

Mask DINO: Towards A Unified Transformer-based Framework for Object Detection and Segmentation

Feng Li^{1,3*†}, Hao Zhang^{1,3*†}, Huaizhe Xu^{1,3}, Shilong Liu^{2,3},
 Lei Zhang^{3‡}, Lionel M. Ni^{1,4}, Heung-Yeung Shum^{1,3}

¹The Hong Kong University of Science and Technology.

²Dept. of CST., BNRist Center, Institute for AI, Tsinghua University.

³International Digital Economy Academy (IDEA).

⁴The Hong Kong University of Science and Technology (Guangzhou).

{fliay, hzhangcx, hxubr}@connect.ust.hk {lius120}@mails.tsinghua.edu.cn {leizhang}@idea.edu.cn {ni, hshum}@ust.hk

Abstract

In this paper we present Mask DINO, a unified object detection and segmentation framework. Mask DINO extends DINO (DETR with Improved Denoising Anchor Boxes) by adding a mask prediction branch which supports all image segmentation tasks (instance, panoptic, and semantic). It makes use of the query embeddings from DINO to dot-product a high-resolution pixel embedding map to predict a set of binary masks. Some key components in DINO are extended for segmentation through a shared architecture and training process. Mask DINO is simple, efficient, and scalable, and it can benefit from joint large-scale detection and segmentation datasets. Our experiments show that Mask DINO significantly outperforms all existing specialized segmentation methods, both on a ResNet-50 backbone and a pre-trained model with SwinL backbone. Notably, Mask DINO establishes the best results to date on instance segmentation (54.5 AP on COCO), panoptic segmentation (59.4 PQ on COCO), and semantic segmentation (60.8 mIoU on ADE20K) among models under one billion parameters. Code is available at <https://github.com/IDEA-Research/MaskDINO>.

1. Introduction

Object detection and image segmentation are fundamental tasks in computer vision. Both tasks are concerned with localizing objects of interest in an image but have different levels of focus. Object detection is to localize objects of interest and predict their bounding boxes and category

labels, whereas image segmentation focuses on pixel-level grouping of different semantics. Moreover, image segmentation encompasses various tasks including instance segmentation, panoptic segmentation, and semantic segmentation with respect to different semantics, e.g., instance or category membership, foreground or background category.

Remarkable progress has been achieved by classical convolution-based algorithms developed for these tasks with specialized architectures, such as Faster RCNN [24] for object detection, Mask RCNN [9] for instance segmentation, and FCN [21] for semantic segmentation. Although these methods are conceptually simple and effective, they are tailored for specialized tasks and lack the generalization ability to address other tasks. The ambition to bridge different tasks gives rise to more advanced methods like HTC [2] for object detection and instance segmentation and Panoptic FPN [14], K-net [33] for instance, panoptic, and semantic segmentation. Task unification not only helps simplify algorithm development but also brings in performance improvement in multiple tasks.

Recently, DETR-like [1] models developed based on Transformers [27] have achieved inspiring progress on many detection and segmentation tasks. As an end-to-end object detector, DETR adopts a set-prediction objective and eliminates hand-crafted modules such as anchor design and non-maximum suppression. Although DETR addresses both the object detection and panoptic segmentation tasks, its segmentation performance is still inferior to classical segmentation models. To improve the detection and segmentation performance of Transformer-based models, researchers have developed specialized models for object detection [15, 18, 32, 35], image segmentation [3, 4, 33], instance segmentation [7], panoptic segmentation [23], and semantic segmentation [12].

Among the efforts to improve object detection, DINO [32] takes advantage of the dynamic anchor box for-

*Equal contribution.

†Work done when Feng Li and Hao Zhang were interns at IDEA.

‡Corresponding author.

mulation from DAB-DETR [18] and query denoising training from DN-DETR [15], and further achieves the SOTA result on the COCO object detection leaderboard for the first time as a DETR-like model. Similarly, for improving image segmentation, MaskFormer [4] and Mask2Former [3] propose to unify different image segmentation tasks using query-based Transformer architectures to perform mask classification. Such methods have achieved remarkable performance improvement on multiple segmentation tasks.

However, in Transformer-based models, the best-performing detection and segmentation models are still not unified, which prevents task and data cooperation between detection and segmentation tasks. As an evidence, in CNN-based models, Mask-RCNN [9] and HTC [2] are still widely acknowledged as unified models that achieve mutual cooperation between detection and segmentation to achieve superior performance than specialized models. Though we believe detection and segmentation can help each other in a unified architecture in Transformer-based models, the results of simply using DINO for segmentation and using Mask2Former for detection indicate that they can not do other tasks well, as shown in Table 1 and 2. Moreover, trivial multi-task training can even hurt the performance of the original tasks. It naturally leads to two questions: 1) *why cannot detection and segmentation tasks help each other in Transformer-based models?* and 2) *is it possible to develop a unified architecture to replace specialized ones?*

To address these problems, we propose Mask DINO, which extends DINO with a mask prediction branch in parallel with DINO’s box prediction branch. Inspired by other unified models [3, 4, 28] for image segmentation, we reuse content query embeddings from DINO to perform mask classification for all segmentation tasks on a high-resolution pixel embedding map (1/4 of the input image resolution) obtained from the backbone and Transformer encoder features. The mask branch predicts binary masks by simply dot-producting each content query embedding with the pixel embedding map. As DINO is a detection model for region-level regression, it is not designed for pixel-level alignment. To better align features between detection and segmentation, we also propose three key components to boost the segmentation performance. First, we propose a unified and enhanced query selection. It utilizes encoder dense prior by predicting masks from the top-ranked tokens to initialize mask queries as anchors. In addition, we observe that pixel-level segmentation is easier to learn in the early stage and propose to use initial masks to enhance boxes, which achieves task cooperation. Second, we propose a unified denoising training for masks to accelerate segmentation training. Third, we use a hybrid bipartite matching for more accurate and consistent matching from ground truth to both boxes and masks.

