

# MP-Former: Mask-Piloted Transformer for Image Segmentation

Hao Zhang<sup>1,4\*†</sup>, Feng Li<sup>1,4\*†</sup>, Huaizhe Xu<sup>1,4</sup>, Shijia Huang<sup>2,4</sup>,  
Shilong Liu<sup>3,4</sup>, Lionel M. Ni<sup>5,1</sup>, Lei Zhang<sup>4‡</sup>

<sup>1</sup>The Hong Kong University of Science and Technology.

<sup>2</sup>The Chinese University of Hong Kong.

<sup>3</sup>Dept. of CST., BNRist Center, Institute for AI, Tsinghua University.

<sup>4</sup>International Digital Economy Academy (IDEA).

<sup>5</sup>The Hong Kong University of Science and Technology (Guangzhou).

## Abstract

We present a mask-piloted Transformer which improves masked-attention in Mask2Former for image segmentation. The improvement is based on our observation that Mask2Former suffers from inconsistent mask predictions between consecutive decoder layers, which leads to inconsistent optimization goals and low utilization of decoder queries. To address this problem, we propose a mask-piloted training approach, which additionally feeds noised ground-truth masks in masked-attention and trains the model to reconstruct the original ones. Compared with the predicted masks used in mask-attention, the ground-truth masks serve as a pilot and effectively alleviate the negative impact of inaccurate mask predictions in Mask2Former. Based on this technique, our MP-Former achieves a remarkable performance improvement on all three image segmentation tasks (instance, panoptic, and semantic), yielding +2.3AP and +1.6mIoU on the Cityscapes instance and semantic segmentation tasks with a ResNet-50 backbone. Our method also significantly speeds up the training, outperforming Mask2Former with half of the number of training epochs on ADE20K with both a ResNet-50 and a Swin-L backbones. Moreover, our method only introduces little computation during training and no extra computation during inference. Our code will be released at <https://github.com/IDEA-Research/MP-Former>.

## 1. Introduction

Image segmentation is a fundamental problem in computer vision which includes semantic, instance, and panoptic

segmentation. Early works design specialized architectures for different tasks, such as mask prediction models [2, 15] for instance segmentation and per-pixel prediction models [25, 28] for semantic segmentation. Panoptic segmentation [18] is a unified task of instance and semantic segmentation. But universal models for panoptic segmentation such as Panoptic FPN [18] and K-net [34] are usually sub-optimal on instance and semantic segmentation compared with specialized models.

Recently, Vision Transformers [29] has achieved tremendous success in many computer vision tasks, such as image classification [11, 24] and object detection [1]. Inspired by DETR’s set prediction mechanism, Mask2Former [7] proposes a unified framework and achieves a remarkable performance on all three segmentation tasks. Similar to DETR, it uses learnable queries in Transformer decoders to probe features and adopts bipartite matching to assign predictions to ground truth (GT) masks. Mask2Former also uses predicted masks in a decoder layer as attention masks for the next decoder layer. The attention masks serve as a guidance to help next layer’s prediction and greatly ease training.

Despite its great success, Mask2Former is still an initial attempt which is mainly based on a vanilla DETR model that has a sub-optimal performance compared to its later variants. For example, in each decoder layer, the predicted masks are built from scratch by dot-producting queries and feature map without performing layer-wise refinement as proposed in deformable DETR [37]. Because each layer build masks from scratch, the masks predicted by a query in different layers may change dramatically. To show this inconsistent predictions among layers, we build two metrics  $mIoU-L^i$  and  $Util^i$  to quantify this problem and analyze its consequences in Section 3. As a result, this problem leads to a low utilization rate of decoder queries, which is especially severe in the first few decoder layers. As the attention masks predicted from an early layer are usually inaccurate, when serving as a guidance for its next layer, they will lead to

\*Equal contribution.

†This work was done when Hao Zhang and Feng Li were interns at IDEA.

‡Corresponding author.

sub-optimal predictions in the next layer.

Since detection is a similar task to segmentation, both aiming to locate instances in images, we seek inspirations from detection to improve segmentation. For example, the recently proposed denoising training [19] for detection has shown very effective to stabilize bipartite matching between training epochs. DINO [33] achieves a new SOTA on COCO detection built upon an improved denoising training method. The key idea in denoising training is to feed noised GT boxes in parallel with learnable queries into Transformer decoders and train the model to denoise and recover the GT boxes.

Inspired by DN-DETR [19], we propose a mask-piloted (MP) training approach. We divide the decoder queries into two parts, an MP part and a matching part. The matching part is the same as in Mask2Former, feeding learnable queries and adopting bipartite matching to assign predictions to GT instances. In the MP part, we feed class embeddings of GT categories as queries and GT masks as attention masks into a Transformer decoder layer and let the model to reconstruct GT masks and labels. To reinforce the effectiveness of our method, we propose to apply MP training for all decoder layers. Moreover, we also add point noises to GT masks and flipping noises to GT class embeddings to force the model to recover GT masks and labels from inaccurate ones.

We summarize our contributions as follows.

1. We propose a mask-piloted training approach to improving masked-attention in Mask2Former, which suffers from inconsistent and inaccurate mask predictions among layers. We also develop techniques including multi-layer mask-piloted training with point noises and label-guided training to further improve the segmentation performance. Our method is computationally efficient, introducing no extra computation during inference and little computation during training.
2. Through analyzing the predictions of Mask2Former, we discover a failure mode of Mask2Former—inconsistent predictions between consecutive layers. We build two novel metrics to measure this failure, which are layer-wise query utilization and mean IoU. We also show that our method improves the two metrics by a large margin, which validates its effectiveness on alleviating inconsistent predictions.
3. When evaluating on three datasets, including ADE20k, Cityscapes, and MS COCO, our MP-Former outperforms Mask2Former by a large margin. Besides improved performance, our method also speeds up training. Our model exceeds Mask2Former on all three segmentation tasks with only half of the number of training steps on ADE20k. Also, our model trained for 36 epochs is comparable with Mask2Former trained for 50 epochs on instance and panoptic segmentation on MS COCO.

