

Real-time Instance Segmentation with Discriminative Orientation Maps

Wentao Du Zhiyu Xiang* Shuya Chen Chengyu Qiao Yiman Chen Tingming Bai
College of Information Science & Electronic Engineering, Zhejiang University

Abstract

Although instance segmentation has made considerable advancement over recent years, it's still a challenge to design high accuracy algorithms with real-time performance. In this paper, we propose a real-time instance segmentation framework termed *OrienMask*. Upon the one-stage object detector YOLOv3, a mask head is added to predict some discriminative orientation maps, which are explicitly defined as spatial offset vectors for both foreground and background pixels. Thanks to the discrimination ability of orientation maps, masks can be recovered without the need for extra foreground segmentation. All instances that match with the same anchor size share a common orientation map. This special sharing strategy reduces the amortized memory utilization for mask predictions but without loss of mask granularity. Given the surviving box predictions after NMS, instance masks can be concurrently constructed from the corresponding orientation maps with low complexity. Owing to the concise design for mask representation and its effective integration with the anchor-based object detector, our method is qualified under real-time conditions while maintaining competitive accuracy. Experiments on COCO benchmark show that *OrienMask* achieves 34.8 mask AP at the speed of 42.7 fps evaluated with a single RTX 2080 Ti. The code is available at <https://github.com/duwt/OrienMask>.

1. Introduction

Instance segmentation aims at pixel-wise predictions for every individual object. It integrates instance-level object detection [28, 26, 19, 16] and pixel-level semantic segmentation [21, 6, 7], formulating a more fine-grained visual perception task. Currently there are two dominating types of solutions, namely detection-based and segmentation-based methods. The former extends an object detector with additional foreground dense predictions while the latter deploys specific per-pixel attributes or embeddings to separate instances of the same category in a bottom-up way.




Figure 1. **Orientation-based Mask Construction.** Those arrowed lines denote densely predicted orientation vectors. Each mask is constructed by gathering all pixels pointing to the central region of the instance in the matched orientation map (white or black).

Both paradigms have apparent drawbacks. Conventional detection-based methods like Mask R-CNN [9] rely on the features pooling operation to project all regions of interest (RoIs) into a fixed size. Since the subsequent mask head should be applied to abundant feature maps of every region proposal, the speed is largely constrained especially when objects densely appear. Moreover, the constant mask resolution brings in unnecessary computation for small objects and loses valuable details for large targets. On the contrary, segmentation-based methods [20, 15] retain the fine-grained appearance and geometry in a pixel-to-pixel manner. They can acquire satisfactory results at elementary scenarios but often fall behind detection-based approaches in accuracy. When the scale of objects varies and the number of categories increases, the generalization of pixel-level clustering adopted in segmentation-based methods is still in doubt.

For the requirement of real-time inference, YOLACT [2] is proposed along with a special mask construction scheme, which linearly combines shared non-local prototypes with instance-wise coefficients. It discards the RoI pooling operation that commonly adopted in earlier detection-based methods and directly assembles masks from fine-grained feature maps. Following this paradigm, an improved approach named BlendMask [4] is put forward. It replaces the 1D instance-specific coefficients with a set of attention maps, which supply additional spatial-adaptive information

*Corresponding author. Email: xiangzy@zju.edu.cn

to enrich fine-grained details of masks. The success of these solutions shows great potential to incorporate informative global features into detection-based methods. However, one noticeable flaw of these approaches lies in the dependence of RoI cropping operations when generating the assembled masks, which may bring in some mask incompleteness due to inaccurate bounding box predictions.

In this work, we attempt to integrate fine-grained expressions with an one-stage detector in another way. To be specific, we focus on compact mask representation and efficient integration with the anchor-based detector YOLOv3 [27] to achieve real-time performance. First, a novel discriminative orientation map is proposed to encode multiple masks independently, where pixels are assigned with centripetal or centrifugal vectors according to their positive or negative labels. This design is totally free from any other semantic segmentation or foreground predictions and is lightweight to decode the complete masks. In addition, considering objects of diverse scales vary the magnitude distributions of orientation vectors, multi-scale design is also taken into account. We assign different orientation maps for instances matching with certain anchor sizes so that the completeness of mask representation is guaranteed. OrienMask merely extends an extra head to the object detector and its function is tightly combined with the box assignment and pre-defined anchor sizes. During inference, for each predicted bounding box, its instance mask can be quickly constructed based on the discriminative vectors in the corresponding orientation map, as shown in Figure 1. This process is simple and direct, consisting of nothing but determining binary labels for all pixels by the spatial destinations that orientation vectors indicate. The main contributions of our work can be summarized as follows:

