

InstanceFormer: An Online Video Instance Segmentation Framework

Rajat Koner^{*1,3}, Tanveer Hannan^{*1,3}, Suprosanna Shit², Sahand Sharifzadeh¹,
Matthias Schubert^{1,3}, Thomas Seidl^{1,3}, Volker Tresp^{1,3}

¹ Ludwig Maximilian University of Munich

² Technical University of Munich

³ MCML

{koner, hannan}@dbs.ifi.lmu.de

Abstract

Recent transformer-based offline video instance segmentation (VIS) approaches achieve encouraging results and significantly outperform online approaches. However, their reliance on the whole video and the immense computational complexity caused by full Spatio-temporal attention limit them in real-life applications such as processing lengthy videos. In this paper, we propose a single-stage transformer-based efficient online VIS framework named InstanceFormer, which is especially suitable for long and challenging videos. We propose three novel components to model short-term and long-term dependency and temporal coherence. First, we propagate the representation, location, and semantic information of prior instances to model short-term changes. Second, we propose a novel memory cross-attention in the decoder, which allows the network to look into earlier instances within a certain temporal window. Finally, we employ a temporal contrastive loss to impose coherence in the representation of an instance across all frames. Memory attention and temporal coherence are particularly beneficial to long-range dependency modeling, including challenging scenarios like occlusion. The proposed InstanceFormer outperforms previous online benchmark methods by a large margin across multiple datasets. Most importantly, InstanceFormer surpasses offline approaches for challenging and long datasets such as YouTube-VIS-2021 and OVIS. Code is available at <https://github.com/rajatkoner08/InstanceFormer>.

Introduction

Video Instance Segmentation (VIS) (Yang et al. 2019) aims to simultaneously classify, segment, and track objects throughout video sequences. Therefore, VIS provides a holistic video understanding for various downstream tasks like autonomous driving and AR/VR applications. A vast amount of literature has been proposed to address VIS and can be grouped into two categories offline and online. Offline methods (Wu et al. 2021; Cheng et al. 2021; Hwang et al. 2021; Fang et al. 2021) process an entire video at once while online methods (Han et al. 2022; Yang et al. 2021b) process it sequentially as a stream.

As in many computer vision applications, such as, object detection (Zhu et al. 2020), object tracking (Meinhardt

^{*}Joint first-authorship and equal contributions.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

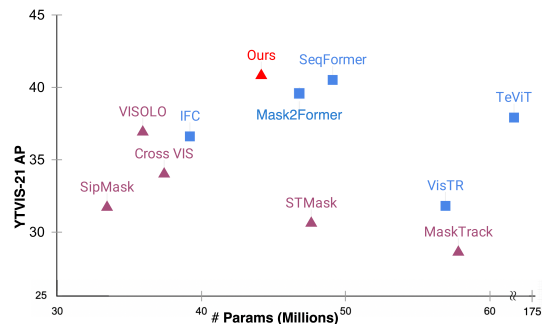


Figure 1: Comparison of performance (AP) vs. model size (params) in YTVIS-21 with existing offline and online models. InstanceFormer outperforms all existing methods with comparable or fewer parameters. See Tab. 2 for references.

et al. 2022), and scene-graph generation (Shit et al. 2022), recent transformer (Vaswani et al. 2017) based methods achieve remarkable success on offline VIS. These offline approaches predominantly employ Spatio-temporal attention (Wang et al. 2021) on the complete video or accumulate all frame-specific instance representations (Wu et al. 2021; Cheng et al. 2021). Thanks to the simultaneous attention over all frames from past and future, offline approaches can classify and segment instances with high precision even under challenging scenarios. However, relying on full videos and having an intractable Spatio-temporal complexity limits the application of offline methods to real-world scenarios where the videos are very long or streamed online. In contrast, online approaches can be employed in real-time and are applicable to long videos because of their sequential frame-by-frame processing. However, current online approaches suffer from a large performance gap compared to offline methods. The performance gap can mostly be attributed to a lack of long-term temporal dependency (Yang et al. 2021b, 2019) or semantic data association (Li et al. 2021; Han et al. 2022). We set three objectives to address these major obstacles for improving online VIS.

In a video, object instances usually change their appearance or location gradually. To explicitly capture the gradual changes in a scene, our **first objective** is to emphasize the immediate past with a robust instance propagation module.