## 2. Related Work

Generally, image segmentation can be divided into three tasks including instance segmentation, semantic segmentation, and panoptic segmentation with respect to different semantics.

**Traditional Segmentation models** Traditionally, researchers develop specialized models and optimization objectives for each task. Instance segmentation is to predict a set of binary masks and their associated categories. Previous methods often predict masks based on bounding boxes produced by detection models. Mask R-CNN [15] builds segmentation upon detection model Faster R-CNN [27] by adding a mask branch in parallel with the detection branch. HTC [2] further proposes to interleave the two branches and add mask information flow to improve segmentation performance. Semantic segmentation focus on category-level semantics without distinguishing instances. Previous models widely formulate it into a per-pixel classification problem. The pioneering work, FCN [25] generates a label for each pixel to solve this problem. Many follow-up works continue with this idea and design more precise pixel-level classification models [3, 4]. Panoptic segmentation [18, 30] is a combination of the two segmentation tasks above to segment both the foreground instances ("thing") and background semantics ("stuff").

**Vision Transformers for segmentation** Transformer-based segmentation models emerge with DETR (DEtection TRANSformer) [1]. Since that, Transformer [29] has been widely used in many detection [1, 19, 23, 33] and segmentation models [7, 8, 32]. Generally, these models adopt a set prediction objective with bipartite matching to query features of interest. Inspiring progress has been made in many specialized query-based architectures for instance segmentation [13], semantic segmentation [32, 35], panoptic segmentation [21]. Notably, some recent works unify the three segmentation tasks into one model [7, 8, 20, 34]. Among them, Mask2Former [7] proposes masked attention and achieves remarkable performance on all three segmentation tasks. Following DETR, Mask2Former is an end-to-end Transformer for image segmentation. It is composed of a backbone, a pixel decoder, and a Transformer decoder. The pixel decoder is an encoder-only deformable Transformer that takes multi-scale image features as input and output processed multi-scale features. The Transformer decoder takes learnable queries as input to attend to the feature maps output by the pixel decoder. Each decoder layer attends to a feature map of one scale following a coarse-to-fine manner. The predicted masks of one decoder layer will feed to the next layer as the attention masks, which enables the following decoder to focus on more meaningful regions. The design of attention masks eases the training and improves performance.

**Ease Training for Vision transformers** Vision transformers have an advantage over the convolution-based methods due

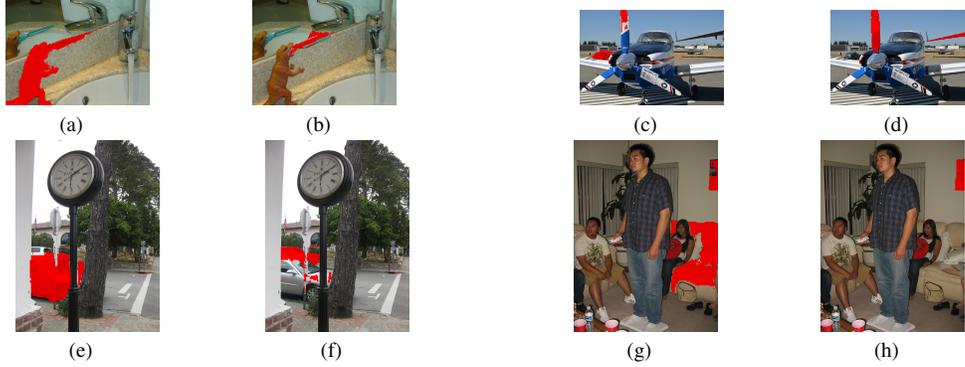


Figure 1. We visualize layer-wise predictions of Mask2Former and show 4 pairs of failure cases. Each pair is the predictions of the same query in adjacent decoder layers. The red regions are the predicted masks. These cases show that the predictions of a query may change dramatically between consecutive layers.

to their global attention—they can see all parts of the image. However, it is hard to train because globally searching for an object is a hard task. This phenomenon exists in both detection and segmentation. In detection, DETR suffers from slow convergence requiring 500 training epochs for convergence. Recently, researchers have dived into the meaning of the learnable queries [23, 26, 31, 37]. They either express the queries as reference points or anchor boxes. [19, 33] proposed to add noised ground truth boxes as positional queries for denoising training and they speed up detection greatly. In segmentation, Mask2Former proposed mask attention which makes training easier and speeds up convergence when compared with MaskFormer. We have found some common points for these works that speed up training for vision Transformers. Firstly, they usually give clearer meaning to learnable queries to reduce ambiguity. Secondly, they give local limitations to cross attention and reduce the scope for the model to search objects. Therefore, we believe that giving guidance to cross attention leads to good segmentation performance. Our method is similar to DN-DETR [19] in introducing GT annotations to the Transformer decoder but ours is totally different from it in both the problems we solve and the method we take. Firstly, our method aims to address the inconsistent prediction problem which leads to the low utilization of queries, while DN-DETR aims to solve unstable matching among epochs. Secondly, we use GT masks as the attention masks, while DN-DETR feeds noised GT boxes as decoder queries. Thirdly, noises are optional in our method, but DN-DETR does not work without noises.

