

The following publication M. Li, S. Li, L. Li and L. Zhang, "Spatial Feature Calibration and Temporal Fusion for Effective One-stage Video Instance Segmentation," 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), Nashville, TN, USA, 2021, pp. 11210-11219 is available at <https://doi.org/10.1109/CVPR46437.2021.01106>.

Spatial Feature Calibration and Temporal Fusion for Effective One-stage Video Instance Segmentation

Minghan Li^{1,2}, Shuai Li^{1,2}, Lida Li¹ and Lei Zhang^{1,2*}

¹The HongKong Polytechnic University, ²DAMO Academy, Alibaba Group
liminghan0330@gmail.com, {csshuaili, cslli, cslzhang}@comp.polyu.edu.hk

Abstract

Modern one-stage video instance segmentation networks suffer from two limitations. First, convolutional features are neither aligned with anchor boxes nor with ground-truth bounding boxes, reducing the mask sensitivity to spatial location. Second, a video is directly divided into individual frames for frame-level instance segmentation, ignoring the temporal correlation between adjacent frames. To address these issues, we propose a simple yet effective one-stage video instance segmentation framework by spatial calibration and temporal fusion, namely STMASK. To ensure spatial feature calibration with ground-truth bounding boxes, we first predict regressed bounding boxes around ground-truth bounding boxes, and extract features from them for frame-level instance segmentation. To further explore temporal correlation among video frames, we aggregate a temporal fusion module to infer instance masks from each frame to its adjacent frames, which helps our framework to handle challenging videos such as motion blur, partial occlusion and unusual object-to-camera poses. Experiments on the YouTube-VIS valid set show that the proposed STMASK with ResNet-50/101 backbone obtains 33.5 % / 36.8 % mask AP, while achieving 28.6 / 23.4 FPS on video instance segmentation. The code is released online <https://github.com/MinghanLi/STMASK>.

1. Introduction

Video instance segmentation aims to obtain the pixel-level segmentation mask for individual instances of all classes over the entire frames of a video, which heavily depends on spatial position-sensitive features to localize frame-level objects and redundant temporal information to track instances across frames. Following object detection and image instance segmentation works, modern video instance segmentation approaches usually adopt the top-down

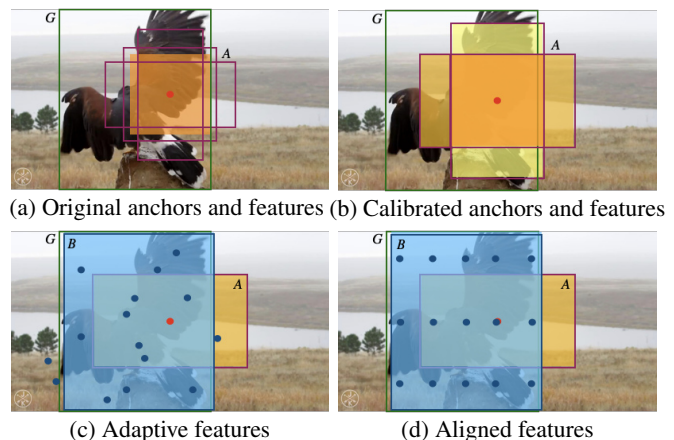


Figure 1. Spatial calibration for anchor boxes and bounding boxes. (a) and (b) display anchors and features in original and calibrated one-stage networks respectively. (c) and (d) demonstrate adaptive and aligned features extracted from predicted bounding boxes. Purple, blue and green rectangles denote anchors, predicted and ground-truth bounding boxes respectively, where coloured areas indicate the receptive field of their convolutional features.

framework of first detecting and segmenting objects frame by frame and then linking instance masks across frames.

Top-down video instance segmentation approaches can be divided into two-stage and one-stage methods. By adding a tracking branch to Mask R-CNN [18], two-stage video instance segmentation methods [43, 2, 16] first predict region-of-interests (RoIs) around ground-truth bounding boxes, and then feed aligned features via RoIPooling [32] or RoIAlign [18] to segment frame-level object masks and to track cross-frame instances. To obtain better location-sensitive features for mask predictor, many spatial feature calibration strategies for RoIs have been continuously proposed in recent years like Deformable RoI [12] and Hybrid Task Cascade [8]. For temporal information exploration, recently proposed MaskProp [2] utilizes temporal features propagated from all frames of a video clip for clip-level instance tracking. Obviously, two-stage methods have recognized the importance of spatial feature localisation and temporal feature tracking for video instance

*Corresponding author. This work is supported by the Hong Kong RGC RIF grant (R5001-18).

segmentation. One-stage instance segmentation networks [11, 42, 5, 35, 6, 24, 25], which focus more on real-time speed, usually employ a fully convolutional network structure to directly predict the final mask for instances. Without RoIs of two-stage methods for localisation, early one-stage methods have to introduce extra position-sensitive information to improve segmentation performance like position-sensitive score maps [11] or semantic features [9]. In image domain, recently proposed anchor-based one-stage methods like Yolact [5] and CondInst [35] decompose instance segmentation as the linear combination between instance-specific mask coefficients and instance-independent prototypes. Furthermore, SipMask [6], introducing a tracking branch in Yolact[5], achieves real-time speed but inferior performance in video instance segmentation task.

Analysing these anchor-based one-stage instance segmentation methods, we observe that, as shown in Fig.1 (a), multiple anchors of different shapes at each spatial position (purple rectangles) share same convolutional features (yellow area), which are neither aligned with pre-defined anchor boxes nor with ground-truth bounding boxes. This fact does violate that instances segmentation is a spatial location-sensitive task. On the other hand, one-stage video instance segmentation methods directly divide a video into separate frames to perform image instance segmentation frame by frame and then track them across frames, which completely ignores high temporal correlation between adjacent frames. This may fail to handle those challenging videos with motion blur, partial occlusion, or unconventional object-to-camera poses. In other words, modern one-stage video instance segmentation methods achieve real-time speed at the cost of discarding spatial feature calibration and temporal feature correlation.

