

RefineMask: Towards High-Quality Instance Segmentation with Fine-Grained Features

Gang Zhang¹ Xin Lu² Jingru Tan³

Jianmin Li¹ Zhaoxiang Zhang⁴ Quanquan Li² Xiaolin Hu^{1*}

¹State Key Laboratory of Intelligent Technology and Systems, Institute for AI, BNRist, Department of Computer Science and Technology, Tsinghua University,

²SenseTime Research, ³Tongji University, ⁴Institute of Automation, CAS & UCAS

Abstract

The two-stage methods for instance segmentation, e.g. Mask R-CNN, have achieved excellent performance recently. However, the segmented masks are still very coarse due to the downsampling operations in both the feature pyramid and the instance-wise pooling process, especially for large objects. In this work, we propose a new method called RefineMask for high-quality instance segmentation of objects and scenes, which incorporates fine-grained features during the instance-wise segmenting process in a multi-stage manner. Through fusing more detailed information stage by stage, RefineMask is able to refine high-quality masks consistently. RefineMask succeeds in segmenting hard cases such as bent parts of objects that are over-smoothed by most previous methods and outputs accurate boundaries. Without bells and whistles, RefineMask yields significant gains of 2.6, 3.4, 3.8 AP over Mask R-CNN on COCO, LVIS, and Cityscapes benchmarks respectively at a small amount of additional computational cost. Furthermore, our single-model result outperforms the winner of the LVIS Challenge 2020 by 1.3 points on the LVIS test-dev set and establishes a new state-of-the-art. Code will be available at <https://github.com/zhanggang001/RefineMask>.

1. Introduction

Instance segmentation [4, 9, 14, 16, 25, 28] is a fundamental but challenging task in computer vision, which aims to assign each pixel into a specific semantic category and differentiate instances in the same category. The recent top-performing methods [4, 9, 14, 16] for instance segmentation generally follow a two-stage paradigm. Taking Mask R-CNN [14] as an example, a powerful detector is first employed to generate high-quality bounding boxes and

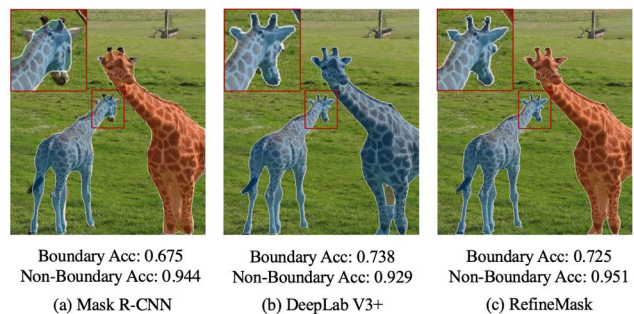


Figure 1: Comparison among three methods for image segmentation: (a) Mask R-CNN, (b) DeepLabV3+, and (c) our proposed **RefineMask**. The segmentation results are indicated by the blue or orange shadows. The boundary accuracy denotes the average accuracy of the predicted masks in boundary regions, and the non-boundary accuracy is defined similarly (the definition of boundary region can be found in Section 3.3). Both accuracies of the three methods are calculated over the entire COCO dataset [20].

then a parallel segmentation branch is introduced to predict binary mask for each instance inside the bounding box. Specifically, in the latter step, a pooling operation, e.g., RoIAlign [14], is used to extract instance features from the feature pyramid [19], then pixel-wise classification is performed based on the output features of the mask head.

Despite the strong abilities provided by the powerful object detector [23] to locate and distinguish instances, Mask R-CNN loses the image details which are indispensable for high-quality instance segmentation task, see Figure 1 (a). The loss of details is mainly due to two factors. Firstly, the features fed into the pooling operation are from multiple levels of the feature pyramid, and the higher-level features usually result in coarser spatial resolution. For these high-level features, it is hard to preserve details when mapping the mask prediction back to input space. Secondly, the pooling operation itself further reduces the spatial size of the features into a smaller size, e.g. 7×7 or 14×14 , which

*Corresponding Author

also causes information loss.

In contrast to instance segmentation, the goal of semantic segmentation is to classify each pixel into a fixed set of categories without differentiating object instances. Since semantic segmentation does not need the extreme high-level features to distinguish large instances, it can make full use of the high-resolution features, *e.g.* P2 in the feature pyramid. Many recently proposed semantic segmentation methods [7, 8, 15, 27] take advantages of high-resolution features to generate high-quality semantic representation and successfully segment sharp object boundaries. These methods have higher prediction accuracy on boundary regions of objects than the two-stage instance segmentation methods, as shown in Figure 1 (b). Moreover, it is obvious that there is no necessity to utilize any instance-wise pooling operation, *e.g.* RoIAlign, to extract instance features in semantic segmentation, further alleviating loss of details.