- We put forward a light and discriminative orientation-based mask representation for real-time instance segmentation. By defining opposite orientation vectors for foreground and background pixels, we are able to effectively encode multiple instance masks in a fine-grained two-channel map without the need for explicit foreground segmentation. In inference, given the target regions of instances, their masks can be easily constructed from orientation maps in parallel.
- To deal with objects with various sizes, we propose an instance grouping mechanism derived from anchor-based detectors. Each group of instances with similar sizes are assigned to share a common class-agnostic orientation map. We also expand the annotated bounding boxes to provide sufficient supervision for background. The enlarged valid training area not only balances the number of positive and negative samples but also helps distinguishing them around the borders.
- We integrate the discriminative orientation maps into a fast anchor-based detector YOLOv3 and implement

the resulting model OrienMask end-to-end. Experiments show that it is able to achieve 34.8 mask AP at a speed of 42.7 fps on COCO benchmark, which is quite competitive among state-of-the-art real-time methods.

2. Related Work

Detect-then-segment Methods With the help of proposals generated by object detectors, detect-then-segment methods first extract reliable RoIs from feature maps, and then obtain fine-grained instance representations. Driven by the success of two-stage detector Faster R-CNN [28], Mask R-CNN [9] adds a branch for mask prediction in parallel with bounding box regression, and employs RoIAlign to fix the misalignment caused by spatial quantization. After that, PANet [18] is presented to strengthen message passing in the bottom-up path and fuse pooling features from all levels. HTC [5] extends Mask R-CNN into a cascade structure by interleaving the mask and box branches while maintaining semantic features fusion. Instead of using the confidence of the detector, Mask Scoring R-CNN [10] predicts an extra score to accurately represent the mask quality. However, due to the heavy computations in the second stage, these approaches can hardly satisfy the requirement of real-time inference.

Detect-and-segment Methods Benefiting from those compact architectures of one-stage object detection [19, 16, 27, 13, 8, 30], detect-and-segment methods customize masks from the global feature maps jointly with implicit instance-specific representations. In YOLACT [2], acknowledged as the milestone of this paradigm, a series of mask coefficients are produced along with the box predictions. Then they are multiplied with a set of high-resolution prototypes to generate instance masks. Chen *et al.* [4] rethink the trade-off between feature resolution and coefficients dimension, and propose BlendMask, which blends some attention maps of each instance with a group of shared bases. Inspired by conditionally parameterized convolutions [38], CondInst [29] predicts instance-aware convolutional kernel weights and applies them on high-resolution feature maps. Thanks to the flexible framework and fine-grained representation, these methods keep a good balance between speed and accuracy. Compared with them, our OrienMask employs an explicit and discriminative features sharing scheme for mask representation rather than implicit parameterized forms, which is more concise and provides strong interpretability.

Other Mask Representations Apart from bounding boxes and foreground probability maps, some compact mask representations also contribute to instance segmentation. For example, Jetley *et al.* [11] deploy an auto-encoder to compress masks into low-dimensional vectors which can be incorporated in a detector. Xu *et al.* [37] describe an instance as a series of inner-center radii and encode them into Chebyshev polynomial coefficients. To achieve better precision,

