

Video Mask Transfinner for High-Quality Video Instance Segmentation

Lei Ke^{1,2}, Henghui Ding¹, Martin Danelljan¹, Yu-Wing Tai³,
Chi-Keung Tang², and Fisher Yu¹

¹ Computer Vision Lab, ETH Zürich

² The Hong Kong University of Science and Technology

³ Kuaishou Technology

<http://vis.xyz/pub/vmt>

Abstract. While Video Instance Segmentation (VIS) has seen rapid progress, current approaches struggle to predict high-quality masks with accurate boundary details. Moreover, the predicted segmentations often fluctuate over time, suggesting that temporal consistency cues are neglected or not fully utilized. In this paper, we set out to tackle these issues, with the aim of achieving highly detailed and more temporally stable mask predictions for VIS. We first propose the Video Mask Transfinner (VMT) method, capable of leveraging fine-grained high-resolution features thanks to a highly efficient video transformer structure. Our VMT detects and groups sparse error-prone spatio-temporal regions of each tracklet in the video segment, which are then refined using both local and instance-level cues. Second, we identify that the coarse boundary annotations of the popular YouTube-VIS dataset constitute a major limiting factor. Based on our VMT architecture, we therefore design an automated annotation refinement approach by iterative training and self-correction. To benchmark high-quality mask predictions for VIS, we introduce the HQ-YTVIS dataset, consisting of a manually re-annotated test set and our automatically refined training data. We compare VMT with the most recent state-of-the-art methods on the HQ-YTVIS, as well as the Youtube-VIS, OVIS and BDD100K MOTS benchmarks. Experimental results clearly demonstrate the efficacy and effectiveness of our method on segmenting complex and dynamic objects, by capturing precise details.

Keywords: video instance segmentation, multiple object tracking and segmentation, video mask transfinner, iterative training, self-correction

1 Introduction

Video Instance Segmentation (VIS) requires tracking and segmenting all objects from a given set of categories. Most recent state-of-the-art methods [11, 14, 35, 36] are transformer-based, using learnable object queries to represent each tracklet in order to predict instance masks for each object. While achieving promising results, their predicted masks suffer from oversmoothed object boundaries and




Fig. 1. Video instance segmentation results by VisTr [35], IFC [14], SeqFormer [36], and VMT (Ours) along with the YTVIS Ground Truth. All methods adopt R101 as backbone. VMT achieves highly accurate boundary details, e.g. at the feet and tail regions of the tiger, even exceeding the quality of the GT annotations.

temporal incoherence, leading to inaccurate mask predictions, as shown in Figure 1. This motivates us to tackle the problem of *high-quality* video instance segmentation, with the aim to achieve accurate boundary details and temporally stable mask predictions.

Although high-resolution instance segmentation [15, 19] has been explored in the image domain, video opens the opportunity to leverage rich temporal information. Multiple temporal views can help to accurately identify object boundaries, and allow the use of correspondences across frames to achieve temporally consistent and robust segmentation. However, high-quality VIS poses major challenges, most importantly: 1) utilizing long-range spatio-temporal cues in the presence of dynamic and fast-moving objects; 2) the large computational and memory costs brought by high-resolution video features for capturing low-level details; 3) how to fuse fine-grained local features and with global instance-aware context for accurate boundary prediction; 4) the inaccurate boundary annotation of existing large-scale datasets [37]. In this work, we set out to address all these challenges, in order to achieve VIS with highly accurate mask boundaries.

We propose Video Mask Transfomer (VMT), an efficient video transformer that performs spatio-temporal segmentation refinement for high-quality VIS. To achieve efficiency, we take inspiration from Ke et al. [15] and identify a set of sparse error-prone regions. However, as illustrated in Figure 2, we detect 3D spatio-temporal points, which are often located along object motion boundaries. These regions are represented as a sequence of quadtree points to encapsulate various spatial and temporal scales. To effectively utilize long-range temporal cues, we group all points and jointly process them using a spatio-temporal refinement transformer. Thus, the input sequence for the transformer contains both detailed spatial and temporal information. To effectively integrate instance-aware global context, besides using the aggregated points as both input queries and keys of the transformer, we design an additional instance guidance layer (IGL). It makes our transformer aware of both local boundary details and global semantic context.

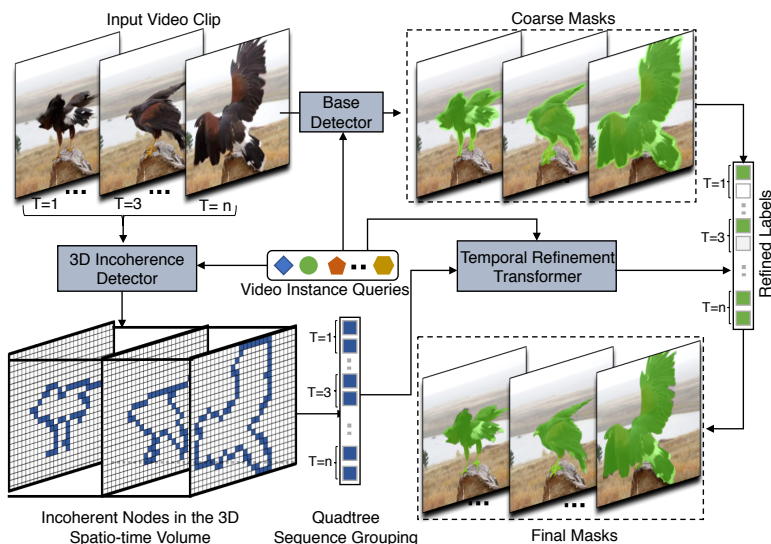


Fig. 2. We propose VMT for high-quality video instance segmentation. It adopts a temporal refinement transformer to jointly correct the 3D error-prone regions in the spatio-temporal volume. We employ VMT for automatically correcting the YTVIS with an iterative training paradigm by taking its annotation as coarse masks input.