### 3. Failure Cases and Solution

**Failure cases:** We view the mask prediction in Mask2Former as a mask refinement process through Transformer decoder layers. By taking predictions from a layer as the attention masks of the next layer, we expected to obtain refined predictions from the next layer. Generally, the predicted masks are assigned to the GT masks by bipartite matching and

refined in a coarse-to-fine procedure. Therefore, keeping refinement consistency in different layers is the key to refining the predicted masks consistently. However, when we visualize the predictions of Mask2Former layer by layer, we found a severe inconsistency problem among the predictions from the same query between consecutive decoder layers. As shown in Fig. 1, the predictions of adjacent layers differ significantly. We define a metric named layer-wise mean Intersection-over-Union (IoU) denoted as  $mIoU-L^i$  to quantify this inconsistency. We denote the predicted masks from the  $i$ -th decoder layer as  $\mathbf{M}^i = \{M_0^i, M_1^i, \dots, M_{N-1}^i\}$ , where  $N$  is the number of decoder queries. Note that  $M^0$  means the predictions before queries enter the first Transformer decoder layer. We have

$$mIoU-L^i = \sum_{n=0}^{N-1} IoU(M_n^{i-1}, M_n^i) \quad (1)$$

As shown in Table 1, the  $mIoU-L$  is quite low, especially for the first few layers, which indicates inconsistent mask prediction results between adjacent layers. Such an inconsistent result means a query may be assigned to different instances (or background for panoptic and semantic segmentation) during layer-by-layer updates. To evaluate the inconsistent matching problem, we define another metric named as layer-wise query utilization denoted as  $Util^i$  for the  $i$ -th layer, which means the ratio of the matched queries in the  $i$ -th layer that matches exactly the same instance as in the last layer. We denote the ground truth objects as  $\mathbf{T} = \{T_0, T_1, T_2, \dots, T_{O-1}\}$  where  $O$  is the number of ground truth objects. After bipartite matching, we compute an index vector  $\mathbf{V}^i = \{V_0^i, V_1^i, \dots, V_{N-1}^i\}$  to store the matching result of layer  $i$  as follows.

$$V_n^i = \begin{cases} o, & \text{if } M_n^i \text{ matches } T_o \\ -1, & \text{if } M_n^i \text{ matches nothing} \end{cases} \quad (2)$$

We define the query utilization of layer  $i$  as the proportion of GT masks that match with the same query in layer  $i$  and

Layer Number		1	2	3	4	5	6	7	8	9
Mask2Former	mIoU-L <sup>i</sup> (%)	23.0	50.3	59.5	69.7	70.1	69.9	78.6	76.3	76.0
	Util <sup>i</sup> (%)	38.1	65.5	70.8	73.5	78.5	81.1	83.2	83.2	100
MP training	mIoU-L <sup>i</sup> (%)	51.5	83.2	86.6	96.2	96.0	93.8	97.6	97.9	96.3
	Util <sup>i*</sup> (%)	94.0	93.8	94.6	93.2	94.0	94.8	94.0	94.0	100
Multi-layer MP	mIoU-L <sup>i</sup> (%)	51.3	87.8	88.8	97.5	97.5	96.2	99.0	98.8	98.2
	Util <sup>i*</sup> (%)	94.2	95.4	97.3	97.6	97.3	97.0	98.2	98.2	100

Table 1. The mIoU-L<sup>i</sup> and Util<sup>i</sup> of Mask2Former and the proposed methods. The mIoU-L<sup>i</sup> for a query is taken as the IoU of the predicted masks in the (i - 1)-th and i-th decoder layers. The mIoU-L<sup>i</sup> takes the average over all queries. Note that the index of decoder layers starts from 1, so "prediction of the 0-th decoder" means the prediction before entering Transformer decoder. The Util<sup>i</sup> is the ratio of GT masks that match with the same query in i-th layer and the last layer. \* means that We test Util<sup>i</sup> with bipartite matching to assign predictions in MP part to GT masks. Note that when directly assigning predictions in MP part to corresponding GT masks with hard assignment, Util<sup>i</sup> is always 100%.

the last layer, which is calculated as

$$\text{Util}^i = \frac{1}{O} \sum_{j=0}^N \mathbb{1}(V_n^i = V_n^l \text{ and } V_n^i \neq -1), \quad (3)$$

where  $\mathbb{1}(\cdot)$  is the indicator function.  $\mathbb{1}(x) = 1$  if  $x$  is true and 0 otherwise. The query utilization of layer  $i$  for the whole data set is averaged over the utilization numbers for all images. Table 1 shows that the query utilizations of the first few layers are too low, which indicate many matched queries in early layers are ended up not being used in the last layer. The low mIoU-L<sup>i</sup> and Util<sup>i</sup> have potential consequences such as instant optimization goals and low training efficiency.

Our observations based on these two metrics motivate us to develop a more effective training method to make the matching more consistent among layers.

**Solution:** To address the inconsistent prediction problem, inspired by DN-DETR, we proposed a mask-piloted training method. We feed GT masks as attention masks and expect these queries with GT masks can focus on their corresponding GT masks and not be distracted by other instances. As shown in row 2 and 3, Table 1, both mIoU-L<sup>i</sup> and Util<sup>i</sup> are improved by a large margin when adopting mask-piloted training. We also find that only applying mask-piloted (MP) training in the first layer is not enough. Fig 3(a)(b) shows that inconsistent prediction exists even when given GT masks. Therefore, we add a GT mask to all decoder layers to reinforce MP training. As shown in Table 1, multi-layer MP training further improves both mIoU-L<sup>i</sup> and Util<sup>i</sup> especially in latter decoder layers.

## 4. MP-Former

### 4.1. Overview

Our method is a training method for improving Mask2Former. Our MP-Former only introduces minor changes in the Mask2Former pipeline during training as shown in the red-line part in Fig. 2. Same as Mask2Former, our model is composed of a backbone, a pixel decoder, and

a Transformer decoder. Our improvements are on the Transformer decoder. We divide decoder queries into two parts, a mask-piloted (MP) part and a matching part. The matching part is the same as Mask2Former. The MP part feeds GT masks and GT class embeddings as input and the predictions are directly assigned to the corresponding GT masks without bipartite matching. Our method includes three components which are multi-layer MP training, mask noises, and label-guided training.