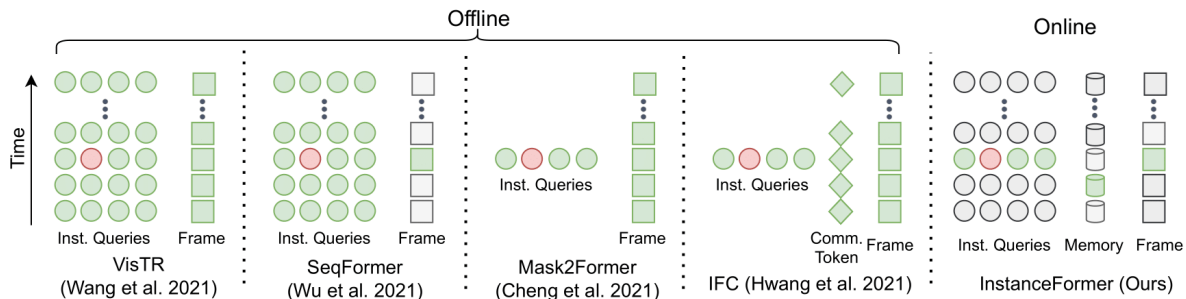


Figure 2: Comparison of attention scheme for a particular instance query (red) for recent transformer-based offline VIS and ours. The instance query receives information from the objects (green) to which it pays attention, whereas the other objects (gray) do not participate in the attention mechanism. Note that InstanceFormer drastically sparsifies Spatio-temporal attention by stressing the valuable past and the current information.

Next, we argue that instead of full Spatio-temporal attention, which requires intractable computational memory for long videos, one could encode the necessary semantic and temporal information in a considerably more compact way. Humans rarely need to memorize every detail in a frame; instead, we extract a high-level compact representation of the recent past and differentially recognize the upcoming video feed. A key component in this process is our working memory. To this end, our **second objective** is to adopt a realization of working memory into VIS. Finally, as a **third objective**, this memory must be temporally coherent to model real-life challenges such as occlusion.

Keeping the above objectives in mind, we propose InstanceFormer, an efficient framework for online VIS based on Deformable-DETR (Zhu et al. 2020). In Deformable-DETR, *instance queries* are responsible for efficient instance representation and learned through sparsified cross-attention from image features by restricting the attention to a set of *reference points*. Hence, we argue that the reference points are excellent markers of compact instance location. Complementary to the reference points, the *class-score* provides semantic information about the instances. Utilizing reference points, class scores, and instance queries, we efficiently establish inter-frame communication via a *prior propagation module* without cumbersome Spatio-temporal attention or additional data-association (Cao et al. 2020; Han et al. 2022). We derive the current frame’s instance queries and reference points from those of the previous frame. Thus, a specific instance query can be used to represent an instance throughout the sequence. In addition, the current frame’s categorization incorporates historical confidence scores to ensure the consistent and reliable classification of instances. Note that TrackFormer (Meinhardt et al. 2022) also handles multi-object tracking with query propagation. However, we significantly differ from them in terms of query initialization, reference point propagation and application scope.

In itself, prior propagation is insufficient for modeling scenarios like occlusions. Therefore, we introduce a memory mechanism that maintains a consistent representation and anticipates possible instance trajectories. We propose a memory queue of fixed temporal width that stores a compact representation of a specified number of past instances. We

exploit this stored memory in an additional cross-attention layer with the instance queries, which enables the current queries to be aware of past changes and anticipate deviations. Further, we aim to inject discriminative features into the instance representation to increase the similarity of the same instances and promote diversity among dissimilar instances across different time points. Hence, we impose supervised contrastive (Khosla et al. 2020) loss in the temporal direction to promote easy distinguishability of the instances.

In summary, our contributions are fourfold:

- We propose InstanceFormer, a simple single-stage transformer-based online VIS framework eliminating the need for an external tracking head or data association.
- We introduce a novel prior propagation module using reference points, class scores, and instance queries to enable efficient communication between consecutive frames.
- We propose a novel memory module to which the current instance queries attend to recollect the recent past. Additionally, temporally contrastive training makes the memory discriminative and easy to identify.
- InstanceFormer sets a new state-of-the-art online VIS by a large margin on YTVIS-19/21 and OVIS datasets. Crucially, InstanceFormer is the first online method that even outperforms the offline methods in the challenging YTVIS-21 (c.f. Fig. 1) and OVIS.

Related Work

Video instance segmentation (Yang et al. 2019) extends the task of image segmentation by tracking each instance throughout the videos. It can be grouped into two primary categories; offline, which processes an entire video simultaneously, and online, which processes a video sequentially.

Offline VIS: Initial attempts for offline VIS, relied on mask propagation (Bertasius et al. 2020; Lin et al. 2021). Inspired by DETR (Carion et al. 2020), there has been a recent surge of offline transformer-based end-to-end VIS frameworks. These approaches exploited instance queries for video-level instance understanding. VisTR (Wang et al. 2021) trivially extended DETR and proposed frame-specific instance queries and their association for the whole video. IFC (Hwang et al. 2021) established an inter-frame communication mechanism using a custom token. Later, Mask2Former