While our VMT already achieves higher segmentation performance, we observed the boundary quality of the YTVIS [37] training annotations to be the next major bottleneck in the strive towards higher-quality mask predictions and evaluation on this popular, large-scale, and highly challenging dataset. Most importantly, we notice that many videos in YTVIS suffer object boundary inflation issues, as shown in Figure 1 and Figure 5. This introduces a learned bias in the trained model and prohibits very accurate evaluation. In fact, high-quality training data for VIS is difficult to obtain since dense pixel-wise annotations are costly for a large number of videos. To address this difficulty, instead of manually relabeling the training data, we design an automatic refinement procedure by employing VMT with iterative training. To self-correct mask annotations of YTVIS, both VMT model and training data are alternately evolved, as in Figure 3. To initialize the training of VMT annotation refinement, we use recently proposed OVIS [28] with better boundary annotations.

To enable benchmarking of high-quality VIS, we introduce the High-Quality YTVIS (HQ-YTVIS) dataset, consisting of our automatically refined training annotations and a manually re-annotated val & test split. Moreover, we propose the Tube-Boundary AP evaluation metric that better focuses on segmentation boundary accuracy as well as tracking ability. With the proposed HQ-YTVIS dataset, we retrain our VMT and several recent VIS baselines [11, 14, 16, 35, 37, 38] using our boundary-accurate annotations, providing a comprehensive comparison with current state-of-the-art. We also compare our VMT with state-of-the-

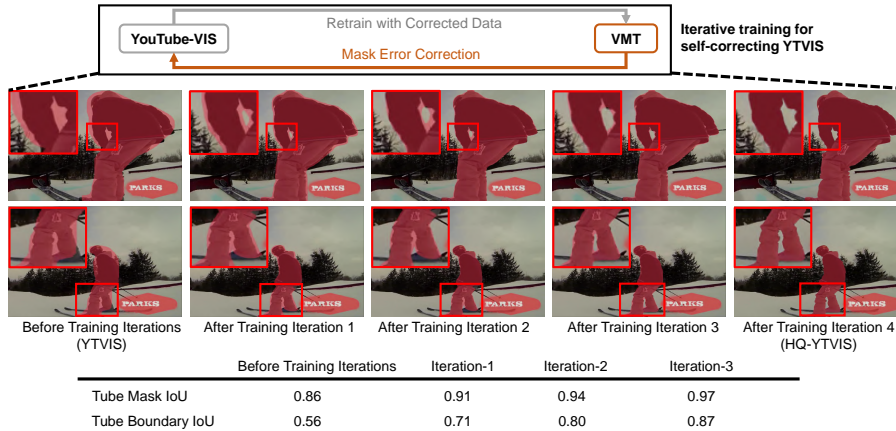


Fig. 3. Illustration and intermediate results visualization for iterative training. We show the mask quality change for the given case when correcting YTVIS coarse labels both qualitatively and quantitatively. The predicted instance masks boundaries by VMT becomes more fine-grained with more correction iterations on the YTVIS.

art methods on the OVIS [28] and BDD MOTS [39] benchmarks with better annotated boundaries. Quantitative and qualitative results on all three benchmarks demonstrate that VMT not only consistently outperforms existing VIS methods, but also predicts masks at much higher resolution size with small additional computation costs to current video transformer-based methods. We hope our VMT and HQ-YTVIS benchmark could facilitate the community in achieving ever more accurate video instance segmentation.

2 Related Work

Video Instance Segmentation (VIS) Extended from image instance segmentation, existing VIS methods can be divided into three categories: two-stage, one-stage, and transformer-based. Earlier methods [3, 21, 37] widely adopted the two-stage Mask R-CNN family [12, 13, 17] by introducing a tracking head for object association. Later works [5, 20, 23] adopted a one-stage instance segmentation framework by using anchor-free detectors [31] and linear combination of mask bases [4]. For longer temporal information modeling [22], CrossVIS [38] proposes instance-to-pixel relation learning and PCAN [16] introduces prototypical cross-attention operations for reading space-time memory. For the transformer-based approach, VisTr [35] first uses vision transformer [6] for VIS, which is then improved by IFC [14] using memory token communication. Seqformer [36] designs query decomposition mechanism. The aforementioned approaches put very limited emphasis on generating very accurate boundary details necessary of high-quality video object masks. In contrast, VMT is the first method targeting for very high-quality video instance segmentation.

Multiple Object Tracking and Segmentation (MOTS) MOTS methods [25, 26, 33] mainly follow the tracking-by-detection paradigm. To utilize temporal features, different from [2, 16] in clustering/grouping spatio-temporal feature, VMT directly detects the sparse error-prone points in the 3D feature space w/o feature compression and yield highly accurate boundary details.

Refinement for Segmentation Existing works [19, 30] on instance segmentation refinement are single-image based and thus neglect temporal information. Most of them adopt convolutional networks [30] or MLPs [19]. The latest image-based method Mask Transfiner [15] detects incoherent regions and adopts quadtree transformer for correcting region errors. Some methods [9, 10, 29, 34, 40] focus on refining semantic segmentation details. However, they apply on images without temporal object associations.

We build VMT based on [15], due to its efficiency and accuracy for single image segmentation. The key design of our VMT lies in leveraging temporal information and multi-view object associations of the input video clip. We explore new ways of using video instance queries to detect 3D incoherent points and correct spatio-temporal segmentation errors. Besides, VMT is also a part of our iterative training and self-correction to construct the HQ-YTVIS benchmark.

Self Training To reduce the expense of large-scale human-annotation on pixels, some semantic segmentation methods produce pseudo labels for unlabeled data using teacher model [7, 42] or data augmentation [43]. Then, their models are jointly trained on both human-labeled and pseudo labels. In contrast, VMT aims at self-correcting the coarsely or wrongly annotated VIS data. Considering that high-quality VIS requires very accurate video mask annotations to reveal object boundary details, our proposed self-correction and iterative training become even more valuable by eliminating such exhaustive manual labeling.