Mask DINO is conceptually simple and easy to implement under the DINO framework. To summarize, our contri-

butions are three-fold. **1)** We develop a unified Transformer-based framework for both object detection and segmentation. As the framework is extended from DINO, by adding a mask prediction branch, it naturally inherits most algorithm improvements in DINO including anchor box-guided cross attention, query selection, denoising training, and even a better representation pre-trained on a large-scale detection dataset. **2)** We demonstrate that detection and segmentation can help each other through a shared architecture design and training method. Especially, detection can significantly help segmentation tasks, even for segmenting background "stuff" categories. Under the same setting with a ResNet-50 backbone, Mask DINO outperforms all existing models compared to DINO (+0.8 AP on COCO detection) and Mask2Former (+2.6 AP, +1.1 PQ, and +1.5 mIoU on COCO instance, COCO panoptic, and ADE20K semantic segmentation). **3)** We also show that, via a unified framework, segmentation can benefit from detection pre-training on a large-scale detection dataset. After detection pre-training on the Objects365 [26] dataset with a SwinL [20] backbone, Mask DINO significantly improves all segmentation tasks and achieves the best results on instance (54.5 AP on COCO), panoptic (59.4 PQ on COCO), and semantic (60.8 mIoU on ADE20K) segmentation among models under one billion parameters.

2. Related Work

Detection: Mainstream detection algorithms have been dominated by convolutional neural network-based frameworks, until recently Transformer-based detectors [1, 15, 18, 32] achieve great progress. DETR [1] is the first end-to-end and query-based Transformer object detector, which adopts a set-prediction objective with bipartite matching. DAB-DETR [18] improves DETR by formulating queries as 4D anchor boxes and refining predictions layer by layer. DN-DETR [15] introduces a denoising training method to accelerate convergence. Based on DAB-DETR and DN-DETR, DINO [32] proposes several new improvements on denoising and anchor refinement and achieves new SOTA results on COCO detection. Despite the inspiring progress, DETR-like detection models are not competitive for segmentation. Vanilla DETR incorporates a segmentation head in its architecture. However, its segmentation performance is inferior to specialized segmentation models and only shows the feasibility of DETR-like detection models to deal with detection and segmentation simultaneously.

Segmentation: Segmentation mainly includes instance, semantic, and panoptic segmentation. Instance segmentation is to predict a mask and its corresponding category for each object instance. Semantic segmentation requires to classify each pixel including the background into different semantic categories. Panoptic segmentation [14] unifies the instance and semantic segmentation tasks and predicts a mask for

each object instance or background segment. In the past few years, researchers have developed specialized architectures for the three tasks. For example, Mask-RCNN [9] and HTC [2] can only deal with instance segmentation because they predict the mask of each instance based on its box prediction. FCN [21] and U-Net [25] can only perform semantic segmentation since they predict one segmentation map based on pixel-wise classification. Although models for panoptic segmentation [13, 30] unifies the above two tasks, they are usually inferior to specialized instance and semantic segmentation models. Until recently, some image segmentation models [3, 4, 33] are developed to unify the three tasks with a universal architecture. For instance, Mask2Former [3] improves MaskFormer [4] by introducing masked-attention to Transformer. Mask2Former has a similar architecture as DETR to probe image features with learnable queries but differs in using a different segmentation branch and some specialized designs for mask prediction. However, while Mask2Former shows a great success in unifying all segmentation tasks, it leaves object detection untouched and our empirical study shows that its specialized architecture design is not suitable for predicting boxes.

Unified Methods: As both object detection and segmentation are concerned with localizing objects, they naturally share common model architectures and visual representations. A unified framework not only helps simplify the algorithm development effort, but also allows to use both detection and segmentation data to improve representation learning. There have been several previous works to unify segmentation and detection tasks, e.g., Mask RCNN [9], HTC [2], and DETR [1]. Mask RCNN extends Faster RCNN and pools image features from Region Of Interest (ROI) proposed by RPN. HTC further proposes an interleaved way of predicting boxes and masks. However, these two models can only perform instance segmentation. DETR predicts boxes and masks together in an end-to-end manner. However, its segmentation performance largely lags behind other models. According to Table 2, adding DETR’s segmentation head to DINO results in inferior instance segmentation results. How to attain mutual assistance between segmentation and detection has long been an important problem to solve.

3. Mask DINO

Mask DINO is an extension of DINO [32]. On top of content query embeddings, DINO has two branches for box prediction and label prediction. The boxes are dynamically updated and used to guide the deformable attention in each Transformer decoder. Mask DINO adds another branch for mask prediction and minimally extends several key components in detection to fit segmentation tasks. To better understand Mask DINO, we start by briefly reviewing DINO and then introduce Mask DINO.

