

LOMM: Latest Object Memory Management for Temporally Consistent Video Instance Segmentation

Seunghun Lee^{1,*} Jiwan Seo¹ Minwoo Choi¹ Kiljoon Han¹
 Jahoon Jeong¹ Zane Durante² Ehsan Adeli^{2,†} Sang Hyun Park¹ Sunghoon Im^{1,†}

¹DGIST, Daegu, Republic of Korea ²Stanford University, Stanford, CA, USA

Abstract

In this paper, we introduce Latest Object Memory (LOM), a system for robustly tracking and continuously updating the latest states of objects by explicitly modeling their presence across video frames. LOM enables consistent tracking and accurate identity management across frames, enhancing both performance and reliability through the video segmentation process. Building upon LOM, we present Latest Object Memory Management (LOMM) for temporally consistent video instance segmentation, significantly improving long-term instance tracking. This enables consistent tracking and accurate identity management across frames, enhancing both performance and reliability through the video segmentation process. Moreover, we introduce Decoupled Object Association (DOA), a strategy that separately handles newly appearing and already existing objects. By leveraging our memory system, DOA accurately assigns object indices, improving matching accuracy and ensuring stable identity consistency, even in dynamic scenes where objects frequently appear and disappear. Extensive experiments and ablation studies demonstrate the superiority of our method over traditional approaches, setting a new state-of-the-art in video instance segmentation. Notably, our LOMM achieves an AP score of 54.0 on YouTube-VIS 2022, a dataset known for its challenging long videos. Project page: [this https URL](https://github.com/seunghunlee/LOMM)

1. Introduction

Video instance segmentation (VIS) is a complex task that requires segmenting, classifying, and tracking objects across video frames [42]. Recent breakthroughs in this field have been significantly propelled by the adoption of query-based segmentation networks [5, 6], which have notably enhanced the precision of instance segmentation. These networks utilize object queries to extract distinctive features for each

object within a frame and then group pixels on the image feature map to delineate object regions. Building on the capabilities of these architectures, the latest VIS approaches [14–16, 20, 23, 36, 46, 47] focus on developing sophisticated methods for instance association, aiming to improve the accuracy and efficiency of tracking objects across a video.

To enhance object association between frames, transformer-based trackers [47, 48] directly align predictions using cross-frame attention, dynamically reconstructing object representations from prior frames. While effective for short-term tracking, these methods heavily rely on previous frame data, which can hinder long-term consistency, as shown in the first row of Fig. 1. To mitigate this issue, some approaches incorporate a memory system, updating object representations based on similarity [19, 36, 46] or momentum mechanisms [11]. However, as shown in the second row of Fig. 1, these methods often struggle due to mixed representations, causing ambiguity in memory-object matching and leading to tracking failures. Even SAM2 [30], which stores not only object features but also spatial feature maps in memory, fails to maintain consistency in long videos, as demonstrated in the third row of Fig. 1. These limitations highlight the need for a more robust approach that consistently updates and associates object representations over long videos. (We refer the reader to the **Supplementary Material** for deeper understanding.)

To tackle the challenges inherent in VIS, we introduce *Latest Object Memory Management (LOMM)* designed for temporally consistent object tracking. As its core, *Latest Object Memory (LOM)* maintains and updates the latest state of each object, using the most recent object state as the key cue for tracking. By explicitly modeling object presence in each frame, our method efficiently handles dynamic scene where objects frequently appear and disappear. This memory-driven approach ensures stable identity management across frames, significantly improving VIS accuracy and reliability. We also introduce an *Decoupled Object Association (DOA)*, a method that leverages our memory system for precise object matching. This strategy separates the object association

* This work was done while visiting Stanford University.

† Corresponding author.



Figure 1. **Qualitative comparisons on a challenging video from YouTube-VIS 2022.** **Yellow text** highlights objects that failed to be tracked. (First row) DVIS++ [48] often fails to track objects that reappear after a temporal gap, as it relies solely on the previous frame. (Second row) CTVIS [46] employs a memory mechanism that aggregates object features across frames based on similarity weights, potentially leading to ID switching. (Third row) SAM2 [30] stores spatial feature maps alongside object features; however, it fails to effectively associate this information, leading to failures in challenging scenarios. (Fourth row) Our approach continuously updates and matches the latest state of each object, enabling robust and consistent tracking.

process into two stages: (1) tracking existing objects using stored memory states to maintain consistent identities and (2) assigning new IDs to newly appeared objects using our Occupancy-guided Hungarian Matching (OHM). By managing unoccupied indices, our method preserves existing object identities while seamlessly integrating new one.

Our approach improves upon traditional cross-frame attention mechanisms [47], which often misassign new objects to background indices, leading to learning conflicts. Instead of relying solely on object features from the previous frame as anchor queries, we incorporate DOA to determine object IDs in the current frame, enabling construction of adaptive anchor queries for more robust matching. By separately handling existing and new objects, our decoupled strategy ensures precise and consistent object alignment across frames, overcoming key limitations of prior methods. Extensive experiments and ablation study demonstrate that the effectiveness of our method, achieving superior performance over existing techniques.

Our contributions can be summarized as follows:

1. We present a new memory-based video instance segmentation, called *Latest Object Memory Management (LOMM)*, which dynamically manages the latest object states for temporally consistent object matching.
2. We introduce *Latest Object Memory (LOM)*, which accurately updates the most recent object features, crucial for precise object matching.

3. We propose *Decoupled Object Association (DOA)*, which matches existing objects in memory and assigns new IDs to newly detected ones.
4. LOMM excels in challenging long videos, achieving state-of-the-art performance in VIS benchmarks.