3 High-Quality Video Instance Segmentation

We tackle the problem of high-quality Video Instance Segmentation (VIS), by proposing an efficient temporal refinement transformer, Video Mask Transfiner (VMT), in Section 3.1. We further introduce a new iterative training paradigm for automatically correcting inaccurate annotations of YTVIS in Section 3.2. To facilitate the research in high-quality VIS, we contribute a large-scale HQ-YTVIS benchmark, and propose the Tube-Boundary AP metric in Section 3.3. The proposed benchmark and metric contribute to existing and future VIS models, with high-quality annotations for both better training and more precise evaluation.

3.1 Video Mask Transfiner

Figure 4 depicts the overall architecture of Video Mask Transfiner (VMT). Our design is inspired by the image-based instance segmentation method Mask Transfiner [15]. This single-image method first detects incoherent regions, where segmentation errors most likely occur in the coarse mask prediction. A quadtree transformer is then used to refine the segmentation in these regions. However, in

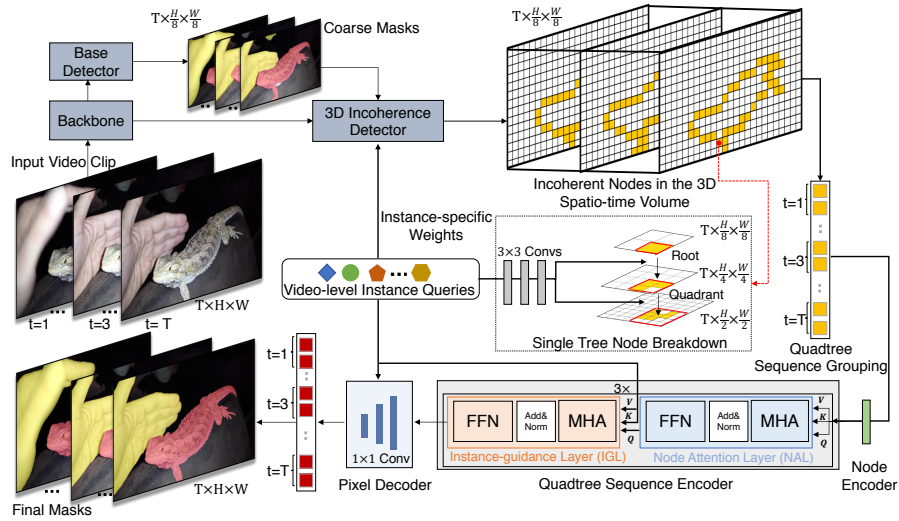


Fig. 4. Our VMT framework. A sequence of quadtrees are first constructed in the spatio-time volume by the 3D incoherence detector. Then, these incoherent nodes are concatenated across frames by Quadtree Sequence Grouping. The produced new spatio-temporal node sequences are corrected by temporal refinement transformer under the guidance of video instance queries with global instance context.

case of video, the usage of temporal information, including object associations between different frames, is not accounted for by Mask Refiner. This limits its segmentation performance in the video domain, leading to temporally incoherent mask results. To effectively and efficiently leverage the high-resolution temporal features, we propose three new components for our VMT: 1) an instance query based 3D incoherent points detector; 2) quadtree sequence grouping for temporal information aggregation; and 3) instance query guided incoherent points segmentation. We will describe each of these key components in this section, after a brief summary of the employed base detector in the following.

Backbone and Base Detector Given a video clip that consists of multiple image frames as input, we first use CNN backbone and transformer encoder [41] to extract feature maps for each frame. Then, we adopt video-level instance queries to detect and segment objects for each frame following [36]. This base detector [36] generates initial coarse mask predictions of the video tracklets at low resolution $T \times \frac{H}{8} \times \frac{W}{8}$, where T , H and W are the length, height and width of the input video clip. Given this input data, our goal is to predict highly accurate video instance segmentation masks at $T \times H \times W$.

Query-based 3D Incoherent Points Detection To detect the incoherent regions in the video clip, where segmentation errors are concentrated, a lightweight

3D incoherent region detector is designed. The detector, which encodes the video-level instance query embedding to generate a set of dynamic convolutional weights, consists of three 3×3 convolution layers with ReLU activations. The predicted instance-specific weights are then convolved with the spatio-temporal feature volume at resolution $T \times \frac{H}{8} \times \frac{W}{8}$, followed by a binary classifier to detect the 3D sparse incoherent tree roots.

We further break down these predicted incoherent points in the 3D volume into each frame. Each point serves as root node in a tree, by branching each node into its four quadrants on the corresponding lower-level frame feature map, which is $2 \times$ higher in resolution. The branching is recursive until reaching the largest feature resolution. We share this 3-layer dynamic instance weights to detect incoherent points for the same video instance across backbone feature sizes at $\{\frac{H}{8} \times \frac{W}{8}, \frac{H}{4} \times \frac{W}{4}, \frac{H}{2} \times \frac{W}{2}\}$, as visualized in Figure 4. This allows VMT to save a huge computational and memory cost, because only a small part of the high-resolution video features are processed, occupying less than 10% of the all the points in the 3D temporal volume. Video-level instance query captures both positional and appearance information for a time sequence of the same instance in a video clip. The instance-specific information are already contained in the correlation weights. Thus, different from [15], instance query-based detection removes the necessity of constructing ROI pooling feature pyramid for each video object. Our 3D incoherent region detector directly operates on the spatio-temporal feature volume from the backbone.