3.1. Preliminaries: DINO

DINO is a typical DETR-like model, which is composed of a backbone, a Transformer encoder, and a Transformer decoder. The framework is shown in Fig. 1 (the blue-shaded part without red lines). Following DAB-DETR [18], DINO formulates each positional query in DETR as a 4D anchor box, which is dynamically updated through each decoder layer. Note that DINO uses multi-scale features with deformable attention [35]. Therefore, the updated anchor boxes are also used to constrain deformable attention in a sparse and soft way. Following DN-DETR [15], DINO adopts denoising training and further develops contrastive denoising to accelerate training convergence. Moreover, DINO proposes a mixed query selection scheme to initialize positional queries in the decoder and a look-forward-twice method to improve box gradient back-propagation.

3.2. Why a universal model has not replaced the specialized models in DETR-like models?

Remarkable progress has been achieved by Transformer-based detectors and segmentation models. For instance, DINO [32] and Mask2Former [3] have achieved the best results on COCO detection and panoptic segmentation, respectively. Inspired by such progress, we attempted to simply extend these specialized models for other tasks but found that the performance of other tasks lagged behind the original ones by a large margin, as shown in Table 1 and 2. It seems that trivial multi-task training even hurts the performance of the original task. However, in convolution-based models, it has shown effective and mutually beneficial to combine detection and instance segmentation tasks. For example, detection models with Mask R-CNN head [9] is still ranked the first on the COCO instance segmentation. We will take DINO and Mask2Former as examples to discuss the challenges in unifying Transformer-based detection and segmentation.

– **What are the differences between specialized detection and segmentation models?** Image segmentation is a pixel-level classification task, while object detection is a region-level regression task. In DETR-based model, the decoder queries are responsible for these tasks. For example, Mask2Former uses such decoder queries to dot-product the high-resolution feature maps to produce segmentation masks, while DINO uses them to regress boxes. However, as such queries in Mask2Former only have to compare per-pixel similarity with the image features, they may not be aware of the region-level position of each instance. On the contrary, queries in DINO are not designed to interact with such low-level features to learn pixel-level representation. Instead, they encode rich positional information and high-level semantics for detection.

– **Why cannot Mask2Former do detection well?** The Transformer decoder of Mask2Former is designed for seg-

Model	Box AP	Mask AP
Mask2Former	46.2*	43.7
Mask2Former + detection head	21.6	41.3
DINO	50.7	–

Table 1. Simply adding a detection head to Mask2Former results in low detection performance. * indicates the boxes are generated from the predicted masks. The generated boxes from Mask2Former are also inferior to DINO (-4.5 AP). The models are trained for 50 epochs.

mentation tasks and does not suit detection for three reasons. First, its **queries** follow the design in DETR [1] without being able to utilize better positional priors as studied in Conditional DETR [22], Anchor DETR [29], and DAB-DETR [18]. For example, its content queries are semantically aligned with the features from the Transformer encoder, whereas its positional queries are just learnable vectors as in vanilla DETR instead of being associated with a single-mode position¹. If we remove its mask branch, it reduces to a variant of DETR [1], whose performance is inferior to recently improved DETR models. Second, Mask2Former adopts **masked attention** (multi-head attention with attention mask) in Transformer decoders. The attention masks predicted from a previous layer are of high resolution and used as hard-constraints for attention computation. They are neither efficient nor flexible for box prediction. Third, Mask2Former cannot explicitly perform **box refinement** layer by layer. Moreover, its coarse-to-fine mask refinement in decoders fails to use multi-scale features from the encoder. As shown in Table 1, the generated box AP from mask is 4.5 AP worse than DINO and trivial multi-task learning by adding a detection head is not working².

–**Why cannot DETR/DINO do segmentation well?** As shown in Table 2, simply **1)** adding DETR’s segmentation head or **2)** adding Mask2Former’s segmentation head result in inferior performance compared to Mask2Former. We analyze the reasons as follows. **The reason for 1)** is that DETR’s segmentation head is not optimal. The vanilla DETR lets each query embedding dot-product with the smallest feature map to compute attention maps and then upsamples them to get the mask predictions. This design lacks an interaction between queries and larger feature maps from the backbone. In addition, the head is too heavy to use mask auxiliary loss for mask refinement. **The reason for 2)** is that features in improved detection models are not aligned with segmentation. For example, DINO inherits many designs from [18, 32, 35] like query formulation, denoising training, and query selection. However, these components are designed to strengthen region-level representation for detection, which is not optimal for segmentation.

¹We refer the interested readers to discussions in Sec. 3 in DAB-DETR [18]

²We also notice there are issues in official Mask2Former Github (<https://github.com/facebookresearch/Mask2Former/issues/43>) that fail to make Mask2Former work well by adding a detection head.

Model	Box AP	Mask AP
DINO + Mask2Former segmentation head	49.9	40.2
DINO + DETR segmentation head (finetune DINO pretrained on COCO detection)	–	35.8
Mask2Former	–	43.7

Table 2. Simply adding a segmentation head to DINO results in low instance segmentation performance. The predicted masks from DINO are also inferior to Mask2Former (-3.5 AP). The models are trained for 50 epochs.

3.3. Our Method: Mask DINO

Mask DINO adopts the same architecture design for detection as in DINO with minimal modifications. In the Transformer decoder, Mask DINO adds a mask branch for segmentation and extends several key components in DINO for segmentation tasks. As shown in Fig. 1, the framework in the blue-shaded part is the original DINO model and the additional design for segmentation is marked with red lines.