2. Related Works

Video Instance Segmentation. VIS primarily learns frame-to-frame feature associations using instance segmentation architectures. The seminal work, MaskTrack R-CNN [42], add a tracking head to Mask R-CNN [13], enhancing instance association. This is advanced by SipMask [2] and CrossVIS [43], which improves temporal associations with cross-frame learning. Additionally, IDOL [36] incorporates contrastive learning into a query-based architecture [52], boosting online method performance. Beyond online methods, VisTR [34] apply DETR [3] for clip-level instance predictions, constrained by dense self-attention. Efficiency improvements are further pursued with IFC [18], introducing a transformer with separate spatial and temporal attentions, and TeViT [44] and Seqformer [35], which adapt vision transformer backbones to enhance temporal associations and video-level instance predictions, respectively. Recently, query-based segmentation networks have become fundamental to contemporary VIS approaches, prominently featuring Mask2Former [6] as a key underlying technology. MinVIS [16] simplifies

tracking by using post-processing based on cosine similarity between object features. VITA [14] enhances this by temporally associating frame-level queries to identify instance prototypes within a video. GenVIS [15] creates a tracking network that operates at the sub-clip level. CTVIS [46] employs contrastive learning across an expanded frame set to achieve detailed frame associations. DVIS [47] introduces a decoupled architecture that segments the processes into distinct tasks of segmentation, tracking, and refinement.

Object Tracking with Memory. Memory-based methods have shown significant advancements in video analysis, particularly in tasks requiring sustained long-term consistency, such as video object segmentation [7, 9, 17, 31, 32, 38, 49], video instance segmentation [15, 22, 36, 42, 46], and video object tracking [1, 10, 11, 27, 39, 45, 50]. Some research has successfully utilized external memory in multi-object tracking scenarios. MaskTrack R-CNN [42] employs an external memory to store predicted instance representations, updating them with new data from the latest frame through a straightforward replacement rule. Building on this, with the introduction of DETR [3], Meinhardt et al. [27] developed TrackFormer, which innovates a *tracking-by-attention* approach. This method uses object tokens to maintain temporal instance memories, enhancing tracking continuity. Recent advancements have further improved long-term consistency by either stacking these tokens in a memory buffer [1] or applying momentum to update the tokens [11]. Despite these advancements, current techniques often fall short in addressing challenges related to the initial appearance and eventual disappearance of objects, especially due to occlusion. Even unified models [33, 37, 40, 41] struggle to handle such cases effectively. This gap indicates significant potential for further improvements in memory-based tracking methods.

3. Preliminary

Transformer-based Tracker. Video instance segmentation (VIS) involves segmenting and tracking objects consistently across video frames. To address this challenge, recent studies [14–16, 23, 46–48] have adopted query-based segmentation network like Mask2Former [6]. The segmentation network \mathcal{S} generates object representations $\tilde{Q}_t \in \mathbb{R}^{N \times C}$, categorical probabilities $P_t \in \mathbb{R}^{N \times (K+1)}$, and segmentation masks $M_t \in \mathbb{R}^{N \times H \times W}$ for each video frame $\{I_t\}_{t=1}^T$ as follows:

$$\left[\tilde{Q}_t, P_t, M_t \right] = \mathcal{S}(I_t), \quad \forall t = \{1, \dots, T\}, \quad (1)$$

where N is a sufficiently large number to detect objects in an image, while H , W , and C denote the height, width of predicted mask and the channel dimensions of the object representations, respectively. The categorical prediction head in the segmentation network classifies objects into K categories or as no-object \emptyset . The class label $c_t \in \mathbb{R}^N$ is determined by applying the argmax operation to the probability matrix P_t




Figure 2. **Limitations of existing trackers.** When a new object appears, there is no reference for assigning its ID, so it must be mapped to an ID previously used for the background. This results in an unnatural learning process where foreground and background features become overly similar, leading to suboptimal performance.

along the $(K + 1)$ dimension.

The principal challenge in VIS lies in maintaining consistent object representations across frames, ensuring that the aligned sequences of predicted objects $Q_t \in \mathbb{R}^{N \times C}$ consistently correspond to the same physical entities throughout the video, as follows:

$$Q_t = \mathcal{T}(Q_{t-1}, \tilde{Q}_t), \quad \forall t \in \{1, \dots, T\}, \quad (2)$$

where \mathcal{T} is a tracking network consisting of multiple transformer blocks and the initial queries Q_0 are initialized as raw features using raw features extracted from the segmentation network, \tilde{Q}_1 , from the first frame. Recent state-of-the-art methods [15, 47, 51] use a transformer-based tracker that encodes object features from each frame, informed by the sequence from previous frames. Cross-attention layers within these transformers align current frame objects with previous ones based on feature similarity, ensuring accurate and consistent tracking across the video.

Discussion. The tracking network described in Eq. (2) encounters two primary challenges. First, it relies on maintaining consistent object indices once they are assigned, but this consistency is not guaranteed. Since object association depends on matching objects across consecutive frames, objects that temporarily disappear and later reappear are at risk of being assigned incorrect indices. Although existing approaches, such as updating long-term memory with momentum [11] or enhancing object-wise similarity [19, 36, 46] aim to mitigate this, they often fall short of reliably preserving object identities over time.

The second challenge relates to the handling of newly appearing objects. The cross-attention mechanism reconstructs objects by comparing them to those already identified. However, new objects lack any prior reference, meaning they should be assigned to previously unallocated indices. For this to work, their features would need to resemble background features at first—an unnatural and impractical constraint, as shown in Fig. 2. This leads to ambiguous feature learning, ultimately degrading the system’s ability to accurately track newly appearing objects. To address these issues, we introduce an Latest Object Memory Management (LOMM) described in Sec. 4.