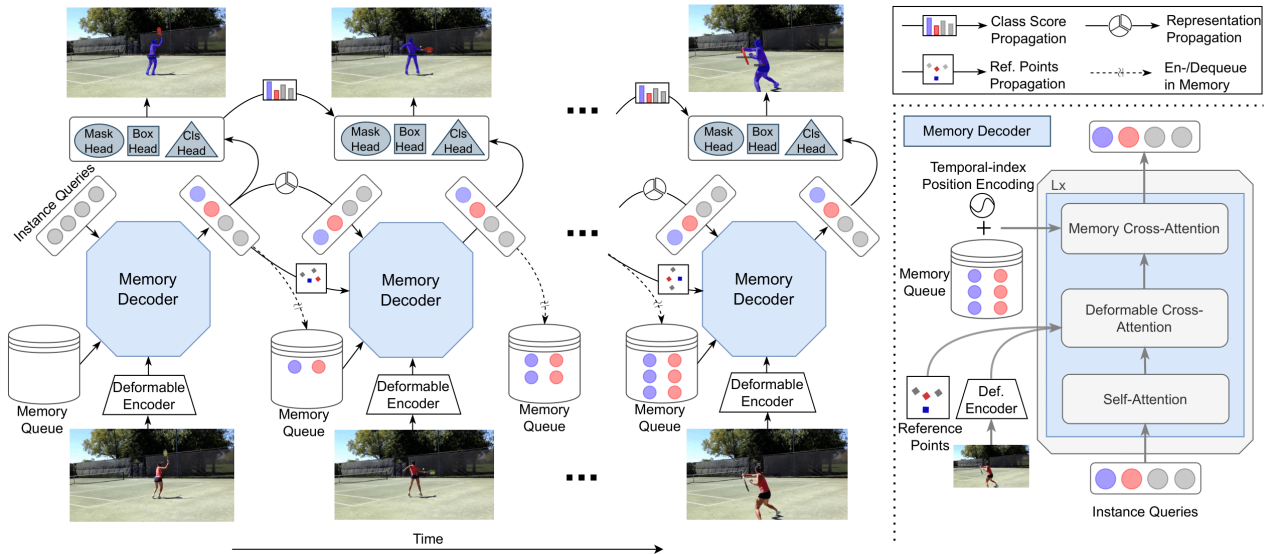


Figure 3: The overall architecture of InstanceFormer. Our key contributions are proposed propagation (representation, reference point, and class distribution) module and the memory decoder with a compact memory queue derived from salient instance queries. For a given frame, a deformable encoder extracts the image feature. In memory decoder, the initial instance queries attend to the image features and memory queue. Thereafter, the learned queries and reference points for deformable attention and class distribution are passed to the next frame. The memory decoder is elaborated on the right inset.

(Cheng et al. 2021), and SeqFormer (Wu et al. 2021) significantly improved the performance across multiple datasets. SeqFormer shares the initial instance query of each frame to learn a video-level instance embedding. Contrary to EfficientVIS’s (Wu et al. 2022) query re-initialization for inter-clip matching and factorized attention for intra-clip Spatio-temporal modeling, we use inter-frame propagation module. Despite their impressive performance, offline methods are not suitable for real-life applications or long video sequences due to their reliance on full videos and substantial memory requirements. Figure 2 compares the attention mechanism of existing transformer-based VIS methods with our proposed InstanceFormer.

Online VIS: In contrast, online methods are more challenging because of their inaccessibility to future frames. In recent time, a large number of online-VIS methods have been proposed that employ sequential frame-by-frame processing. (Yang et al. 2019) uses Mask-RCNN with a tracking head to assign instances to candidate boxes. (Yang et al. 2021b) proposes a cross-over learning scheme over the past temporal domain. QueryInst (Fang et al. 2021) uses query-based frame-wise segmentation and a tracking head. The current state-of-the-art (SOTA) online VIS method is VISOLO (Han et al. 2022) which uses grid memory to store information from earlier frames. Contrary to HITF (Li et al. 2022), we do not need IoU based instance matching and space-time attention on multiple past frames. However, the current online methods, including VISOLO, suffer from a large performance gap compared to the offline ones due to inadequate exploration of trajectories and representations of past frames. In this work, we aim to strengthen the performance of online VIS by propagating prior instance charac-

teristics over short- and long-term intervals.

Method

The overall architecture of InstanceFormer is shown in Fig. 3. Our method inherits the deformable encoder from Deformable-DETR (Zhu et al. 2020) and employs a novel prior propagation module followed by the proposed contrastively trained memory decoder. Finally, it has three heads (Wu et al. 2021): classification, box, and segmentation. It sequentially processes each frame in an online fashion. In the following, we will discuss each module in detail.

Deformable Encoder

Given an input video frame $x_t \in \mathbb{R}^{H \times W \times 3}$ at time t , of height H and width W , a feature extractor, e.g., a ResNet-50 (He et al. 2016) provides a multi-scale spatial feature map. We apply a projection layer (1×1 convolution with channel dimension of $C = 256$) on the extracted features. The projection layer transforms the feature map from different scales to a fixed size. Next, the projected features are flattened, and a fixed positional encoding is added. Then a series of deformable attention layers¹ processes the sequence to obtain the final feature map f_t .

Prior Propagation Module

The deformable decoder learns an instance representation by restricting attention to a few reference points close to the object instead of the whole image. Reference points are learnable and conditioned on their respective instance query. This

¹One can find details of deformable attention and reference points in the supplementary.