3.4. Segmentation branch

Following other unified models [3, 4, 28] for image segmentation, we perform mask classification for all segmentation tasks. Note that DINO is not designed for pixel-level alignment as its positional queries are formulated as anchor boxes and its content queries are used to predict box offset and class membership. To perform mask classification, we adopt a key idea from Mask2Former [3] to construct a pixel embedding map which is obtained from the backbone and Transformer encoder features. As shown in Fig. 1, the pixel embedding map is obtained by fusing the 1/4 resolution feature map C_b from the backbone with an upsampled 1/8 resolution feature map C_e from the Transformer encoder. Then we dot-product each content query embedding q_c from the decoder with the pixel embedding map to obtain an output mask m .

$$m = q_c \otimes \mathcal{M}(\mathcal{T}(C_b) + \mathcal{F}(C_e)), \quad (1)$$

where \mathcal{M} is the segmentation head, \mathcal{T} is a convolutional layer to map the channel dimension to the Transformer hidden dimension, and \mathcal{F} is a simple interpolation function to perform 2x upsampling of C_e . This segmentation branch is conceptually simple and easy to implement in the DINO framework, as shown in Fig. 1.

3.5. Unified and Enhanced Query Selection

Unified query selection for mask: Query selection has been widely used in traditional two-stage models [24] and many DETR-like models [32, 35] to improve detection performance. We further improve the query selection scheme in Mask DINO for segmentation tasks.

The encoder output features contain dense features, which can serve as better priors for the decoder. Therefore, we adopt three prediction heads (classification, detection, and segmentation) in the encoder output. Note that the three

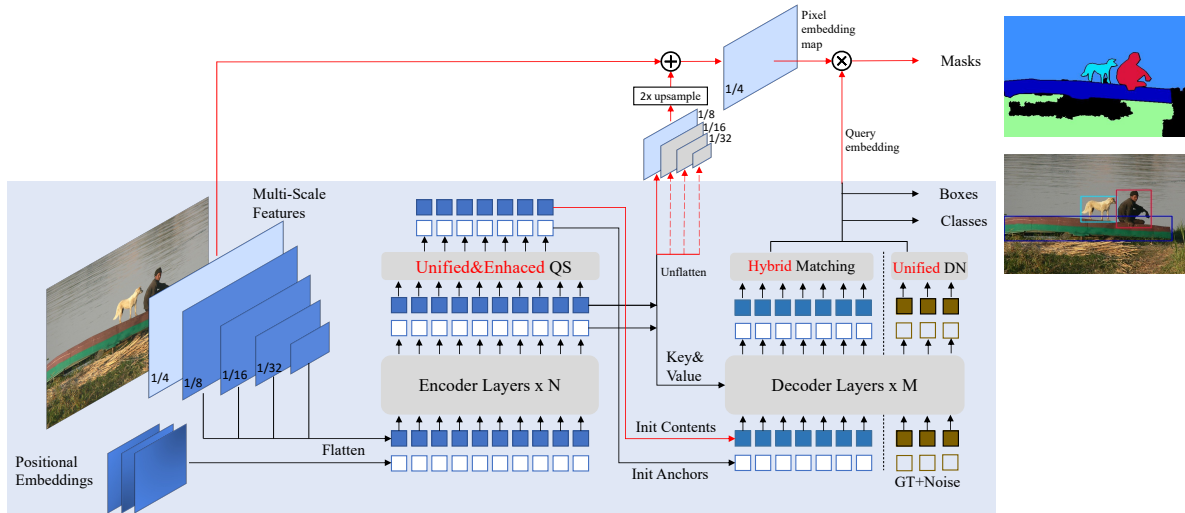


Figure 1. The framework of Mask DINO, which is based on DINO (the blue-shaded part) with extensions (the red part) for segmentation tasks. 'QS' and 'DN' are short for query selection and denoising training, respectively.

heads are identical to the decoder heads. The classification score of each token is considered as the confidence to select top-ranked features and feed them to the decoder as content queries. The selected features also regress boxes and dot-product with the high-resolution feature map to predict masks. The predicted boxes and masks will be supervised by the ground truth and are considered as initial anchors for the decoder. Note that we initialize both the content and anchor box queries in Mask DINO whereas DINO only initializes anchor box queries.

Mask-enhanced anchor box initialization: As summarized in Sec 3.2, image segmentation is a pixel-level classification task while object detection is a region-level position regression task. Therefore, compared to detection, though segmentation is a more difficult task with fine-granularity, it is easier to learn in the initial stage. For example, masks are predicted by dot-producting queries with the high-resolution feature map, which only needs to compare per-pixel semantic similarity. However, detection requires to directly regress the box coordinates in an image. Therefore, in the initial stage after unified query selection, mask prediction is much more accurate than box (the qualitative AP comparison between mask prediction and box prediction in different stages is also shown in Table 8 and 9). Therefore, after unified query selection, we derive boxes from the predicted masks as better anchor box initialization for the decoder. By this effective task cooperation, the enhanced box initialization can bring in a large improvement to the detection performance.

3.6. Segmentation Micro Design

Unified denoising for mask: Query denoising in object detection has shown effective [15, 32] to accelerate convergence and improve performance. It adds noises to ground-

truth boxes and labels and feed them to the Transformer decoder as noised positional queries and content queries. The model is trained to reconstruct ground truth objects given their noised versions. We also extend this technique to segmentation tasks. As masks can be viewed as a more fine-grained representation of boxes, box and mask are naturally connected. Therefore, we can treat boxes as a noised version of masks, and train the model to predict masks given boxes as a denoising task. The given boxes for mask prediction are also randomly noised for more efficient mask denoising training. The detailed noise and its hyperparameters used in our model are shown in Appendix ??.