To address the issues, we propose a simple yet effective one-stage video instance segmentation framework, named STMask. Firstly, we design a feature calibration strategy for anchor boxes and ground-truth bounding boxes to obtain more precise spatial features. Specifically, as shown in Fig.1 (b), to enable each anchor box can extract its own specific features, we first design multiple convolutional kernels at each spatial position, and then generate anchors according to the receptive field of these convolutional kernels. To improve feature presentation for objects segmentation and tracking, we first predict regressed bounding boxes around ground-truth bounding boxes by regression branch, and then extract features from them to segment and track instances. As shown in Fig.1 (c) and (d), we provide two strategies to extract features from regressed bounding boxes, including adaptive features by a single 1×1 convolutional layer and aligned features by mathematical derivation. Finally, we explore temporal correlation between video frames by adding a temporal fusion module to infer instance masks from adjacent frames, thereby improv-

ing the performance of objects detection, segmentation and tracking for those challenging videos.

2. Related Work

Video instance segmentation directly benefits from advances in image instance segmentation and video object detection field. Thus this section consists of three parts.

Image Instance Segmentation. Existing image instance segmentation methods either follow bottom-up or top-down paradigms. The bottom-up methods [9, 36] widely employ multiple stages to first perform semantic segmentation and then identify the specific location of each instance by boundary detection [20], pixel clustering [25], or pixels embedding loss [30, 22, 31], position-sensitive pooling [11, 24]. The top-down instance segmentation approaches first predict bounding boxes by object detectors and then perform mask segmentation within the predicted boxes. Mask R-CNN [18] extends Faster R-CNN [32] by adding a mask segmentation branch on each Region of Interest (RoI), where RoIAlign operator is introduced to obtain ‘re-pool’ features of each proposal for better mask prediction. Follow-up works try to improve its accuracy by aligning spatial features [12], enriching the FPN features [28] or addressing the incompatibility between a mask’s confidence score and its localization accuracy [19]. One-stage instance segmentation approaches [42, 34, 7, 45, 38, 5, 40, 21] recently are proposed to keep the trade-off between speed and performance. Yolact-based methods [5, 35, 6] break up instance segmentation into two parallel subtasks: generating a set of instance-independent prototypes and predicting instance-specific mask coefficients.

Video Object Detection. For handling challenging videos such as motion blur and occlusion, optical flow [46, 39], correlation operation [15], deformable convolutional networks [3, 41] and relation networks [14, 10] are the popular technologies to propagate or align features across frames. In addition, methods like [13, 33] try to utilise semantic similarity between frames to assist object detection in videos.

Video Instance Segmentation. To jointly perform the detection, segmentation and tracking tasks simultaneously, most of video instance segmentation methods extend Mask R-CNN [18] by adding a new branch for tracking. For example, MaskTrack R-CNN [43] predicts an extra embedding vector for each instance and use an external memory to store them for tracking across frames. MaskProp [2] introduces a mask propagation branch to propagate frame-level instance masks from each video frame to all the other frames in a video clip and then match clip-level instance masks for tracking, providing state-of-the-art instance segmentation performance and quite limited speed. Besides, [26] proposes a modified variation autoencoder (VAE) architecture built on the top of Mask R-CNN [18]. Recently,




Figure 2. An overview architecture of STMask, including frame-level instance segmentation (blue area) and cross-frame instance segmentation (yellow area). Specifically, we first use Yolact with spatial calibration to obtain frame-level instance masks in time $t - 1$ and t respectively, then feed FPN features of two adjacent frames and the output of correlation operation to TemporalNet to infer the displacement of instances from time $t - 1$ to t , and finally merge frame-level and cross-frame instance masks to obtain final instance masks.

SipMask [6] also introduces a tracking branch (same as [43]) in the one-stage image instance segmentation network Yolact [5] to obtain inferior performance but real-time speed. Besides, inspired by Guided Anchor [38, 23], SipMask further aligns feature with regressed bounding boxes to improve feature representation for classification and mask coefficients generation. Existing spatial feature calibration for one-stage approaches only focus on the misalignment between features and regressed bounding boxes.

3. STMask

The overall architecture of STMask shown in the Fig. 2 consists of frame-level instance segmentation with spatial calibration and cross-frame instance segmentation by temporal fusion module.

3.1. Spatial Calibration

Our goal is to align features with anchors and ground-truth bounding boxes respectively for one-stage instance segmentation methods. For spatial feature calibration on anchors, multiple convolutional kernels shown in Fig. 3 are introduced to alleviate the misalignment between features and anchors. For spatial feature calibration on bounding boxes, shown in the pink rectangle of Fig. 2, we first individually predict regressed bounding boxes around ground-truth bounding boxes, and then extract features from those regressed bounding boxes to classify, segment and track instances.

3.1.1 Feature Calibration for Anchors (FCA)

One-stage anchor-based object detectors usually sample a large number of regions in the input image, determine whether these regions contain objects of interest, and adjust the edges of the regions so as to more accurately predict the ground-truth bounding boxes of the targets. In general, the sliding-window fashion to generate anchors is the most popular method, which generates multiple boxes with different scales and aspect ratios while centring on each pixel. As shown in Fig. 1(a), for multiple anchors with different shapes, one-stage detectors directly employ a 3×3 convolution on the central point of anchors to extract features. In practice, the receptive field of convolution should be positively related to the size of anchor boxes. For example, those larger anchor boxes should have larger receptive fields, while those smaller boxes should have smaller ones. To address this issue, thus, we adopt multiple convolution kernels with different aspect ratios on each FPN layer. For example, we replace the single 3×3 convolution with three new aspect ratios of convolutions, 3×3 , 3×5 , 5×3 respectively. The FCA architecture implemented on the bounding box regression branch is shown in Fig. 3. To further ensure the calibration between convolutional features and anchors, we keep the scales unchanged while changing the anchor aspect ratios from $[1, 1/2, 2]$ to $[1, 3/5, 5/3]$, which are same as the aspect ratios of convolutions. The simple feature calibration for anchors maintains position sensitivity for object detection and segmentation.