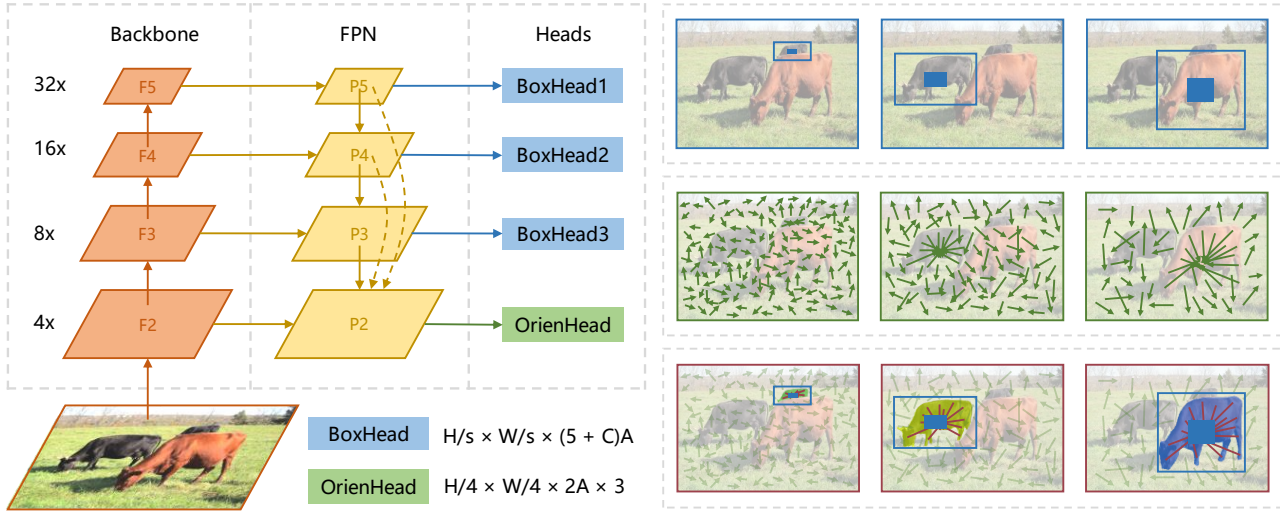


Figure 2. **OrienMask Architecture.** Left: The network is built upon YOLOv3 with an additional head to predict different orientation maps for each anchor size. H and W are height and width of the input image while s denotes the output stride. There are C categories and A anchors per spatial position. Pyramid features fed to OrienHead are strengthened in our improved version as dashed lines indicate. Right: There are three instances matching with different anchor sizes. The bounding boxes in the first row are predicted by BoxHeads while the orientation maps in the second row are predicted by OrienHead. Each pair in the same column individually determines an instance mask.

polar centerness and polar IoU loss are proposed in PolarMask [36]. Peng *et al.* [25] implement the snake algorithm in a learning-based form and the circular convolution is proposed to iteratively regress sampling points towards contour positions. Serving as effective representation, pixel offset and its variants are popular in separating instances within the segmented foreground. Uhrig *et al.* [31] deploy template matching by predicted depth classes and discrete directions to assign pixels of the same semantic label to different instance centers. Box2Pix [32] matches the foreground pixels with predicted box centers according to offset vectors. Similarly, Li *et al.* [14] merge foreground pixels based on adaptive voting zones around detection centers. Neven *et al.* [22] propose a joint optimization scheme for class-specific seed and sigma maps along with dense offset vectors. Given the learned clustering bandwidth, masks are sequentially recovered. PersonLab [24] utilizes the short-range and mid-range offset vectors to decode person poses, and then clusters foreground pixels by long-range offsets. Novotny *et al.* [23] propose a semi-convolutional operator, which adds coordinates to part of learned embedding. PointGroup [12] extends offset descriptor to 3D instance segmentation, where points of the same category are grouped step by step. Our method also draws inspiration from spatial offset descriptor. However, unlike above methods which mostly use spatial offset to assign segmented foreground pixels to instances, our orientation map is self-discriminative. It is able to filter out background regions and separate instances at the same time. To achieve this goal, special valid training areas are defined and distinctive orientation vectors for both positive and negative samples are considered. Besides, these fine-

grained orientation maps are tightly bonded with anchors of the detector, which retains the mask completeness of different sizes and eases the regression.

3. OrienMask

3.1. Overall Architecture

The network architecture of OrienMask is mainly built upon anchor-based detector YOLOv3 [27] with the backbone Darknet-53. As illustrated in Figure 2, we add an extra OrienHead to predict orientation maps, which is the key part of the whole framework. The obtained bounding boxes and orientation maps are combined for mask construction.

YOLOv3 deploys 9 bounding box priors and evenly assigns them across 3 scales. Given an image with height H and width W as input, the BoxHead at the scale of output stride s produces $H/s \times W/s \times A$ bounding box predictions, where A denotes the number of anchors per grid cell. OrienHead takes the largest feature maps P2 after feature pyramid network (FPN) as input and then predicts $3A$ orientation maps of fixed resolution $H/4 \times W/4$ with two channels, matching with $3A$ different anchor sizes respectively. Given that handling feature maps of high resolutions is time-consuming, our OrienHead is designed to be lightweight. It is interleaved with three 3×3 and 1×1 convolution layers of input channels 128 and 256 alternatively. After the standard non-maximum suppression, each bounding box is paired with an orientation map according to its anchor size. As demonstrated in the right part of Figure 2, all pixels whose orientation vectors end within a contracted bounding box form a foreground mask.




Figure 3. **Orientation Maps.** For each instance, pixels covered by the mask are positive samples whose orientations are defined as vectors pointing to the base position. Remaining pixels inside the valid training area serve as negative samples, which should point to the border of the valid training area in centrifugal directions.

3.2. Orientation-based Mask Representation

An orientation map $O^* \in \mathbb{R}^{H \times W \times 2}$ stores pixel-wise spatial offsets in horizontal and vertical directions for multiple instances. Some pivotal concepts at the stage of training are illustrated in Figure 3 and will be explained in detail.

Valid Training Area We first expand regions enclosed by annotated bounding boxes to form valid training areas. Pixels outside of any expanded region are ignored during training, which means they are not involved in the loss calculation. All remaining valid pixels are divided into two parts, namely positive and negative samples, based on whether they are covered by instance masks or not. Since the number of positive samples is constant and more negative samples will be counted when the valid training area expands, a proper expand ratio should be determined to balance the number of positive and negative samples. Meanwhile, this expansion also provides sufficient guidance for distinguishing pixels nearby the instance borders.