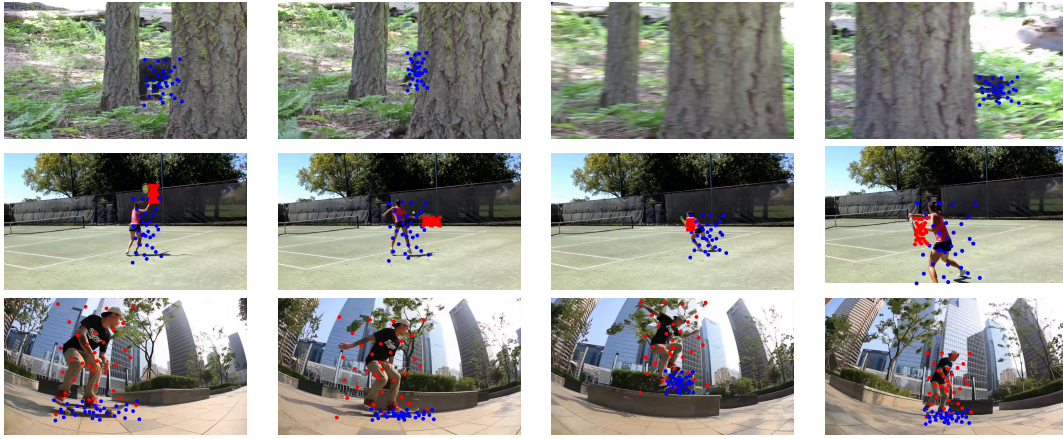


Figure 4: The reference points of the last layer of the decoder are overlaid on the image sequences, which show the compact propagation of object locations through reference points. Note that this propagation handles considerable occlusion and motion.

makes the attention sparse, computationally fast, and improves the overall object detection. Instances in a video often change their appearance and trajectory gradually over time. Hence, once an instance is localized, it is easier to calculate the offset from the earlier location than from scratch. Therefore to compute the current frame, we take advantage of instance representations, reference points, and class scores from the previous frame.

Representation Propagation: Let q_t^i be the final latent representation of the i^{th} instance query at time t . We propagate this information to the next frame at $t + 1$ by initializing its query weights with q_t^i . During the training, we only call the Hungarian matcher (Carion et al. 2020) in the presence of new ground truth instances to uniquely assign them to corresponding queries. Since we keep the index of matched queries the same between the previous frame and the current one, we do not need to solve the matching for every frame. Consequently, the query indices act as tracking ids for the instances, eliminating the need for any explicit tracking (Yang et al. 2019) or data association (Yang et al. 2021b).

Location Propagation: Representation propagation is helpful for high-level instance continuation, but it alone may not be sufficient for consistent low-level location and trajectory modeling. Hence, we explicitly propagate the prior location in terms of reference points of deformable attention. The reference points are dedicated markers of important spatial regions, including object boundaries or key points to which the instance queries attend to. Thus, the reference points are arguably the most representative feature for a specific instance as shown in figure 4. Note that InstanceFormer is the first method propagating reference points instead of bounding box (Meinhardt et al. 2022) or segmentation masks (Bertasius et al. 2020). As reference points are learnable, this propagation is flexible and easy to adapt under incremental changes. Thus, temporal continuation of reference points is analogous to the implicit trajectory, or motion prediction (Yang et al. 2021a; Hannan et al. 2022). At time t for the i^{th}

instance, reference points (ref_t^i) are computed as:

$$ref_t^i = \begin{cases} \sigma(W_{ref}(q_t^i)) & \text{if } t = 0 \\ \sigma(\sigma(W_{ref}(q_t^i)) \times ref_{t-1}^i) & \text{if } t > 0, \end{cases} \quad (1)$$

where W_{ref} , a linear layer, predicts reference points for each query, σ is the sigmoid function that projects a reference point to a normalized 2D image coordinate. Using the previous location as a starting point, the network learns to determine the offset for the current frame. As a result, the deformable cross attention benefits from the past location and helps to maintain a consistent instance trajectory.

Class Prior: Despite the representation and location prior, variation in appearance across frames makes the final instance classification inconsistent. To smooth out the probability of the class across frames, we propagate the class probability of past frames as a weighted sum. At time t the class probability c_t^i of the i^{th} instance is calculated as,

$$c_t^i = \hat{c}_t^i \times \text{Softmax}(\sigma(T_{cls}([c_f^i]_{f=t-d}^{t-1}))), \quad (2)$$

where the classification head consists of W_{cls} , temporal linear layer (T_{cls}), and a *sigmoid* activation. We concatenate the class distribution of the previous d frames, process them through a linear layer T_{cls} , and apply a *Softmax* over the temporal dimension. Eq. 2 essentially gathers class distribution over a short temporal domain and amplifies the most consistent class score in the current frame.

Altogether, our prior propagation provides a frame-to-frame inductive bias to replace the cumbersome Spatio-temporal attention and eliminates external tracking. This, in turn, reduces network overhead by computing the offset from the earlier frame. At the same time, it strengthens the temporal agreement between frames in the presence of gradual changes.

