-VIS 2021 [42], and OVIS [30] datasets. YouTube-VIS 2019 is the first and largest dataset for video instance segmentation. It contains 2,238 training, 302 validation, and 343 test high-resolution YouTube video clips, with an average video duration of 4.61s. YouTube-VIS 2021 is an extended version of YouTube-VIS 2019. Both datasets have 40 object categories, but the category label set is slightly different. OVIS dataset is a relatively new and challenging dataset. It consists of 607 training videos, 140 validation videos, and 154 test videos. Compared with YouTube-VIS, its videos are much longer and last 12.77s on average, and more importantly, it contains much more videos that record objects with severe occlusion, complex motion patterns, and rapid deformation. All these features make OVIS an ideal dataset to evaluate and analyze different methods. We report standard metrics such as AP, AP₅₀, AP₇₅, AR₁, and AR₁₀. IoU threshold is used during evaluation.

Table 1. Oracle experiments on association quality. Frame oracle means gt instance ID is provided (same to Fig. 1). We set the clip length of IFC to 10. AP is reported.

Dataset	Method	Publish	Predicted	Frame Oracle
YouTube-VIS	CrossVIS	ICCV 2021	43.4	52.8
	IFC	NeurIPS 2021	46.8	50.1
OVIS	CrossVIS	ICCV 2021	10.1	29.9
	IFC	NeurIPS 2021	8.7	25.1

Table 2. Oracle experiments on clip length for offline models. AP is reported.

Dataset	Method	Oracle Type	Clip Length (frame)					
			1	3	5	10	20	30
YouTube-VIS	IFC	frame	48.1	49.5	49.7	50.1	51.3	51.8
		clip	48.1	49.1	49.0	48.9	49.7	50.1
	SeqFormer	frame	57.9	58.2	57.9	57.4	57.6	57.0
		clip	57.9	53.8	53.14	52.4	50.2	50.2
OVIS	IFC	frame	24.6	25.8	25.9	25.1	24.2	22.8
		clip	24.6	23.6	22.5	18.8	13.5	10.5
	SeqFormer	frame	32.0	32.3	31.8	31.5	30.3	29.8
		clip	32.0	28.3	25.0	18.9	11.5	8.9

4.2 Analysis of Current SOTA VIS Models

Since no annotation for the validation set is available, we split the original training set into custom training split and validation split. All models are trained on the training split and evaluated on the validation split. YouTube-VIS 2019 and OVIS are used. We analyze the results as follows:

Performance gaps between online and offline models: First, we compare the mask quality of two recent online and offline methods in Table. 1. They are both published in 2021 thus can be considered as work in the same period. When ground-truth instance ID is provided (the column of ‘frame oracle’), CrossVIS outperforms IFC on both YouTube-VIS and OVIS. However, when the instance ID is not provided (the column of ‘predicted’), the methods are required to match the results, and the performance of CrossVIS drops dramatically by 9.4 AP while the performance of IFC drops by 3.3 AP on YouTube-VIS, leading to the poor performance of CrossVIS. Offline methods enable the model to match predicted masks by itself and avoid using hand-designed association modules. It works well on simple datasets and benefits current offline models, but it still fails on challenging datasets like OVIS.

Analysis of current offline models: In Table. 2, we give detailed analyses for SOTA offline models, hoping to provide insights for future research. Compared with online methods, offline models in theory have two inherent advantages:

First, as we mentioned above, it avoids hand-designed association. However, this step is very sensitive to the occlusion and the complexity of videos, especially when the clip becomes longer. For example, when clip length is equal to 5, the