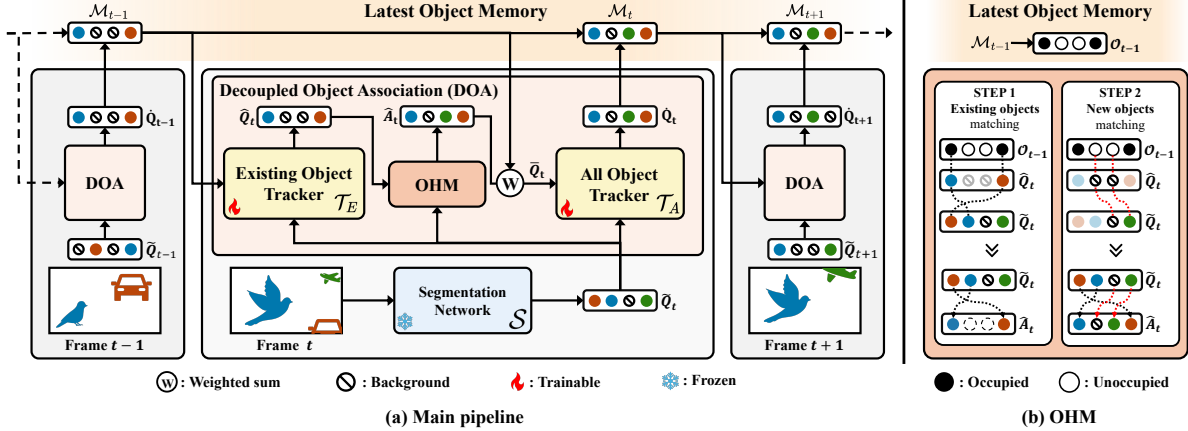


Figure 3. **Overall pipeline of our LOMM.** (a) For each frame, the existing object tracker \mathcal{T}_E utilizes \mathcal{M}_{t-1} from the Latest Object Memory and \tilde{Q}_t from the segmentation network \mathcal{S} to predict aligned features \hat{Q}_t for matched existing objects. The current frame’s existing objects \tilde{Q}_t and all objects \hat{Q}_t are matched through Occupancy-guided Hungarian Matching (OHM), ensuring accurate ID assignment for potential new objects and generating aligned object features \hat{A}_t . For robust tracking, an adaptive query \tilde{Q}_t is created by blending past information \mathcal{M}_{t-1} with the current frame’s aligned objects \hat{A}_t . This adaptive query \tilde{Q}_t and the current frame’s objects \hat{Q}_t are processed by \mathcal{T}_A to obtain the final tracking result \hat{Q}_t . (b) OHM assigns newly appeared objects to unoccupied slots using the occupancy prior \mathcal{O}_{t-1} . To ensure precise matching, OHM first assigns IDs to existing objects at occupied indices. Then, it allocates IDs to newly appeared objects at unoccupied indices, producing the aligned object features \hat{A}_t as output.

4. LOMM

As illustrated in Fig. 3, our model consists of two tracking networks, \mathcal{T}_E for existing objects and \mathcal{T}_A for all objects including potential new objects. These two networks align the object features \tilde{Q} produced by the pretrained segmentation network \mathcal{S} . Our approach incorporates two principal components to enhance tracking precision and consistency: a Latest Object Memory in Sec. 4.1 and a Decoupled Object Association strategy in Sec. 4.2. These sections detail how they contribute to maintaining object continuity and seamlessly integrating potential new objects into the scene.

4.1. Latest Object Memory

To enhance tracking in video instance segmentation (VIS), we first introduce Latest Object Memory, a crucial component of our LOMM framework. This memory system, denoted as \mathcal{M}_t , continuously updates object features online at each time step t . Specifically, it accumulates the aligned object features $\hat{Q}_t \in \mathbb{R}^{N \times C}$ over time as follows:

$$\mathcal{M}_t = (1 - p_t)\mathcal{M}_{t-1} + p_t\hat{Q}_t, \quad \forall t = \{1, \dots, T\}, \quad (3)$$

where p_t is the foreground probability for each object index. This probability is computed by summing the categorical probabilities $\{P_t^k\}_{k=1}^K$ across all object classes, excluding the no-object probability $P_t^{(K+1)}$. Our memory is initialized at \mathcal{M}_0 with \hat{Q}_1 and the aligned queries \tilde{Q}_t are defined in Eq. (6). This dynamic update process ensures consistent and robust maintenance of object information throughout the video sequence. By adaptively incorporating changes from the current frame, the memory effectively tracks and updates

object representations as they appear, move, and potentially disappear across the video. Unlike traditional memory update methods that might prioritize background information, our method focuses on foreground objects. This prioritization is crucial in dynamic video environments where objects frequently appear and disappear, ensuring that our system maintains accurate and reliable object information over time.

4.2. Decoupled Object Association

Based on our observation described in Sec. 3, we design a Decoupled Object Association strategy that consists of the following steps: (1) existing object tracking, (2) index assignment to new objects, and (3) object alignment with an adaptive anchor query. To avoid unnatural associations that can arise from directly tracking all objects observed in the instance queries \tilde{Q}_t , we initially focus on matching all instances to the existing objects recorded in the memory \mathcal{M}_{t-1} as follows:

$$\hat{Q}_t = \mathcal{T}_E(\mathcal{M}_{t-1}, \tilde{Q}_t), \quad (4)$$

where \hat{Q}_t represents aligned features of matched existing objects, and \mathcal{T}_E denotes the existing object tracker. This mechanism specifically focuses on identifying existing objects in the current frame, thereby minimizing ambiguity in the object association process. It is important to note that \hat{Q}_t excludes any objects that have newly appeared at time t . This ensures that the system accurately tracks and updates only those objects that persist across frames, without being confounded by newly introduced elements.

When new objects appear, they should be assigned previously unallocated indices for accurate tracking. To manage

the allocation state of each object, we define occupancy statuses \mathcal{O} as follows:

$$\mathcal{O}_t = \mathcal{O}_{t-1} \vee \mathcal{O}_t, \quad \mathcal{O}_t^n = \begin{cases} 0 & \text{if } c_t^n = \emptyset, \\ 1 & \text{otherwise} \end{cases}, \quad (5)$$

where c_t^n represents the predicted class, and \emptyset denotes the no-object class. \mathcal{O}_0 is initialized as a zero vector of length N , and each slot n is marked as occupied once a foreground object has been assigned to it at least once.

In our framework, the queries from segmentation network \tilde{Q}_t encompasses both existing and newly appeared objects, whereas the queries \hat{Q}_t contains only the existing objects at occupied indices. By matching these two object features, appropriate indices for the new objects can be identified. To facilitate this, we introduce an occupancy-guided Hungarian matching mechanism. This ensures that newly appeared objects are not erroneously matched to indices that are already occupied \mathcal{O}_{t-1} from the previous time ($t-1$). The matching starts by aligning objects corresponding to the occupied indices in \hat{Q}_t , followed by the remaining objects.