### 4.2. Multi-layer MP training

In order to address the inconsistent prediction problem, we propose a multi-layer mask-piloted (MP) training method. In the MP part, we feed an additional set of queries and masks into Transformer where the queries and masks are GT class embeddings and GT masks. On the loss side, we independently assign the outputs of the MP part and the matching part with the GT instances. The predictions of the MP part are directly assigned to their corresponding GT masks and those in the matching part are assigned with bipartite graph matching. The losses for both parts are the same, following the loss design in Mask2Former. We also propose to add GT masks in multiple layers and show that this can further improve the performance with experiments in Table 8. Since we follow Mask2Former to use feature maps of different resolutions in different decoder layers, we interpolate GT masks into different resolutions when applied to different layers.

### 4.3. Noised masks

Besides feeding GT masks to alleviate the inconsistency problem, we further propose to use noised GT masks to make mask refinement more robust in each decoder. The intuition is simple, as Mask2former refines the mask from an inaccurate mask predicted from the previous decoder. GT masks without noise may be too easy for the task which prevents further refinement. Therefore, we feed noised masks to the decoder and train the model to reconstruct the original ones, which is similar to DN-DETR [19] for box denoising

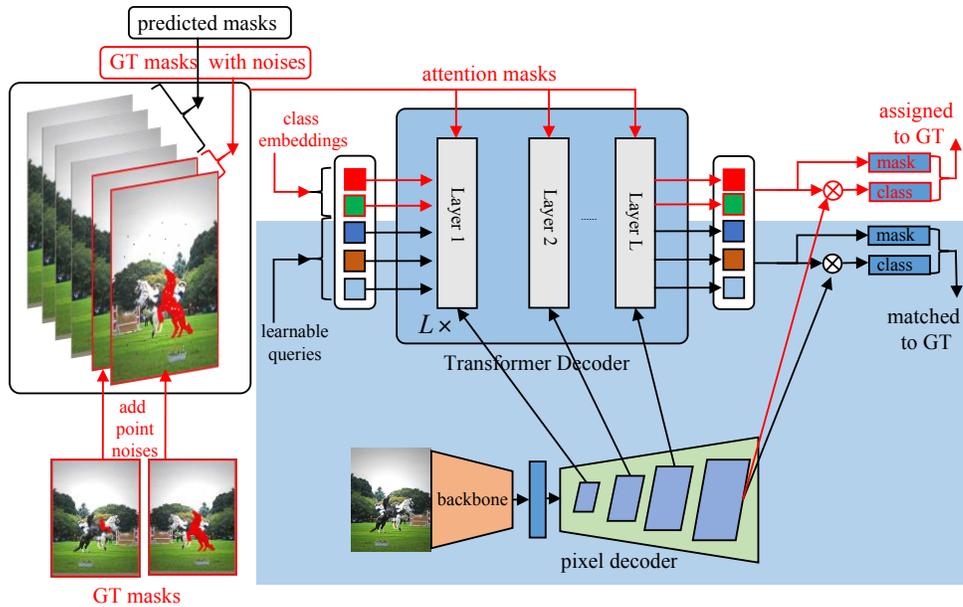


Figure 2. The architecture of our method is the same as Mask2Former (the blue-shaded part), which consists of a backbone, a pixel decoder, and a Transformer decoder. The difference is that we feed extra queries and attention masks which are called the MP part to the Transformer decoder (red-line part in the figure). The MP part contains GT masks as attention masks and GT class embeddings as queries. We feed GT masks into the MP part of all decoder layers. We also add point noises to GT masks and flipping noises to class embeddings which can further improve the performance. Note that this architecture is just for training. In the inference time, the red-line part does not exist, and thus, our pipeline is exactly the same as Mask2Former.

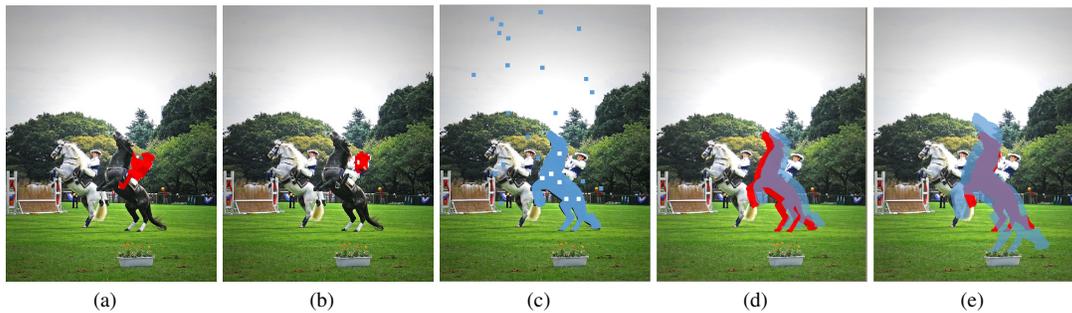


Figure 3. (a)The red region is a given GT mask in first decoder layer which covers the man riding the black horse. (b) The predicted mask of the first decoder layer which only covers the head of the man. The predicted mask will be used as the attention mask of the second layer and may mislead the second layer. (c)(d)(e) are demonstrations of point, shift and scale noises. The blue regions are noised masks and the red regions are the GT masks.

method	backbone	panoptic			instance				semantic mIoU
		PQ	$AP_{pan}^{Th}$	$mIoU_{pan}$	AP	$AP^S$	$AP^M$	$AP^L$	
MaskFormer [8]	R50	34.7	-	-	-	-	-	-	-
Panoptic-DeepLab [6]	SWideRNet [5]	37.9*	-	50.0*	-	-	-	-	-
Mask2Former [7]	R50	39.7	26.5	46.1	26.4	10.4	28.9	43.1	47.2
<b>MP-Former</b>	R50	<b>40.8(+1.1)</b>	<b>27.1</b>	<b>48.3(+2.2)</b>	<b>28.0(+1.6)</b>	10.5	<b>30.7(+1.8)</b>	<b>44.6</b>	<b>48.1</b>
MaskFormer [8]	Swin-L <sup>†</sup>	-	-	-	-	-	-	-	54.1
FaPN-MaskFormer [8,17]	Swin-L <sup>†</sup>	-	-	-	-	-	-	-	55.2
Mask2Former [7]	Swin-L <sup>†</sup>	48.1	34.2	54.5	34.9	16.3	40.0	54.7	56.1
<b>MP-Former</b>	Swin-L <sup>†</sup>	<b>49.4(+1.3)</b>	<b>35.2</b>	<b>56.0</b>	<b>35.3</b>	16.3	<b>40.7</b>	<b>55.5</b>	<b>56.9</b>