Our main idea in this work is to perform instance segmentation by keeping the strong ability of current two-stage methods for distinguishing instances and supplementing the lost details with fine-grained features during the instance-wise segmenting process. To achieve this goal, we propose a new framework named RefineMask. RefineMask builds a new semantic head on the highest resolution feature map in the feature pyramid to generate fine-grained semantic features. These fine-grained features are used to supplement the lost details in the instance-wise segmenting process. Different from Mask R-CNN, RefineMask uses a multi-stage refinement strategy in the mask head. Specifically, after the RoI-Align operation, it gradually up-scales the prediction size and incorporates the fine-grained features to alleviate the loss of details for high-quality instance mask prediction. Moreover, RefineMask uses a boundary-aware refinement strategy to focus on the boundary regions for predicting more accurate boundaries. Through fusing more fine-grained features iteratively and focusing on the boundary regions explicitly, RefineMask is able to consistently refine higher quality masks. As shown in Figure 1 (c), RefineMask outputs much higher quality segmentation results than Mask R-CNN and obtains comparable details as state-of-the-art semantic segmentation methods, especially in hard regions such as object boundaries.

We evaluated our method on different datasets for instance segmentation and achieved significant improvements consistently. Without bells and whistles, RefineMask outperformed Mask R-CNN by 2.6, 3.4 and 3.8 points on COCO, LVIS and Cityscapes validation sets respectively. Evaluated under the finer LVIS annotations, RefineMask trained on COCO achieved 4.1 points AP improvements over Mask R-CNN, and the gain reached 6.0 points for large objects. Furthermore, our single-model result surpassed the winner of the LVIS Challenge 2020 [24] by 1.3 points on the LVIS test-dev set and established a new state-of-the-art.

2. Related Work

Instance segmentation. The dominant methods for instance segmentation [4, 9, 14, 16] often utilize a powerful detector to generate bounding boxes and then categorize each pixel inside the bounding box as a foreground or background pixel. However, as these methods rely on a pooling operation, *e.g.* RoIAlign [14], to extract canonical-size features from the feature pyramid [19] for each instance, which loses many details, it’s hard for the segmenters to predict high quality instance masks, especially for large objects. In this work, we perform instance segmentation by incorporating information extracted from the fine-grained features stage by stage to supplement the lost details.

Multi-stage refinement. Multi-stage refinement is widely used to improve performance in object detection [26, 30] and image segmentation [4, 16, 17, 22]. Cascade R-CNN [30] uses a sequence of detectors to regress precise bounding boxes. Deep Layer Cascade [17] treats a single deep model as a cascade of several sub-models to promote semantic segmentation. HTC [4] designs an intertwined cascade mask head to boost performance for object detection and instance segmentation, but with too much computational cost. PointRend [16] performs point-based predictions at the blurred areas iteratively for high-quality image segmentation. Different from PointRend, we refine entire objects together. SharpMask [22] shares similar motivation with us, but it focuses on the object proposal generation.

Methods using semantic segmentation. Semantic segmentation is also utilized as a supplement to instance segmentation in some recent methods [1, 3, 4, 6]. HTC [4] introduces a semantic branch to provide spatial contexts for both object detection and instance segmentation. Mask Lab [6] combines the semantic branch with a direction branch to distinguish instances inside the bounding box. HRNet [27] generates high-resolution representation via a well-designed structure, which can also benefit instance segmentation. However, all these methods still suffer from the loss of details, as they incorporate the semantic features before the instance-wise pooling operations. In this work, we supplement the lost details iteratively with the fine-grained semantic features in the refinement process.

Boundary-aware segmentation. Obtaining sharp boundaries is essential for high-quality image segmentation, and lots of work [9, 13, 18, 29] make attempt on this. Poly-Transform [18] utilizes a deforming network to transform the polygons generated by the results of existing segmentation methods to better fit the object boundaries. SegFix [29] presents a post-processing scheme to improve the boundary quality by replacing the originally unreliable predictions of boundary pixels with the predictions of interior pixels. BMask R-CNN [9] enhances the mask features by introducing an additional contour supervision to better align with