Figure 3. Bounding Box Regression Architecture. Taking the regression branch as example, we design multiple convolutions for calibrating features with pre-defined anchor boxes, where $a=3$ is the number of anchors.

3.1.2 Feature Calibration for Bounding Boxes (FCB)

After the feature calibration with anchors in last part, features for classification, mask coefficients and embedding vectors of tracking branch still fail to align with the regressed bounding boxes. To further address the issue, we adopt the prediction architecture of first predicting regressed bounding boxes and then extracting features from the regressed bounding boxes to segment and track objects.

Let $\{(P_i, G_i)\}_{i=1, \dots, N}$ denote the pairs of anchor boxes and ground-truth bounding boxes. Note that subscript i is only added when necessary. $P = (P_x, P_y, P_w, P_h)$ specifies the pixel coordinates of the centre and the width and height of the anchor box. The same goes for ground-truth bounding box $G = (G_x, G_y, G_w, G_h)$. The bounding-box regression aims to learn a transformation d that maps an anchor box to its ground-truth bounding box:

$$d = [d_x, d_y, d_w, d_h]. \quad (1)$$

After that, its regressed bounding box B_i can be calculated by applying the transformation [17]

$$\begin{aligned} B_x &= P_w d_x + P_x, & B_y &= P_h d_y + P_y; \\ B_w &= P_w \exp(d_w), & B_h &= P_h \exp(d_h). \end{aligned} \quad (2)$$

Based on the transformation between anchor boxes and predicted bounding boxes, we can introduce a 2D deformable convolution to calibrate convolutional features from anchor boxes to predicted bounding boxes. For easier understanding, we take a 3×3 convolution with dilation 1 as an example to explicate the process of spatial feature calibration for bounding boxes. For each location p_0 on the output map \mathbf{g} , the 2D deformable convolution can be formulated as

$$\mathbf{g}(p_0) = \sum_{p_n \in \mathcal{R}} w(p_n) \cdot \mathbf{f}(p_0 + p_n + \Delta p_n), \quad (3)$$

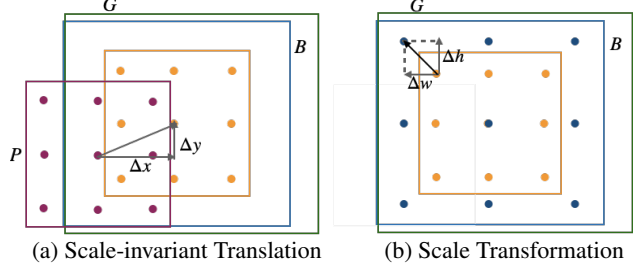


Figure 4. Illustration of aligned features from anchor boxes to predicted bounding boxes.

where Δp_n is the offsets on the position $p_n \in \mathcal{R}$, which is a point of the grid

$$\mathcal{R} = \left\{ \begin{array}{ccc} (-1, -1) & (-1, 0) & (-1, 1) \\ (0, -1) & (0, 0) & (0, 1) \\ (1, -1) & (1, 0) & (1, 1) \end{array} \right\}, \quad (4)$$

To extract convolutional features from predicted bounding boxes, the offsets \mathcal{O} should be dominated by the transformation d . Thus, this paper provides two strategies to obtain offsets: adding a single 1×1 convolution layer to predict adaptive offsets, or directly deriving the aligned offsets through mathematical geometric knowledge.

Adaptive Features on Bounding Boxes. Inspired by the anchor-guided feature adaption module [38], we also transform the feature at each individual location based on the underlying anchor transformation:

$$\mathcal{O} = \mathcal{N}_O(d), \quad (5)$$

where \mathcal{N}_O is a 1×1 convolutional layer to predict offsets according to the bounding-box regression transformation d . When the predicted offsets $\mathcal{O} \in \mathbb{R}^{2 \times 3 \times 3}$ is fed into a deformable convolutional layer to produce adaptive features shown in blue points in Fig. 1(c). For such regression-dependent offsets, each regressed bounding box can learn its own adaptive features to perform further objects classification, segmentation and tracking.

Aligned Features on Bounding Boxes. For the regression transformation d in terms of four functions in Eq. 1, the first two specify a scale-invariant translation of the centre of anchor box, while the second two specify log-space translations of the width and height of anchor box. According to the four parameters of transformation d , the derivation process of generating the offsets can also be divided into two steps: *scale-invariant translation* and *scale transformation*, shown in Fig. 4 (a) and (b) respectively. On one hand, Fig. 4 (a) demonstrates that all sampling points on the grid \mathcal{R} of the anchor box have the same scale-invariant translation as the centre point. On the other hand, Fig. 4 (b) shows that the absolute scale transformation of the width and height on the grid \mathcal{R} is related to their own coordinate position. Overall, the derived offsets for all points on the grid should be

$$\mathcal{O} = (\Delta y, \Delta x)I + (\Delta h, \Delta w)\mathcal{R}. \quad (6)$$

where I is a matrix with all elements as 1, $\Delta x, \Delta y, \Delta h$ and Δw are listed below:

$$\Delta x = k_w d_x, \Delta w = \exp(d_w) - 1, \quad (7)$$

$$\Delta y = k_h d_y, \Delta h = \exp(d_h) - 1. \quad (8)$$

where $k = (k_w, k_h)$ is the kernel size on the width and the height. Due to the limit of space, the detailed formula derivation process is provided in the **supplementary materials**. In fact, the mathematical derivation for the offset in this part is equivalent to a special case of RoIAlign operation in two-stage networks, where each bin only takes the central point as its output value.