To ensure robust object association, we combine the aligned object queries \hat{A}_t from Alg. 1 with the latest object queries stored in \mathcal{M}_{t-1} to construct an adaptive anchor query for \mathcal{T}_A . Then, we reconstruct the current object features \hat{Q}_t using the adaptive anchor query through the alignment network \mathcal{T}_A as follows:

$$\hat{Q}_t = \mathcal{T}_A(\bar{Q}, \tilde{Q}_t), \quad \bar{Q}_t = \text{Adp}(\hat{A}_t, \mathcal{M}_{t-1}) \quad (6)$$

where $\text{Adp}(A, B) = \text{sim}(A, B) \cdot A + (1 - \text{sim}(A, B)) \cdot B$ and $\text{sim}(\cdot, \cdot)$ measures cosine similarity.

4.3. Training

Early training. In the initial stages of training, the quality of representations, such as \hat{Q}_t , can be poor, making effective object matching as outlined in Alg. 1 challenging. To improve early training outcomes and provide a more stable foundation for learning, we adopt the following approach:

$$\begin{aligned} \bar{Q}_t &= \text{Adp}(\tilde{Q}_t^*, \mathcal{M}_{t-1}), \\ \tilde{Q}_t^* &= \text{Hungarian}(\tilde{Q}_{t-1}^*, \tilde{Q}_t), \end{aligned} \quad (7)$$

where ‘‘Hungarian’’ refers to the Hungarian matching algorithm [21], used to improve the initial alignment of object representations between frames. The initial queries, \tilde{Q}_0^* , are initialized as raw features extracted from the segmentation network, \tilde{Q}_1 , in the first frame.

The initial outputs from the tracking networks \mathcal{T}_E and \mathcal{T}_A often exhibit considerable noise, which can impede the accuracy of the tracking process. To mitigate this, we adopt the approach described in [47] for ground truth assignment.

Algorithm 1: Occupancy-guided Hungarian Matching

- 1 **Input:** Existing objects $\hat{Q}_t \in \mathbb{R}^{N \times C}$, Current objects $\tilde{Q}_t \in \mathbb{R}^{N \times C}$, Occupancy statuses $\mathcal{O}_{t-1} \in \mathbb{B}^N$
 - 2 **Output:** Aligned object queries $\hat{A}_t \in \mathbb{R}^{N \times C}$
 - 3 **Initialize:** $\hat{A}_t \leftarrow$ empty matrix of size $N \times C$
 - 4 **Get index sets:**

$$\mathcal{I}_o \leftarrow \{i \mid \mathcal{O}_{t-1}^i = 1, i \in \{1, \dots, N\}\}$$

$$\mathcal{I}_u \leftarrow \{i \mid \mathcal{O}_{t-1}^i = 0, i \in \{1, \dots, N\}\}$$
 - 5 **Get object features for each index set:**

$$f_o \leftarrow \hat{Q}_t[\mathcal{I}_o], \quad f_u \leftarrow \hat{Q}_t[\mathcal{I}_u]$$
 - 6 **Find index set for occupied objects:**

$$\mathcal{J}_o \leftarrow \text{Hungarian}(f_o, \tilde{Q}_t)$$
 - 7 **Get unmatched features in \tilde{Q}_t :**

$$\tilde{f}_u \leftarrow \tilde{Q}_t \setminus \{\tilde{Q}_t[\mathcal{J}_o]\}$$
 - 8 **Find index set for unoccupied objects:**

$$\mathcal{J}_u \leftarrow \text{Hungarian}(f_u, \tilde{f}_u)$$
 - 9 **Assign aligned object query:**

$$\hat{A}_t[\mathcal{I}_o] \leftarrow \tilde{Q}_t[\mathcal{J}_o], \quad \hat{A}_t[\mathcal{I}_u] \leftarrow \tilde{Q}_t[\mathcal{J}_u]$$
 - 11 **Return** \hat{A}_t
-

Specifically, predictions \hat{y} from \tilde{Q}_t^* are used for assigning ground truth via Hungarian matching. This method is strategically implemented during the first half of the total training iterations. This phased application allows the model to adapt incrementally to the task’s complexity, enhancing the quality of the training representations as the process evolves. Such a staged training approach not only stabilizes the learning curve but also significantly improves alignment and tracking accuracy over time (see Tab. 3-(f)).

Training loss. In our approach, we utilize a comprehensive loss function aligned with those [4, 23, 47]. This function incorporates categorical cross-entropy, binary cross-entropy, and dice losses, which are pivotal for effectively training our model. We specifically focus on optimizing the networks \mathcal{T}_E and \mathcal{T}_A , while keeping other parameters static. We use both the actual ground truth y_t and a modified version, \hat{y}_t^n , which adapts the indices by assigning the no-object label \emptyset and zero mask where new objects appear. This modification ensures that \hat{y}_t reflects only predictions for existing objects, thereby enhancing the relevance and accuracy of the training process. The structure of our primary training loss is computed as follows:

$$\begin{aligned} \mathcal{L}_{\text{Track}} &= \sum_{t=1}^T \sum_{n=1}^{NGT} \left(\mathcal{L}(\hat{y}_t^n, \hat{y}_t^{\hat{\sigma}(n)}) + \mathcal{L}(y_t^n, \hat{y}_t^{\hat{\sigma}(n)}) \right), \\ \hat{\sigma} &= \arg \min_{\sigma \in \mathfrak{S}_N} \sum_{n=1}^{NGT} \mathcal{L}_{\text{Match}}(y_{f(n)}^n, \hat{y}_{f(n)}^{\sigma(n)}), \end{aligned} \quad (8)$$

where \hat{y}_t and \hat{y}_t are the predictions from the feature representations \tilde{Q}_t and \hat{Q}_t , respectively. \mathfrak{S}_N represents a permutation of N elements, and $\mathcal{L}_{\text{Match}}$ denotes a pair-wise matching cost

Table 1. Comparison on YouTube-VIS validation sets. * denotes offline methods. We employ a temporal refiner [47] for offline model.