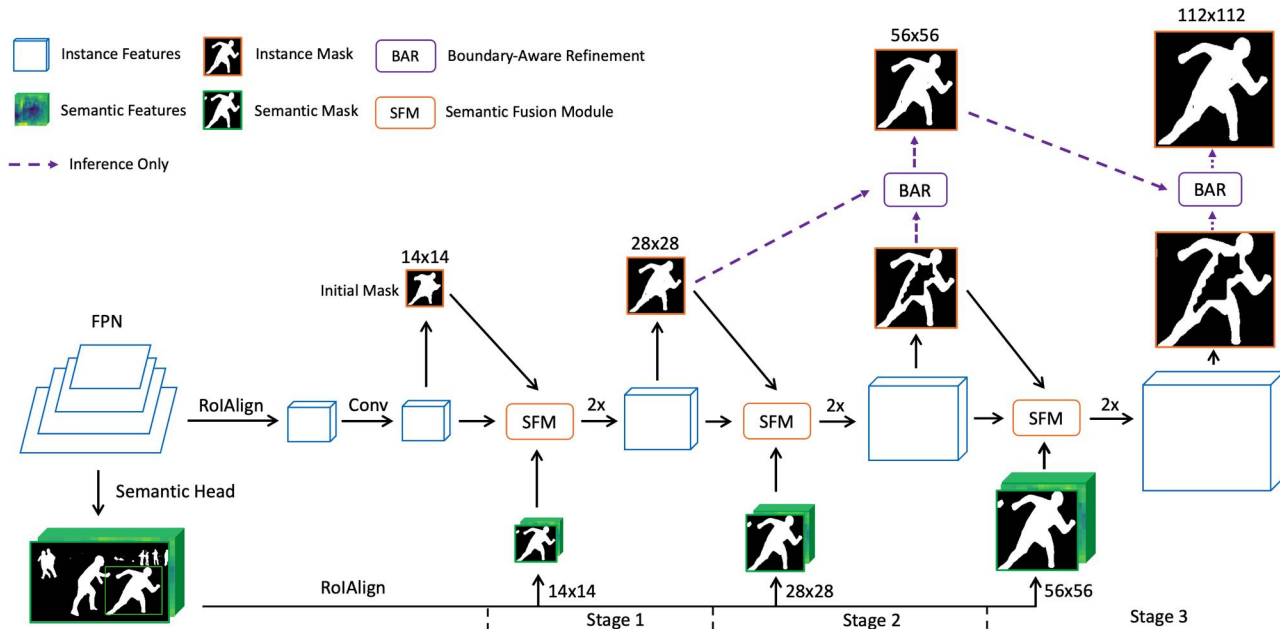


Figure 2: **Framework of RefineMask.** Based on FPN [19], a mask head (the middle row) parallel to the detection head (omitted for clarity) is introduced to perform instance segmentation in a multi-stage manner, and a semantic head is attached to P2 to generate fine-grained features. Each stage has a Semantic Fusion Module (SFM) to fuse the instance features obtained from its preceding stage and the semantic features pooled from the output of the semantic head, also receiving the instance mask and the semantic mask as a guide. Moreover, a boundary-aware refinement (BAR) strategy is proposed to focus on the boundary regions for predicting more accurate boundaries in later stages.

object boundaries. Unlike existing methods, we facilitate the instance boundaries through focusing on the boundary regions explicitly in later stages of the refinement process.

3. RefineMask

An overview of RefineMask is shown in Figure 2. Based on the powerful detector FPN [19], RefineMask relies on two small network modules, i.e. the semantic head and the mask head (the middle row, right of the FPN), to perform high-quality instance segmentation.

The semantic head takes the highest resolution feature map from the feature pyramid as input and performs semantic segmentation. The output of the semantic head keeps the same resolution as the input without using spatial compression operations such as downsampling. The fine-grained features (see the definition in Section 3.1) generated by the semantic head are utilized to facilitate instance segmentation in the mask head.

The mask head performs instance segmentation in a multi-stage fashion. At each stage, the mask head incorporates the semantic features and the semantic mask extracted from the fine-grained features and increases the spatial size of the features for finer instance mask prediction. In addition to that, a boundary-aware refinement strategy is proposed in the mask head to explicitly focus on the boundary

regions for predicting more crisp boundaries. We delve into details of each component in the following subsections.

3.1. Semantic Head

The semantic head is a fully convolutional neural network attached to P2 (the highest resolution feature map of FPN). It consists of four convolutional layers to extract the *semantic features* of the entire image, and a binary classifier to predict the probability of each pixel belonging to the foreground. Under the supervision of binary cross-entropy loss, it predicts a high-resolution *semantic mask* for the entire image. We define *fine-grained features* as the union of the semantic features and the semantic mask. These fine-grained features are further used to supplement the lost details in the mask head for high-quality mask prediction.

3.2. Mask Head

The mask head is a fully convolutional instance segmentation branch. In the mask head, features extracted by the 14×14 RoIAlign operation are first fed into two 3×3 convolutional layers to generate *instance features*. After that, a 1×1 convolutional layer is adopted to predict the *instance mask* like Mask R-CNN but the spatial size of the mask is only 14×14 . This coarse mask is served as an initial mask for later refinement stages.

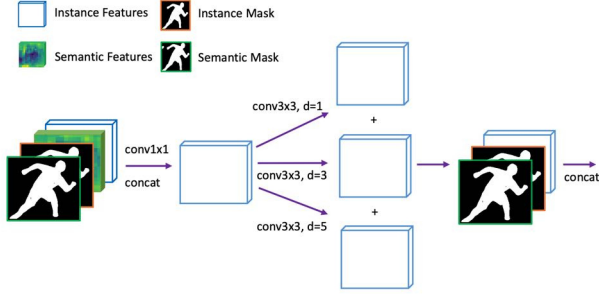


Figure 3: **Semantic Fusion Module.** Four input parts are first concatenated and compressed by a 1×1 convolutional layer, then three parallel convolutional layers with different dilations are used to fuse the features and masks.