Table 2. **Image segmentation results on ADE20K val1.** MP-Former improves the performance on all three segmentation tasks with R50 and Swin-L as backbone. All metrics are evaluated with *single-scale* inference. Backbones pre-trained on ImageNet-22K are marked with <sup>†</sup>. - means the results are not reported for the methods.

training. We have tried three types of noise which are point noise, shift noise, and scale noise as shown in Fig. 3. For

point noise, we randomly sample some points on the mask and flip them from 1 to 0 or from 0 to 1. The number of

method	backbone	panoptic			instance		semantic
		PQ	AP <sub>pan</sub> <sup>Th</sup>	mIoU <sub>pan</sub>	AP	AP50	mIoU
Panoptic-DeepLab [6]	R50	60.3	32.1	78.7	-	-	-
Mask2Former [7]	R50	62.1	37.3	77.5	37.4	61.9	79.4
<b>MP-Former</b>	R50	<b>62.7</b>	<b>39.3(+2.0)</b>	<b>79.6(+2.1)</b>	<b>39.7(+2.3)</b>	<b>65.4(+3.5)</b>	<b>81.0(+1.6)</b>
Mask2Former [7]	Swin-L <sup>†</sup>	66.6	43.6	82.9	43.7	71.4	83.3
<b>MP-Former</b>	Swin-L <sup>†</sup>	<b>67.5</b>	<b>45.8(+2.2)</b>	<b>83.5</b>	<b>44.9(+1.2)</b>	<b>72.4</b>	<b>83.9</b>

Table 3. **Image segmentation results on Cityscapes v1.1.** We report results on all three segmentation tasks with R50 or Swin-L as backbone. All metrics are evaluated with *single-scale* inference. Backbones pre-trained on ImageNet-22K are marked with <sup>†</sup>

sampled points is proportional to the area of a GT mask. For shift noise, we randomly shift a mask instance vertically and horizontally while keeping the center of the instance within its GT bounding box. For the scaling noise, we enlarge or shrink the instance by a random ratio. A comparison of three types of noises is shown in Table 8. Through experiments, we find point noises improve the performance while shifting and scaling noises do not work for our method. Therefore, we adopt point noises in our implementation. We use a noise ratio  $\lambda_p$  to control how many noise points are added. Denote a binary GT mask as  $M = [m_{i,j}]_{H \times W}$  where  $m_{i,j} = 1$  means the pixel belongs to the mask. The number of noise points is randomly sampled from 0 to  $\lambda_p \times \sum_{i,j} m_{i,j}$ . Since we apply mask-piloted training to all decoder layers, we independently add different noises to GT masks in different layers.

#### 4.4. Label-guided training

We use the class embeddings of GT categories as queries in the MP part. The reason why we use class embeddings is that queries will dot-product with image features and an intuitive way to distinguish instances is to use their categories. Note that we only use class embeddings in the first layer and do not replace predicted queries with class embeddings in subsequent layers. We also have a classification loss in the MP part. Since the classification task becomes very easy given GT class embeddings, we add flipping noises in class embeddings. We randomly sample GT class embeddings and let them randomly flip to the class embeddings of other categories. For example,  $E_0, E_1, \dots, E_K$  denotes the class embeddings of  $K$  categories. An image contains  $N$  instances which belong to Categories  $C^0, C^1, \dots, C^N$ . So their default class embeddings should be  $E_{C^0}, E_{C^1}, \dots, E_{C^N}$ . To add flipping noise, we enumerate all  $N$  instances. For  $i$ -th instance, we generate a random number  $r_i$  in  $[0, 1]$  to decide whether to flip the class embedding  $E_{C^i}$  or not. If  $r_i < \lambda_l$ , we randomly sample another class embedding in  $E_0, E_1, \dots, E_K$  to replace  $E_{C^i}$ .  $\lambda_l$  denotes the proportion of the sampled class embeddings.

## 5. Experiments

We conduct extensive experiments to compare our method with Mask2Former and other segmentation methods

on all three segmentation tasks on three datasets to show the effectiveness of our method. We also conduct several ablation studies to show the effectiveness of each component in our method.

### 5.1. Implementation details

**Model:** Since our method is an improved training method based on Mask2Former, we adopt an identical model setting as Mask2Former except for the mask-piloted (MP) part. In the MP part, we adopt 100 GT masks with point noises. We adopt noise ratios  $\lambda_p = 0.2$  and  $\lambda_l = 0.2$  for point noises on masks and flipping noises on class embeddings. We adopt ResNet50 [16] and Swin-L [24] when evaluating our method on ADE20K and Cityscapes. On COCO2017, we use ResNet50 and ResNet101 as the backbone.

**Dataset and metrics:** We evaluate MP-Former on three challenging datasets: Cityscapes [9, 10] and ADE20K [36] for all three segmentation tasks and COCO 2017 [22] dataset for panoptic segmentation and instance segmentation. For instance segmentation, we evaluate the mask AP [22] on instances that are denoted as "thing" categories in datasets. For semantic segmentation, we evaluate mean Intersection-over-Union (mIoU) over all categories including both foreground and background. For panoptic segmentation, we evaluate the results with the panoptic quality (PQ) metric [18].

### 5.2. Main Results

We trained on three challenging datasets ADE20K, Cityscapes, and MS COCO with different backbones. Our method outperforms Mask2Former on all three segmentation tasks (instance, panoptic, and semantic) under the same setting, which demonstrates its effectiveness and generalization ability to different datasets.

