

## Vision Transformers Are Good Mask Auto-Labelers

Shiyi Lan<sup>1</sup> Xitong Yang<sup>2</sup> Zhiding Yu<sup>1</sup> Zuxuan Wu<sup>3</sup> Jose M. Alvarez<sup>1</sup> Anima Anandkumar<sup>1,4</sup>  
<sup>1</sup>NVIDIA <sup>2</sup>Meta AI, FAIR <sup>3</sup>Fudan University <sup>4</sup>Caltech

<https://github.com/NVlabs/mask-auto-labeler>




Figure 1. Examples of mask pseudo-labels generated by Mask Auto-Labeler on COCO. Only human-annotated bounding boxes are used as supervision during training to obtain these results.

### Abstract

We propose Mask Auto-Labeler (MAL), a high-quality Transformer-based mask auto-labeling framework for instance segmentation using only box annotations. MAL takes box-cropped images as inputs and conditionally generates their mask pseudo-labels. We show that Vision Transformers are good mask auto-labelers. Our method significantly reduces the gap between auto-labeling and human annotation regarding mask quality. Instance segmentation models trained using the MAL-generated masks can nearly match the performance of their fully-supervised counterparts, retaining up to 97.4% performance of fully supervised models. The best model achieves 44.1% mAP on COCO instance segmentation (test-dev 2017), outperforming state-of-the-art box-supervised methods by significant margins. Qualitative results indicate that masks produced by MAL are, in some cases, even better than human annotations.

### 1. Introduction

Computer vision has seen significant progress over the last decade. Tasks such as instance segmentation have made it possible to localize and segment objects with pixel-level accuracy. However, these tasks rely heavily on expensive human mask annotations. For instance, when creat-

ing the COCO dataset, about 55k worker hours were spent on masks, which takes about 79% of the total annotation time [1]. Moreover, humans also make mistakes. Human annotations are often misaligned with actual object boundaries. On complicated objects, human annotation quality tends to drop significantly if there is no quality control. Due to the expensive cost and difficulty of quality control, some other large-scale detection datasets such as Open Images [2] and Objects365 [3], only contain partial or even no instance segmentation labels.

In light of these limitations, there is an increasing interest in pursuing box-supervised instance segmentation, where the goal is to predict object masks from bounding box supervision directly. Recent box-supervised instance segmentation methods [4–8] have shown promising performance. The emergence of these methods challenges the long-held belief that mask annotations are needed to train instance segmentation models. However, there is still a non-negligible gap between state-of-the-art approaches and their fully-supervised oracles.

**Our contributions:** To address box-supervised instance segmentation, we introduce a two-phase framework consisting of a mask auto-labeling phase and an instance segmentation training phase (see Fig. 2). We propose a Transformer-based mask auto-labeling framework, Mask Auto-Labeler (MAL), that takes Region-of-interest (RoI) images as inputs

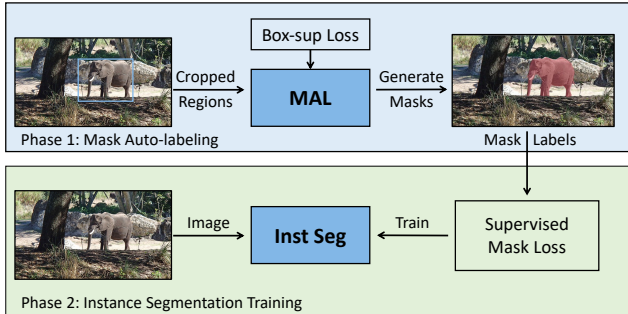


Figure 2. An overview of the two-phase framework of box-supervised instance segmentation. For the first phase, we train Mask Auto-Labeler using box supervision and conditionally generate masks of the cropped regions in training images (top). We then train the instance segmentation models using the generated masks (bottom).

and conditionally generates high-quality masks (demonstrated in Fig. 1) within the box. Our contributions can be summarized as follows:

- Our two-phase framework presents a versatile design compatible with any instance segmentation architecture. Unlike existing methods, our framework is simple and agnostic to instance segmentation module designs.
- We show that Vision Transformers (ViTs) used as image encoders yield surprisingly strong auto-labeling results. We also demonstrate that some specific designs in MAL, such as our attention-based decoder, multiple-instance learning with box expansion, and class-agnostic training, crucial for strong auto-labeling performance. Thanks to these components, MAL sometimes even surpasses humans in annotation quality.
- Using MAL-generated masks for training, instance segmentation models achieve up to 97.4% of their fully supervised performance on COCO and LVIS. Our result significantly narrows down the gap between box-supervised and fully supervised approaches. We also demonstrate the outstanding open-vocabulary generalization of MAL by labeling novel categories not seen during training.

Our method outperforms all the existing state-of-the-art box-supervised instance segmentation methods by large margins. This might be attributed to good representations of ViTs and their emerging properties such as meaningful grouping [9], where we observe that the attention to objects might benefit our task significantly (demonstrated in Fig. 6). We also hypothesize that our class-agnostic training design enables MAL to focus on learning general grouping instead of focusing on category information. Our strong results pave the way to remove the need for expensive human annotation for instance segmentation in real-world settings.

## 2. Related work

### 2.1. Vision Transformers

Transformers were initially proposed in natural language processing [10]. Vision Transformers [11] (ViTs) later emerged as highly competitive visual recognition models that use multi-head self-attention (MHSA) instead of convolutions as the basic building block. These models are recently marked by their competitive performance in many visual recognition tasks [12]. We broadly categorize existing ViTs into two classes: plain ViTs, and hierarchical ViTs.

**Standard Vision Transformers.** Standard ViTs [11] are the first vision transformers. Standard ViTs have the simplest structures, which consist of a tokenization embedding layer followed by a sequence of MHSA layers. However, global MHSA layers can be heavy and usually face significant optimization issues. To improve their performance, many designs and training recipes are proposed to train ViTs in data-efficient manners [9, 13–19].

**Hierarchical Vision Transformers.** Hierarchical Vision Transformers [12, 20–22] are pyramid-shaped architectures that aim to benefit other tasks besides image classification with their multi-scale designs. On top of plain ViTs, these ViTs [20, 21] separate their multi-head self-attention layers into hierarchical stages. Between the stages, there are spatial reduction layers, such as max-pooling layers. These architectures are usually mixed with convolutional layers [23] and often adopt efficient self-attention designs to deal with long sequence lengths.

### 2.2. Instance segmentation

Instance segmentation is a visual recognition task that predicts the bounding boxes and masks of objects.

**Fully supervised instance segmentation.** In this setting, both bounding boxes and instance-level masks are provided as the supervision signals. Early works [24–27] follow a two-stage architecture that generates box proposals or segmentation proposals in the first stage and then produces the final segmentation and classification information in the second stage. Later, instance segmentation models are broadly divided into two categories: some continue the spirit of the two-stage design and extend it to multi-stage architectures [28, 29]. Others simplify the architecture and propose one-stage instance segmentation, e.g., YOLACT [30], SOLO [31, 32], CondInst [33], PolarMask [34, 35]. Recently, DETR and Deformable DETR [36, 37] show great potential of query-based approaches in object detection. Then, methods like MaxDeepLab [38], MaskFormer [39], PanopticSegFormer [40], Mask2Former [41] and Mask DINO [42] are introduced along this line and have pushed the boundary of instance segmentation. On the other hand, the instance segmentation also benefits from more power-