Figure 4: **An illustration of object boundary region.** d_{ij} is defined as the Euclidean distance from pixel p_{ij} to its nearest pixel on the mask contour (the white line).

Multi-stage refinement. After the above process, we obtain a coarse instance mask. Next, a multi-stage refinement procedure is proposed to refine the mask in an iterative manner. Inputs of each stage consists of four parts, *i.e.* the instance features and the instance mask obtained from its preceding stage, the semantic features and the semantic mask pooled from the output of the semantic head. A Semantic Fusion Module (SFM) is proposed to integrate these inputs and the fused features are then up-scaled to higher spatial size (see details later). The mask head runs this refinement procedure iteratively and outputs a high-quality instance mask of resolution up to 112×112 .

Note that before being up-scaled to higher spatial size, the fused features in SFM are compressed with a 1×1 convolutional layer to halve its channels. Therefore, although the spatial size of features grows larger and larger, the additional computational cost introduced is quite low.

Semantic Fusion Module. In order to better integrate the fine-grained features, we design a simple fusion module called *Semantic Fusion Module (SFM)* to make sure that each neuron in the mask head perceive its surrounding context, as shown in Figure 3. It concatenates four input parts of each stage mentioned above, following a 1×1 convolutional layer to fuse these features and reduce the channel di-

mension. After that, three parallel 3×3 convolutional layers with different dilations are used to fuse information around a single neuron while keeping the local details. Finally, the instance mask and the semantic mask are concatenated with the fused features again, as a guide for later prediction.

3.3. Boundary-Aware Refinement

For the purpose of predicting accurate boundaries, we propose a boundary-aware refinement strategy to focus on the boundary regions.

Definition of boundary region. Let M^k denotes the binary instance mask of stage k , and the spatial size of the mask can be formulated as $14 \cdot 2^k \times 14 \cdot 2^k$, where $k=1, 2, 3$ (Figure 2). The boundary region of M^k is defined as the region consisting of pixels whose distance to the mask contour is less than \hat{d} pixels. We introduce a binary mask B^k to represent the boundary region of M^k , and B^k can be formulated as follows:

$$B^k(i, j) = \begin{cases} 1, & \text{if } d_{ij} \leq \hat{d} \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

where (i, j) denotes position of the pixel p_{ij} in M^k , and d_{ij} is the Euclidean distance from pixel p_{ij} to its nearest pixel on the mask contour. An illustration is shown in Figure 4. For efficient implementation, we design a convolutional operator to approximate the calculation of boundary regions (details can be found in Appendix). As objects have different scales, we first resize the instance mask into a fixed size, *e.g.* 28×28 in stage 1 and 56×56 in stage 2, and then calculate the boundary mask.

Training. The first stage predicts a complete instance mask with a size of 28×28 . In the two subsequent stages whose output sizes are 56×56 and 112×112 , only certain boundary regions are trained with supervised signals. These regions are determined by both the ground-truth mask and the predicted mask of its *preceding stage*:

$$R^k = f_{\text{up}}(B_G^{k-1} \vee B_P^{k-1}) \quad (2)$$

where f_{up} denotes the bilinear upsampling operation with scale factor of 2, B_G^{k-1} denotes the boundary region of the ground-truth mask in stage $k-1$, B_P^{k-1} denotes the boundary region of the predicted mask in stage $k-1$, and \vee denotes the union of above two boundary regions. The training loss \mathcal{L}^k for the k -th stage ($k=2, 3$) with an output size of $S_k \times S_k$ can be defined as follows:

$$\mathcal{L}^k = \frac{1}{\delta_n} \sum_{n=0}^{N-1} \sum_{i=0}^{S_k-1} \sum_{j=0}^{S_k-1} R_{nij}^k \cdot l_{nij} \quad (3)$$

$$\delta_n = \sum_{n=0}^{N-1} \sum_{i=0}^{S_k-1} \sum_{j=0}^{S_k-1} R_{nij} \quad (4)$$

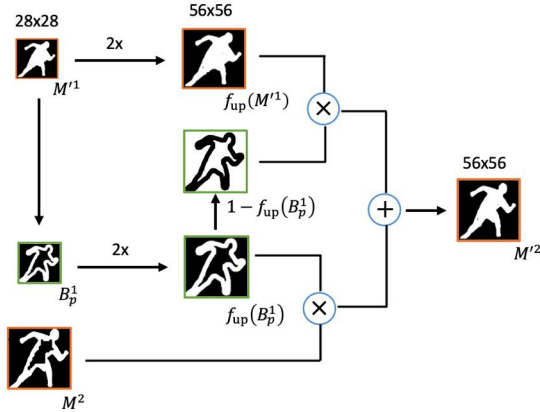


Figure 5: **The inference process of boundary-aware refinement (the second stage).** M'^1 and M'^2 are inputs of the BAR module, B_p^1 is the boundary region of M'^1 , and M'^2 is the output mask of the second stage. The two operators with \times and $+$ denote pixel-wise multiplication and pixel-wise addition respectively.

where N is the number of instances, l_{nij} is a binary cross-entropy loss at position (i, j) for instance n .