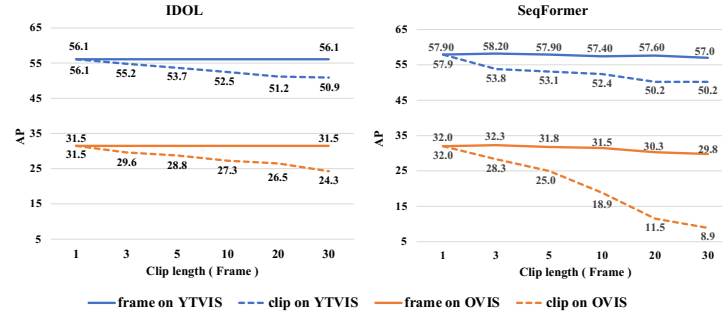


Fig. 3. Oracle experiments on IDOL and SeqFormer.

black-box association degrades the performance of IFC and SeqFormer on OVIS by 3.4 AP (25.9 AP *vs.* 22.5 AP) and 6.8 AP (31.8 AP *vs.* 25.0 AP), respectively. When clip length is set to 30, the performance drops by 12.3 AP (22.8 AP *vs.* 10.5 AP) and 20.9 AP (29.8 AP *vs.* 8.9 AP), respectively. What’s more, even a clip length of 30 still doesn’t meet the requirement of real-world application. Clip matching is still inevitable, and it further decreases the overall performance.

Another inherent advantage of offline models is the ability to use multiple frames for instance segmentation, which provides more information to handle occlusion and optimize the results. However, current models still fail to fully utilize this feature. Currently, it only works for simple videos: compared with per-frame segmentation (clip length=1), it improves the mask quality by at most 3.7 AP for IFC (when clip length=30) and 0.3 AP for SeqFormer (when clip length=3). When testing on OVIS, multiple frames segmentation only improves the mask quality by at most 1.3 AP for IFC (clip length=5) and 0.3 AP for SeqFormer (clip length=3), and even degrades the performance by 1.8 AP for IFC and 2.2 AP for SeqFormer when clip size becomes longer (clip length=30). What’s more, the improvement is even less obvious in practice when no ground-truth instance ID is provided, even for simple videos, due to the association problem. It only improves the performance of IFC on YouTube-VIS by 2.0 AP, but degrades the performance in all the other experiments.

To prove the effectiveness of our method, we further analyze IDOL with oracle experiments in Fig. 3. Since IDOL is an online model, the results of frame oracles with different clip lengths are the same as per-frame segmentation oracle results. For clip oracles, IDOL is required to do association within the clips by itself. Compared with SeqFormer, the gaps between frame oracles and clip oracles of IDOL are much smaller on OVIS, proving that IDOL performs a more robust association between frames than the offline model on challenging datasets.

4.3 Main Results

We compare IDOL against current online and offline SOTA methods on YouTube-VIS 2019, YouTube-VIS 2021, and OVIS validation sets. The results are reported

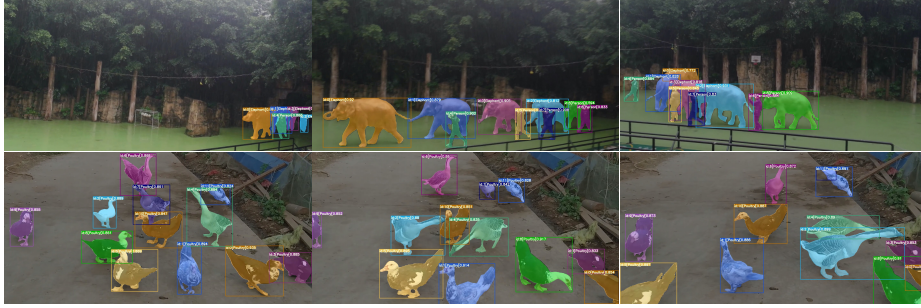


Fig. 4. Qualitative results on the OVIS validation dataset. IDOL achieves very good results on complex scenes. Please refer to the supplementary for more qualitative results and comparison with other methods.