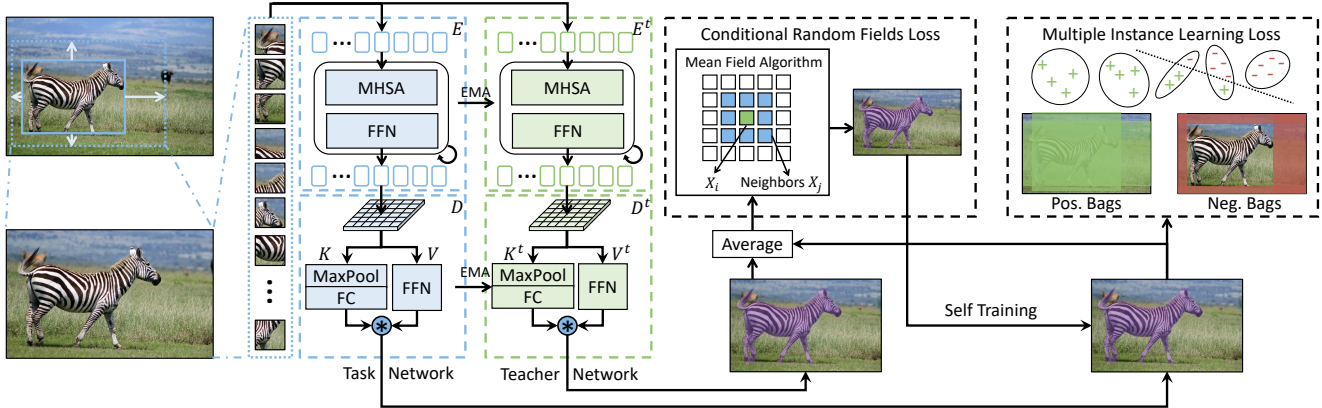


Figure 3. Overview of MAL architecture. We visualize the architecture of Mask Auto-Labeler. Mask Auto-Labeler takes cropped images as inputs. Mask Auto-Labeler consists of two symmetric networks, **Task Network** and **Teacher Network**. Each network contains the image encoder  $E$  (or  $E^t$ ), and the mask decoder  $D$  (or  $D^t$ ). We use the exponential moving average (EMA) to update the weights of the teacher network. We apply multiple instance learning (MIL) loss and conditional random fields (CRFs) loss. The CRF loss takes the average mask predictions of the teacher network and the task network to make the training more stable and generate refined masks for self-training.

ful backbone designs, such as Swin Transformers [12, 22], ViTDet [43], and ConvNeXt [44].

**Weakly supervised instance segmentation.** There are two main styles of weakly supervised instance segmentation: learning with image-level and box-level labels. The former uses image-level class information to perform instance segmentation [45–49], while the latter uses box-supervision. Hsu et al. [4] leverages the tight-box priors. Later, Box-Inst [5] proposes to leverage color smoothness to improve accuracy. Besides that, DiscoBox [7] proposes to leverage both color smoothness and inter-image correspondence for the task. Other follow-ups [6, 8] also leverage tight-box priors and color smoothness priors.

### 2.3. Deep learning interpretation

The interest in a deeper understanding of deep networks has inspired many works to study the interpretation of deep neural networks. For example, Class Activation Map (CAM) [50] and Grad-CAM [51] visualize the emerging localization during image classification training of convolutional neural networks (CNNs). This ability has also inspired much weakly-supervised localization and shows deep connections to general weakly-supervised learning, which partly motivates our decoder design in this paper. DINO [9] further shows that meaning visual grouping emerges during self-supervised learning with ViTs. In addition, FAN [52] shows that such emerging properties in ViTs are linked to their robustness.

## 3. Method

Our work differs from previous box-supervised instance segmentation frameworks [4–8] that simultaneously learns detection and instance segmentation. We leverage a two-

phase framework as visualized in Fig. 2, which allows us to have a network focused on generating mask pseudo-labels in phase 1, and another network focused on learning instance segmentation [24, 28, 41, 43] in phase 2. Our proposed auto-labeling framework is used in phase 1 to generate high-quality mask pseudo-labels.