Inference. For each instance, the first stage outputs a coarse and complete mask M'^1 with size 28×28 , also generating its boundary mask B_p^1 (omitted in Figure 2). The rule to generate the finer and complete instance mask M'^k (the final output of stage k) can be formulated as follows:

$$M'^1 = M^1 \quad (5)$$

$$M'^k = f_{up}(B_p^{k-1}) \otimes M'^k + (1 - f_{up}(B_p^{k-1})) \otimes f_{up}(M'^{k-1}) \quad (6)$$

where \otimes denotes pixel-wise multiplication. Figure 5 displays the inference process of the second stage. We repeat this process until getting the finest mask.

3.4. Implementation Details

We adopted Mask R-CNN [14] as our baseline, and replaced the default mask head with our proposed multi-stage refinement head, and there were three refinement stages in the mask head by default.

All hyper-parameters were kept the same as Mask R-CNN implemented in MMDetection [5] except the new designed parts. We used the loss defined in Equation 3 for the last two refinement stages. For the semantic head and the other mask prediction stages, we used the average binary cross-entropy loss. Losses for the other parts, including the RPN and the detection head, were kept the same as Mask R-CNN. Loss weights for the initial mask prediction stage and the three refinement stages were set as 0.25, 0.5, 0.75 and 1.0 respectively. To balance the losses between the detection head and the mask head, the loss weight for

the detection head was set as 2.0, including both the classification and regression loss. \hat{d} was set as 2 for training and 1 for inference. In addition, all models presented in the ablation experiments were trained with $1 \times$ learning schedule. No data augmentations except standard horizontal flipping were used unless otherwise stated.

4. Experiments

We performed extensive experiments on three standard instance segmentation datasets: COCO [20], LVIS [12] and Cityscapes [10]. For all three datasets, we used the standard mask AP metric [20] as the evaluation metric.

COCO has 80 categories with instance-level annotations. Our models were trained on *train2017*. Following [16], we also report AP*, which evaluates the COCO categories using LVIS annotations. As the LVIS annotations have significantly higher quality than COCO, it can better reflect improvements in mask quality. Note that the results for AP* were from the same models trained on COCO.

LVIS is a long-tail instance segmentation dataset consisting of 1203 categories, having more than 2 million high-quality instance mask annotations. It contains about 100k, 20k, 20k images for training, validation and test respectively.

Cityscapes is a real-world dataset that consists of 2975, 500, 1525 images with resolution of 2048×1024 for training, validation and test respectively. It contains 8 semantic categories for instance segmentation task.

4.1. Main Results

We first evaluated our RefineMask with different backbones and different learning schedules on COCO *val2017* (Table 1). RefineMask outperformed Mask R-CNN by a large margin under various configurations. Without bells and whistles, RefineMask achieved 2.6 points AP improvements over the Mask R-CNN baseline using ResNet-50-FPN as the backbone network. When evaluating the same COCO results using the 80 categories subset of LVIS annotations, RefineMask surpassed Mask R-CNN by 4.1 points AP*. This better demonstrates the effectiveness of RefineMask in predicting high-quality instance masks.

We also present runtime of RefineMask in Table 2. Compared to Mask R-CNN, RefineMask achieved significant improvement at a small amount of extra computational cost. Comparisons with PointRend [16] and HTC [4] showed the superiority of RefineMask in both accuracy and speed.

4.2. Ablations Experiments

We conducted extensive ablation experiments on COCO *val2017* to analyze RefineMask.

Different number of stages. We compared models with different number of stages. For each additional stage, the

Method	Backbone	Schedule	AP	AP*	AP _S *	AP _M *	AP _L *
Mask R-CNN	R50-FPN	1×	34.7	36.8	22.6	43.7	52.0
RefineMask	R50-FPN	1×	37.3	40.9	24.1	48.8	58.0
Mask R-CNN	R50-FPN	2×	35.4	37.7	22.8	44.7	53.8
RefineMask	R50-FPN	2×	37.9	41.5	24.5	48.7	59.8
Mask R-CNN	R101-FPN	1×	36.1	38.4	22.8	46.0	54.6
RefineMask	R101-FPN	1×	38.6	41.7	24.9	49.5	59.9
Mask R-CNN	R101-FPN	2×	36.6	39.3	23.5	46.8	56.6
RefineMask	R101-FPN	2×	38.8	42.3	24.7	50.2	61.7
Mask R-CNN [‡]	X101-FPN	3×	39.4	41.8	27.2	49.0	57.7
RefineMask [‡]	X101-FPN	3×	41.5	45.3	28.6	53.1	62.8

Table 1: **Comparison with Mask R-CNN on COCO val2017.** Models with [‡] were trained with 3× schedule using multi-scale training with shorter side range [640, 800].