Orientation Vectors To easily distinguish positive and negative samples in mask construction process, the orientation vectors of them are defined to point at opposite directions. To be specific, a base position is first specified for each instance and the bounding box centroid is a preferable choice in our experiments. The positive samples on the orientation map are defined pointing to the base position while the negative ones should point to the boarder of the valid training area in centrifugal directions. Denoting the base position as b , the target orientation vector o_i^* for pixel i at location p_i can be expressed as

$$o_i^* = \begin{cases} b - p_i, & \text{if positive,} \\ (\alpha_i - 1)(p_i - b), & \text{if negative,} \end{cases} \quad (1)$$

$$\alpha_i = 1 / \max\left(\frac{p_i^x - b^x}{v^l - b^x}, \frac{p_i^x - b^x}{v^r - b^x}, \frac{p_i^y - b^y}{v^t - b^y}, \frac{p_i^y - b^y}{v^b - b^y}\right). \quad (2)$$

The superscripts x and y denote horizontal and vertical directions respectively. Elements in (v^l, v^r, v^t, v^b) represent the leftmost, rightmost, topmost, and bottommost coordinates of the valid training area. All positive samples have the highest priority to override other orientations so that the mask completeness is preserved as much as possible. If a negative sample is overlapped by multiple valid training areas, we simply take the average to avoid ambiguity.

Since orientation vectors are locally defined as spatial offsets pointing to some neighboring destinations, they vary smoothly in both directions within the interior places. The harmony may be slightly disturbed when valid training areas of different instances overlap but the overall unity is still maintained. For each pair of neighboring pixels locating across the mask boundary, the positive is pulled to the base position while the negative is pushed towards outside to the expanded boarder. Thus one direction of gradients at these positions equals to the distance between the base position and the corresponding boarder of valid training area, which is significantly larger than one pixel difference in other interior places. Our experiments in Section 4.3 do prove that the learned orientation maps retain this property, which is helpful to accurately delineate instance masks.

Instance Grouping Although simply stacking all instance masks onto a two-channel orientation map is fascinating for real-time considerations, it could fail in handling some instance overlapping scenarios, such as a person wearing a tie. To alleviate this problem, we introduce an instance grouping mechanism. Noticing that YOLOv3 assigns objects to different anchor sizes based on the intersection over unions with those bounding box priors, we naturally transfer this assignment to our mask representation. To be specific, instances are divided into several groups according to the anchor sizes that they are matched, and each group of instance masks are assigned to an independent orientation map.

Besides solving the overlapping problem caused by objects of different aspect ratios or scales, the instance grouping mechanism has more additional advantages. Since bigger objects often require larger receptive field, this arrangement is beneficial to adapt each orientation map to appropriate scale. Meanwhile, the orientation vectors of grouped instances can be normalized within a small interval so that the magnitude distribution of each group does not vary significantly, which is good for the network training. Moreover, our design also complies with the observation that an image may contain many small instances but only a few large ones. Hence it can preserve as many objects as possible.

3.3. Mask Construction

The mask construction process involves two elements: a predicted bounding box B and an orientation map O . Recall that each bounding box prediction has an anchor size and each anchor size is associated with an orientation map.

Therefore, each bounding box is bound to be matched with an orientation map. Supposing that B and O have been paired, we take the centroid of B as the base position according to the definition in Eq. (1). Then a rectangular target region centered on the base position is defined, whose size is proportional to the width and height of B . If we denote the base position as b and the size of bounding box as s , the constructed mask M can be expressed as

$$M = (|O_x + P_x - b_x| < \tau \cdot s_x) \cap (|O_y + P_y - b_y| < \tau \cdot s_y). \quad (3)$$

Here $P \in \mathbb{R}^{H \times W \times 2}$ stores the coordinates of each pixel, and τ is a contracting factor to define the target region. In a nutshell, if an orientation vector points to somewhere within certain realms around the base position, the corresponding pixel belongs to the foreground of the given instance. This simple procedure requires no more than point-wise arithmetic or logic operations. For all bounding boxes that survive NMS, the mask construction process can be easily executed in parallel without foreground segmentation. Furthermore, it does not require any RoI cropping operations and directly considers the whole orientation map, resulting little precision leakage caused by inaccurate box predictions.

3.4. Loss Function

The loss function consists of two components, providing supervisions for object detection and orientation maps respectively. It can be formulated as

$$\mathcal{L} = \mathcal{L}_{det} + \lambda \mathcal{L}_{orien}, \quad (4)$$

where λ is a hyper-parameter to balance these two terms. \mathcal{L}_{det} is completely copied from official YOLOv3 without any tricks in the literature.