3.2. Temporal Fusion Module

Compared with image instance detection and segmentation, video instance segmentation task faces more challenges such as partial occlusion, unusual view, motion blur. To address the issue, thus, we further set up a temporal fusion module for cross-frame bounding boxes regression and mask segmentation. The yellow area of Fig. 2 shows the schematic diagram of temporal fusion module between two adjacent frames, which concatenate features of FPN between two adjacent frames to infer the displacement of instances from the previous frame to the current frame, thereby obtaining cross-frame instance masks from the previous frame to the current frame. Afterthat, we only need to merge frame-level detected instance masks and cross-frame tracked instance masks to get the final instances masks. Such a double guarantee architecture does improve the accuracy of detection and segmentation in video domain.

Frame-level Instance Segmentation. Same as Yolact [5], STMask decomposes instance segmentation into the linear combination between instance-independent prototypes $P \in R^{H' \times W' \times k}$ and instance-specific mask coefficients $C \in R^{k \times n}$. Essentially, the process of learning prototypes P is equivalent to learn the online basis set, also called dictionary [29], when each instance i can find a linear combination $C_i \in R^{k \times 1}$ of a 'few' atoms from prototypes that is 'close' to the ground-truth mask $M_i \in R^{H' \times W'}$. The each instance segmentation can be implemented efficiently using a single matrix multiplication and the activation function (sigmoid function σ):

$$M_i = \text{Crop}(\sigma(PC_i), B_i). \quad (9)$$

where the final masks M_i is also cropped by the predicted bounding box B_i .

Cross-frame Instance Segmentation. Given a pair of frames I^{t-1}, I^t from a video $\mathcal{V} \in R^{H \times W \times V}$, the goal of proposed temporal fusion module is to track those object masks from time $t-1$ to t , noted as $M^{t-1,t}$. Since the prototypes are instance-independent, according to Eq. (9), the temporal fusion module intuitively needs to predict the




Figure 5. Temporal Network Architecture.

mask coefficients $C^{t-1,t}$ and bounding boxes $B^{t-1,t}$ of instances from time $t-1$ to t .

Inspired by RoI tracking process in video object detection task [15], the temporal fusion module of STMask should include a bounding boxes regressor and a mask coefficient predictor. To be specific, as shown in the yellow area of Fig. 2, we first adopt a correlation operation on FPN features of $x^{t-1} \in R^{H' \times W' \times f}$ and $x^t \in R^{H' \times W' \times f}$ to embed motion information of instances, noted as $x_{corr}^{t-1,t} \in R^{H' \times W' \times d^2}$, where d is the side length of local square. Then the concatenation of features $\{x^{t-1}, x^t, x_{corr}^{t-1,t}\}$ is transported to the temporal network, shown in Fig. 5, to infer the displacement of bounding boxes $d^{t-1,t}$ and mask coefficients $\Delta C^{t-1,t}$ of instances from time $t-1$ to t , where

$$d^{t-1,t} = \{d_x^{t-1,t}, d_y^{t-1,t}, d_w^{t-1,t}, d_h^{t-1,t}\}. \quad (10)$$

Thus, the cross-frame bounding boxes $B^{t-1,t}$ from time $t-1$ to t can be formulated as

$$B_x^{t-1,t} = B_w^{t-1} d_x^{t-1,t} + B_x^{t-1}, B_y^{t-1,t} = B_h^{t-1} d_y^{t-1,t} + B_y^{t-1}, \quad (11)$$

$$B_w^{t-1,t} = B_w^{t-1} \exp(d_w^{t-1,t}), B_h^{t-1,t} = B_h^{t-1} \exp(d_h^{t-1,t}). \quad (12)$$

where the predicted bounding box coordinates B^{t-1} acts as an anchor. On the other hand, due to the similarity of the prototypes between adjacent frames, the mask coefficients of each instance on different frames should also be 'close' to each other. Based on the mask coefficients C^{t-1} in the $t-1$ frame, the cross-frame mask coefficients $C^{t-1,t}$ inferred from time $t-1$ to t can be obtained by

$$C^{t-1,t} = C^{t-1} + \Delta C^{t-1,t}. \quad (13)$$

Finally, the instance masks $M^{t-1,t}$ inferred from time $t-1$ to t can be conducted by

$$M^{t-1,t} = \text{Crop}(\sigma(P^t C^{t-1,t}), B^{t-1,t}). \quad (14)$$

Merge Frame-level and Cross-frame Instance Segmentation. We denote the set of video-level instance IDs as \mathcal{Y} , which is incrementally built, and N_t is the number of instances predicted by frame-level instance segmentation of Eq. (9) at time t . At the beginning, we assign IDs

$\mathcal{Y} = \{1, \dots, N_1\}$ for all frame-level object segmentations in the first frame of a video. For following video frames, we first perform frame-level instance segmentations of Eq. (9) to predict frame-level instance masks M^t , and then execute the temporal fusion module of Eq. (14) to infer cross-frame instance masks $M^{t-1,t}$ from time $t-1$ to t . Obviously, the cross-frame instance masks $M^{t-1,t}$ are naturally given the same embedding vectors E^{t-1} and the same instance IDs as M^{t-1} . Thus, the matching scores between frame-level masks M^t and cross-frame masks $M^{t-1,t}$ consist of two components: the cosine similarity of embedding vectors and mask IoU:

$$s_{ij}^t = \alpha \frac{E_i^t \cdot E_j^{t-1}}{\|E_i^t\| \|E_j^{t-1}\|} + \beta \text{MIoU}(M_i^t, M_j^{t-1,t}) \quad (15)$$