Model	AP	AP*	Runtime (fps)
Mask R-CNN	34.7	36.8	15.7
PointRend	35.6	38.7	11.4
HTC	37.4	40.7	4.4
RefineMask	37.3	40.9	11.4

Table 2: **Efficiency of RefineMask.** We reimplemented PointRend [16] and HTC [4] for fair comparison. All models were trained with 1× schedule using R50-FPN as the backbone network. The inference time was measured on a single Tesla V100 GPU.

Stages	Output size	AP	AP*	AP _S *	AP _M *	AP _L *
1	28×28	35.7	38.3	23.1	45.6	54.4
2	56×56	36.6	40.3	23.3	47.9	57.3
3	112×112	37.3	40.9	24.1	48.8	58.0
4	224×224	37.1	41.0	24.3	48.5	58.7

Table 3: **Different number of stages.** Model with more refinement stages has larger output size.

output size is twice larger than its preceding stage. The results are shown in Table 3. Models with more stages (also larger output size) obtained significant performance improvements and the large objects benefited most from that. The model with three stages and the model with four stages obtained comparable performance, but the former one obviously had lower computational cost.

Effectiveness of SFM. We compared the SFM with other three fusion modules (Table 4). Our proposed SFM performed much better than the other three fusion modules, as it has the largest receptive field and can better capture the context information for each single neuron.

Effectiveness of fine-grained features. We analyzed effectiveness of the fine-grained features by removing the semantic head in RefineMask. The results are shown in Table 5. With fine-grained features, RefineMask brought an improvement of 1.6 points AP and 2.5 points AP*, which indicates the usefulness of fine-grained features for high-quality mask prediction. In addition, objects with large

Fusion module	AP	AP*	AP _S *	AP _M *	AP _L *
1 single 1×1 Conv	35.8	38.7	22.5	46.1	55.3
1 single 3×3 Conv	36.3	39.9	23.4	47.1	57.2
3 parallel 3×3 Convs	36.7	39.9	23.4	47.2	57.6
SFM	37.3	40.9	24.1	48.8	58.0

Table 4: **Different fusion modules.** Comparison of different designs of fusion module.

Semantic head	AP	AP*	AP _S *	AP _M *	AP _L *
✓	35.7	38.4	23.1	45.7	54.3
✓	37.3	40.9	24.1	48.8	58.0

Table 5: **Effectiveness of fine-grained features.** Without the semantic head, the mask head only relies on the instance features to predict instance masks.

Multi-stage refinement	AP	AP*	AP _S *	AP _M *	AP _L *
✓	36.3	39.5	23.1	46.9	56.2
✓	37.3	40.9	24.1	48.8	58.0

Table 6: **Effectiveness of multi-stage refinement.** Only the last stage was supervised without multi-stage refinement.

BAR	AP	AP*	AP _S *	AP _M *	AP _L *
✓	36.9	40.0	23.6	47.4	56.8
✓	37.3	40.9	24.1	48.8	58.0

Table 7: **Effectiveness of boundary-aware refinement.** BAR denotes boundary-aware refinement. Each stage predicted a complete mask for each instance without boundary-aware refinement.

scales benefited more from fine-grained features, indicated by the large AP_L* improvement (54.3→58.0).

Effectiveness of multi-stage refinement. We designed an ablation experiment to prove the necessity of multi-stage refinement by removing supervision of all previous stages. The result is shown in Table 6. With multi-stage refinement, RefineMask improved the AP* by 1.4 points and the AP_L* by 1.8 points, indicating that multi-stage refinement is important to predict high-quality instance masks.

Method	Backbone	AP _{dev}	AP*	AP _S *	AP _M *	AP _L *	Runtime (fps)
BMask R-CNN [9]	R50-FPN	35.9	-	-	-	-	-
Mask R-CNN [†]		35.7	37.7	22.8	44.7	53.8	15.7
PointRend [†]		36.8	39.7	22.9	46.7	57.5	11.4
HTC [4]		38.4	-	-	-	-	4.4
HTC [†]		38.6	41.3	25.0	48.5	58.8	4.4
RefineMask		38.2	41.5	24.5	48.7	59.8	11.4
RefineMask[‡]		40.2	43.4	27.5	50.6	60.7	11.4
Mask R-CNN [14]		R101-FPN	35.7	-	-	-	-
BMask R-CNN [9]	37.7		-	-	-	-	-
Mask R-CNN [†]	37.1		39.3	23.5	46.8	56.6	13.5
PointRend [†]	38.2		41.4	24.7	49.0	59.8	10.0
HTC [4]	39.7		-	-	-	-	3.9
HTC [†]	39.7		42.5	26.2	50.4	60.4	3.9
RefineMask	39.4		42.3	24.7	50.2	61.7	9.6
RefineMask[‡]	41.2		44.6	27.7	52.5	63.2	9.6
Mask R-CNN [14]	X101-FPN	37.1	-	-	-	-	-
Mask R-CNN ^{†‡}		39.6	41.8	27.2	49.0	57.7	8.0
PointRend ^{†‡}		41.1	44.4	27.8	52.0	62.0	6.7
HTC [4]		41.2	-	-	-	-	3.4
HTC [†]		41.3	44.1	27.2	51.9	61.5	3.4
RefineMask		41.0	43.6	25.8	51.8	62.2	6.6
RefineMask[‡]		41.8	45.3	28.6	53.1	62.8	6.6