Quadtree Sequence Grouping After detecting 3D incoherent points, we build a sequence of quadtree points within the video clip, each of which resides in a single frame. To effectively utilize the temporal information across frames, VMT groups together all the tree nodes from all frames of the quadtree sequence, and concatenate them in the token dimension for the transformer. The resulting new sequence is the input for the temporal refinement transformer, which contains tree nodes across both spatial and temporal scales, thus encapsulating both detailed spatial and temporal information. We study the influence of different video clip lengths in Table 1, which reveals that the input sequence from longer video clips with more diverse and rich information boosts the accuracy of temporal segmentation refinement.

Instance Query Guided Temporal Refinement For segmenting the newly formed incoherent sequence above, instead of solely leveraging the incoherent points as both input queries and keys [15], our Node Attention Layer (NAL) utilizes video-level instance queries as additional semantic guidance. In Figure 4, to inject each point with instance-specific information, we introduce the Instance Guidance Layer (IGL) after each NAL in a level-wise manner. IGL uses incoherent points only as queries, and adopts the video-level instance embedding as the keys and values. This helps our temporal refinement transformer be aware of both local boundary details and global instance-level context, thus better separating incoherent points among different foreground instances. Besides, we add

a low-level RGB feature embedding, produced by a network consisting of three 3×3 Conv. layers directly operating on the image. This further encapsulates fine-grained object edge details as input to the node encoder. Finally, the output is sent into the dynamic pixel decoder for final prediction.

3.2 Iterative Training Paradigm for Self-correcting YTVIS

We observed the boundary annotation quality of the YTVIS dataset to be an important bottleneck when aiming to learn highly accurate segmentation masks. We show the inaccurate and coarse boundary annotations of YTVIS in Figure 5, Figure 1 and the supplemental video. In particular, we randomly sample 200 videos from the original YTVIS annotations, and find around 28% of the cases suffer from the boundary inflation problem, where a halo about 5 pixels is around the real object contour. These coarse annotations may due to small number of selected polygon points during instance labeling, which introduces a severe bias in the training, leading to inaccurate boundary prediction. Based on VMT, we therefore design a method for automatic annotation refinement, and apply it to correct the inaccurate annotations of YTVIS. The core idea is to take the coarse mask annotations from HQ-YTVIS as input and alternate between refining the training data and training the model to achieve gradually improved annotations.

At the beginning, to equip VMT with initial boundary correction ability, we pretrain VMT on the better annotated OVIS dataset as the first iteration, which has similar data categories and sources as YTVIS. We train the temporal refinement transformer of VMT in a class-agnostic way, leveraging only the incoherent points and video-level instance queries as the input. To simulate various shapes and output of inaccurate segmentation, we degrade the video mask annotations of OVIS [28] by subsampling the boundary regions followed by random dilations and erosions. Examples of such degraded masks are in the supplemental file. VMT is trained to correct the errors in the ground-truth incoherent regions, and we further enlarge the regions by dilating 3 pixels to introduce both the diversity and the balance of foreground and background pixels ratio in this region.

After training on OVIS, we employ the trained VMT to correct the mask boundary annotations of YTVIS, where the mask annotations of YTVIS are regarded as the coarse mask inputs. We only correct the mask labels when the confidence of the most likely predicted class (foreground or background) is larger than 0.65. Then, we obtain a corrected version of YTVIS and use this new corrected YTVIS data to retrain the temporal refinement transformer of VMT as the 2nd iteration. We iterate this process until the model performance on the manually labeled validation set reaches saturation, requiring 4 iterations. We illustrate the iterative training process and show the intermediate visualizations in Figure 3. After each iteration, the produced annotations masks of YTVIS become more fine-grained until final convergence. We compare the training results using different iterated versions of the YTVIS data, and evaluate their performance on the human-relabeled *val* set in Table 3.




Fig. 5. Masks quality comparisons between YTVIS [37] and HQ-YTVIS annotations.

3.3 The HQ-YTVIS Benchmark

To facilitate the research in high-quality VIS, we further contribute a new benchmark HQ-YTVIS and design a new evaluation metric Tube-Boundary AP.

HQ-YTVIS To construct the HQ-YTVIS, we first randomly re-split the original YTVIS training set (2238 videos) with coarse mask boundary annotations into train (1678 videos, 75%), val (280 videos, 12.5%) and test (280 videos, 12.5%) subsets following the splitting ratios in YTVIS. Then, the masks annotations on the train subset is self-corrected automatically by VMT using iterative training as described in Section 3.2. The smaller set of validation and test videos are carefully relabeled by human annotators to ensure high mask boundary quality. Figure 5 shows the mask annotation differences of the same image from the training set between HQ-YTVIS and YTVIS. HQ-YTVIS has much more accurate object boundary annotations. We retrained VMT and all baselines [11, 14, 16, 35, 37, 38] on HQ-YTVIS from scratch, and compare the results with those obtained by training them on the original YTVIS annotations with the same set of images. We conduct quantitative results comparisons results in Table 4, which clearly shows the advantage brought by HQ-YTVIS. We also include the relevant qualitative comparisons in the Supp. file. We hope HQ-YTVIS can serve a new and more accurate benchmark to facilitate future development of VIS methods aiming at higher mask quality.

Tube-Boundary AP We propose a new segmentation measure Tube-Boundary AP for high-quality video instance segmentation. The standard tube mask AP in [37] is biased towards object interior pixels [8, 19], thus falling short of revealing motion boundary errors, especially for large moving objects. Given a sequence

of GT masks $G_{b\dots e}^i$ for instance i , a sequence detected masks $P_{\hat{b}\dots\hat{e}}^j$ for predicted instance j , we extend frame index b and \hat{b} to 1, e and \hat{e} to T for temporal length alignment using empty masks. Tube-Boundary AP (AP^B) is computed as,

$$\text{AP}^B(i, j) = \frac{\sum_{t=1}^{t=T} |(G_t^i \cap g_t^i) \cap (P_t^j \cap p_t^j)|}{\sum_{t=1}^{t=T} |(G_t^i \cap g_t^i) \cup (P_t^j \cap p_t^j)|} \quad (1)$$

where spatio-temporal boundary regions g and p are respectively the sequential set of all pixels within d pixels distance from the contours of $G_{b\dots e}^i$ and $P_{\hat{b}\dots\hat{e}}^j$ in the video clip. By definition, Tube-Boundary AP not only focuses on the boundary quality of the objects, but also considers spatio-temporal consistency between the predicted and ground truth object masks. For example, detected object masks with frequent id switches will lead to a low IoU value.