**Performance on ADE20K.** On ADE20K, we report our performance trained for 160K steps following Mask2Former. As shown in Table 2, our method exceeds Mask2Former and other baselines on all metrics with both R50 and Swin-L. Notably, MP-Former achieves +1.6AP on instance segmentation with a R50 backbone and +1.3PQ on panoptic segmentation with a Swin-L backbone.

**Performance on Cityscapes.** On Cityscapes, we train the model with 90K training steps following Mask2Former. As shown in Table 3, our method consistently outperforms Mask2Former on all three segmentation tasks across differ-

method	backbone	query type	epochs	AP	AP <sup>S</sup>	AP <sup>M</sup>	AP <sup>L</sup>	#params.	FLOPs	fps
MaskFormer [8]	R50	100 queries	300	34.0	16.4	37.8	54.2	45M	181G	19.2
Mask R-CNN [15]	R50	dense anchors	36	37.2	18.6	39.5	53.3	44M	201G	15.2
Mask R-CNN [12, 14, 15]	R50	dense anchors	400	42.5	23.8	45.0	60.0	46M	358G	10.3
Mask2Former [7]	R50	100 queries	50	43.7	23.4	47.2	64.8	44M	226G	9.7
<b>MP-Former</b>	R50	100 queries	50	<b>44.8(+1.1)</b>	<b>24.7</b>	<b>48.1</b>	<b>65.5</b>	44M	226G	9.7
Mask R-CNN [15]	R101	dense anchors	36	38.6	19.5	41.3	55.3	63M	266G	10.8
Mask R-CNN [12, 14, 15]	R101	dense anchors	400	43.7	<b>24.6</b>	46.4	61.8	65M	423G	8.6
Mask2Former [7]	R101	100 queries	50	44.2	23.8	47.7	<b>66.7</b>	63M	293G	7.8
<b>MP-Former</b>	R101	100 queries	50	<b>45.1(+0.9)</b>	<b>24.6</b>	<b>48.7</b>	66.2	63M	293G	7.8
QueryInst [13]	Swin-L <sup>†</sup>	300 queries	50	48.9	30.8	52.6	68.3	-	-	3.3
Swin-HTC++ [2, 24]	Swin-L <sup>†</sup>	dense anchors	72	49.5	31.0	52.4	67.2	284M	1470G	-
Mask2Former [7]	Swin-L <sup>†</sup>	200 queries	100	50.1	29.9	53.9	72.1	216M	868G	4.0
<b>MP-Former (ours)</b>	Swin-L <sup>†</sup>	200 queries	72	<b>50.8(+0.7)</b>	30.3	<b>54.5</b>	<b>72.6</b>	216M	868G	4.0

Table 4. **Instance segmentation on COCO val2017 with 80 categories.** MP-Former outperforms strong baseline Mask2Former without extra computation cost in inference time. For a fair comparison, we only consider single-scale inference and models trained using only COCO train2017 set data. Backbones pre-trained on ImageNet-22K are marked with <sup>†</sup>

method	backbone	query type	epochs	PQ	PQ <sup>Th</sup>	PQ <sup>St</sup>	AP <sup>Th</sup> <sub>pan</sub>	mIoU <sub>pan</sub>	#params.	FLOPs	fps
DETR [1]	R50	100 queries	500+25	43.4	48.2	36.3	31.1	-	-	-	-
MaskFormer [8]	R50	100 queries	300	46.5	51.0	39.8	33.0	57.8	45M	181G	17.6
Mask2Former [7]	R50	100 queries	50	51.4*	57.2	42.7	41.9	61.5	44M	226G	8.6
<b>MP-Former (ours)</b>	R50	100 queries	50	<b>52.2(+0.8)</b>	<b>58.3</b>	<b>42.9</b>	<b>42.5</b>	<b>61.6</b>	44M	226G	8.6
DETR [1]	R101	100 queries	500+25	45.1	50.5	37.0	33.0	-	-	-	-
MaskFormer [8]	R101	100 queries	300	47.6	52.5	40.3	34.1	59.3	64M	248G	14.0
Mask2Former [7]	R101	100 queries	50	52.6	58.5	43.7	42.6	<b>62.4</b>	63M	293G	7.2
<b>MP-Former (ours)</b>	R101	100 queries	50	<b>52.9(+0.3)</b>	<b>58.7</b>	<b>44.0</b>	<b>43.0</b>	61.8	63M	293G	7.2
Max-DeepLab [30]	Max-L	128 queries	216	51.1	57.0	42.2	-	-	451M	3692G	-
MaskFormer [8]	Swin-L <sup>†</sup>	100 queries	300	52.7	58.5	44.0	40.1	64.8	212M	792G	5.2
K-Net [34]	Swin-L <sup>†</sup>	100 queries	36	54.6	60.2	46.0	-	-	-	-	-
Mask2Former	Swin-L <sup>†</sup>	200 queries	100	57.8	64.2	48.1	48.6	67.4	216M	868G	4.0
<b>MP-Former</b>	Swin-L <sup>†</sup>	200 queries	72	<b>58.1(+0.3)</b>	<b>64.4</b>	<b>48.4</b>	<b>48.9</b>	<b>67.6</b>	216M	868G	4.0

Table 5. **Panoptic segmentation on COCO panoptic val2017 with 133 categories.** MP-Former outperforms Mask2Former on COCO panoptic segmentation with both R50 and R101 as backbone. Results that are implemented are marked with \*.

ent backbones. MP-Former also scales well when adopting Swin-L as the backbone. Especially, our method outperforms Mask2Former by +2.3AP on instance segmentation and +1.6mIoU on semantic segmentation with R50.