in Tables 3, 4, and 5, respectively. We compare the methods with different backbones [14, 24] for a fair comparison. Notably, our method significantly surpasses all previous online methods by at least 10.1 AP. In addition, we also outperform all previous offline methods under all evaluation metrics when training on the same data. More importantly, our method achieves an overall first place [39] in the YouTube-VIS Challenge 2022, which proves our SOTA performance. In addition, our method only decreases the inference speed of the adopted instance segmentation pipeline by 1.1 FPS on an RTX-2080Ti, which proves our efficiency. In general, our method is simple and very effective compared with all baseline methods. Qualitative results on sample videos of the challenging OVIS dataset are shown in Fig. 4. More qualitative results can be found in the supplementary. We analyze the performance in detail as follows:

YouTube-VIS 2019: It is the most commonly used dataset. We reported the results in Table 3. When using the same backbones, IDOL significantly outperforms previous online methods by 10.1 AP (on ResNet-50) and 10.6 AP (on ResNet-101). Compared with offline methods, we achieve better performance when no extra data is used for training, surpassing previous methods by 1.3 AP and 2.2 AP with ResNet-50 and Swin-L as the backbone, respectively. For the result with superscript “†”, we randomly crop images from COCO twice to form pseudo key-reference frame pairs to pre-train the contrastive embedding part of our model before training on real video datasets. It improves the performance by 3.1 AP on ResNet-50 and 0.7 AP on Swin-L backbone, outperforming previous SOTA models comprehensively.

YouTube-VIS 2021: It is an extended version of YouTube-VIS 2019. It contains more videos with a larger number of instances and frames. As shown in Table 4, we achieve 43.9 AP with a ResNet-50 backbone, surpassing the previous best online method and offline method by 9.7 AP and 3.4 AP, respectively.

OVIS: As mentioned before, OVIS contains long video sequences with heavy occlusion and complex motion, thus it is extremely difficult for all algorithms and exceeds the capability of offline methods due to the limit of computational

Table 3. Comparison on YouTube-VIS 2019 val set. The best results with the same backbone are in **bold** and second underline. ‘V’ means only YouTube-VIS training set is used. ‘V+I’ means synthetic videos from COCO with overlapping categories are also used for joint training. Note that offline models take advantage of larger batch sizes thus having much bigger FPS, while online models handle video frame by frame. The result with superscript “†” is obtained by pre-training on COCO pseudo key-reference frame pairs, and resolution of 480p is used during inference.

Backbone	Method	Type	FPS	Data	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
ResNet-50	MaskTrack R-CNN [45]	online	20.0	V	30.3	51.1	32.6	31.0	35.5
	SipMask [4]	online	30.0	V	33.7	54.1	35.8	35.4	40.1
	CompFeat [8]	online	-	V	35.3	56.0	38.6	33.1	40.3
	CrossVIS [46]	online	39.8	V	36.3	56.8	38.9	35.6	40.7
	PCAN [18]	online	-	V	36.1	54.9	39.4	36.3	41.6
	STEm-Seg [1]	offline	7.0	V+I	30.6	50.7	33.5	37.6	37.1
	VisTR [37]	offline	69.9	V	36.2	59.8	36.9	37.2	42.4
	MaskProp [3]	offline	-	V	40.0	-	42.9	-	-
	Propose-Reduce [21]	offline	-	V+I	40.4	63.0	43.8	41.1	49.7
	IFC [16]	offline	107.1	V	42.8	65.8	46.8	43.8	51.2
	SeqFormer [40]	offline	72.3	V	45.1	66.9	50.5	<u>45.6</u>	54.6
	SeqFormer [40]	offline	72.3	V+I	<u>47.4</u>	69.8	51.8	45.5	54.8
	IDOL(ours)	online	30.6	V	46.4	<u>70.7</u>	<u>51.9</u>	44.8	<u>54.9</u>
	IDOL(ours)[†]	online	30.6	V	49.5	74.0	52.9	47.7	58.7
ResNet-101	MaskTrack R-CNN [45]	online	-	V	31.8	53.0	33.6	33.2	37.6
	CrossVIS [46]	online	35.6	V	36.6	57.3	39.7	36.0	42.0
	PCAN [18]	online	-	V	37.6	57.2	41.3	37.2	43.9
	STEm-Seg [1]	offline	-	V+I	34.6	55.8	37.9	34.4	41.6
	VisTR [37]	offline	57.7	V	40.1	64.0	45.0	38.3	44.9
	MaskProp [3]	offline	-	V	42.5	-	45.6	-	-
	Propose-Reduce [21]	offline	-	V+I	43.8	65.5	47.4	43.0	53.2
	IFC [16]	offline	89.4	V	44.6	69.2	49.5	44.0	52.1
	SeqFormer [40]	offline	64.6	V+I	<u>49.0</u>	71.1	<u>55.7</u>	<u>46.8</u>	<u>56.9</u>
	IDOL(ours)	online	26.0	V	48.2	73.6	52.5	45.6	55.5
	IDOL(ours)[†]	online	26.0	V	50.1	<u>73.1</u>	56.1	47.0	57.9
Swin-L	SeqFormer [40]	offline	27.7	V+I	59.3	82.1	66.4	51.7	64.4
	IDOL(ours)	online	17.6	V	<u>61.5</u>	<u>84.2</u>	69.3	<u>53.3</u>	<u>65.6</u>
	IDOL(ours)[†]	online	17.6	V	62.2	86.5	<u>69.2</u>	54.6	68.1