where $i \in \{1, \dots, N_t\}, j \in \mathcal{Y}$, and α, β are hyperparameters to balance the effect of different components. For each frame-level instance mask M_i^t , let $s_i^t = \max_{j \in \mathcal{Y}} s_{ij}^t$ be the maximum score among all instances IDs $j \in \mathcal{Y}$. If s_i^t is greater than a certain threshold ϵ , the frame-level instance mask M_i^t will be assigned the instance ID with the highest score. Otherwise, it will be considered as a new one and will be assigned an instance ID $|\mathcal{Y}| + 1$. Overall, the frame-level instance masks M_i^t will be assigned instance IDs by:

$$y_i^t = \begin{cases} \arg \max_{j \in \mathcal{Y}} s_{ij}^t, & \text{if } s_i^t > \epsilon, \\ |\mathcal{Y}| + 1, & \text{otherwise.} \end{cases} \quad (16)$$

We denote $\mathcal{Y}^t = \{y_i^t | i \in \{1, \dots, N_t\}\}$ as the set including all frame-level instance IDs at time t . If there are instance IDs that do not appear in the set of frame-level instance IDs \mathcal{Y}^t but appear in the set of cross-frame instance IDs \mathcal{Y} , in other words,

$$\mathcal{Y} - \mathcal{Y} \cap \mathcal{Y}^t \neq \emptyset,$$

these cross-frame instance masks $M_{\mathcal{Y} - \mathcal{Y} \cap \mathcal{Y}^t}^{t-1,t}$ will be supplemented as the missing instances in the t frame. Finally, the merged masks of all instances in the t -th frame should be

$$\bar{M}^t = M^t \cup M_{\mathcal{Y} - \mathcal{Y} \cap \mathcal{Y}^t}^{t-1,t}. \quad (17)$$

4. Experimental Results

Training. We conduct video instance segmentation experiments on YouTube-VIS [43] dataset using the standard metrics. In training, the YouTube-VIS train set is split as two subsets: train-sub and valid-sub set. One is used for training and the other is used as the valid set during training. We train all models on pairs of frames, where the second frame in a pair is randomly selected with a time gap $\delta \in [-5, 5]$ relative to the first frame, using pre-trained models on MS COCO dataset [27]. Similar as MaskTrack R-CNN[43], we use the input size 360×640 for training. All model is trained with batch size 16 for 160k iterations and divide the

learning rate at 66K and 133K. Training takes 1-2 days on four NVIDIA 2080Ti.

Inference. Given an input image, we forward it through the frame-level network to obtain the outputs including classification scores, bounding box, embedding vectors, mask coefficients and the prototypes. Box-based Non-maximum suppression (NMS) with the threshold being 0.5 is used to remove duplicated detections and then the top 100 bounding-boxes are used to compute frame-level masks. After that, we forward the features of the previous frame and the current frames through temporal fusion module to first infer masks from the previous frame to current frame, then assign IDs for frame-level masks of the current frame and supplement the missing masks. We take the hyperparameters as $\alpha = 1, \beta = 1$ in temporal fusion module.

4.1. Design Choices

In this section, we discuss the design choices of spatial calibration and temporal fusion module for our proposed one-stage video instance segmentation framework STMASK. Similar as MaskTrack R-CNN [43], we extend Yolact [5] by adding a tracking branch to predict embedding vectors for tracking. For this part, we adopt ResNet101 with deformable convolution layers (interval=3) as the backbone of Yolact [5]. Note that all experiments do not include other improvements in Yolact++ [4] such as more anchors and semantic segmentation loss.

Design Choices of FCA. Our baseline Yolact [5] has three anchors with aspect ratios [1, 1/2, 2] and a 3×3 convolution in prediction head. As shown in Table 1 (1st row), it achieves 28.9% in mask AP. After adding FCA strategy, the setting of aspect ratios with smaller convolutions $\{3 \times 3, 3 \times 2, 2 \times 3\}$ still brings the gain of 1% mask AP, and the setting with larger convolutions of aspect ratios $\{3 \times 3, 3 \times 5, 5 \times 3\}$ brings a significant improvement of 1.9% mask AP. Note that, for fair comparison, we set the kernel size to nearly 3×3 to ensure that the performance gain does come from feature calibration on anchors rather than a larger receptive field. For all experiments below, we choose the latter setting by default.

Design choices of FCB. To analyse the impact of the feature calibration for regression bounding boxes (FCB) on classification, mask coefficients and embedding vectors of tracking branch, we discuss some possible combinations among them on the prediction head. As shown in Table 2, our baseline with FCA (1st row) achieves 30.8% mask AP. Overall, all experiments adding FCB with adaptive and aligned features strategy on all possible combination among branches of the prediction head obtain significantly improved results on all metrics, around 1-3%. Besides, all experiments with aligned features calibration is relatively lower than that of adaptive features, because the latter may be more adaptable than the former. By default, FCB (ada)

Table 1. Design choices for FCA on YouTube-VIS valid set.

Aspect ratios	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
{3:3, 3:3, 3:3}	28.9	45.8	30.1	9.8	26.9	40.3
{3:3, 3:2, 2:3}	29.9	48.9	31.2	8.0	27.8	40.3
{3:3, 3:5, 5:3}	30.8	49.8	31.0	11.8	29.7	44.6

Table 2. Design choices for FCB with adaptive and aligned features on YouTube-VIS valid set.

Baseline+FCA	+Class	+Mask	+Track	AP(ada)	AP(ali)
✓				30.8	30.8
✓	✓	✓	✓	33.7	31.7
✓	✓	✓		34.4	31.7
✓	✓		✓	34.3	32.4
✓		✓	✓	32.5	33.1

Table 3. Design choices for temporal fusion module.