Hybrid matching: Mask DINO, as in some traditional models [2, 9], predicts boxes and masks with two parallel heads in a loosely coupled manner. Hence the two heads can predict a pair of box and mask that are inconsistent with each other. To address this issue, in addition to the original box and classification loss in bipartite matching, we add a mask prediction loss to encourage more accurate and consistent matching results for one query. Therefore, the matching cost becomes $\lambda_{cls}\mathcal{L}_{cls} + \lambda_{box}\mathcal{L}_{box} + \lambda_{mask}\mathcal{L}_{mask}$, where \mathcal{L}_{cls} , \mathcal{L}_{box} , and \mathcal{L}_{mask} are the classification, box, and mask loss and λ are their corresponding weights. The detailed losses used in our model and their corresponding weights are shown in Appendix ??.

4. Experiments

We conduct extensive experiments and compare with several specialized models for four popular tasks including object detection, instance, panoptic, and semantic segmentation on COCO [17], ADE20K [34], and Cityscapes [5]. For all experiments, we use batch size 16 and A100 GPUs with

Model	Epochs	Query type	Mask AP	Box AP	AP_{50}^{mask}	AP_{75}^{mask}	AP_S^{mask}	AP_M^{mask}	AP_L^{mask}	GFLOPS	Params	FPS
ResNet-50 backbone												
Mask-RCNN [6, 8, 9]	400	Dense anchors	42.5	48.2	—	—	23.8	45.0	60.0	207	40M	10.3
HTC [2]	36	Dense anchors	39.7	44.9	61.4	43.1	22.6	42.2	50.6	441	80M	5
QueryInst [7]	36	300 queries	40.6	45.6	63.0	44.0	23.4	42.5	52.8	—	—	7.0
DINO-4scale [32]	36	900 queries	—	<u>50.9</u>	—	—	—	—	—	245	47M	19.6
Mask2Former [3]	12	100 queries	38.7	—	59.8	41.2	18.2	41.5	59.8	226	44M	8.2
Mask DINO (ours)	12	300 queries	41.4 ^(+2.7)	45.7	62.9	44.6	21.1	44.2	61.4	286	52M	14.8
Mask DINO (ours)	24	300 queries	44.2 ^(+0.5)	48.4	66.6	47.9	23.9	47.0	64.0	286	52M	14.8
Mask2Former* [3]	50	100 queries	43.7	46.2 [†]	66.0	46.9	23.4	47.2	64.8	226	44M	8.2
Mask DINO (ours)	50	100 queries	45.4	49.8	67.9	49.3	25.2	48.3	65.8	280	52M	15.2
Mask DINO (ours)	50	300 queries	46.0	50.5	68.9	50.3	26.0 ^(+2.6)	49.3 ^(+2.1)	65.5 ^(+0.7)	286	52M	14.8
Mask DINO [‡] (ours)	50	300 queries	46.3 ^(+2.6)	51.7 ^(+0.8)	69.0	50.7	26.1 ^(+2.7)	49.3 ^(+2.1)	66.1 ^(+1.3)	286	52M	14.2
SwinL backbone												
HTC++ [2, 20]	72	dense anchors	49.5	57.1	—	—	31.0	52.4	67.2	1470	284M	—
Mask2Former [3]	100	200 queries	<u>50.1</u>	—	—	—	29.9	53.9	72.1	868	216M	4.0
DINO [32]	36	900 queries	—	<u>58.5</u>	—	—	—	—	—	1285	217M	8.1
Mask2Former	100	200 queries	50.1	—	—	—	29.9	53.9	72.1	868	216M	4.0
Mask DINO (ours)	50	300 queries	52.1	58.3	76.5	57.6	32.9	55.4	72.5	1326	223M	6.1
Mask DINO [‡] (ours)	50	300 queries	52.3 ^(+2.2)	59.0 ^(+0.5)	76.6	57.8	33.1	55.4	72.6	1326	223M	5.6

Table 3. Results for Mask DINO and other object detection and instance segmentation models with ResNet-50 and SwinL backbone on COCO val2017 without extra data or tricks. Following DINO [32], we use ResNet-50 with four feature scales by default, and use five scales under large models with a SwinL backbone. We follow the common practice in DETR-like models to use 300 queries. * Mask2Former using 300 queries is not listed as its performance will degenerate when using 300 queries. † indicates the box AP is derived from mask prediction. ‡ we use the proposed mask-enhanced box initialization to further improve detection performance. We test the FPS and GFLOPS of Mask2Former and Mask DINO on the A100 GPU using detectron2.

Model	Epochs	Query type	PQ	PQ^{Th}	PQ^{St}	Box AP_{pan}^{Th}	Mask AP_{pan}^{Th}
ResNet-50 backbone							
DETR [1]	500 + 25	100 queries	43.4	48.2	36	—	31.1
Panoptic Segformer [16]	24	353 queries	49.6	54.4	42.4	—	41.7
Mask2Former* [3]	50	100 queries	51.9/51.5 [†]	57.7	43.0	—	41.7
Mask DINO (ours)	50	100 queries	52.3	58.3	43.2	47.7	43.7
Mask DINO (ours)	50	300 queries	53.0 ^(+1.1)	59.1 ^(+1.4)	43.9 ^(+0.9)	48.8	44.3 ^(+2.6)
Mask DINO (ours)	24	300 queries	51.5	57.3	42.6	46.4	42.8
Mask2Former [3]	12	100 queries	46.9	52.5	38.4	—	37.2
Panoptic Segformer [16]	12	353 queries	48.0	52.3	41.5	—	—
Mask DINO (ours)	12	300 queries	49.0 ^(+1.0)	54.8	40.2	43.2	40.4 ^(+3.2)
SwinL backbone							
Mask2Former [3]	100	100 queries	57.8	64.2	48.1	—	48.6
OneFormer [11]	100	150 queries	57.9	64.4	48.0	—	49.0
Mask DINO (ours)	50	300 queries	58.3 ^(+0.5)	65.1	48.0	56.2	50.6 ^(+2.0)

Table 4. Results for Mask DINO and other panoptic segmentation models with a ResNet-50 backbone on COCO val2017. * Mask2Former using 300 queries is not listed as its performance will degenerate when using 300 queries. † Our reproduced result.