resources. STEm-Seg [1] is the only offline method that can be directly evaluated on OVIS since its design enables it to run in a nearly online manner. To compare with the SOTA offline methods (*e.g.* IFC [16] and SeqFormer [40]), we split the video into short clips and apply the clipping matching method provided in IFC on these two methods (SeqFormer doesn’t provide its matching method). The results are provided in Table 5. Note that the previous best method only gains 15.4 AP on the validation set. IDOL with the same ResNet-50 backbone achieves 2× performance and gains 30.2 AP, surpassing the previous method by 14.8 AP. What’s more, when using a stronger backbone (*i.e.*, Swin-L) to extract better features, IDOL achieves the state-of-the-art performance of 42.6 AP, which is a huge improvement over previous best results.

Table 4. Comparison on YouTube-VIS 2021 val set. Best in **bold**, second underline.

Backbone	Method	Type	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
ResNet-50	MaskTrack R-CNN [45]	online	28.6	48.9	29.6	26.5	33.8
	SipMask [4]	online	31.7	52.5	34.0	30.8	37.8
	STMask [20]	online	31.1	50.4	33.5	26.9	35.6
	CrossVIS [46]	online	34.2	54.4	37.9	30.4	38.2
	IFC [16]	offline	36.6	57.9	39.3	-	-
	SeqFormer [40]	offline	<u>40.5</u>	<u>62.4</u>	<u>43.7</u>	<u>36.1</u>	<u>48.1</u>
	IDOL(ours)	online	43.9	68.0	49.6	38.0	50.9
Swin-L	SeqFormer [40]	offline	51.8	74.6	58.2	42.8	58.1
	IDOL(ours)	online	56.1	80.8	63.5	45.0	60.1

Table 5. Comparison on OVIS 2021 val set. Best in **bold** and second underline. The results with superscript “†” are not reported in [16, 40]. Videos in OVIS are much longer than those in YTVIS, offline models are unable to take the entire video as input due to the limit of computational resources. We split the video into clips of length 10 and perform clip matching provided by [16] on overlapping frames to get the final results.