Layers	d	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
P3	13	33.8	51.3	35.8	32.1	39.6
P4	11	35.0	53.0	38.7	32.8	40.1

and FCB (ali) denote FCB with adaptive features on all three branches, and that with aligned features on mask co-efficient and tracking branches, respectively.

Design Choices of Temporal Fusion Module. We first conduct experiments on different FPN layers and the side length of local square for the correlation operation. As shown in Table 3, the P4 FPN layer with the side length 11 obtains the better performance for temporal fusion module.

4.2. Ablation Study

We conduct the ablation study on YouTube-VIS valid set using Yolact with ResNet101-DCN backbone[4] as our baseline. For fair comparison, the baseline of all experiments adopts three anchors and does not include semantic segmentation loss. Table 4 shows the impact of progressively integrating our different components: feature calibration with anchor boxes (FCA), adaptive / aligned feature calibration with anchor boxes (FCB), and temporal fusion module (TF) to the baseline. The four components progressively increase segmentation performance to 36.8 % mask AP. Besides, to better explore the individual improvements for mask segmentation brought by each single component, we also conduct the experiments adding the different components individually to the baseline shown in Table 5, where FCA, FCB(ada), FCB(ali) and TF component respectively contribute to the gains of 1.9 % , 2.3 % , 2.4 % and 5.2 % mask AP. Among these components, the FCB(ada) and TF provide the most improvement in accuracy over the baseline. These results suggest that each of our components individually contributes towards improving the final performance.

We also perform more ablation study for STMask on COCO dataset to verify the effectiveness of spatial feature

Table 4. Ablation study on YouTube-VIS valid: progressively integrating different components into the baseline.

Baseline+FCA	+FCB	+TF	AP	AP ₅₀	AP ₇₅	FPS
			(ada / ali)			
✓			28.9	45.8	30.1	31.2
✓	✓		30.8	49.8	31.0	29.5
✓	✓	✓ /	34.4	53.0	34.6	28.3
✓	✓	/ ✓	33.1	51.4	35.2	26.7
✓	✓	✓ /	36.8	56.8	38.0	23.4
✓	✓	/ ✓	36.3	55.2	39.9	22.1

Table 5. Ablation study on YouTube-VIS valid: individually integrating different components into the baseline.

Methods	Impro.	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
Baseline	-	28.9	45.8	30.1	9.8	26.9	40.3
+ FCA	+1.9	30.8	49.8	31.0	11.8	29.7	44.6
+ FCB(ada)	+2.3	31.2	52.1	31.2	11.1	32.8	43.6
+ FCB(ali)	+2.4	31.3	50.0	32.8	10.3	31.7	43.3
+ TF	+5.2	34.1	51.9	36.0	11.4	30.3	46.5

calibration on image instance segmentation task, which are provided in the **supplementary material**.

4.3. Mask Results

We compare the proposed STMask with state-of-the-arts on YouTube-VIS valid set. The results are shown in Table 6. The ResNet50 backbone is used by the competing methods. The speed is computed on a single 2080Ti GPU.

From Table 6, we can see that among existing fast video object segmentation (VOS) methods, OSMN [44] with ‘track-by-detect’ and FEELVOS [37] achieve relatively lower mask AP scores. The bottle-up method STEm-Seg with ResNet 50 and ResNet 101 backbone reaches 30.6 % and 34.6 % mask AP respectively. The recently introduced two-stage methods MaskTrack R-CNN [43] and MaskProp [2] obtain mask AP scores of 30.3 % and 40.0 %, respectively. However, MaskProp requires a video clip (more than 13 frames) to process an image, resulting in very slow speed. For the one-stage method using ResNet50 Backbone, SipMask [6] and our proposed STMask obtain 32.5 % and 33.5 % mask AP with a speed of 30.0 and 28.6 FPS respectively. In addition, STMask with ResNet101-DCN backbone achieves 36.8 % and 36.3 % mask AP on adaptive and aligned FCB settings, respectively, without any other extra tricks. Although STMask has relatively lower mask AP than MaskProp, it does achieve a better trade-off between accuracy and speed.

Figs. 6 and 7 visualise the instance segmentation masks of our STMask on challenging videos with small objects, occlusion and uncommon camera-to-object view. We can see that the temporal fusion module can indeed reduce the missed detection across frames, compared with frame-level instance segmentation.

Table 6. Quantitative performance comparison on YouTube-VIS valid set, where a, b, c, and d respectively refer to semantic segmentation, high-resolution mask refinement, temporal information, and multi-scale training.

Type	Methods	Backbone	Frames	Anchors	Others	FPS	AP	AP ₅₀	AP ₇₅	AR ₁	AR ₁₀
VOS baseline	OSMN [44]	R50-FPN	3	3	-	-	27.5	45.1	29.1	28.6	33.1
	FEELVOS [37]	R50-FPN	3	3	-	-	26.9	42.0	29.7	29.9	33.4
	STEm-Seg [1]	R50-FPN	-	-	-	-	30.6	50.7	33.5	31.6	37.1
	STEm-Seg [1]	R101-FPN	-	-	-	2.1	34.6	55.8	37.9	34.4	41.6
Two-stage methods	MaskTrack [43]	R50-FPN	2	3	-	18.4	30.3	51.1	32.6	31.0	35.5
	MaskProp [2]	R50-FPN	13	3	b,c,d	-	40.0	-	42.9	-	-
	MaskProp [2]	R101-FPN	13	3	b,c,d	-	42.5	-	45.6	-	-
One-stage methods	SipMask++ [6]	R50-DCN	2	-	-	30.0	32.5	53.0	33.3	33.5	38.9
	SipMask++ [6]	R101-DCN	2	-	-	27.8	35.0	56.1	35.2	36.0	41.2
	STMask (ada)	R50-DCN	2	9	c	28.6	33.5	52.1	36.9	31.1	39.2
	STMask (ada)	R101-DCN	2	3	c	23.4	36.8	56.8	38.0	34.8	41.8
	STMask (ali)	R101-DCN	2	3	c	22.1	36.3	55.2	39.9	33.7	42.0