Table 8: **Comparisons of single-model results on COCO val2017 and test-dev.** [†] denotes our implementation. AP_{dev} denotes the evaluation results on *test-dev*, and AP* denotes the evaluation results on COCO val2017 using the 80 categories subset of LVIS. Note that HTC employs the same detection head as Cascade R-CNN [30] and uses the extra COCO-Stuff [2] annotations. Models with [‡] were trained with 3× schedule using multi-scale training with shorter side range [640, 800].

Effectiveness of boundary-aware refinement. We also conducted an ablation experiment to analyze effectiveness of the boundary-aware refinement. Without boundary-aware refinement, each stage of RefineMask predicts complete masks. Here we only used the output from the last stage, which had the largest output size. The results are shown in Table 7. With boundary-aware refinement, RefineMask improved the AP* by 0.9 points. Although the output size becomes higher in later stages, pixels far away from the object contour cannot benefit more from this, as these pixels generally belongs to the easy non-boundary regions.

4.3. Comparison with Previous Methods

We present single-model results on both COCO val2017 (AP*) and test-dev (AP_{dev}) in Table 8 to compare RefineMask with previous methods. Compared with Mask R-CNN, RefineMask improved the performance by 3.8, 3.0, 3.5 points AP* with ResNet-50-FPN, ResNet-101-FPN and ResNext-101-FPN as the backbone network respectively. Without utilizing the more powerful detection head in Cascade R-CNN [30] and the extra COCO-Stuff [2] data annotations, RefineMask still achieved comparable performance as HTC, with a much faster inference time. RefineMask also outperformed the PointRend [16] under various configurations with comparable speed. As [16] did not release the experimental results on COCO test-dev, we reimplemented

Method	Backbone	AP	AP _r	AP _c	AP _f
Mask R-CNN	R50-FPN	22.1	10.1	21.7	30.0
RefineMask	R50-FPN	25.5	14.2	24.3	31.7
		+3.4	+4.1	+2.6	+1.7

Table 9: **Results on LVISv1.0 validation set.** All Models were trained with 1× schedule and the hyper-parameters were kept the same as [12].

it following the same configurations except that we used the pytorch-style ImageNet [11] pretrained models and did not use multi-scale data augmentation during training for fair comparison unless otherwise stated. Our reimplemented version had similar performance with [16].

4.4. Experiments on LVIS

We first evaluated our RefineMask on the LVIS validation set. To save training memory, we replaced the default class-specific classifier in the last stage with a class-agnostic classifier. The results are shown in Table 9. RefineMask achieved 3.4 points AP improvement compared with the Mask R-CNN baseline, which is larger than that on COCO due to the finer annotations.

We then managed to obtain a good result on the LVIS test-dev dataset based on RefineMask. Specifically, we combined the mask head of the model presented in [24] with our multi-stage refinement head. Following the rule of the




Figure 6: COCO example result tuples from Mask R-CNN [14] (the top row), RefineMask (the middle row) and the ground truth masks (the bottom row). RefineMask predicted masks with substantially higher quality than Mask R-CNN, even better than the ground truth masks, especially on the sharp boundaries.

Method	AP	AP _r	AP _c	AP _f
IvisTraveler [24]	41.2	31.9	40.4	46.4
IvisTraveler + RefineMask	42.5	33.5	41.5	47.7
	+1.3	+1.6	+1.1	+1.3

Table 10: Results on LVISv1.0 test-dev set. The first row is result of the winner of the LVIS challenge 2020.

Method	Backbone	AP	AP _S	AP _M	AP _L
PointRend [16]	R50-FPN	35.8	-	-	-
Mask R-CNN	R50-FPN	33.8	12.0	31.5	51.8
RefineMask	R50-FPN	37.6	14.6	34.0	58.1
		+3.8	+2.6	+2.5	+6.3

Table 11: Results on Cityscapes validation set. All models were trained on the fine annotations with 64 epochs, using multi-scale training with shorter side range [800, 1024].

LVIS Challenge 2020, we did not use any extra data with human labels. Our final single-model result outperformed the winner of LVIS challenge 2020 by 1.3 points (Table 10).