Memory Attention

Although prior propagation is sufficient for simple videos, it suffers in complex scenarios and accumulates errors from

past frames. For example, in the case of occlusion, a heavy reliance on the immediate prior may lose its effectiveness. Hence, an instance may lose its trace, resulting in performance drop. To alleviate this problem in an online fashion, we introduce a memory module to recollect salient instances from past frames. This recollection is achieved via our proposed instance query to memory cross-attention. Consequently, it allows an instance query to look into its past and correct any possible mistake. Memory in this case consists of memory tokens $([m_f^i]_{i,f})$ where i and f denote the instance index and relative temporal positions. The memory token m_f^i is defined as:

$$m_f^i = W_{m_proj}(sg(q_f^i, b_f^i, c_f^i)), \quad (3)$$

where b_{t-1}^i is box location, sg is the stop-gradient and $W_{m_proj} \in \mathbb{R}^C$ is a linear projection. To get a compact temporal perspective, we take $k(k \ll N)$ number of instances for each of the past $d(d \ll T)$ frames. Note that the number of tokens and memory size is much smaller than the number of instance queries (N) and the video length (T), respectively. Thus, having a small number of instances with discriminative features from past frames helps the network to learn continual temporal dependencies while keeping the additional computation in check. During training, we keep the matched query instances in the memory, and during inference, we keep the top- k query instances based on the class confidence. Additionally, we introduce a custom *temporal-index* embedding for memory tokens. This temporal-index embedding contains two parts: first, an index position embedding (Carion et al. 2020) of the instance queries at a given frame and second, a temporal 1D sinusoidal (Vaswani et al. 2017) embedding of memory tokens relative to the current frame. Our memory module maintains a fixed-size temporal queue of d frames with k number of instances per frame. As time proceeds, the oldest memory is replaced with the newest one through an enqueue/dequeue operation. One can express query to memory cross attention as

$$q_t^i = \text{Attn}(q_t^i, M_t), \quad (4)$$

where M_t is flattened memory tokens across all frames that serve as key and value to attention. The ordering of memory attention in the decoder is a design choice. Hierarchically, memory attention enables a higher order information gateway to assist the output of the deformable cross-attention through past instance cues. Therefore, we select the order of different attention layers in the decoder as 1) self, 2) deformable-cross, and 3) memory attention.

Temporal Contrastive Loss

Like DETR, we use the Hungarian algorithm to uniquely assign one instance query to a ground truth upon the appearance of any new instance. Thus, a particular query remains responsible for a specific instance throughout the video.

In addition to classification loss (\mathcal{L}_{cls}), box loss (\mathcal{L}_{box}) and cross-entropy mask loss (\mathcal{L}_{mask}) as in Deformable-DETR, we introduce a temporal contrastive loss (TCL) as shown in Fig. 5. As mentioned in Section , our third objective is to get temporally coherent embeddings of recurring instances while being discriminative to other instances

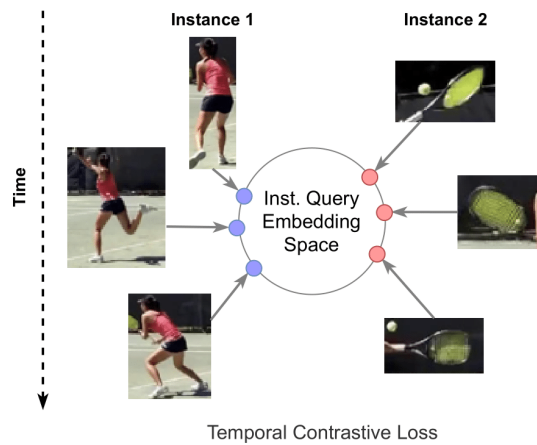


Figure 5: Temporal contrastive loss considers different appearances of an instance (anchor) throughout time as positive and all other instances as negative against the anchor.

across frames. Such a coherent embedding allows memory attention to effectively identify instances in scenarios like reappearance from occlusion. We take inspiration from (Khosla et al. 2020) and adopt the supervised contrastive loss to the temporal domain. In loss computation, we consider the embedding of an instance at different time points as a positively augmented pair, while the remaining combinations constitute negative pairs. Therefore, for the i^{th} instance at time t , \mathcal{L}_{tcl}^i can be expressed as

$$\mathcal{L}_{tcl}^i = - \sum_{f=t-d}^{t-1} \log \frac{\exp(q_t^i \cdot q_f^i / \tau)}{\sum_{j=0, j \neq i}^k \exp(q_t^i \cdot q_f^j / \tau)} \quad (5)$$

where τ is a temperature parameter. Finally, our joint training loss with their respective weights ($[\lambda_i]_{i=1}^4$) is:

$$\mathcal{L} = \lambda_1 \mathcal{L}_{cls} + \lambda_2 \mathcal{L}_{box} + \lambda_3 \mathcal{L}_{mask} + \lambda_4 \mathcal{L}_{tcl}. \quad (6)$$