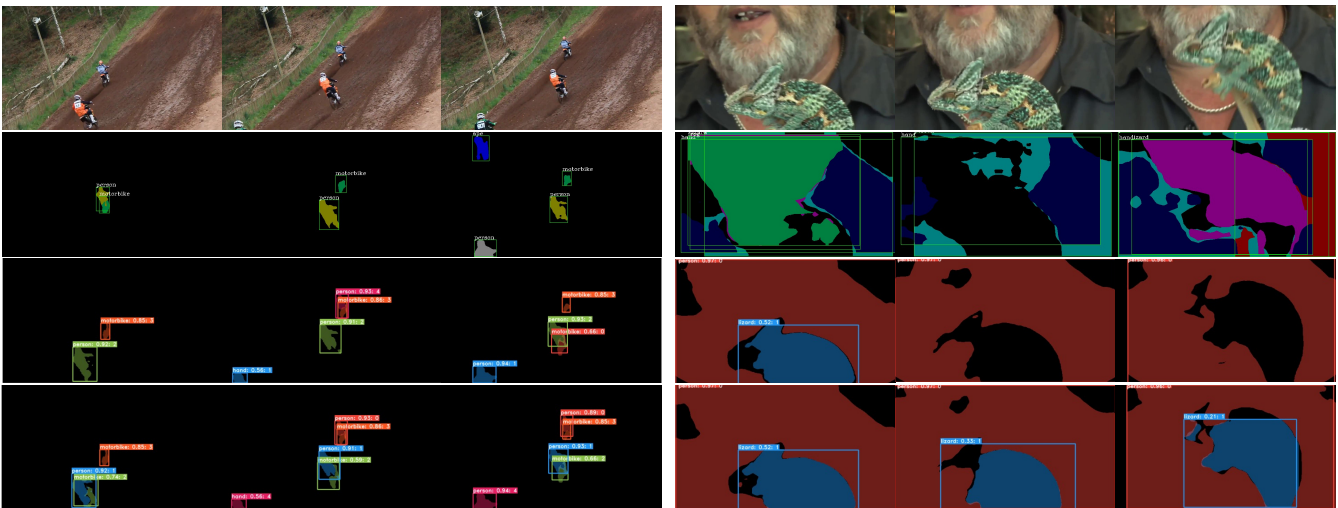


Figure 6. Segmentation visual comparison on two videos with small objects and occlusion. The first two rows show the original frames and the instance masks produced by MaskTrack R-CNN [43] respectively, while the last two rows are the results obtained by our STMask with only spatial calibration and with both spatial calibration and temporal fusion module, respectively.




Figure 7. Segmentation visual comparison on a video with uncommon camera-to-object view.

5. Conclusion

We observe that one-stage instance segmentation approaches underestimate the importance of spatial feature calibration and temporal redundancy information between video frames for video instance segmentation. To address the issue, we first propose a simple spatial feature calibration to detect and segment object masks frame-by-frame, and further introduce a temporal fusion module to track instance across video frames to effectively reduce missed instances on challenging videos like motion blur, partial occlusion and unusual object-to-camera poses. Overall, on YouTube-VIS valid set, our proposed STMask with ResNet-50/101 backbone obtains 33.5% / 36.8 % mask AP, while achieving 28.6 / 23.4 FPS on video instance segmentation, which does keep the trade-off between accuracy and speed.