Method	Backbone	YouTube-VIS 2019					YouTube-VIS 2021					YouTube-VIS 2022				
		AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
MinVIS [16]	R50	47.4	49.0	52.1	45.7	55.7	44.2	66.0	48.1	39.2	51.7	23.3	47.9	19.3	20.2	28.0
VITA* [14]	R50	49.8	72.6	54.5	49.4	61.0	45.7	67.4	49.5	40.9	53.6	32.6	53.9	39.3	30.3	42.6
GenVIS [15]	R50	50.0	71.5	54.6	49.5	59.7	47.1	67.5	51.5	41.6	54.7	37.5	61.6	41.5	32.6	42.2
DVIS [47]	R50	51.2	73.8	57.1	47.2	59.3	46.4	68.4	49.6	39.7	53.5	31.6	52.5	37.0	30.1	36.3
TCOVIS [23]	R50	52.3	73.5	57.6	49.8	60.2	49.5	71.2	53.8	41.3	55.9	38.6	59.4	41.6	32.8	46.7
DVIS-DAQ [51]	R50	55.2	78.7	61.9	50.6	63.7	50.4	72.4	55.0	41.8	57.6	34.6	-	35.5	-	41.1
DVIS++ [48]	R50	55.5	80.2	60.1	51.1	62.6	50.0	72.2	54.5	42.8	58.4	37.2	57.4	40.7	31.8	44.6
Ours	R50	55.7	79.8	61.4	51.3	62.7	50.7	72.9	56.9	43.7	58.4	41.1	62.4	46.2	35.8	47.5
MinVIS [16]	Swin-L	61.6	83.3	68.6	54.8	66.6	55.3	76.6	62.0	45.9	60.8	33.1	54.8	33.7	29.5	36.6
VITA* [14]	Swin-L	63.0	86.9	67.9	56.3	68.1	57.5	80.6	61.0	47.7	62.6	41.1	63.0	44.0	39.3	44.3
DVIS [47]	Swin-L	63.9	87.2	70.4	56.2	69.0	58.7	80.4	66.6	47.5	64.6	39.9	58.2	42.6	33.5	44.9
GenVIS [15]	Swin-L	64.0	84.9	68.3	56.1	69.4	59.6	80.9	65.8	48.7	65.0	45.1	69.1	47.3	39.8	48.5
TCOVIS [23]	Swin-L	64.1	86.6	69.5	55.8	69.0	61.3	82.9	68.0	48.6	65.1	51.0	73.0	53.5	41.7	56.5
DVIS++ [48]	ViT-L	67.7	88.8	75.3	57.9	73.7	62.3	82.7	70.2	49.5	68.0	37.5	53.7	39.4	32.4	43.5
DVIS-DAQ [51]	ViT-L	68.3	88.5	76.1	58.0	73.5	62.4	83.6	70.8	49.1	68.0	42.0	-	43.0	-	48.4
DVIS++* [48]	ViT-L	68.3	90.3	76.1	57.8	73.4	63.9	86.7	71.5	48.8	69.5	50.9	75.7	52.8	40.6	55.8
Ours	ViT-L	69.1	89.3	76.5	58.1	73.5	65.0	86.0	72.7	49.6	69.1	48.2	70.5	53.2	40.7	52.6
Ours*	ViT-L	70.1	90.7	77.7	58.5	74.8	66.2	88.6	74.9	49.9	70.6	54.0	77.1	57.9	43.2	58.8

[4]. $f(n)$ is computed specifically for the frame in which each n -th object first appears, following the approach used in [47]. This ensures that our model specifically learns from relevant, existing object features and avoids any confusion from non-object areas or noise.

To ensure that occupied indices in \tilde{Q} consistently represent the same objects across frames and that unoccupied indices reflect changes, we implement a similarity-based loss function as follows:

$$\mathcal{L}_{\text{Sim}} = \frac{1}{TN_{GT}} \sum_{t=2}^T \sum_{n=1}^{N_{GT}} \left(\text{sim} \left(\hat{A}_t^{\sigma(n)}, \mathcal{M}_{t-1}^{\sigma(n)} \right) - O_{t-1}^{\sigma(n)} \right)^2. \quad (9)$$

Our model is jointly trained using an objective function that combines the tracking loss and the similarity loss, with a balance determined by the weight λ_{Sim} :

$$\mathcal{L}_{\text{Total}} = \mathcal{L}_{\text{Track}} + \lambda_{\text{Sim}} \mathcal{L}_{\text{Sim}}. \quad (10)$$

5. Experiments

5.1. Implementation Details

We evaluate the performance of our LOMM using standard benchmark datasets: YouTubeVIS datasets (2019, 2021, 2022) [42] and OVIS [29]. For our segmentation network, we employ the Mask2Former architecture [6] equipped with three distinct backbone encoders: ResNet-50 [12], Swin-L [26], ViT-L and ViT-H [8]. All backbones are initialized with parameters pre-trained on COCO [25]. Our tracking framework integrates two networks \mathcal{T}_E and \mathcal{T}_A , each comprising three transformer blocks and enhanced with a referring cross-attention layer [47] for improved accuracy. Our tracking networks are trained with all other parameters frozen as previous studies [23, 47]. We empirically set λ_{Sim} as 1.0. Further details are described in the supplementary materials.

Table 2. Comparisons on OVIS validation sets.

Method	Backbone	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
MinVIS[16]	Swin-L	39.4	61.5	41.3	18.1	43.3
IDOL[36]	Swin-L	40.0	63.1	40.5	17.6	46.4
MDQE[24]	Swin-L	41.0	67.9	42.7	18.3	45.2
NOVIS[28]	Swin-L	43.0	66.9	44.5	18.9	46.3
GenVIS[15]	Swin-L	45.2	69.1	48.4	19.1	48.6
DVIS[47]	Swin-L	45.9	71.1	48.3	18.5	51.5
TCOVIS[23]	Swin-L	46.7	70.9	49.5	19.1	50.8
CTVIS[46]	Swin-L	46.9	71.5	47.5	19.1	52.1
Ours	Swin-L	47.8	73.6	51.4	19.2	52.4
DVIS++[48]	ViT-L	49.6	72.5	55.0	20.8	54.6
Ours	ViT-L	51.7	73.9	57.5	21.1	56.2
UNINEXT[41]	ViT-H	49.0	72.5	52.2	-	-
Ours	ViT-H	52.9	76.1	55.3	20.1	57.4

5.2. Main Results

Following the standard evaluation metrics, Average Precision (AP) and Average Recall (AR), we benchmark the performance of LOMM against the current state-of-the-art methods in video instance segmentation.