4 Experiments

4.1 Experimental Setup

HQ-YTVIS & YTVIS We conduct experiments on YTVIS [37] and our HQ-YTVIS datasets. YTVIS contains 2,883 videos with 131k annotated object instances belonging to 40 categories. We identify its inaccurate mask boundary issues in Figure 5 and Section 3.2, which influences both model training and accuracy in testing evaluation. For HQ-YTVIS, we split the original YTVIS training set (2238 videos) into a new *train* (1678 videos, 75%), *val* (280 videos 12.5%) and *test* (280 videos 12.5%) sets following the ratios in YTVIS. The masks annotations on the *train* subset of HQ-YTVIS is self-corrected by VMT, while the smaller sets of *val* and *test* are carefully relabeled by human annotators to ensure high mask boundary quality. We employ both the standard tube mask AP^M in [37] and our Tube-Boundary AP^B as evaluation metrics.

OVIS We also report results on OVIS [28], a recently proposed VIS benchmark on occlusion learning. OVIS has better-annotated boundaries for instance masks with 607, 140 and 154 videos for train, valid and test respectively.

BDD100K MOTS We further train and evaluate Video Mask Transfmer on the large-scale BDD100K [39] MOTS, which is a self-driving benchmark with high-quality instance masks. It contains 154 videos (30,817 images) for training, 32 videos (6,475 images) for validation, and 37 videos (7,484 images) for testing.

4.2 Implementation Details

Video Mask Transfmer is implemented on the query-based detector [41], and employ [36] to provide coarse mask predictions for video instances. For the temporal refinement transformer, we adopt 3 multi-head attention layers, setting the hidden dimension to 64 and using 4 attention heads. The instance queries are shared between temporal refinement transformer with the base object detector.

Table 1. Quadtree sequence grouping (QSG) across frames in varying video clip lengths on HQ-YTVIS *val* set.

Length	QSG	AP ^B	AP ₅₀ ^B	AR ₁ ^B	AP ^M	AP ₅₀ ^M	AR ₁ ^M
1		26.1	59.8	23.8	44.5	64.2	40.1
5		30.2	63.3	29.6	47.5	69.8	43.5
5	✓	31.4	64.2	30.7	48.2	70.5	44.1
10		31.2	64.1	30.3	48.9	70.6	44.3
10	✓	32.5	65.3	31.2	49.6	71.2	44.9
All		32.3	66.0	30.6	49.7	71.8	45.5
All	✓	33.7	67.2	31.8	50.5	72.4	46.2

Table 2. Ablation on 3D incoherent region detector, and refinement region types comparison on HQ-YTVIS validation set. IQ: Instance Query.

Region Type	Incoherence Detector	AP ^B	AP ₅₀ ^B	AR ₁ ^B	AP ^M	AP ₅₀ ^M	AR ₁ ^M
Detected Object Boundary	FCN	31.8	65.4	30.5	48.7	71.0	45.0
	IQ (Frame-level)	31.3	64.8	29.9	48.1	69.9	44.3
	IQ (Video-level)	32.8	66.2	31.0	49.8	71.6	45.7
3D Incoherent Region	FCN	32.2	65.1	30.7	49.1	70.9	45.3
	IQ (Frame-level)	31.8	65.2	30.5	48.9	70.6	45.1
	IQ (Video-level)	33.7	67.2	31.8	50.5	72.4	46.2

During training, we follow the setting in [36] and use video clips consisting of 5 frames and sample them from the whole video. We train VMT for 12 epochs and use AdamW [24] as optimizer, with initial learning rate set to 2e-4. Our VMT executes at 8.2 FPS on Swin-L backbone. The learning rate is decayed at the 5th and 11th epochs by factor of 0.1. More details are in the Supp. file.

4.3 Ablation Experiments

We conduct detailed ablation studies for VMT using ResNet-101 as backbone on HQ-YTVIS and OVIS *val* sets. We analyze the impact of each proposed component. Besides, we study the effect of iterative training for self-correcting YTVIS, and compare the same models trained on our HQ-YTVIS vs. YTVIS.

Effect of the Quadtree Sequence Grouping Table 1 analyzes the influence of video clip lengths to the Quadtree Sequence Grouping (QSG). It reveals that the longer video clips with richer temporal amount indeed brings more performance gain to our VMT. When we increase the tube length from 1 to all frames in the video, a remarkable gain in tube boundary AP^B from 26.1 to 33.7 is achieved. This demonstrate that our approach effectively leverages temporal information, since a tube length 1 performs independent prediction for each frame. Moreover, models **w/o** QSG are refining the inherent points in each frame separately as [15]. The multiple boundary view of the same object brings an gain in temporal refinement for over 1.0 AP^B.

Ablation on the 3D Incoherence Detector We study the design choices of our 3D incoherence detector in Table 2. We compare fixed FCN and dynamic FCN (three 3×3 Convs) with weights produced by frame-level or video-level instance queries used in [36]. Video-level instance queries achieve the highest AP^B, improving 1.9 point compared to the frame-level queries, which shows the effect of temporally aggregated video-level instance information. We also compare 3D incoherent regions with detected object mask boundaries, where the 3D incoherent regions achieves 0.9 AP^B gain.