References

- [1] Ali Athar, Sabarinath Mahadevan, Aljosa Osep, Laura Leal-Taixé, and Bastian Leibe. Stem-seg: Spatio-temporal embeddings for instance segmentation in videos. In *European Conference on Computer Vision*, pages 158–177. Springer, 2020. 8
- [2] Gedas Bertasius and Lorenzo Torresani. Classifying, segmenting, and tracking object instances in video with mask propagation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9739–9748, 2020. 1, 2, 7, 8
- [3] Gedas Bertasius, Lorenzo Torresani, and Jianbo Shi. Object detection in video with spatiotemporal sampling networks. In *Eur. Conf. Comput. Vis.*, pages 331–346, 2018. 2
- [4] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact++: Better real-time instance segmentation. *arXiv preprint arXiv:1912.06218*, 2019. 6, 7
- [5] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact: Real-time instance segmentation. In *Int. Conf. Comput. Vis.*, pages 9157–9166, 2019. 2, 3, 5, 6
- [6] Jiale Cao, Rao Muhammad Anwer, Hisham Cholakkal, Fahad Shahbaz Khan, Yanwei Pang, and Ling Shao. SipMask: Spatial information preservation for fast image and video instance segmentation. *arXiv preprint arXiv:2007.14772*, 2020. 2, 3, 7, 8
- [7] Hao Chen, Kunyang Sun, Zhi Tian, Chunhua Shen, Yongming Huang, and Youliang Yan. BlendMask: Top-down meets bottom-up for instance segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8573–8581, 2020. 2
- [8] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiao-xiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4974–4983, 2019. 1
- [9] Liang-Chieh Chen, Alexander Hermans, George Papandreou, Florian Schroff, Peng Wang, and Hartwig Adam. MaskLab: Instance segmentation by refining object detection with semantic and direction features. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4013–4022, 2018. 2
- [10] Yihong Chen, Yue Cao, Han Hu, and Liwei Wang. Memory enhanced global-local aggregation for video object detection. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10337–10346, 2020. 2
- [11] Jifeng Dai, Kaiming He, and Jian Sun. Instance-aware semantic segmentation via multi-task network cascades. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3150–3158, 2016. 2
- [12] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Int. Conf. Comput. Vis.*, pages 764–773, 2017. 1, 2
- [13] Hanming Deng, Yang Hua, Tao Song, Zongpu Zhang, Zhen-gui Xue, Ruhui Ma, Neil Robertson, and Haibing Guan. Object guided external memory network for video object detection. In *Int. Conf. Comput. Vis.*, pages 6678–6687, 2019. 2
- [14] Jiajun Deng, Yingwei Pan, Ting Yao, Wengang Zhou, Houqiang Li, and Tao Mei. Relation distillation networks for video object detection. In *Int. Conf. Comput. Vis.*, pages 7023–7032, 2019. 2
- [15] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *Int. Conf. Comput. Vis.*, pages 3038–3046, 2017. 2, 5
- [16] Qianyu Feng, Zongxin Yang, Peike Li, Yunchao Wei, and Yi Yang. Dual embedding learning for video instance segmentation. In *Int. Conf. Comput. Vis. Worksh.*, pages 0–0, 2019. 1
- [17] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 580–587, 2014. 4
- [18] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Int. Conf. Comput. Vis.*, pages 2961–2969, 2017. 1, 2
- [19] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring R-CNN. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6409–6418, 2019. 2
- [20] Alexander Kirillov, Evgeny Levinkov, Bjoern Andres, Bogdan Savchynskyy, and Carsten Rother. InstanceCut: from edges to instances with multicut. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5008–5017, 2017. 2
- [21] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. PointRend: Image segmentation as rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9799–9808, 2020. 2
- [22] Shu Kong and Charless C Fowlkes. Recurrent pixel embedding for instance grouping. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9018–9028, 2018. 2
- [23] Shuai Li, Lingxiao Yang, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Dynamic anchor feature selection for single-shot object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6609–6618, 2019. 3
- [24] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, and Yichen Wei. Fully convolutional instance-aware semantic segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2359–2367, 2017. 2
- [25] Xiaodan Liang, Liang Lin, Yunchao Wei, Xiaohui Shen, Jianchao Yang, and Shuicheng Yan. Proposal-free network for instance-level object segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 40(12):2978–2991, 2017. 2
- [26] Chung-Ching Lin, Ying Hung, Rogerio Feris, and Linglin He. Video instance segmentation tracking with a modified vae architecture. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 13147–13157, 2020. 2
- [27] 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 *Eur. Conf. Comput. Vis.*, pages 740–755. Springer, 2014. 6
- [28] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8759–8768, 2018. 2

- [29] Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online dictionary learning for sparse coding. pages 689–696, 2009. [5](#)
- [30] Davy Neven, Bert De Brabandere, Marc Proesmans, and Luc Van Gool. Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 8837–8845, 2019. [2](#)
- [31] Alejandro Newell, Zhiao Huang, and Jia Deng. Associative embedding: End-to-end learning for joint detection and grouping. In *Adv. Neural Inform. Process. Syst.*, pages 2277–2287, 2017. [2](#)
- [32] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Adv. Neural Inform. Process. Syst.*, pages 91–99, 2015. [1](#), [2](#)
- [33] Mykhailo Shvets, Wei Liu, and Alexander C Berg. Leveraging long-range temporal relationships between proposals for video object detection. In *Int. Conf. Comput. Vis.*, pages 9756–9764, 2019. [2](#)
- [34] Konstantin Sofiiuk, Olga Barinova, and Anton Konushin. AdaptIS: Adaptive instance selection network. In *Int. Conf. Comput. Vis.*, pages 7355–7363, 2019. [2](#)
- [35] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. *arXiv preprint arXiv:2003.05664*, 2020. [2](#)
- [36] Jonas Uhrig, Eike Rehder, Björn Fröhlich, Uwe Franke, and Thomas Brox. Box2Pix: Single-shot instance segmentation by assigning pixels to object boxes. In *IEEE Intell. Veh. Symp.*, pages 292–299. IEEE, 2018. [2](#)
- [37] Paul Voigtlaender, Yuning Chai, Florian Schroff, Hartwig Adam, Bastian Leibe, and Liang-Chieh Chen. FEELVOS: Fast end-to-end embedding learning for video object segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9481–9490, 2019. [7](#), [8](#)
- [38] Jiaqi Wang, Kai Chen, Shuo Yang, Chen Change Loy, and Dahua Lin. Region proposal by guided anchoring. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2965–2974, 2019. [2](#), [3](#), [4](#)
- [39] Shiyao Wang, Yucong Zhou, Junjie Yan, and Zhidong Deng. Fully motion-aware network for video object detection. In *Eur. Conf. Comput. Vis.*, pages 542–557, 2018. [2](#)
- [40] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. SOLO: Segmenting objects by locations. *arXiv preprint arXiv:1912.04488*, 2019. [2](#)
- [41] Fanyi Xiao and Yong Jae Lee. Video object detection with an aligned spatial-temporal memory. In *Eur. Conf. Comput. Vis.*, pages 485–501, 2018. [2](#)
- [42] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. PolarMask: Single shot instance segmentation with polar representation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 12193–12202, 2020. [2](#)
- [43] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *Int. Conf. Comput. Vis.*, pages 5188–5197, 2019. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [44] Linjie Yang, Yanran Wang, Xuehan Xiong, Jianchao Yang, and Aggelos K Katsaggelos. Efficient video object segmentation via network modulation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6499–6507, 2018. [7](#), [8](#)
- [45] Rufeng Zhang, Zhi Tian, Chunhua Shen, Mingyu You, and Youliang Yan. Mask encoding for single shot instance segmentation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 10226–10235, 2020. [2](#)
- [46] Xizhou Zhu, Yujie Wang, Jifeng Dai, Lu Yuan, and Yichen Wei. Flow-guided feature aggregation for video object detection. In *Int. Conf. Comput. Vis.*, pages 408–417, 2017. [2](#)