Results on Youtube-VIS. We compare LOMM with leading methods on the YouTube-VIS (YTVIS) datasets. The performance metrics, detailed in Tab. 1, show that LOMM outperforms state-of-the-art method, DVIS++, by achieving higher AP scores: +1.8 AP on YTVIS19, +2.2 AP on YTVIS21, and +3.1 AP on YTVIS22. Notably, LOMM also surpasses the performance of DVIS-DAQ, which addresses newly emerging and disappearing objects, under the same online setting equipped with a ViT-L backbone, achieving margins of +0.8 AP, +2.6 AP, and +6.2 AP on YTVIS19, YTVIS21, and YTVIS22, respectively. This significant improvement, particularly on longer video sequences, underscores LOMM’s ability to maintain long-term consistency effectively. This is largely attributed to its innovative latest object memory mechanism that adapts dynamically to complex video contexts.




Figure 4. **VIS results from our model on challenging videos in OVIS dataset.** (First and Second row) In both videos, the objects are successfully tracked without any ID switch even after complete occlusion. (Third row) The model robustly tracks a newly appeared object throughout the entire video, even in a crowded scene.

Results on OVIS. Our results on the OVIS benchmark are detailed in Tab. 2, where LOMM exhibits remarkable superiority over existing models. Notably, both with the ViT-L and ViT-H backbone, our model sets a new state-of-the-art by outperforming CTVIS by +0.9 AP, DVIS++ by +2.1 AP, and UNINEXT by +3.9 AP. The strong performance across these configurations, particularly in a dataset characterized by frequent occlusions and dynamic object appearances, highlights the efficacy of LOMM’s latest-state-aware memory and object association techniques in maintaining accurate and consistent tracking under challenging conditions. As shown in Fig. 4, our model demonstrates a clear advantage in complex scenes, accurately predicting object trajectories even amid severe occlusions and multiple interacting objects.

5.3. Ablation Study

To verify the effectiveness of our method on long videos, we conduct an ablation study on the YouTube-VIS-2022 dataset. We use a ResNet-50 backbone for all experiments and adopt the AP score as the evaluation metric.

Latest Object Memory. We demonstrate the effectiveness of our Latest Object Memory through the results in Tab. 3-(a) and (b), which show a notable improvement by +7.2 and +5.3 AP compared to the baseline MinVIS [16] and DVIS [47], respectively. The conventional memory systems struggle to effectively update recent object information, resulting in suboptimal tracking performance. (Refer to our supplementary materials for a more detailed analysis.) This underscores the effectiveness of our approach, con-

firmed that updating memory with foreground probability significantly enhances the model’s accuracy and reliability, especially in challenging scenarios where precise object association is essential.

Decoupled Object Association. Tab. 3-(a), (c), (d), and (e) demonstrates the effectiveness of the Decoupled Object Association, which consists of two trackers, T_E and T_A , and Occupancy-guided Hungarian Matching. Applying DOA to the baseline DVIS [47] results in +3.3 AP improvement, and combining it with LOM achieves a significant synergy, leading to a +7.2 AP gain. The aligned object queries \hat{A}_t , obtained by aligning the segmentation network’s output through Occupancy-guided Hungarian Matching (OHM), achieve a strong performance of 39.1 AP. Further refining the current frame’s object features \hat{A} using past information \mathcal{M}_{t-1} via \mathcal{T}_A results in the refined object features \hat{Q} , achieving an even higher score of 41.1 AP. Using either the current frame’s object \hat{A}_t or the past memory \mathcal{M}_{t-1} alone as the anchor query for \mathcal{T}_A results in suboptimal performance, yielding 39.3 AP and 39.2 AP, respectively. Notably, OHM without occupancy guidance still achieves a competitive 40.3 AP, but incorporating occupancy information enables more robust matching, further boosting performance. These results highlight the effectiveness of our matching strategy in consistently tracking both existing and newly appeared objects.

Early training strategy. Although LOMM demonstrates robust performance, the initial stages of training the tracker can be particularly challenging due to the scarcity of learned

Table 3. Ablation studies on each component of LOMM. All experiments report AP scores on the YouTube-VIS 2022 dataset using a ResNet-50 backbone. (a) Effect of Latest Object Memory (LOM) and Decoupled Object Association (DOA). (b) Performance comparison of different memory systems across baselines. (c) Comparison between prediction from \hat{A} and its refined output \hat{Q} via \mathcal{T}_A . (d) Performance comparison using \hat{A}_t (current frame), \mathcal{M}_{t-1} (past frames), and \bar{Q}_t (blended adaptive anchor) as anchors for \mathcal{T}_A . (e) Comparison between OHM and standard Hungarian matching. (f) Effect of early training.

(a) LOM, DOA			(b) Memory type					(c) Necessity of \mathcal{T}_A	
LOM	DOA	AP	Baseline	Memory type			Prediction from	AP	
		33.9		\times	Similarity	Momentum	LOM		
	✓	36.2	MinVIS [16]	26.8	28.5	30.9	34.0	\hat{A}	39.1
✓		39.2	DVIS [47]	33.9	35.8	35.2	39.2	\hat{Q}	41.1
✓	✓	41.1							

(d) Anchor query for \mathcal{T}_A				(e) Occupancy-guided Hungarian Matching		(f) Early training			
Metric	Anchor query			Metric	Occupancy awareness		Metric	Early training	
	\hat{A}_t	\mathcal{M}_{t-1}	\bar{Q}_t		\times	✓		\times	✓
AP	39.3	39.2	41.1	AP	40.3	41.1	AP	40.0	41.1

Table 4. Trade-off between computational cost and performance.