Backbone	Method	Type	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
ResNet-50	MaskTrack R-CNN [45]	online	10.8	25.3	8.5	7.9	14.9
	SipMask [4]	online	10.2	24.7	7.8	7.9	15.8
	CMaskTrack R-CNN [30]	online	<u>15.4</u>	<u>33.9</u>	13.1	9.3	20.0
	CrossVIS [46]	online	14.9	32.7	12.1	10.3	19.8
	STEm-Seg [1]	offline	13.8	32.1	11.9	9.1	20.0
	IFC† [16]	offline	13.1	27.8	11.6	9.4	23.9
	SeqFormer† [40]	offline	15.1	31.9	<u>13.8</u>	<u>10.4</u>	<u>27.1</u>
	IDOL(ours)	online	30.2	51.3	30.0	15.0	37.5
Swin-L	IDOL(ours)	online	42.6	65.7	45.2	17.9	49.6

4.4 Ablation Study

In this section, we conduct extensive ablation experiments to study the importance of the core factors of our method on YouTube-VIS 2019. In the “Contrastive” setting, we use box IoU between predictions and ground truth for positive and negative embeddings selection following [29]. We further evaluate our optimal transport method to dynamically select positive and negative embeddings, termed as “OT”. For ablation study on inference strategy, “multi-cues” setting combines semantic consistency, spatial correlation, detection confidence and appearance similarity together to perform association following [45, 46]. Our association strategy is termed “embedding”.

Contrastive Training. To evaluate the importance of our contrastive embeddings, we apply the same association method on the embeddings predicted by ID Head and contrastive head. As shown in Table 6, contrastive training only improves 1.5 AP with the multi-cues association but improves 8.9 AP when it comes to the embedding-based association. Furthermore, optimal transport matching (OT) improves the results by 2.3 AP, which indicates that the choice

Table 6. Ablation study on contrastive learning and inference strategies on YTVIS.

Training			Inference		AP	AP ₇₅	AR ₁
ID Head	Contrastive	OT	Matching	Temporal			
✓	-	-	multi-cues	-	30.3	30.5	31.5
✓	-	-	embeddings	-	33.6	36.8	38.7
✓	-	-	embeddings	✓	34.8	37.6	39.4
-	✓	-	multi-cues	-	31.8	35.3	31.9
-	✓	-	embeddings	-	42.5	45.7	42.2
-	✓	✓	embeddings	-	44.8	49.9	43.4
-	✓	✓	embeddings	✓	46.4	51.9	44.8

of positive and negative embeddings plays an important role in learning discriminative embeddings. OT provides a better selection of positive and negative embeddings during training, and thus improves the quality of the embedding. We show the visualization of positive and negative embeddings selected by these two strategies in supplementary.

Association Strategy and Temporally Weighted Softmax. As shown in Table 6, compared with “multi-cues”, our embedding association strategy takes advantage of the discriminative embedding learned by contrastive learning, and improves the AP from 31.8 to 42.5 on YouTube-VIS. In addition, when temporally weighted softmax is added, it can be further improved by 1.6 AP. Utilizing information and priors from multiple previous frames improve robustness of association. Considering the problem of false positives and disappearing-and-reappearing that the online model needs to deal with, we believe this strategy helps maintain temporal consistency. We provide more visualization results, detailed analysis, and additional ablation experiments on OVIS in supplementary.

5 Conclusions

Online video instance segmentation methods have their inherent advantage in handling long/ongoing videos, but they are inferior to the offline models in performance. In this work, we aim to bridge the performance gap. We first deeply analyze the current online and offline models and find that the gap mainly comes from the error-prone association between frames. Based on this observation, we propose IDOL, which enables models to learn more discriminative and robust instance features for VIS tasks. It significantly outperforms all online and offline methods and achieves new SOTA on three benchmarks. We believe our insights on VIS methods will inspire future work in both online and offline methods.

Acknowledgements This work was supported by NSF 1763705. We thank the reviewers for their efforts and valuable feedback to improve our work.