Effect of Iterative Training In Table 3, we compare MaskTrack [37], SeqFormer [36] and VMT for correcting coarse masks of YTVIS in the iterative training. We observe that the improvement scales after each iteration of MaskTrack and SeqFormer on HQ-YTVIS *val* is minor, where the boundary quality AP^B after the 3rd iteration are still coarse (around 60.0 using **GT** object classes,

Table 3. Comparison on iterative training. Models after each correction is evaluated on HQ-YTVIS *val* by taking **GT** classes, ids and coarse masks as input.




Table 4. Training on YTVIS vs. HQ-YTVIS with the same images from scratch. We evaluate the trained models on HQ-YTVIS and OVIS *val* sets.

Method	YTVIS	HQ-YTVIS	HQ-YTVIS			OVIS	
			AP^B	AP_{50}^B	AP^M	AP^M	AP_{50}^M
MaskTrack [37]	✓		19.8	48.9	40.2	9.3	24.2
MaskTrack [37]		✓	21.7	50.5	41.1	10.5	25.1
SeqFormer [36]	✓		28.9	64.2	48.6	13.8	32.1
SeqFormer [36]		✓	31.0	66.1	50.5	15.2	33.7
VMT (Ours)	✓		30.5	64.7	48.9	15.9	33.8
VMT (Ours)		✓	33.7	67.2	50.5	17.1	35.0

identities and corresponding coarse masks). In contrast, VMT achieves consistent and large mask quality improvements after three training iterations, which reveals the design advantages of our temporal refinement transformer.

Training on YTVIS vs. HQ-YTVIS In Table 4, we evaluate the performance of three different approaches when training on either YTVIS or HQ-YTVIS. We train MaskTrack [37], SeqFormer [36] and our VMT from scratch with the same set of images. We use HQ-YTVIS and OVIS for evaluation due to the better annotated mask boundaries. For evaluation on OVIS, we train the mask heads of all these methods in a class-agnostic way, and fix the model weights of the mask head when finetuning them on OVIS for object detection and tracking parts. All three methods trained using HQ-YTVIS obtain consistent and large performance gain of over 2.0 AP^B on the manually labeled HQ-YTVIS *val* set, and over 1.0 AP^M on the OVIS *val* set. This shows our self-corrected HQ-YTVIS dataset consistently improves existing VIS methods for segmentation quality, without overfitting to the specific dataset.

Temporal Attention Visualization In Figure 6, we visualize the temporal attention distribution for incoherent nodes in a video-clip of length 5. The attention weights are extracted from the last NAL of the refinement transformer. For the sampled point R1 at T=3, it attends more to the feet regions of the giraffe with semantic correspondence in both the current and neighboring frames. Also, the attention weights for the temporally farther frames are smaller.

4.4 Comparison with State-of-the-art Methods

We compare VMT with the state-of-the-art methods on the benchmarks HQ-YTVIS, YTVIS, OVIS and BDD100K MOTS. Note that we only conduct iterative training when producing the training annotations of HQ-YTVIS. When retraining VMT and all other baselines on the HQ-YTVIS benchmark, all methods are trained from scratch and only once on the same data for fair comparison.

HQ-YTVIS & YTVIS Table 5 compares VMT with state-of-the-art instance segmentation methods on both HQ-YTVIS and YTVIS benchmarks. VMT achieves consistent performance advantages on different backbones, showing its effectiveness by surpassing SeqFormer [36] by around 2.8 AP_{75}^B on HQ-YTVIS using ResNet-50. As in Figure 5 and Sec 3.2, the mask boundary an-




Fig. 6. Temporal attention visualizations on the sparse incoherent regions for a video clip of length 5. The sampled red node R1 attends more to the feet regions of the giraffe with semantic correspondence in both the current and neighboring frames. The top 10 attended incoherent node regions are marked in yellow.

Table 5. Comparison with state-of-the-art methods on HQ-YTVIS test set and YTVIS [37] validation set. All methods, including VMT, are retrained on HQ-YTVIS and YTVIS training sets respectively from scratch for fair comparisons. Results are reported in terms of Tube-Mask AP^M [37] and our Tube-boundary AP^B . VMT predicts mask at output sizes $16\times$ larger than SeqFormer [36]. The advantage of VMT is not fully revealed on YTVIS due to its inaccurate and coarse boundary annotation.

Method	Backbone	Params	HQ-YTVIS						YTVIS		
			AP^B	AP_{75}^B	AR_1^B	AP^M	AP_{75}^M	AR_1^M	AP^M	AP_{75}^M	AR_1^M
MaskTrack [37]	R50	58.1M	19.8	10.6	21.1	38.8	48.6	40.3	30.3	32.6	31.0
CrossVIS [38]	R50	37.5M	23.6	16.2	24.9	43.0	52.3	44.0	36.3	38.9	35.6
VisTr [35]	R50	57.2M	24.0	16.3	25.1	43.3	52.9	44.5	36.2	36.9	37.2
PCAN [16]	R50	36.9M	23.9	16.1	25.2	42.2	51.8	43.9	36.1	39.4	36.3
IFC [14]	R50	39.3M	26.5	19.6	27.5	46.6	51.5	46.9	42.8	46.8	43.8
SeqFormer [36]	R50	49.3M	28.6	21.4	29.3	48.5	52.2	48.5	47.4	51.8	45.5
VMT (Ours)	R50	51.5M	30.7	24.2	31.5	50.5	54.5	50.2	47.9	52.0	45.8
MaskTrack [37]	R101	77.2M	21.7	13.1	22.8	41.1	49.1	41.7	31.8	33.6	33.2
CrossVIS [38]	R101	56.6M	24.5	19.7	26.5	44.1	52.6	44.5	36.6	39.7	36.0
VisTr [35]	R101	76.3M	25.1	20.5	27.7	45.3	53.2	45.1	40.1	45.0	38.3
PCAN [16]	R101	54.8M	24.8	20.1	27.0	44.0	52.1	44.3	37.6	41.3	37.2
IFC [14]	R101	58.3M	27.2	23.6	28.3	48.2	51.8	47.6	44.6	49.5	44.0
SeqFormer [36]	R101	68.4M	30.7	27.3	30.1	49.0	52.3	46.6	49.0	55.7	46.8
VMT (Ours)	R101	70.5M	32.5	28.9	32.6	51.2	55.1	49.3	49.4	56.4	46.7
SeqFormer [36]	Swin-L	220.0M	43.3	41.5	41.6	63.7	69.7	58.7	59.3	66.4	51.7
VMT (Ours)	Swin-L	222.3M	44.8	43.4	43.0	64.8	70.1	59.3	59.7	66.7	52.0