With regard to \mathcal{L}_{orien} , we compute smooth-l1 loss per pixel in valid training areas, and then take the average over positive and negative samples respectively. In addition, we multiply them by the number of instances N_{inst} in accordance with \mathcal{L}_{det} . The complete expression is written as

$$\begin{aligned} \mathcal{L}_{orien} = & \frac{N_{inst}}{N_{pos}} \sum_i \mathbb{1}_i^{\mathcal{P}} f_{smooth,l1}\left(\frac{o_i}{a}, \frac{o_i^*}{a}\right) \\ & + \frac{N_{inst}}{N_{neg}} \sum_i \mathbb{1}_i^{\mathcal{N}} f_{smooth,l1}\left(\frac{o_i}{a}, \frac{o_i^*}{a}\right), \end{aligned} \quad (5)$$

where $\mathbb{1}^{\mathcal{P}}$ and $\mathbb{1}^{\mathcal{N}}$ are the indicator functions for positive and negative samples. o_i denotes the orientation vector predicted at pixel i while o_i^* is the corresponding ground truth. To lift numerical stability and reduce the variance across different scales, o_i and o_i^* are normalized by their anchor size a instead of directly measured in pixels. \mathcal{L}_{orien} for each scale is calculated separately and finally summed up. Though the orientation maps produced by the network is

one fourth dimensions of the input image, we calculate and aggregate losses after upsampling them to the full resolution. Likewise, bilinear interpolation is also implemented before feeding orientation maps to construct masks.

4. Experiments

We conduct experiments on the challenging MS COCO dataset [17] and evaluate the predictions with standard metrics. Following the common practice, all models are trained with 118k images of `train2017` and tested on 5k images of `val2017` or 20k images of `test-dev` subset.

Training Details For the network structure, we retain the official implementation of YOLOv3 and extend a fully convolutional OrientHead as described above. The backbone Darknet-53 is initialized with a pretrained detector and the network is trained end-to-end. We employ stochastic gradient descent (SGD) optimizer with momentum 0.9 and weight decay 0.0005. The batch size is 16 and the synchronized batch normalization is used in our ultimate model but not in ablation study. The initial learning rate is 0.001 and divided by 10 at iterations 520k and 660k respectively. All models are trained for 100 epochs with the input resolution 544×544 . Multiple data augmentations are applied, such as color jitter, random resizing, and horizontal flipping.

Inference Details Similar to YOLACT [2], input images are directly resized to 544×544 without test-time augmentation. The inference speed is evaluated on RTX 2080 Ti by default and measured with frames per second (FPS).

4.1. Ablation Study

In our ablation experiments, implementations of the object detector is fixed. We adjust other hyper-parameters of our method to obtain the best configurations. To integrate OrientHead with the detector more tightly, some additional refinements will also be applied to the base model.

Valid Training Area For orientation maps, the definition of negative samples is closely related to the boarder of valid training area. Given that the size of each valid training area is proportional to its bounding box, we vary the expand ratio r from 1.0 to 1.6 with the stride 0.2. As the experimental results shown in Table 1, our model achieves best performance when $r = 1.2$ and either smaller or larger expand ratio brings some drops in AP. We notice the amount of negative samples and their magnitudes of orientation vectors increase simultaneously as valid training areas expand. A moderate expand ratio keeps a balance between the number of positive and negative samples while maintains enough differentiation around boundaries. It also keeps a proper numeric distribution of negative samples. These two aspects both contribute to better convergence of our model.

Orientation Loss Weight In the loss function of Eq. (4), \mathcal{L}_{det} and \mathcal{L}_{orien} are used to provide box-level and pixel-

r	AP	AP ₅₀	AP ₇₅	AP ^{bb}
1.0	32.0	53.4	33.1	37.0
1.2	32.5	53.9	33.6	37.4
1.4	32.3	53.8	33.0	37.4
1.6	31.9	53.6	32.5	37.1

Table 1. **Valid Training Area.** An appropriate expand ratio r balances the amounts of positive and negative samples while casting enough supervision to distinguish them.

λ	AP	AP ₅₀	AP ₇₅	AP ^{bb}
5	31.0	52.7	31.3	36.7
10	31.8	53.5	32.5	36.8
15	32.1	53.4	33.3	37.1
20	32.5	53.9	33.6	37.4

Table 2. **Orientation Loss Weight.** As the orientation loss weight λ increases, the mask AP performance of our model is gradually boosted while the box-level AP metric remains relatively stable.

level supervisions respectively. In order to associate these two terms together, we explore the orientation loss weight λ from 5 to 20 and obtain the results in Table 2. The mask AP metrics gradually increase as λ becomes larger. However, we find the training process gets unstable and the performance saturates when applying weights larger than 20. Hence $\lambda = 20$ is adopted in our subsequent experiments. Moreover, we observe that the adjustment of the orientation loss weight does not disturb the box-level performance remarkably and even has a bit positive effects, which indicates OrientHead maintains the network stability to some extent.