**Performance on COCO.** On COCO val2017, we follow Mask2Former’s setting to train for 50epochs. As shown in Table 4 and Table 5, MP-Former outperforms both Mask2Former and classical methods such as Mask R-CNN [15] on instance segmentation and panoptic segmentation. For instance segmentation, our method achieves +1.1AP with R50, +0.9AP with R101 and +0.7 with Swin-L [24]. Note that we only train MP-Former for 72 epochs with Swin-L. On panoptic segmentation, our method achieves +0.8PQ with R50 and +0.3PQ with R101. Notably, our method is only a pluggable training method that will not affect the model architecture or inference speed. In addition, it only introduces negligible parameters (class embeddings) during training.

### 5.3. Speed up training

In Table 6, we show our MP-Former can significantly speed up training on ADE20k and COCO. We do not report the results on Cityscapes because this dataset is relatively

small and only needs 90k steps to converge in Mask2Former, which is not necessary for additional speedup. The results show that our method trained with 36 epochs achieved comparable results with Mask2Former trained with 50 epochs with both R50 and R101 as the backbone on both panoptic and instance segmentation. Moreover, our method trained with half number of training steps (80K steps) achieves comparable results with Mask2Former trained with 160K steps. Note that we use the same batch size as Mask2Former. We also report the clock time. The results show that our method reduces the actual training time by large. For a fair comparison, we use the same number of GPUs to run MP-Former and Mask2Former. Because there may exist differences in speed for different machines, we run our method and Mask2Former on the same machine to test the running time.

### 5.4. Ablation Study

We conduct an ablation study on the COCO instance segmentation task with R50 as backbone.

**Effectiveness of each component:** We show the effectiveness of each version of our method trained for 12 epochs and 50 epochs in Table 7, where we show the result of raw Mask2Former and the result after each version of our method

dataset	method	backbone	epochs*	panoptic				instance					semantic	
				PQ	AP <sub>pan</sub> <sup>Th</sup>	mIoU <sub>pan</sub>	clock time/h	AP	AP <sup>S</sup>	AP <sup>M</sup>	AP <sup>L</sup>	clock time/h	mIoU	clock time/h
ADE20k	Mask2Former	R50	160k	39.7	26.5	46.1	52.8	26.4	10.4	28.9	43.1	40.3	47.2	21.3
	MP-Former	R50	<b>80k</b>	39.7	<b>27.0</b>	<b>48.4</b>	34.2	<b>27.4</b>	10.2	<b>30.0</b>	<b>44.5</b>	25.5	<b>47.4</b>	12.3
	Mask2Former	Swin-L <sup>†</sup>	160k	48.1	34.2	54.5	63.2	34.9	16.3	40.0	54.7	51.4	56.1	39.8
	MP-Former	Swin-L <sup>†</sup>	<b>80k</b>	<b>48.5</b>	<b>34.7</b>	<b>56.1</b>	45.7	<b>35.2</b>	15.8	39.6	<b>55.6</b>	26.7	<b>56.9</b>	21.5
COCO	Mask2Former	R50	50	51.4*	<b>41.9</b>	61.5	77.9	43.7	23.4	47.2	64.8	62.5	-	-
	MP-Former	R50	<b>36</b>	<b>51.5</b>	41.8	<b>61.6</b>	59.5	<b>44.1</b>	<b>24.1</b>	<b>47.4</b>	<b>65.3</b>	50.8	-	-
	Mask2Former	R101	50	52.6	42.6	62.4	113.1	44.2	<b>23.8</b>	47.7	<b>66.7</b>	86.0	-	-
	MP-Former	R101	<b>36</b>	52.5	42.6	62.0	84.4	44.2	23.2	<b>47.9</b>	66.6	71.5	-	-

Table 6. **Image segmentation results on ADE20K va1. and COCO va12017** MP-Former speeds up training by large. On ADE20K, MP-Former trained for 80K steps can achieve comparable result with Mask2Former trained for 160K steps with both R50 and Swin-L as backbone. On COCO, MP-Former trained for 36 epochs can achieve comparable result with Mask2Former trained for 50 epochs with both R50 and R101 as backbone. \* means that we adopt number of steps instead of epochs on ADE20k following Mask2Former. We report only single-scale inference results. Backbones pre-trained on ImageNet-22K are marked with <sup>†</sup>

#Row	12 ep (AP mask)	50 ep (AP mask)	Layers for MP training	ep	AP mask	Feature map	ep	AP mask
1. Mask2Former [7]	38.7	43.7	First layer	12	39.4	Coarse to fine	12	40.2
2. Row1+MP training on 1st layer	39.4	44.2	First three layers	12	39.6(+0.2)	All large	12	40.0
3. Row2+Label noises	39.6	44.2	All layers	12	40.2(+0.8)	<b>Assignment in Mask-Piloted part</b>		
4. Row3+Point noises	39.9	44.6	<b>Noise Type</b>			Match	12	39.4
5. Ours (Row4+All-layer MP)	<b>40.2</b>	<b>44.8</b>	No noise	12	39.6	Hard assignment	12	39.6
			Point noise	12	<b>39.9</b>			
			Shifting noise	12	39.6			
			Scaling noise	12	39.4			

Table 7. The effectiveness of each component of our method when trained for 12 epochs and 50 epochs. The table shows that each component of our model effectively improves the results. Label noise does not work when trained for 50 epochs. Especially, our method without noises can improve the results by +0.9AP. With multi-layer MP training, our method achieves +1.5AP on COCO instance segmentation with R50 backbone trained for 12 epochs. "ep" is short for epochs.

Table 8. Top 3 lines shows the effectiveness of multi-layer Mask-Piloted (MP) training. The bottom lines of the table shows that shifting and scaling noise do not work for our method. "ep" is short for epochs.

Table 9. Some hyperparameter tuning: 1.coarse-to-fine manner works better for our method than all-large feature maps. 2.Hard assignment lead to better results than matching for MP part queries. "ep" is short for epochs.

is adopted. Without adding noises, our method achieves an improvement of +0.9AP. With multi-layer MP-training, our method achieves +1.5AP. We also find that label noise does not work when trained for 50 epochs. Note that except for the last line of the table, other lines of our method apply MP training on the first 3 layers.