40GB memory. We use a ResNet-50 [10] and a SwinL [20] backbone for our main results and SOTA model. Under ResNet-50, we use 4 A100 GPUs for all tasks without extra data. The implementation details are in Appendix ??.

4.1. Main Results

Instance segmentation and object detection. In Table 3, we compare Mask DINO with other instance segmentation and object detection models. Mask DINO outperforms both the specialized models such as Mask2Former [3] and DINO [32] and hybrid models such as HTC [2] under the same setting. Especially, the instance segmentation results surpass the strong baseline Mask2Former by a large margin (+2.7 AP and +2.6 AP) on the 12-epoch and 50-epoch settings. Moreover, Mask DINO significantly improves the convergence speed, outperforming Mask2Former with less

than half training epochs (44.2 AP in 24 epochs). In addition, after using mask-enhanced box initialization, our detection performance has been significantly improved (+1.2 AP), which even outperforms DINO by **0.8** AP. These results indicate that task unification is beneficial. Without bells and whistles, we achieve the best detection and instance segmentation performance among DETR-like model with a SwinL backbone without extra data.

Panoptic segmentation. We compare Mask DINO with other models in Table 4. Mask DINO outperforms all previous best models on both the 12-epoch and 50-epoch settings by **1.0** PQ and **1.1** PQ, respectively. This indicates Mask DINO has the advantages of both faster convergence and superior performance. One interesting observation is that we outperform Mask2Former [3] in terms of both PQ^{Th} and PQ^{St} . However, instead of using dense and hard-

Model	Iterations	Crop size	mIoU (mean)	mIoU (high)	mIoU (reported)
Mask2Former [3]	160k	512	46.1	46.5	47.2
Mask DINO (ours)	160k	512	47.7 ^(+1.6)	48.7 ^(+2.2)	48.7 ^(+1.6)

Table 5. Results for Mask DINO and Mask2Former with 100 queries using a ResNet-50 backbone on ADE20K val. We found the performance variance on this dataset is high and run three times to report both the mean and highest results for both models.

Model	Iterations	mIoU (mean)	mIoU (high)	mIoU (reported)
Mask2Former [3]	90k	78.7	79.0	79.4
Mask DINO (ours)	90k	79.8 ^(+1.1)	80.0 ^(+1.0)	80.0 ^(+0.6)

Table 6. Results for Mask DINO and Mask2Former with 100 queries using a ResNet-50 backbone on Cityscapes val. We found the performance variance on this dataset is high and run three times to report both the mean and highest results for both models.

Method	Params	Backbone	Backbone Pre-training Dataset	Detection Pre-training Dataset	val w/o TTA w/ TTA	
Instance segmentation on COCO					AP	
Mask2Former [3]	216M	SwinL	IN-22K-14M	–	50.1	–
Soft Teacher [31]	284M	SwinL	IN-22K-14M	O365	51.9	52.5
SwinV2-G-HTC++ [19]	3.0B	SwinV2-G	IN-22K-ext-70M [19]	O365	53.4	53.7
MasK DINO(Ours)	223M	SwinL	IN-22K-14M	–	52.6	–
MasK DINO(Ours)	223M	SwinL	IN-22K-14M	O365	54.5 ^(+1.1)	–
Panoptic segmentation on COCO					PQ	
Panoptic SegFormer [16]	–M	SwinL	IN-22K-14M	–	55.8	–
Mask2Former [3]	216M	SwinL	IN-22K-14M	–	57.8	–
MasK DINO (ours)	223M	SwinL	IN-22K-14M	–	58.4 ^(+0.6)	–
MasK DINO (ours)	223M	SwinL	IN-22K-14M	O365	59.4 ^(+1.6)	–
Semantic segmentation on ADE20K					mIoU	
Mask2Former [3]	215M	SwinL	IN-22K-14M	–	56.1	57.3
SeMask-L MSFaPN-Mask2Former [12]	–M	SwinL-FaPN	IN-22K-14M	–	–	58.2
SwinV2-G-UperNet [19]	3.0B	SwinV2-G	IN-22K-ext-70M [19]	–	59.3	59.9
MasK DINO (ours)	223M	SwinL	IN-22K-14M	–	56.6	–
MasK DINO (ours)	223M	SwinL	IN-22K-14M	O365	59.5	60.8 ^(+0.9)

Table 7. Comparison of the SOTA models on three segmentation tasks. Mask DINO outperforms all existing models. "TTA" means test-time-augmentation. "O365" denotes the Objects365 [26] dataset.

constrained masked attention, we predict boxes and then use them in deformable attention to extract query features. Therefore, our box-oriented deformable attention also works well with "stuff" categories, which makes our unified model simple and efficient. In addition, we improve the mask AP_{pan}^{Th} by **2.6** to 44.3 AP, which is 0.6 higher than the specialized instance segmentation model Mask2Former (43.7 AP).