Orientation Target Area For any bounding box that survives NMS, the mask is simply constructed by collecting all pixels whose orientation vectors point to somewhere close to its base position, without any other operations like RoI cropping. The bounding box is contracted by a scale factor to form an orientation target area so that it is compatible with objects of different aspect ratios or scales. Here we choose contraction ratio τ from 0.4 to 0.8. From Table 3 we find the performance is sensitive to the contraction ratio and the best performance is obtained when $\tau = 0.6$. More specific parameter tuning for each group of instances is optional to achieve even higher AP.

Other Improvements We further explore some measures to better integrate OrientHead with the detector and improve the overall performance. These measures are applied step by step and the results are illustrated in Table 4. For the orientation definition, we initially choose the grid centers as base positions, which keeps consistent with the rule of box regression in the detector. Since the box centroid is more accurate to locate an instance, we adopt it as the base position. This refinement improves the mask AP metric from 32.5 to

τ	AP	AP ₅₀	AP ₇₅
0.4	30.4	50.7	31.3
0.5	32.2	52.9	33.2
0.6	32.5	53.9	33.6
0.7	31.7	53.8	32.4
0.8	30.1	52.8	29.9

Table 3. **Orientation Target Area.** The contraction ratio τ for orientation target area has strong effects on overall performance.

Method	AP	AP ₅₀	AP ₇₅
base	32.5	53.9	33.6
+ box centroid	33.3	54.8	34.4
+ larger anchors	33.8	55.2	35.2
+ fpn plus	34.1	55.7	35.4

Table 4. **Other Improvements.** ‘+ box centroid’ means replacing grid centers with box centroids as base positions for orientation maps. ‘+ larger anchors’ inherits anchor settings from YOLOv4 which have about 1.2 times the sizes of the original version. ‘+ fpn plus’ merges pyramid features of P3, P4, and P5 to generate P2 for better multi-scale expressions in OrientHead.

33.3. Then we are inspired by the next generation of YOLO framework [1] to adopt larger anchors, which are proved more suitable for the given input resolution. It should be noted that no extra training tricks are taken and we still maintain other settings as pure YOLOv3. Benefiting from the larger anchors, the performance is boosted by 0.5 AP. In earlier experiments, we follow the standard FPN to produce P2 for OrientHead but the predicted orientation maps are related with box predictions from multiple scales. To closely associate these two outputs, we merge multi-scale pyramid features to predict orientation maps while keeping the same network structure in subsequent layers, as dashed lines in Figure 2 indicate. The resulting model surpasses the previous one again with negligible extra computing cost.

4.2. Comparison with State-of-the-art Methods

We first evaluate our OrientMask on the canonical COCO `test-dev` benchmark and select a series of representative frameworks for comparison. From the quantitative results displayed in Table 5, we find OrientMask is both faster and more accurate when compared against those having similar input resolutions like YOLACT and CenterMask, and some methods that aim at simplified mask representation like PolarMask and MEInst. We admit that OrientMask falls behind some non-real-time approaches with either higher input image resolutions or more complicated pipelines. When considering the twice or even three times faster in inference, it seems reasonable to accept some sacrifice in accuracy.

Since real-time instance segmentation is the main mo-

Method	Backbone	Size	FPS	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
Mask R-CNN* [9]	ResNet-50	800×	18.5	37.5	59.3	40.2	21.1	39.6	48.3
BlendMask [4]	ResNet-50	800×	17.6	37.0	58.9	39.7	17.3	39.4	52.5
CondInst [29]	ResNet-50	800×	17.9	37.8	59.1	40.5	21.0	40.3	48.7
SOLO [33]	ResNet-50	800×	11.2	36.8	58.6	39.0	15.9	39.5	52.1
PolarMask [36]	ResNet-101	800×	12.3*	32.1	53.7	33.1	14.7	33.8	45.3
MEInst [39]	ResNet-101	800×	12.8*	33.9	56.2	35.4	19.8	36.1	42.3
YOLACT [2]	ResNet-101	550	38.5	29.8	48.5	31.2	9.9	31.3	47.7
CenterMask [35]	DLA-34	512	25.2*	33.1	53.8	34.9	13.4	35.7	48.8
OrienMask	Darknet-53	544	42.7	34.8	56.7	36.4	16.0	38.2	47.8

Table 5. **Quantitative results** on COCO `test-dev`. We compare OrienMask with some typical frameworks. The speed is evaluated on 1k samples with the same platform except for those marked with ‘*’ whose statistics are inferred from their publications. Mask R-CNN model comes from Detectron2. The input size is square by default and the notation ended with ‘×’ means the length of shorter side.