Experiments

Datasets and Metrics: We evaluate InstanceFormer on three challenging datasets OVIS (Qi et al. 2021), YTVIS-19 (Yang et al. 2019), and YTVIS-21 (Xu et al. 2021). The recently proposed OVIS data set is more complex than the YTVIS data sets as it primarily contains long video sequences with high percentages of occlusion and a large number of objects. It contains 296k instance masks (2x of YTVIS-19) of 25 semantic categories, 5.8 instances per video (3.4x of YTVIS-19), and the longest sequence has more than 500 frames (13.5x longer than the same of YTVIS-19). Our second dataset, YTVIS-19, is the most popular VIS dataset, with 3, 859 videos, 40 different object categories, and an average length of 27 frames. Our third dataset, YTVIS-21, contains 33% more videos and almost double the amount of annotation compared to the 2019 version, making it more challenging than its predecessor. We use standard metrics for YTVIS and OVIS, namely Average Precision (AP) and Average Recall (AR).




Figure 6: Qualitative examples of InstanceFormer on the YTVIS-19 validation set. It includes occlusion and different poses.

Methods	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
STEm-Seg †	13.8	32.1	11.9	9.1	20.0
TeVIT †	17.4	34.9	15.0	11.2	21.8
SipMask	10.3	25.4	7.8	7.9	15.8
MaskTrack	10.9	26.0	8.1	8.3	15.2
CrossVIS	14.9	32.7	12.1	10.3	19.8
CMaskTrack	15.4	33.9	13.1	9.3	20.0
Ours	20.0	40.7	18.1	12.0	27.1

Table 1: Comparisons on OVIS dataset without feature calibration. † denotes the near-online methods.

Implementation Details: InstanceFormer is based on Deformable-DETR (Zhu et al. 2020) and employs a similar architecture with 6 layers of encoder and decoder. We first pre-train our network on COCO (Lin et al. 2014) for 12 epochs. Following SeqFormer and STEm-Seg, the pre-trained network is used for training both YTVIS and OVIS along with COCO to prevent over-fitting. We evaluate our model on the official validation split. In our experiments, we have set the size of memory queue $d = 4$ and number of memory token per frame to 10. We train our network using the AdamW optimizer on 4 NVIDIA RTX A6000 GPUs with a batch size of 4 and learning rate of $1e^{-4}$ for 16 epochs. Ablation on these hyper-parameters can be found in Sec. . More details about other model parameters and optimization can be found in the supplementary.

Results

Here, we analyze the quantitative and qualitative findings of our experiments across multiple dataset.

OVIS: Quantitatively, Table 1 shows the excellent performance of InstanceFormer against the previous online methods. Additionally, due to the longer sequences, most current offline counterparts (Wu et al. 2021; Cheng et al. 2021) cannot be applied due to their quadratic memory requirement. The competitive edge of InstanceFormer is also reflected in outperforming recent near online ones (Athar et al. 2020; Thawakar et al. 2022). In OVIS, occlusions are either long-term or short-term. Long-term occlusions involve the disap-

pearance and reappearance of instances that one can solve efficiently with discriminative visual features. However, the more abundant short-term occlusion is far more challenging because similar-looking objects cross each other, resulting in id switches. Recent online approaches heavily rely on visual (Yang et al. 2021b) cues and lack trajectory information which considerably weakens the performance in this scenario. In contrast, our implicit trajectory modeling through reference point propagation coupled with memory attention is more effective in mitigating this problem.

YTVIS-19: Table 2 compares InstanceFormer with state-of-the-art online and offline methods. InstanceFormer significantly outperforms all previous online methods by at least 8% AP. This trend is consistent in other metric like AP₅₀, AP₇₅, AR₁ and AR₁₀. Since YTVIS-19 contains mostly short videos with gradual changes, our proposed prior propagation convincingly supplies a continuous stream of information from past to present. Specifically, the reference point propagation glues consecutive frames in a sustained temporal trajectory. The benefit of such past-to-current communication, for the first time, bridges the gap between online and offline methods. Notably, InstanceFormer outperforms most offline methods while being marginally behind the recent benchmark like SeqFormer. For a fair comparison, we mention the backbone of different competing methods. Figure 6 shows typical qualitative examples.

YTVIS-21: Similar to YTVIS-19, InstanceFormer not only maintains its advantage over all online methods but also conquers the gap with the offline ones and outperforms them. Table 2 shows, InstanceFormer offers 4% and 0.2% higher AP than the closest online VISOLO and offline Mask2Former respectively. In YTVIS-21 increased number of instances per frame demands solving local interaction among instances apart from delineating their correct mask. A global offline model easily handles this interaction and delineation for a small number of instances. However, our explicit emphasis on local interaction reduces the complexity for a large number of instances.