Semantic segmentation. In Table 5 and 6, we show the performance of semantic segmentation with a ResNet-50 backbone. We use 100 queries for these small datasets. We outperform Mask2Former on both ADE20K and Cityscapes by 1.6 and 0.6 mIoU on the reported performance.

4.2. Comparison with SOTA Models

In Table 7, we compare Mask DINO with SOTA models on three image segmentation tasks to show its scalability. We use the SwinL [20] backbone and pre-train DINO on the Objects365 [26] detection dataset. Even without using extra data, we outperform Mask2Former on all three tasks, especially on instance segmentation (**+2.5 AP**). As Mask DINO is an extension of DINO, the pre-trained DINO model can be used to fine-tune Mask DINO for segmentation tasks. After fine-tuning Mask DINO on the corresponding tasks, we

achieve the best results on instance (**54.5 AP**), panoptic (**59.4 PQ**), and semantic (**60.8 mIoU**) segmentation among model under one billion parameters. Compared to SwinV2-G [19], we significantly reduce the model size to 1/15 and backbone pre-training dataset to 1/5. Our detection pre-training also significantly helps all segmentation tasks including panoptic and semantic with "stuff" categories. However, previous specialized segmentation models such as Mask2Former can not use detection datasets and adding a detection head to it results in poor performance as shown in Table 1, which severely limits the data scalability. By unifying four tasks in one model, we only need to pre-train one model on a large-scale dataset and finetune on all tasks for 10 to 20 epochs (Mask2Former needs 100 epochs), which is more computationally efficient and simpler in model design.

4.3. Ablation Studies

We conduct ablation studies using a ResNet-50 backbone to analyze Mask DINO on COCO val₂₀₁₇. Unless otherwise stated, our experiments are based on object detection and instance segmentation without *Mask-enhanced anchor box initialization*.

Query selection. Table 8 shows the results of our query selection for instance segmentation, where we additionally pro-

test layer#	Mask DINO	Mask2Former
layer 0	39.6 (+38.5)	1.1
layer 3	44.0	42.3
layer 6	45.9	43.3
layer 9	46.0	43.7

Table 8. Effectiveness of our query selection for mask initialization. We evaluate the instance segmentation performance from different decoder layers in the same model after training for 50 epochs.

Tasks		Box AP	Mask AP
Box	Mask		
✓		50.1	—
	✓	—	43.3
✓	✓	50.5 (+0.4)	46.0 (+2.7)

Table 11. Task comparison under the 50-epoch setting. We train the same Mask DINO with different tasks and validate that box and mask can achieve mutual cooperation.

	Box AP	Mask AP
Mask DINO (ours)	45.7	41.4
– Mask-enhanced anchor box initialization	44.5(-1.2)	41.4
– Unified query selection for masks	43.6(-2.1)	40.3(-1.1)
– Unified denoising for masks	44.4(-1.3)	40.3 (-1.1)
– Hybrid matching	44.9(-0.8)	40.5 (-0.9)
– remove all the above	41.7 (-4.0)	38.5 (-2.7)

Table 14. Effectiveness of the proposed components under the 12-epoch setting.

vide the performance of different decoder layers in one single model. Mask2Former also predicts the masks of learnable queries as initial region proposals. However, their performance lags behind Mask DINO by a large margin (**-38.5 AP**). With our effective query selection scheme, the mask performance achieves **39.6 AP** without using the decoder. In addition, our mask performance at layer six is already comparable to the final results with 9 layers. In Table 9, we show that in query selection the predicted box is inferior to mask, which indicates segmentation is easier to learn in the initial stage. Therefore, our proposed mask-enhanced box initialization enhances boxes with masks in query selection to provide better anchor boxes (+15.6 AP) for the decoder, which results in **+1.2 AP** improvement in the final detection performance.

Feature scales. Mask2Former [3] shows that concatenating multi-scale features as input to Transformer decoder layers does not improve the segmentation performance. However, in Table 10, Mask DINO shows that using more feature scales in the decoder consistently improves the performance.

Object detection and segmentation help each other. To validate task cooperation in Mask DINO, we use the same model but train different tasks and report the 12 epoch and 50 epoch results. As shown in Table 11, only training one task will lead to a performance drop. Although only training object detection results in faster convergence in the early stage for box prediction, the final performance is still inferior to training both tasks together.

Decoder layer number. In DINO, increasing the decoder layer number to nine will decrease the performance of box. In Table 12, the result indicates that increasing the number of

layer#	w/o ME		w ME	
	Box	Mask	Box	Mask
layer 0	39.6	<u>25.6</u>	39.8	41.2 (+15.6)
layer 9	46.0	<u>50.5</u>	46.3	51.7 (+1.2)

Table 9. Comparison of our model with and without Mask-enhanced anchor box initialization (ME). ME enhances anchor box initialization and improves final detection performance. Trained for 50 epochs.

Decoder layer#	Box AP	Mask AP
3	43.1	40.7
6	44.3	41.1
9	44.5	41.4
12	44.8	41.1

Table 12. Decoder layer number comparison under the 12-epoch setting. Mask DINO benefits from more decoders, while DINO’s performance will decrease with 9 decoders.

Feature scale	box AP	mask AP
single scale(1/8)	45.8	45.1
3 scales	50.5	45.8
4 scales	50.5	46.0

Table 10. Comparison of multi-scale features for Transformer decoder under the 50-epoch setting. Both detection and segmentation benefit from more feature scales.