notation in YTVIS is less accurate. Therefore, the advantages brought by our approach are not fully revealed on this dataset. Yet, VMT exceeds SeqFormer by about 0.5 AP^M on YTVIS with ResNet-50 with higher mask quality as in Figure 7. Moreover, masks predicted by our approach are $16\times$ larger than those of SeqFormer, while only increasing negligible amount of the model parameters.

OVIS The results of OVIS dataset are reported in Table 6, where VMT achieves the best mask AP 19.8 using Swin-L backbone, improving 1.9 point compared to the baseline SeqFormer [36].




Fig. 7. Seqformer (1st row) vs. ours (2nd row) on YTVIS, in terms of mask quality & temporal consistency. Please refer to the Supp. file for more video results comparisons.

Table 6. Comparison with state-of-the-art on the OVIS validation set.

Method	Backbone	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
MaskTrack [37]	R50	10.8	25.3	8.5	7.9	14.9
SipMask [5]	R50	10.2	24.7	7.8	7.9	15.8
CrossVIS [38]	R50	14.9	32.7	12.1	10.3	19.8
STMask [20]	R50	15.4	33.8	12.5	8.9	21.3
CMTrack RCNN [28]	R50	15.4	33.9	13.1	9.3	20.0
SeqFormer [36]	R50	15.6	34.3	12.1	9.6	21.8
VMT (Ours)	R50	16.9	36.4	13.7	10.4	22.7
SeqFormer [36]	R101	16.2	35.1	13.1	10.3	22.9
VMT (Ours)	R101	17.8	35.7	15.4	10.5	23.8
SeqFormer [36]	Swin-L	17.9	35.6	15.6	10.7	24.1
VMT (Ours)	Swin-L	19.8	39.6	17.2	11.2	26.3

Table 7. State-of-the-art comparison on the BDD100K segmentation tracking validation set using ResNet-50. I: ImageNet. C: COCO. S: Cityscapes. B: BDD100K.

Method	Pretrained	mMOTSA \uparrow	mMOTSP \uparrow	mIDF \uparrow	ID sw. \downarrow	mAP \uparrow
SortIoU	I, C, S	10.3	59.9	21.8	15951	22.2
MaskTrack [32]	I, C, S	12.3	59.9	26.2	9116	22.0
STEm-Seg [2]	I, C, S	12.2	58.2	25.4	8732	21.8
QDTrack [27]	I, C, S	22.5	59.6	40.8	1340	22.4
QDTrack-fix [27]	I, B	23.5	66.3	44.5	973	25.5
PCAN [16]	I, B	27.4	66.7	45.1	876	26.6
VMT (Ours)	I, B	28.7	67.3	45.7	825	28.3

BDD100K MOTS Table 7 shows results on BDD100K MOTS, where Mask Transfinner obtains the highest mMOTSA of 28.7 and outperforms the PCAN [16] by 1.3 points by sharing the same object detection tracking heads. The large gain reveals the high quality of temporal masks prediction by VMT.

5 Conclusion

We present Video Mask Transfinner, the first high-quality video instance segmentation method. Enabled by the efficient video transformer design, VMT utilizes the high-resolution spatio-temporal features for temporal mask refinement and achieves large boundary and mask AP gains on the HQ-YTVIS, OVIS, and BDD100K. To refine the coarse annotation of YTVIS, we design an iterative training paradigm and adopt VMT to correct the annotations errors of the training data instead of tedious manual relabeling. We build the new HQ-YTVIS benchmark with more accurate mask boundary annotations than YTVIS, and introduce Tube Boundary AP for accurate performance measure. We believe our method, the new benchmark HQ-YTVIS and evaluation metric will facilitate future video instance segmentation works on improving their mask quality and benefit real-world applications such as video editing [1, 18].

¹ This work is supported in part by the Research Grant Council of the Hong Kong SAR under grant no. 16201420 and Kuaishou Technology.