**Apply MP training on multiple layers:** In Table 8, we show the effectiveness of applying GT masks on multiple layers. The results indicate that more layers of GT masks can lead to better performance. By adding noised GT masks in all layers, MP-Former achieves +0.8AP.

**Add noises to GT masks:** In the last four rows of Table 8, we show the effectiveness of adding different types of noises. Note that without any noise, MP-Former achieves 39.6AP which is corresponding to the 4th row in Table 7. Shifting noise and scaling noise do not work in our method. Point noise achieved +0.3 extra growth.

**Some hyper-parameter tuning:** Mask2Former adopts coarse-to-fine feature maps in decoder layers. We try to adopt the largest feature map (1/8 feature map) in all decoder layers. Table 9 shows that a coarse-to-fine manner is better than using the largest feature map in all layers for our MP-Former. We also try to use matching to assign predictions to GT instances in the MP part. The result shows that matching can also improve the result compared with raw Mask2Former but the hard assignment is better.

## 6. Conclusion

In this paper, we have analyzed failure cases of Mask2Former and found a problem of inconsistent predictions between adjacent layers which potentially leads to inconstant optimization goals and low utilization of queries and therefore limits the performance of Mask2Former. In order to address the problem, we developed a new MP-Former, whose key components include multi-layer mask-piloted training with point noises and label-guided training. Our experimental results showed that the proposed training method effectively mitigates the inconsistent prediction problem. When evaluating on three challenging datasets (ADE20K, Citescapes, and MS COCO), our method achieved improvement by a large margin on all three image segmentation tasks (semantic, instance, and panoptic). In addition, our method only introduces little computation during training and no extra computation in inference time. Finally, our method can also significantly speed up training. Especially, when trained with half number of steps, our method exceeds Mask2Former on all three segmentation tasks on ADE20K with both R50 and Swin-L as the backbones.

**Limitations:** We evaluated MP-Former on Mask2Former, but our method may be useful to some other models such as MaskFormer and HTC. We will take it as our future work.

## References

- [1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 1, 2, 7
- [2] 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4974–4983, 2019. 1, 2, 7
- [3] 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. 2
- [4] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 2
- [5] Liang-Chieh Chen, Huiyu Wang, and Siyuan Qiao. Scaling wide residual networks for panoptic segmentation. *arXiv preprint arXiv:2011.11675*, 2020. 5
- [6] Bowen Cheng, Maxwell D Collins, Yukun Zhu, Ting Liu, Thomas S Huang, Hartwig Adam, and Liang-Chieh Chen. Panoptic-deeplab: A simple, strong, and fast baseline for bottom-up panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12475–12485, 2020. 5, 6
- [7] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. 2022. 1, 2, 5, 6, 7, 8
- [8] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. 2021. 2, 5, 7
- [9] 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 *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 6
- [10] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Scharwächter, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset. In *CVPR Workshop on The Future of Datasets in Vision*, 2015. 6
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [12] Xianzhi Du, Barret Zoph, Wei-Chih Hung, and Tsung-Yi Lin. Simple training strategies and model scaling for object detection. *arXiv preprint arXiv:2107.00057*, 2021. 7
- [13] Yuxin Fang, Shusheng Yang, Xinggang Wang, Yu Li, Chen Fang, Ying Shan, Bin Feng, and Wenyu Liu. Instances as queries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6910–6919, 2021. 2, 7
- [14] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2918–2928, 2021. 7
- [15] 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. 1, 2, 7
- [16] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 6
- [17] Shihua Huang, Zhichao Lu, Ran Cheng, and Cheng He. Fapn: Feature-aligned pyramid network for dense image prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 864–873, 2021. 5
- [18] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9404–9413, 2019. 1, 2, 6
- [19] Feng Li, Hao Zhang, Shilong Liu, Jian Guo, Lionel M Ni, and Lei Zhang. Dn-detr: Accelerate detr training by introducing query denoising. *arXiv preprint arXiv:2203.01305*, 2022. 2, 3, 4
- [20] Feng Li, Hao Zhang, Shilong Liu, Lei Zhang, Lionel M Ni, Heung-Yeung Shum, et al. Mask dino: Towards a unified transformer-based framework for object detection and segmentation. *arXiv preprint arXiv:2206.02777*, 2022. 2
- [21] Zhiqi Li, Wenhai Wang, Enze Xie, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, Tong Lu, and Ping Luo. Panoptic segformer. *arXiv preprint arXiv:2109.03814*, 2021. 2
- [22] 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. 6
- [23] Shilong Liu, Feng Li, Hao Zhang, Xiao Yang, Xianbiao Qi, Hang Su, Jun Zhu, and Lei Zhang. Dab-detr: Dynamic anchor boxes are better queries for detr. *arXiv preprint arXiv:2201.12329*, 2022. 2, 3
- [24] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 1, 6, 7
- [25] 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. 1, 2
- [26] Depu Meng, Xiaokang Chen, Zejia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Con-

- ditional detr for fast training convergence. *arXiv preprint arXiv:2108.06152*, 2021. [3](#)
- [27] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [2](#)
- [28] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [1](#)
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [1](#), [2](#)
- [30] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5463–5474, 2021. [2](#), [7](#)
- [31] Yingming Wang, Xiangyu Zhang, Tong Yang, and Jian Sun. Anchor detr: Query design for transformer-based detector. *arXiv preprint arXiv:2109.07107*, 2021. [3](#)
- [32] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34, 2021. [2](#)
- [33] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel M Ni, and Heung-Yeung Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. *arXiv preprint arXiv:2203.03605*, 2022. [2](#), [3](#)
- [34] Wenwei Zhang, Jiangmiao Pang, Kai Chen, and Chen Change Loy. K-net: Towards unified image segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. [1](#), [2](#), [7](#)
- [35] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6881–6890, 2021. [2](#)
- [36] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019. [6](#)
- [37] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR 2021: The Ninth International Conference on Learning Representations*, 2021. [1](#), [3](#)