Matching		Box AP	Mask AP
Box	Mask		
✓		44.4	40.5
	✓	40.2	38.4
✓	✓	44.5	41.4

Table 13. Matching method comparison under the 12-epoch setting. We train both tasks together but use different matching methods to verify the effectiveness of hybrid matching.

decoder layers will contribute to both detection and segmentation in Mask DINO. We hypothesize that the multi-task training become more complex and require more decoders to learn the needed mapping function.

Matching. In Table 13, we show that only using boxes or masks to perform bipartite matching is not optimal in Mask DINO. A unified matching objective makes the optimization more consistent. In addition, after removing all the proposed components, both detection and segmentation performance drop by a large margin. This result indicates that if we trivially add detection and segmentation tasks in one DETR-based model, the features are not aligned for detection and segmentation tasks to achieve mutual cooperation.

We also present visualization analysis in Appendix ??.

5. Conclusion

In this paper, we have presented Mask DINO as a unified Transformer-based framework for both object detection and image segmentation. Conceptually, Mask DINO is a natural extension of DINO from detection to segmentation with minimal modifications on some key components. Mask DINO outperforms previous specialized models and achieves the best results on all three segmentation tasks (instance, panoptic, and semantic) among models under one billion parameters. Moreover, Mask DINO shows that detection and segmentation can help each other in query-based models. In particular, Mask DINO enables semantic and panoptic segmentation to benefit from a better visual representation pre-trained on a large-scale detection dataset. We hope Mask DINO can provide insights for enabling task cooperation and data cooperation towards designing a universal model for more vision tasks.

Limitations: Different segmentation tasks fail to achieve mutual assistance in Mask DINO in COCO panoptic segmentation. For example, in COCO panoptic segmentation, the mask AP still lags behind the model only trained with instances.

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](#), [3](#), [4](#), [6](#)
- [2] 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/CVF Conference on Computer Vision and Pattern Recognition*, pages 4974–4983, 2019. [1](#), [2](#), [3](#), [5](#), [6](#)
- [3] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention Mask Transformer for Universal Image Segmentation. 2022. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [4] Bowen Cheng, Alexander G. Schwing, and Alexander Kirillov. Per-Pixel Classification is Not All You Need for Semantic Segmentation. 2021. [1](#), [2](#), [3](#), [4](#)
- [5] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. [5](#)
- [6] 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. [6](#)
- [7] Yuxin Fang, Shusheng Yang, Xinggong 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. [1](#), [6](#)
- [8] 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. [6](#)
- [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. [1](#), [2](#), [3](#), [5](#), [6](#)
- [10] 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](#)
- [11] Jitesh Jain, Jiachen Li, MangTik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. OneFormer: One Transformer to Rule Universal Image Segmentation. *arXiv preprint arXiv:2211.06220*, 2022. [6](#)
- [12] Jitesh Jain, Anukriti Singh, Nikita Orlov, Zilong Huang, Jiachen Li, Steven Walton, and Humphrey Shi. SeMask: Semantically Masked Transformers for Semantic Segmentation. *arXiv preprint arXiv:2112.12782*, 2021. [1](#), [7](#)
- [13] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6399–6408, 2019. [3](#)
- [14] 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](#)
- [15] 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. [1](#), [2](#), [3](#), [5](#)
- [16] 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. [6](#), [7](#)
- [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. [5](#)
- [18] 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. [1](#), [2](#), [3](#), [4](#)
- [19] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin Transformer V2: Scaling Up Capacity and Resolution. *arXiv preprint arXiv:2111.09883*, 2021. [7](#)
- [20] 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. [2](#), [6](#), [7](#)
- [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. [1](#), [3](#)
- [22] Depu Meng, Xiaokang Chen, ZeJia Fan, Gang Zeng, Houqiang Li, Yuhui Yuan, Lei Sun, and Jingdong Wang. Conditional DETR for Fast Training Convergence. *arXiv preprint arXiv:2108.06152*, 2021. [4](#)
- [23] Zipeng Qin, Jianbo Liu, Xiaolin Zhang, Maoqing Tian, Aojun Zhou, Shuai Yi, and Hongsheng Li. Pyramid Fusion Transformer for Semantic Segmentation. *arXiv preprint arXiv:2201.04019*, 2022. [1](#)
- [24] 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. [1](#), [4](#)
- [25] 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. [3](#)
- [26] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8430–8439, 2019. [2](#), [7](#)

- [27] 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](#)
- [28] 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](#), [4](#)
- [29] Yingming Wang, Xiangyu Zhang, Tong Yang, and Jian Sun. Anchor detr: Query design for transformer-based detector. *arXiv preprint arXiv:2109.07107*, 2021. [4](#)
- [30] Yuwen Xiong, Renjie Liao, Hengshuang Zhao, Rui Hu, Min Bai, Ersin Yumer, and Raquel Urtasun. Upsnet: A unified panoptic segmentation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8818–8826, 2019. [3](#)
- [31] Mengde Xu, Zheng Zhang, Han Hu, Jianfeng Wang, Lijuan Wang, Fangyun Wei, Xiang Bai, and Zicheng Liu. End-to-end semi-supervised object detection with soft teacher. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3060–3069, 2021. [7](#)
- [32] 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. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [33] 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](#), [3](#)
- [34] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 633–641, 2017. [5](#)
- [35] 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](#), [4](#)