References

1. Athar, A., Mahadevan, S., Osep, A., Leal-Taixé, L., Leibe, B.: Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In: ECCV (2020) 2, 7, 12, 13
2. Bergmann, P., Meinhardt, T., Leal-Taixe, L.: Tracking without bells and whistles. In: ICCV (2019) 4
3. Bertasius, G., Torresani, L.: Classifying, segmenting, and tracking object instances in video with mask propagation. In: CVPR (2020) 2, 4, 7, 12
4. Cao, J., Anwer, R.M., Cholakal, H., Khan, F.S., Pang, Y., Shao, L.: Sipmask: Spatial information preservation for fast image and video instance segmentation. In: ECCV (2020) 2, 4, 12, 13
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020) 5
6. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020) 5
7. Dendorfer, P., Osep, A., Milan, A., Schindler, K., Cremers, D., Reid, I., Roth, S., Leal-Taixé, L.: Motchallenge: A benchmark for single-camera multiple target tracking. *International Journal of Computer Vision* 129(4), 845–881 (2021) 4
8. Fu, Y., Yang, L., Liu, D., Huang, T.S., Shi, H.: Compfeat: Comprehensive feature aggregation for video instance segmentation. arXiv preprint arXiv:2012.03400 (2020) 2, 4, 12
9. Ge, Z., Liu, S., Li, Z., Yoshie, O., Sun, J.: Ota: Optimal transport assignment for object detection. In: CVPR (2021) 6, 7
10. Ge, Z., Liu, S., Wang, F., Li, Z., Sun, J.: Yolox: Exceeding yolo series in 2021. arXiv preprint arXiv:2107.08430 (2021) 6, 7
11. Han, T., Xie, W., Zisserman, A.: Self-supervised co-training for video representation learning. In: NeurIPS (2020) 5
12. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: CVPR (2020) 5
13. He, K., Gkioxari, G., Dollar, P., Girshick, R.: Mask R-CNN. In: ICCV (2017) 4, 5
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016) 11
15. Henaff, O.: Data-efficient image recognition with contrastive predictive coding. In: ICML (2020) 5
16. Hwang, S., Heo, M., Oh, S.W., Kim, S.J.: Video instance segmentation using inter-frame communication transformers. In: NeurIPS (2021) 2, 5, 7, 12, 13
17. Jiang, Z., Gu, Z., Peng, J., Zhou, H., Liu, L., Wang, Y., Tai, Y., Wang, C., Zhang, L.: Stc: Spatio-temporal contrastive learning for video instance segmentation. arXiv preprint arXiv:2202.03747 (2022) 4
18. Ke, L., Li, X., Danelljan, M., Tai, Y.W., Tang, C.K., Yu, F.: Prototypical cross-attention networks for multiple object tracking and segmentation. In: NeurIPS (2021) 2, 4, 12
19. Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. In: NeurIPS (2020) 5
20. Li, M., Li, S., Li, L., Zhang, L.: Spatial feature calibration and temporal fusion for effective one-stage video instance segmentation. In: CVPR (2021) 2, 4, 13
21. Lin, H., Wu, R., Liu, S., Lu, J., Jia, J.: Video instance segmentation with a propose-reduce paradigm. In: ICCV (2021) 2, 4, 12
22. Lin, T.Y., Dollar, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017) 5

23. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (2017) [6](#)
24. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021) [11](#)
25. Meinhardt, T., Kirillov, A., Leal-Taixe, L., Feichtenhofer, C.: TrackFormer: Multi-object tracking with transformers. arXiv preprint arXiv:2101.02702 (2021) [4](#)
26. Miech, A., Alayrac, J.B., Smaira, L., Laptev, I., Sivic, J., Zisserman, A.: End-to-end learning of visual representations from uncurated instructional videos. In: CVPR (2020) [5](#)
27. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: 2016 fourth international conference on 3D vision (3DV) (2016) [6](#)
28. Van den Oord, A., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv e-prints pp. arXiv-1807 (2018) [5](#)
29. Pang, J., Qiu, L., Li, X., Chen, H., Li, Q., Darrell, T., Yu, F.: Quasi-dense similarity learning for multiple object tracking. In: CVPR (2021) [3](#), [4](#), [5](#), [7](#), [13](#)
30. Qi, J., Gao, Y., Hu, Y., Wang, X., Liu, X., Bai, X., Belongie, S., Yuille, A., Torr, P.H., Bai, S.: Occluded video instance segmentation: A benchmark. International Journal of Computer Vision pp. 1–18 (2022) [3](#), [8](#), [13](#)
31. Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: CVPR (2019) [6](#)
32. Sun, P., Jiang, Y., Zhang, R., Xie, E., Cao, J., Hu, X., Kong, T., Yuan, Z., Wang, C., Luo, P.: TransTrack: Multiple-object tracking with transformer. arXiv preprint arXiv:2012.15460 (2020) [4](#)
33. Sun, Y., Cheng, C., Zhang, Y., Zhang, C., Zheng, L., Wang, Z., Wei, Y.: Circle loss: A unified perspective of pair similarity optimization. In: CVPR (2020) [5](#)
34. Tian, Z., Shen, C., Chen, H.: Conditional convolutions for instance segmentation. In: ECCV (2020) [5](#)
35. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: NeurIPS (2017) [4](#)
36. Voigtlaender, P., Krause, M., Osep, A., Luiten, J., Sekar, B.B.G., Geiger, A., Leibe, B.: Mots: Multi-object tracking and segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7942–7951 (2019) [4](#)
37. Wang, Y., Xu, Z., Wang, X., Shen, C., Cheng, B., Shen, H., Xia, H.: End-to-end video instance segmentation with transformers. In: CVPR (2021) [2](#), [4](#), [12](#)
38. Wu, J., Yarram, S., Liang, H., Lan, T., Yuan, J., Eledath, J., Medioni, G.: Efficient video instance segmentation via tracklet query and proposal. In: CVPR (2022) [2](#)
39. Wu, J., Bai, X., Jiang, Y., Liu, Q., Yuan, Z., Bai, S.: 1st place solution for youtubevos challenge 2022: Video instance segmentation [11](#)
40. Wu, J., Jiang, Y., Bai, S., Zhang, W., Bai, X.: Seqformer: Sequential transformer for video instance segmentation. In: ECCV (2022) [2](#), [5](#), [12](#), [13](#)
41. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via non-parametric instance discrimination. In: CVPR (2018) [5](#)
42. Xu, N., Yang, L., Yang, J., Yue, D., Fan, Y., Liang, Y., Huang, T.S.: Youtubevis dataset 2021 version. <https://youtube-vos.org/dataset/vis/> [8](#)
43. Xu, Z., Yang, W., Zhang, W., Tan, X., Huang, H., Huang, L.: Segment as points for efficient and effective online multi-object tracking and segmentation. TPAMI (2021) [4](#)

44. Yan, B., Jiang, Y., Sun, P., Wang, D., Yuan, Z., Luo, P., Lu, H.: Towards grand unification of object tracking. arXiv preprint arXiv:2207.07078 (2022) 4
45. Yang, L., Fan, Y., Xu, N.: Video instance segmentation. In: ICCV (2019) 1, 2, 3, 4, 5, 7, 8, 12, 13
46. Yang, S., Fang, Y., Wang, X., Li, Y., Fang, C., Shan, Y., Feng, B., Liu, W.: Crossover learning for fast online video instance segmentation. In: ICCV (2021) 2, 3, 4, 5, 7, 12, 13
47. Zhang, Y., Sun, P., Jiang, Y., Yu, D., Yuan, Z., Luo, P., Liu, W., Wang, X.: Bytetrack: Multi-object tracking by associating every detection box. arXiv preprint arXiv:2110.06864 (2021) 4
48. Zhang, Y., Wang, C., Wang, X., Zeng, W., Liu, W.: FairMOT: On the fairness of detection and re-identification in multiple object tracking. IJCV (2021) 4
49. Zhou, X., Koltun, V., Krähenbühl, P.: Tracking objects as points. In: ECCV (2020) 4
50. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. In: ICLR (2021) 5