Tracker	Params (M) ↓	Time (ms) ↓	AP (YTVIS22) ↑
DVIS [47]	9.68	91.26	33.9
DVIS++ [48]	26.78	<u>97.73</u>	<u>37.2</u>
DVIS-DAQ [51]	18.62	209.15	34.6
Ours	<u>9.88</u>	101.16	41.1

information. To address this, we adopt a strategy where training initially relies on the outputs from the pre-trained segmentation network \mathcal{S} for the first half of the training iterations. This approach facilitates a more stable and informed beginning to the learning process. The effectiveness of this strategy compared to conventional methods is detailed in Tab. 3-(f), showcasing the advantages of integrating the pre-trained knowledge at the early stages.

5.4. Computational Cost

We introduce a novel tracking approach based on a decoupled architecture and compare its computational cost with existing methods (DVIS, DVIS++, and DVIS-DAQ) in Tab. 4. All models share the same segmentation network, with difference arising only in their tracking networks. The baseline model, DVIS, has a tracker with 9.68M parameters, while DVIS++ doubles the number of channels in its attention operations, increasing its parameter count to 26.78M. Similarly, DVIS-DAQ doubles the number of transformer blocks, resulting in 18.62M parameters. In contrast, our method maintains a nearly identical parameter count to the baseline at 9.88M.

Inference speed (measured in ms/frame) is evaluated using the R50 backbone on an RTX 2080 Ti GPU. Compared to DVIS and DVIS++, our method requires approximately 10ms and 3ms more per frame, respectively, while achieving substantial performance gains of +7.2 AP and +3.9 AP. Compared to DVIS-DAQ, our method is approximately twice as fast while obtaining +6.5 AP. These results highlight that

our method achieves an effective balances performance and computational efficiency.

6. Conclusion

In this paper, we present Latest Object Memory Management (LOMM) for temporally consistent video instance segmentation, a robust method designed to maintain temporal consistency in VIS. At the heart of LOMM is the Latest Object Memory, which utilizes the most recent instance features to effectively discern and track objects over time. The memory is essential for correctly assigning indices to new objects and distinguishing them from existing ones. We also propose the Decoupled Object Association strategy, integrating our object memory with occupancy information to ensure precise object matching. This approach divides the object association process into two phases: continuously tracking existing objects to preserve their identities and aligning all current frame instances via an adaptive anchor query. This ensures both stable and accurate tracking. Extensive experiments confirm that our LOMM achieves top performance across the most standardized VIS benchmarks.

Acknowledgments

This work was supported by the Ministry of Trade, Industry and Energy (MOTIE), South Korea, through the Technology Innovation Program under Grant RS-2024-00445759, for the development of navigation technology utilizing visual information based on vision-language models for understanding dynamic environments in nonlearned spaces), the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00210908) and the Institute of Information & Communications Technology Planning & Evaluation(IITP) grant funded by the Korea government(MSIT) (No. RS-2025-02219277, AI Star Fellowship Support Project(DGIST)).

References

- [1] Jiarui Cai, Mingze Xu, Wei Li, Yuanjun Xiong, Wei Xia, Zhuowen Tu, and Stefano Soatto. Memot: Multi-object tracking with memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8090–8100, 2022. 3
- [2] 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. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, pages 1–18. Springer, 2020. 2
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 2, 3
- [4] Bowen Cheng, Anwesa Choudhuri, Ishan Misra, Alexander Kirillov, Rohit Girdhar, and Alexander G Schwing. Mask2former for video instance segmentation. *arXiv preprint arXiv:2112.10764*, 2021. 5, 6
- [5] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in neural information processing systems*, 34: 17864–17875, 2021. 1
- [6] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 1, 2, 3, 6
- [7] Ho Kei Cheng and Alexander G Schwing. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European Conference on Computer Vision*, pages 640–658. Springer, 2022. 3
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 6
- [9] Kevin Duarte, Yogesh S Rawat, and Mubarak Shah. Capsulevos: Semi-supervised video object segmentation using capsule routing. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8480–8489, 2019. 3
- [10] Zhihong Fu, Qingjie Liu, Zehua Fu, and Yunhong Wang. Stm-track: Template-free visual tracking with space-time memory networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13774–13783, 2021. 3
- [11] Ruopeng Gao and Limin Wang. Memotr: Long-term memory-augmented transformer for multi-object tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9901–9910, 2023. 1, 3
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6
- [13] 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
- [14] Miran Heo, Sukjun Hwang, Seoung Wug Oh, Joon-Young Lee, and Seon Joo Kim. Vita: Video instance segmentation via object token association. *Advances in Neural Information Processing Systems*, 35:23109–23120, 2022. 1, 3, 6
- [15] Miran Heo, Sukjun Hwang, Jeongseok Hyun, Hanjung Kim, Seoung Wug Oh, Joon-Young Lee, and Seon Joo Kim. A generalized framework for video instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14623–14632, 2023. 3, 6
- [16] De-An Huang, Zhiding Yu, and Anima Anandkumar. Minvis: A minimal video instance segmentation framework without video-based training. *Advances in Neural Information Processing Systems*, 35:31265–31277, 2022. 1, 2, 3, 6, 7, 8
- [17] Xuhua Huang, Jiarui Xu, Yu-Wing Tai, and Chi-Keung Tang. Fast video object segmentation with temporal aggregation network and dynamic template matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8879–8889, 2020. 3
- [18] Sukjun Hwang, Miran Heo, Seoung Wug Oh, and Seon Joo Kim. Video instance segmentation using inter-frame communication transformers. *Advances in Neural Information Processing Systems*, 34:13352–13363, 2021. 2
- [19] Hanjung Kim, Jaehyun Kang, Miran Heo, Sukjun Hwang, Seoung Wug Oh, and Seon Joo Kim. Visage: Video instance segmentation with appearance-guided enhancement. *arXiv preprint arXiv:2312.04885*, 2023. 1, 3
- [20] Hojin Kim, Seunghun Lee, Hyeon Kang, and Sunghoon Im. Offline-to-online knowledge distillation for video instance segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 159–168, 2024. 1
- [21] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [22] Seunghun Lee, Jiwan Seo, Kiljoon Han, Minwoo Choi, and Sunghoon Im. Context-aware video instance segmentation. *arXiv preprint arXiv:2407.03010*, 2024. 3
- [23] Junlong Li, Bingyao Yu, Yongming Rao, Jie Zhou, and Jiwen Lu. Tcavis: Temporally consistent online video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1097–1107, 2023. 1, 3, 5, 6
- [24] Minghan Li, Shuai Li, Wangmeng Xiang, and Lei Zhang. Mdqe: Mining discriminative query embeddings to segment occluded instances on challenging videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10524–10533, 2023. 6
- [25] 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 *Computer Vision—ECCV 2014: 13th European Conference*,