References

1. Alldieck, T., Magnor, M., Xu, W., Theobalt, C., Pons-Moll, G.: Video based reconstruction of 3d people models. In: CVPR (2018) [14](#)
2. Athar, A., Mahadevan, S., Ošep, A., Leal-Taixé, L., Leibe, B.: Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In: ECCV (2020) [5](#), [14](#)
3. Bertasius, G., Torresani, L.: Classifying, segmenting, and tracking object instances in video with mask propagation. In: CVPR (2020) [4](#)
4. Bolya, D., Zhou, C., Xiao, F., Lee, Y.J.: Yolact: Real-time instance segmentation. In: ICCV (2019) [4](#)
5. Cao, J., Anwer, R.M., Cholakkal, H., Khan, F.S., Pang, Y., Shao, L.: Sipmask: Spatial information preservation for fast image and video instance segmentation. In: ECCV (2020) [4](#), [14](#)
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020) [4](#)
7. Chen, L.C., Lopes, R.G., Cheng, B., Collins, M.D., Cubuk, E.D., Zoph, B., Adam, H., Shlens, J.: Naive-student: Leveraging semi-supervised learning in video sequences for urban scene segmentation. In: ECCV (2020) [5](#)
8. Cheng, B., Girshick, R., Dollár, P., Berg, A.C., Kirillov, A.: Boundary iou: Improving object-centric image segmentation evaluation. In: CVPR (2021) [9](#)
9. Cheng, H.K., Chung, J., Tai, Y.W., Tang, C.K.: Cascadepsp: Toward class-agnostic and very high-resolution segmentation via global and local refinement. In: CVPR (2020) [5](#)
10. Ding, H., Jiang, X., Liu, A.Q., Thalmann, N.M., Wang, G.: Boundary-aware feature propagation for scene segmentation. In: ICCV (2019) [5](#)
11. Fang, Y., Yang, S., Wang, X., Li, Y., Fang, C., Shan, Y., Feng, B., Liu, W.: Instances as queries. In: ICCV (2021) [1](#), [3](#), [9](#)
12. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: ICCV (2017) [4](#)
13. Huang, Z., Huang, L., Gong, Y., Huang, C., Wang, X.: Mask scoring r-cnn. In: CVPR (2019) [4](#)
14. Hwang, S., Heo, M., Oh, S.W., Kim, S.J.: Video instance segmentation using inter-frame communication transformers. In: NeurIPS (2021) [1](#), [2](#), [3](#), [4](#), [9](#), [13](#)
15. Ke, L., Danelljan, M., Li, X., Tai, Y.W., Tang, C.K., Yu, F.: Mask transfinner for high-quality instance segmentation. In: CVPR (2022) [2](#), [5](#), [7](#), [11](#)
16. 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) [3](#), [4](#), [5](#), [9](#), [13](#), [14](#)
17. Ke, L., Tai, Y.W., Tang, C.K.: Deep occlusion-aware instance segmentation with overlapping bilayers. In: CVPR (2021) [4](#)
18. Ke, L., Tai, Y.W., Tang, C.K.: Occlusion-aware video object inpainting. In: ICCV (2021) [14](#)
19. Kirillov, A., Wu, Y., He, K., Girshick, R.: Pointrend: Image segmentation as rendering. In: CVPR (2020) [2](#), [5](#), [9](#)
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) [4](#), [14](#)
21. Lin, C.C., Hung, Y., Feris, R., He, L.: Video instance segmentation tracking with a modified vae architecture. In: CVPR (2020) [4](#)
22. Lin, H., Wu, R., Liu, S., Lu, J., Jia, J.: Video instance segmentation with a propose-reduce paradigm. In: ICCV (2021) [4](#)

23. Liu, D., Cui, Y., Tan, W., Chen, Y.: Sg-net: Spatial granularity network for one-stage video instance segmentation. In: CVPR (2021) [4](#)
24. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: ICLR (2019) [11](#)
25. Meinhardt, T., Kirillov, A., Leal-Taixe, L., Feichtenhofer, C.: Trackformer: Multi-object tracking with transformers. arXiv preprint arXiv:2101.02702 (2021) [5](#)
26. Milan, A., Leal-Taixé, L., Schindler, K., Reid, I.: Joint tracking and segmentation of multiple targets. In: CVPR (2015) [5](#)
27. 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) [14](#)
28. Qi, J., Gao, Y., Hu, Y., Wang, X., Liu, X., Bai, X., Belongie, S., Yuille, A., Torr, P., Bai, S.: Occluded video instance segmentation. arXiv preprint arXiv:2102.01558 (2021) [3](#), [4](#), [8](#), [10](#), [14](#)
29. Takikawa, T., Acuna, D., Jampani, V., Fidler, S.: Gated-scnn: Gated shape cnns for semantic segmentation. In: ICCV (2019) [5](#)
30. Tang, C., Chen, H., Li, X., Li, J., Zhang, Z., Hu, X.: Look closer to segment better: Boundary patch refinement for instance segmentation. In: CVPR (2021) [5](#)
31. Tian, Z., Shen, C., Chen, H., He, T.: Fcos: Fully convolutional one-stage object detection. In: ICCV (2019) [4](#)
32. Voigtlaender, P., Chai, Y., Schroff, F., Adam, H., Leibe, B., Chen, L.C.: Feelvos: Fast end-to-end embedding learning for video object segmentation. In: CVPR (2019) [14](#)
33. Voigtlaender, P., Krause, M., Osep, A., Luiten, J., Sekar, B.B.G., Geiger, A., Leibe, B.: Mots: Multi-object tracking and segmentation. In: CVPR (2019) [5](#)
34. Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., et al.: Deep high-resolution representation learning for visual recognition. TPAMI (2020) [5](#)
35. 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) [1](#), [2](#), [3](#), [4](#), [9](#), [13](#)
36. Wu, J., Jiang, Y., Zhang, W., Bai, X., Bai, S.: Seqformer: a frustratingly simple model for video instance segmentation. arXiv preprint arXiv:2112.08275 (2021) [1](#), [2](#), [4](#), [6](#), [10](#), [11](#), [12](#), [13](#), [14](#)
37. Yang, L., Fan, Y., Xu, N.: Video instance segmentation. In: ICCV (2019) [2](#), [3](#), [4](#), [9](#), [10](#), [11](#), [12](#), [13](#), [14](#)
38. 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) [3](#), [4](#), [9](#), [13](#), [14](#)
39. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: CVPR (2020) [4](#), [10](#)
40. Yuan, Y., Xie, J., Chen, X., Wang, J.: Segfix: Model-agnostic boundary refinement for segmentation. In: ECCV. Springer (2020) [5](#)
41. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020) [6](#), [10](#)
42. Zhu, Y., Zhang, Z., Wu, C., Zhang, Z., He, T., Zhang, H., Manmatha, R., Li, M., Smola, A.: Improving semantic segmentation via self-training. arXiv preprint arXiv:2004.14960 (2020) [5](#)
43. Zou, Y., Zhang, Z., Zhang, H., Li, C.L., Bian, X., Huang, J.B., Pfister, T.: Pseudoseg: Designing pseudo labels for semantic segmentation. In: International Conference on Learning Representations (ICLR) (2021) [5](#)