4.5. Experiments on Cityscapes

We also evaluated RefineMask on Cityscapes (Table 11). As the experimental results on Cityscapes have high variance, we report the median of three runs. RefineMask obtained larger improvements on this dataset than on both COCO and LVIS. This further demonstrates the superiority of our approach for high-quality mask prediction.

4.6. Qualitative results.

We show some visualization examples from COCO in Figure 6. RefineMask predicted masks with substantially higher quality than Mask R-CNN, especially for the hard regions, such as the tail of the airplane (the second column), the feet of the horse (the third column), and so on. As the COCO annotations are much coarser than LVIS and Cityscapes, RefineMask even yielded more accurate masks than the ground truth masks, which further indicates the necessity of a more stringent evaluation metric (AP*).

5. Conclusion

In this work, we propose a multi-stage framework called RefineMask towards high-quality instance segmentation. Previous two stage methods are struggling to predict accurate masks, as they depend on the pooling operation, *e.g.* RoIAlign, to extract instance features from the feature pyramid, which loses details required for predicting crisp boundaries. To alleviate this problem, RefineMask refines instance masks by incorporating fine-grained features iteratively during the instance-wise segmentation process. We believe RefineMask can serve as a strong baseline for high-quality instance segmentation.

Acknowledgement. This work was supported in part by the National Key Research and Development Program of China (No. 2017YFA0700904), and the National Natural Science Foundation of China (Nos. 61836014, U19B2034, 62061136001 and 61620106010).

References

- [1] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. YOLACT: Real-time instance segmentation. In *ICCV*, 2019. 2
- [2] Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. COCO-Stuff: Thing and stuff classes in context. In *CVPR*, 2018. 7
- [3] Hao Chen, Kunyang Sun, Zhi Tian, Chunhua Shen, Yongming Huang, and Youliang Yan. BlendMask: Top-down meets bottom-up for instance segmentation. In *CVPR*, 2020. 2
- [4] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. Hybrid task cascade for instance segmentation. In *CVPR*, 2019. 1, 2, 5, 6, 7
- [5] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open MMLab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5
- [6] 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 *CVPR*, 2018. 2
- [7] 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. *TPAMI*, 2017. 2
- [8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *ECCV*, 2018. 2
- [9] Tianheng Cheng, Xinggang Wang, Lichao Huang, and Wenyu Liu. Boundary-preserving Mask R-CNN. In *ECCV*, 2020. 1, 2, 7
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 5
- [11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009. 7
- [12] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 5, 7
- [13] Zeeshan Hayder, Xuming He, and Mathieu Salzmann. Boundary-aware instance segmentation. In *CVPR*, 2017. 2
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollar, and Ross Girshick. Mask R-CNN. In *ICCV*, 2017. 1, 2, 5, 7, 8
- [15] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollar. Panoptic feature pyramid networks. In *CVPR*, 2019. 2
- [16] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. PointRend: Image segmentation as rendering. In *CVPR*, 2020. 1, 2, 5, 6, 7, 8
- [17] Xiaoxiao Li, Ziwei Liu, Ping Luo, Change Loy Chen, and Xiaoou Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *CVPR*, 2017. 2
- [18] Justin Liang, Namdar Homayounfar, Wei-Chiu Ma, Yuwen Xiong, Rui Hu, and Raquel Urtasun. PolyTransform: Deep polygon transformer for instance segmentation. In *CVPR*, 2020. 2
- [19] Tsung-Yi Lin, Piotr Dollar, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 1, 2, 3
- [20] 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 *ECCV*, 2014. 1, 5
- [21] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *CVPR*, 2016. 10
- [22] Pedro O. Pinheiro, Tsung-Yi Lin, Ronan Collobert, and Piotr Dollár. Learning to refine object segments. In *ECCV*, 2016. 2
- [23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *TPAMI*, 2015. 1
- [24] Jingru Tan, Gang Zhang, Hanming Deng, Changbao Wang, Lewei Lu, Quanquan Li, and Jifeng Dai. 1st place solution of lvis challenge 2020: A good box is not a guarantee of a good mask. *arXiv preprint arXiv:2009.01559*, 2020. 2, 7, 8
- [25] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *ECCV*, 2020. 1
- [26] Thang Vu, Hyunjun Jang, Trung X Pham, and Chang D Yoo. Cascade RPN: Delving into high-quality region proposal network with adaptive convolution. In *NeurIPS*, 2019. 2
- [27] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2019. 2
- [28] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. SOLO: Segmenting objects by locations. In *ECCV*, 2020. 1
- [29] Yuhui Yuan Yuan, Jingyi Xie, Xilin Chen, and Jingdong Wang. SegFix: Model-agnostic boundary refinement for segmentation. In *ECCV*, 2020. 2
- [30] Cai Zhaowei and Vasconcelos Nuno. Cascade R-CNN: Delving into high quality object detection. In *CVPR*, 2018. 2, 7