Method	Backbone	Size	Space	FPS	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
YOLACT [2]	Darknet-53	550	4.52M	45.9	28.9	46.9	30.3	9.8	30.9	47.3
YOLACT++ [3]	ResNet-50	550	9.04M	40.8	33.7	52.7	35.5	11.9	36.6	54.6
BlendMask [4]	ResNet-50	550×	18.44M	35.6	34.5	54.7	36.5	14.4	37.7	52.1
MEInst [39]	ResNet-50	512	1.41M	28.1	32.2	53.9	33.0	13.9	34.4	48.7
SOLO-Lite [33]	ResNet-50	512×	2.95M	29.7	33.0	53.9	34.9	11.5	37.0	51.5
SOLOv2-Lite [34]	ResNet-50	448×	7.25M	39.6	33.7	53.3	35.6	11.3	36.9	55.4
OrienMask	Darknet-53	544	1.27M	41.9	34.5	56.0	35.8	16.8	38.5	49.1

Table 6. **Real-time instance segmentation** on COCO `val2017`. With source codes and trained weights released by their authors, all methods are evaluated in our platform. ‘Space’ means the memory occupation of feature maps in megabytes that used to construct masks.

tivation of our work, we further compare OrienMask with state-of-the-art methods capable of real-time inference on COCO `val2017`. All models adopt relatively shallow backbones along with small input resolutions, which makes the comparison fair and persuasive. As shown in Table 6, Our method surpasses YOLACT with 5.6 AP at the cost of 4.0 fps slower. Except for that, OrienMask serves as the leading method in speed comparison and outperforms most counterparts in the mask AP metric. It reaches a good balance between efficiency and accuracy. We also calculate the memory occupation of top feature maps that used for mask construction, and record the results in the ‘space’ column. The input resolution of all methods is assumed to be fixed as 544×544 for brevity. The statistics manifest that our method occupies the least memory resources to construct masks, which proves its success in reducing redundancy while maintaining good mask quality.

4.3. Discussions

In this subsection, we analyze some underlying properties of our method from a qualitative view. The advantages and limitations are both covered.

Orientation Maps We pick two predicted orientation maps to unearth the mechanism of our mask representation. As shown in Figure 4, the attention of each orientation map is

grabbed by objects with the specific anchor size. Two kids and a skateboard are assigned to two orientation maps based on their sizes instead of categories. We display the gradient maps of both directions and their pixel-wise sum, which exhibit the obvious difference around instance boundaries. It can also be observed that two components of gradient maps concentrate on different regions of instances, *i.e.*, the left and right parts are highlighted in G_x while the top and bottom parts are accentuated in G_y . Combining these two complementary orientation maps together, the complete contours of objects can be depicted and then all inner pixels for each instance are safely gathered. These visualized patterns along with high-quality predicted masks verify the effectiveness of our orientation-based mask representation.

Qualitative Results As displayed in Figure 5, our method performs well in separating adjacent instances and precisely delineating their masks. Getting rid of RoI cropping and directly collecting foreground pixels according to the vectors in orientation maps, our mask construction procedure has much tolerance to inaccurate bounding box predictions. Meanwhile, our method also performs well in some complex object overlapping scenarios, especially when one or more small objects locates on a large one. This is illustrated in several images of Figure 5, such as persons wearing ties, food placed on the table, baseball gloves at the front of peo-




Figure 4. **Orientation maps (O) and their gradients (G)** for different anchor sizes at two directions (x and y), where the brighter colors mean the larger signed values. Masks (M) constructed by each orientation map are filled with distinct colors and displayed separately.




Figure 5. **Qualitative results** of OrienMask on COCO *test-dev*. Predictions with confidence greater than 0.3 are displayed. Our model is able to handle most complex scenarios with satisfactory precision. Two typical failure cases are also shown in the last column.

ple and so on. Thanks to the instance grouping mechanism, objects matched with different anchor sizes do not disturb each other and their masks are completely preserved. Moreover, we do not introduce any pixel-level categorical information for OrienHead. The class-agnostic orientation maps are proved to be qualified for recovering masks with satisfactory quality, no matter what categories they belong to.

Failure Cases Although OrienMask works well for most cases, we observe two typical failures as shown in the last column of Figure 5. The first case appears when two instances with the same category and similar size heavily overlap. Orientation maps cannot distinguish them because pixels of both masks point to almost the same base position. The second failure happens due to the severe confrontation of some background pixels between instances, especially when the base positions locate close to their mask boundaries. For example, two giraffes in the bottom right corner of Figure 5 both tend to push the intermediate part outwards, which accidentally makes some background pixels intrude into another target region by mistake. Overall, these errors caused by the incompleteness of mask representation are unusual and only occur in limited cases.

5. Conclusion

In this work, a real-time instance segmentation framework termed OrienMask is proposed, which integrates discriminative orientation maps with an anchor-based detector. Apart from those centripetal vectors for foreground pixels, we further consider negative samples in orientation maps so that both background filtering and instances separation can be accomplished at the same time. An instance grouping mechanism is also presented and each orientation map specializes in grouped objects with the same anchor size. Given the target regions indicated by predicted boxes, masks can be efficiently constructed from corresponding orientation maps, without the need for explicit foreground predictions. Experiments on COCO show that the proposed OrienMask can reach competitive accuracy under real-time conditions.