Method	Backbone	YTVIS-19 (Yang et al. 2019)					YTVIS-21 (Xu et al. 2021)					Param (M)	FPS	
		AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀			
Offline	MaskProp	R50	40.0	—	42.9	—	—	—	—	—	—	—	—	—
	VisTR	R50	36.2	59.2	36.9	37.2	42.4	31.8	51.7	34.5	29.7	36.9	57.2	69.7
	EfficientVIS*	R50	37.9	59.7	43.0	40.3	46.6	34.0	57.5	37.3	33.8	42.5	—	36.0
	IFC	R50	41.2	65.1	44.6	42.3	49.6	36.6	57.9	39.3	—	—	39.3	107.1
	SeqFormer*	R50	47.4	69.8	51.8	45.5	54.8	40.5	62.5	43.6	36.2	48.0	49.3	—
	Mask2Former	R50	46.4	68.0	50.0	—	—	40.6	60.9	41.8	—	—	45.0	—
	TeViT	—	46.6	71.3	51.6	44.9	54.3	37.9	61.2	42.1	35.1	44.6	172.3	68.9
Online	M-RCNN	R50	30.3	51.1	32.6	31.0	35.5	28.6	48.9	29.6	26.5	33.8	58.1	26.1
	SipMask	R50	33.7	54.1	35.8	35.4	40.1	31.7	52.5	34.0	30.8	37.8	33.5	35.5
	SG-Net	R50	36.3	57.1	39.6	35.9	43.0	—	—	—	—	—	—	19.8
	CompFeat	R50	35.3	56.0	38.6	33.1	40.3	—	—	—	—	—	—	—
	CrossVIS	R50	36.3	56.8	38.9	35.6	40.7	34.2	54.4	37.9	30.4	38.2	37.5	25.6
	CrossVIS	R101	36.6	57.3	39.7	36.0	42.0	—	—	—	—	—	—	23.3
	STMASK	R50	33.5	52.1	36.9	31.1	39.2	30.6	49.4	32.0	26.4	36.0	47.8	26.1
	STMASK	R101	36.8	56.8	38.0	34.8	41.8	—	—	—	—	—	—	22.4
	VISOLO*	R50	38.6	56.3	43.7	35.7	42.5	36.9	54.7	40.2	30.6	40.9	35.0	40.0
	HITF	R50	41.3	61.5	43.5	42.7	47.8	35.8	56.3	39.1	33.6	40.3	—	—
	Ours*	R50	45.6	68.6	49.6	42.1	53.5	40.8	62.4	43.7	36.1	48.1	44.3	25.0

Table 2: Quantitative evaluation on YTVIS-19/21 validation set. FPS is reported on RTX 2080 ti. * denotes usage of the Augmented COCO (Athar et al. 2020).

Ablation	Query	Ref.	Class	AP	FPS
Prior Only	✓			14.5	27.0
	✓	✓		16.5	27.0
	✓	✓	✓	17.1	27.0
	Prior	Mem.	TCL		
Prior + Mem.	✓	✓		18.8	25.0
	✓	✓	✓	20.0	25.0

Table 3: Impact of the proposed prior propagation and memory module in the OVIS dataset.

Ablation

We perform ablations focusing on two perspectives: first, the effectiveness of the proposed prior propagation and memory module. Second is the influence of the number of memory tokens and past frames. All experiments are conducted on the OVIS dataset with a R50 backbone.

Prior Propagation and Memory: First, we inspect the contribution coming from different components of prior propagation in Table 3. Reference point propagation improves 2% AP over the baseline model with only query propagation. Further, class scores improve 0.6% AP. This shows our joint representation, location, and semantic propagation facilitate seamless information integration from the past frame with no need for an extra tracking module.

Subsequently, we probe the efficacy of our memory attention and TCL in table 3 given the full prior module. Memory attention improves 1.7% AP while TCL further boost 1.2% AP. Hence, memory plays a critical role in solving occlusion, and TCL assists it through instance-wise temporally coherent and discriminative cross-instance embedding.

Number of Tokens and Frames in Memory: Table 4 illustrates the effect of memory token count per frame. We find optimal performance with 10 tokens per frame. De-

# token/frame	5	10	15	20
AP	15.1	20.0	17.1	14.5
# frames	2	3	4	8
AP	19.0	19.5	20.0	16.9

Table 4: Ablation on the number of memory tokens and length of temporal memory frame in the OVIS dataset.

creasing the number will not cover the average count of instances per frame. However, an increase will introduce background instances. Further, we notice a memory queue size of four is adequate to retrieve essential representation from near history. It suggests importance of near history for short-term occlusion and preventing error propagation.

Conclusion

In this work, we proposed InstanceFormer, a single-stage online VIS framework. It specifically tackles real-life challenges such as long sequences and short-term occlusion. Further it simplifies the VIS pipeline by eliminating external data association or tracking. Thanks to our novel prior propagation and memory attention module, it can precisely track and segment instances by accumulating representation, trajectory, and semantic information from the past. Our proposed approach not only outperforms all existing online methods but also establishes a new state-of-the-art performance on challenging YTVIS-21 and OVIS datasets by outperforming existing offline and online methods. In summary, our carefully engineered adaptation of transformers pioneers high-performance online VIS while keeping the computation tractability in check for long video sequences. We hope to accelerate and direct future research towards sparse and efficient video-level interaction to improve VIS and beyond.