- Zurich, Switzerland, September 6-12, 2014, *Proceedings, Part V 13*, pages 740–755. Springer, 2014. 6
- [26] 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. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 6
- [27] Tim Meinhardt, Alexander Kirillov, Laura Leal-Taixe, and Christoph Feichtenhofer. Trackformer: Multi-object tracking with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8844–8854, 2022. 3
- [28] Tim Meinhardt, Matt Feiszli, Yuchen Fan, Laura Leal-Taixe, and Rakesh Ranjan. Novis: A case for end-to-end near-online video instance segmentation. *arXiv preprint arXiv:2308.15266*, 2023. 6
- [29] Jiyang Qi, Yan Gao, Yao Hu, Xinggang Wang, Xiaoyu Liu, Xiang Bai, Serge Belongie, Alan Yuille, Philip HS Torr, and Song Bai. Occluded video instance segmentation: A benchmark. *International Journal of Computer Vision*, 130(8): 2022–2039, 2022. 6
- [30] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 1, 2
- [31] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Learning video object segmentation with visual memory. In *Proceedings of the IEEE international conference on computer vision*, pages 4481–4490, 2017. 3
- [32] Carles Ventura, Miriam Bellver, Andreu Girbau, Amaia Salvador, Ferran Marques, and Xavier Giro-i Nieto. Rvos: End-to-end recurrent network for video object segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5277–5286, 2019. 3
- [33] Junke Wang, Dongdong Chen, Zuxuan Wu, Chong Luo, Xiyang Dai, Lu Yuan, and Yu-Gang Jiang. Omnitracker: Unifying object tracking by tracking-with-detection. *arXiv preprint arXiv:2303.12079*, 2023. 3
- [34] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8741–8750, 2021. 2
- [35] Junfeng Wu, Yi Jiang, Song Bai, Wenqing Zhang, and Xiang Bai. Seqformer: Sequential transformer for video instance segmentation. In *European Conference on Computer Vision*, pages 553–569. Springer, 2022. 2
- [36] Junfeng Wu, Qihao Liu, Yi Jiang, Song Bai, Alan Yuille, and Xiang Bai. In defense of online models for video instance segmentation. In *European Conference on Computer Vision*, pages 588–605. Springer, 2022. 1, 2, 3, 6
- [37] Junfeng Wu, Yi Jiang, Qihao Liu, Zehuan Yuan, Xiang Bai, and Song Bai. General object foundation model for images and videos at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3783–3795, 2024. 3
- [38] Ning Xu, Linjie Yang, Yuchen Fan, Dingcheng Yue, Yuchen Liang, Jianchao Yang, and Thomas Huang. Youtube-vos: A large-scale video object segmentation benchmark. *arXiv preprint arXiv:1809.03327*, 2018. 3
- [39] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10448–10457, 2021. 3
- [40] Bin Yan, Yi Jiang, Peize Sun, Dong Wang, Zehuan Yuan, Ping Luo, and Huchuan Lu. Towards grand unification of object tracking. In *European conference on computer vision*, pages 733–751. Springer, 2022. 3
- [41] Bin Yan, Yi Jiang, Jiannan Wu, Dong Wang, Ping Luo, Zehuan Yuan, and Huchuan Lu. Universal instance perception as object discovery and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15325–15336, 2023. 3, 6
- [42] 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. 1, 2, 3, 6
- [43] Shusheng Yang, Yuxin Fang, Xinggang Wang, Yu Li, Chen Fang, Ying Shan, Bin Feng, and Wenyu Liu. Crossover learning for fast online video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8043–8052, 2021. 2
- [44] Shusheng Yang, Xinggang Wang, Yu Li, Yuxin Fang, Jiemin Fang, Wenyu Liu, Xun Zhao, and Ying Shan. Temporally efficient vision transformer for video instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2885–2895, 2022. 2
- [45] Tianyu Yang and Antoni B Chan. Learning dynamic memory networks for object tracking. In *Proceedings of the European conference on computer vision (ECCV)*, pages 152–167, 2018. 3
- [46] Kaining Ying, Qing Zhong, Weian Mao, Zhenhua Wang, Hao Chen, Lin Yuanbo Wu, Yifan Liu, Chengxiang Fan, Yunzhi Zhuge, and Chunhua Shen. Ctviz: Consistent training for online video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 899–908, 2023. 1, 2, 3, 6
- [47] Tao Zhang, Xingye Tian, Yu Wu, Shunping Ji, Xuebo Wang, Yuan Zhang, and Pengfei Wan. Dvis: Decoupled video instance segmentation framework. *arXiv preprint arXiv:2306.03413*, 2023. 1, 2, 3, 5, 6, 7, 8
- [48] Tao Zhang, Xingye Tian, Yikang Zhou, Shunping Ji, Xuebo Wang, Xin Tao, Yuan Zhang, Pengfei Wan, Zhongyuan Wang, and Yu Wu. Dvis++: Improved decoupled framework for universal video segmentation. *arXiv preprint arXiv:2312.13305*, 2023. 1, 2, 3, 6, 8
- [49] Yizhuo Zhang, Zhirong Wu, Houwen Peng, and Stephen Lin. A transductive approach for video object segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6949–6958, 2020. 3
- [50] Zixu Zhao, Jiase Wang, Max Horn, Yizhuo Ding, Tong He, Zechen Bai, Dominik Zietlow, Carl-Johann Simon-Gabriel, Bing Shuai, Zhuowen Tu, et al. Object-centric multiple object

tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16601–16611, 2023.

[3](#)

- [51] Yikang Zhou, Tao Zhang, Shunping Ji, Shuicheng Yan, and Xiangtai Li. Dvis-daq: Improving video segmentation via dynamic anchor queries. *arXiv preprint arXiv:2404.00086*, 2024. [3](#), [6](#), [8](#)
- [52] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [2](#)