Acknowledgment

This work is supported by National Natural Science Foundation of China-Zhejiang Joint Fund for the Integration of Industrialization and Informatization (U1709214), and Key Research & Development Plan of Zhejiang Province (2021C01196).

References

- [1] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.
- [2] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact: Real-time instance segmentation. In *Proceedings of the IEEE international conference on computer vision*, pages 9157–9166, 2019.
- [3] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact++: Better real-time instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [4] Hao Chen, Kunyang Sun, Zhi Tian, Chunhua Shen, Yongming Huang, and Youliang Yan. Blendmask: Top-down meets bottom-up for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8573–8581, 2020.
- [5] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, et al. Hybrid task cascade for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4974–4983, 2019.
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017.
- [7] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [8] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6569–6578, 2019.
- [9] 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.
- [10] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask scoring r-cnn. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6409–6418, 2019.
- [11] Saumya Jetley, Michael Sapienza, Stuart Golodetz, and Philip HS Torr. Straight to shapes: real-time detection of encoded shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6550–6559, 2017.
- [12] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4867–4876, 2020.
- [13] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 734–750, 2018.
- [14] Yuezun Li, Xiao Bian, Ming-ching Chang, Longyin Wen, and Siwei Lyu. Pixel offset regression (por) for single-shot instance segmentation. In *2018 15th IEEE International Conference on Advanced Video and Signal Based Surveillance (AVSS)*, pages 1–6. IEEE, 2018.
- [15] Xiaodan Liang, Liang Lin, Yunchao Wei, Xiaohui Shen, Jianchao Yang, and Shuicheng Yan. Proposal-free network for instance-level object segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2978–2991, 2017.
- [16] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [17] 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 *European conference on computer vision*, pages 740–755. Springer, 2014.
- [18] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8759–8768, 2018.
- [19] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [20] Yiding Liu, Siyu Yang, Bin Li, Wengang Zhou, Jizheng Xu, Houqiang Li, and Yan Lu. Affinity derivation and graph merge for instance segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 686–703, 2018.
- [21] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015.
- [22] Davy Neven, Bert De Brabandere, Marc Proesmans, and Luc Van Gool. Instance segmentation by jointly optimizing spatial embeddings and clustering bandwidth. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8837–8845, 2019.
- [23] David Novotny, Samuel Albanie, Diane Larlus, and Andrea Vedaldi. Semi-convolutional operators for instance segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 86–102, 2018.
- [24] George Papandreou, Tyler Zhu, Liang-Chieh Chen, Spyros Gidaris, Jonathan Tompson, and Kevin Murphy. Personlab: Person pose estimation and instance segmentation with a bottom-up, part-based, geometric embedding model. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 269–286, 2018.
- [25] Sida Peng, Wen Jiang, Huaijin Pi, Xiuli Li, Hujun Bao, and Xiaowei Zhou. Deep snake for real-time instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8533–8542, 2020.
- [26] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object de-

- tection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [27] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [28] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015.
- [29] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. *arXiv preprint arXiv:2003.05664*, 2020.
- [30] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. Fcos: Fully convolutional one-stage object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 9627–9636, 2019.
- [31] Jonas Uhrig, Marius Cordts, Uwe Franke, and Thomas Brox. Pixel-level encoding and depth layering for instance-level semantic labeling. In *German Conference on Pattern Recognition*, pages 14–25. Springer, 2016.
- [32] Jonas Uhrig, Eike Rehder, Björn Fröhlich, Uwe Franke, and Thomas Brox. Box2pix: Single-shot instance segmentation by assigning pixels to object boxes. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 292–299. IEEE, 2018.
- [33] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. Solo: Segmenting objects by locations. *arXiv preprint arXiv:1912.04488*, 2019.
- [34] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic, faster and stronger. *arXiv preprint arXiv:2003.10152*, 2020.
- [35] Yuqing Wang, Zhaoliang Xu, Hao Shen, Baoshan Cheng, and Lirong Yang. Centermask: single shot instance segmentation with point representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9313–9321, 2020.
- [36] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12193–12202, 2020.
- [37] Wenqiang Xu, Haiyang Wang, Fubo Qi, and Cewu Lu. Explicit shape encoding for real-time instance segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5168–5177, 2019.
- [38] Brandon Yang, Gabriel Bender, Quoc V Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. In *Advances in Neural Information Processing Systems*, pages 1307–1318, 2019.
- [39] Rufeng Zhang, Zhi Tian, Chunhua Shen, Mingyu You, and Youliang Yan. Mask encoding for single shot instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10226–10235, 2020.