References

- Athar, A.; Mahadevan, S.; Osep, A.; Leal-Taixé, L.; and Leibe, B. 2020. Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In *European Conference on Computer Vision*, 158–177. Springer.
- Bertasius, G.; et al. 2020. Classifying, segmenting, and tracking object instances in video with mask propagation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9739–9748.
- Cao, J.; Anwer, R. M.; Cholakkal, H.; Khan, F. S.; Pang, Y.; and Shao, L. 2020. Sipmask: Spatial information preservation for fast image and video instance segmentation. In *European Conference on Computer Vision*, 1–18. Springer.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Cheng, B.; Choudhuri, A.; Misra, I.; Kirillov, A.; Girdhar, R.; and Schwing, A. G. 2021. Mask2former for video instance segmentation. *arXiv preprint arXiv:2112.10764*.
- Fang, Y.; Yang, S.; Wang, X.; Li, Y.; Fang, C.; Shan, Y.; Feng, B.; and Liu, W. 2021. Instances as queries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6910–6919.
- Han, S. H.; Hwang, S.; Oh, S. W.; Park, Y.; Kim, H.; Kim, M.-J.; and Kim, S. J. 2022. VISOLO: Grid-Based Space-Time Aggregation for Efficient Online Video Instance Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2896–2905.
- Hannan, T.; Koner, R.; Kobold, J.; and Schubert, M. 2022. Box Supervised Video Segmentation Proposal Network. *arXiv preprint arXiv:2202.07025*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hwang, S.; Heo, M.; Oh, S. W.; and Kim, S. J. 2021. Video instance segmentation using inter-frame communication transformers. *Advances in Neural Information Processing Systems*, 34: 13352–13363.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. *Advances in Neural Information Processing Systems*, 33: 18661–18673.
- Li, M.; Li, S.; Li, L.; and Zhang, L. 2021. Spatial feature calibration and temporal fusion for effective one-stage video instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11215–11224.
- Li, X.; Wang, J.; Li, X.; and Lu, Y. 2022. Hybrid instance-aware temporal fusion for online video instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 1429–1437.
- Lin, H.; Wu, R.; Liu, S.; Lu, J.; and Jia, J. 2021. Video instance segmentation with a propose-reduce paradigm. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1739–1748.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *European conference on computer vision*, 740–755. Springer.
- Meinhardt, T.; Kirillov, A.; Leal-Taixe, L.; and Feichtenhofer, C. 2022. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8844–8854.
- Qi, J.; Gao, Y.; Hu, Y.; Wang, X.; Liu, X.; Bai, X.; Belongie, S.; Yuille, A.; Torr, P. H.; and Bai, S. 2021. Occluded video instance segmentation. *arXiv preprint arXiv:2102.01558*.
- Shit, S.; Koner, R.; Wittmann, B.; Paetzold, J.; Ezhov, I.; Li, H.; Pan, J.; Sharifzadeh, S.; Kaissis, G.; Tresp, V.; et al. 2022. Relationformer: A Unified Framework for Image-to-Graph Generation. *arXiv preprint arXiv:2203.10202*.
- Thawakar, O.; Narayan, S.; Cao, J.; Cholakkal, H.; Anwer, R. M.; Khan, M. H.; Khan, S.; Felsberg, M.; and Khan, F. S. 2022. Video Instance Segmentation via Multi-scale Spatio-temporal Split Attention Transformer. *arXiv preprint arXiv:2203.13253*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, Y.; Xu, Z.; Wang, X.; Shen, C.; Cheng, B.; Shen, H.; and Xia, H. 2021. End-to-end video instance segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8741–8750.
- Wu, J.; Jiang, Y.; Zhang, W.; Bai, X.; and Bai, S. 2021. SeqFormer: a Frustratingly Simple Model for Video Instance Segmentation. *arXiv preprint arXiv:2112.08275*.
- Wu, J.; Yarram, S.; Liang, H.; Lan, T.; Yuan, J.; Eledath, J.; and Medioni, G. 2022. Efficient video instance segmentation via tracklet query and proposal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 959–968.
- Xu, N.; Yang, L.; Yang, J.; Yue, D.; Fan, Y.; Liang, Y.; and Huang, T. S. 2021. YouTube-VIS Dataset 2021 Version. <https://youtube-vos.org/dataset/vi>. Accessed: 2022-01-01.
- Yang, C.; Lamdouar, H.; Lu, E.; Zisserman, A.; and Xie, W. 2021a. Self-supervised video object segmentation by motion grouping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7177–7188.
- Yang, L.; et al. 2019. Video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5188–5197.
- Yang, S.; Fang, Y.; Wang, X.; Li, Y.; Fang, C.; Shan, Y.; Feng, B.; and Liu, W. 2021b. Crossover learning for fast online video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8043–8052.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.