We propose this two-phase framework because it brings the following benefits:

- We can relax the learning constraints in phase 1 and focus only on mask pseudo-labels. Therefore, in this phase, we can take Region-of-interest (RoI) images instead of untrimmed images as inputs. This change allows us to use a higher resolution for small objects and a strong training technique mentioned in Sec. 3.1, which helps improve the mask quality.
- We can leverage different image encoders and mask decoders in phases 1 and 2 to achieve higher performance. We empirically found that phases 1 and 2 favor different architectures for the image encoders and mask decoders. See the ablation study in Tab. 3 and 4.
- We can use MAL-generated masks to directly train the most fully supervised instance segmentation models in phase 2. This makes our approach more flexible than previous architecture-specific box-supervised instance segmentation approaches [4–8].

As phase 2 follows the previous standard pipelines, which do not need to be re-introduced here, we focus on introducing phase 1 (MAL) in the following subsections.

### 3.1. RoI input generation

Most box-supervised instance segmentation approaches [4–7] are trained using the entire images.

However, we find that using RoI images might have more benefits in box-supervised instance segmentation. Moreover, we compare two intuitive sampling strategies of RoI images to obtain foreground and background pixels and explain the better strategy, box expansion, in detail.

**Benefits of using RoI inputs.** There are two advantages of using RoI images for inputs. First, using the RoI images as inputs is naturally good for handling small objects because no matter how small the objects are, the RoI images are enlarged to avoid the issues caused by low resolution. Secondly, having RoI inputs allows MAL to focus on learning segmentation and avoid being distracted from learning other complicated tasks, e.g., object detection.

**RoI sampling strategy.** The sampling strategy should ensure both positive and negative pixels are included. We present two straightforward sampling strategies:

- The first strategy is to use bounding boxes to crop the images for positive inputs. We crop the images using randomly generated boxes containing only background pixels for negative inputs. MAL does not generate good mask pseudo-labels with cropping strategy. We observe that the networks tend to learn the trivial solution (all pixels are predicted as either foreground or background).
- The second is to expand the bounding boxes randomly and include background pixels, where negative bags are chosen from the expanded rows and columns. We visualize how we define positive/negative bags in Fig. 3 and explain the detail in Sec. 3.3. **This detailed design is critical to make MAL work** as it prevents MAL from learning trivial solutions. Without this design, the generated masks tend to fill the entire bounding box.

**Box expansion specifics.** Given an untrimmed image  $I^u \in \mathbb{R}^{C \times H^u \times W^u}$  and the bounding box  $\mathbf{b} = (x_0, y_0, x_1, y_1)$  indicating the x, y coordinates of the top-left corners and the bottom-right corners. To obtain background pixels, we randomly expand the bounding box  $\mathbf{b}$  to  $\mathbf{b}' = (x_c + \beta_x(x_0 - x_c), y_c + \beta'_x(y_0 - y_c), x_c + \beta_y(x_1 - x_c), y_c + \beta'_y(y_1 - y_c))$ , where  $x_c = (x_0 + x_1)/2$ ,  $y_c = (y_0 + y_1)/2$ . To generate random values of  $\beta_x, \beta'_x, \beta_y, \beta'_y$ , we randomly generate  $\theta_x, \theta_y \in [0, \theta]$  for x- and y-direction, where  $\theta$  is the upper bound of box expansion rate. Next, we randomly generate  $\beta_x \in [0, \theta_x]$  and  $\beta_y \in [0, \theta_y]$ . In the end, we assign  $\beta'_x$  as  $\theta_x - \beta_x$  and  $\beta'_y$  as  $\theta_y - \beta_y$ . Finally, we use  $\mathbf{b}'$  to crop the image and obtain trimmed image  $I^t$ . We conduct the ablation study for  $\theta$  in Tab. 5. At last, We resize the trimmed image  $I^t$  to the size of  $C \times H^c \times W^c$  as the input image  $I^c$ .

### 3.2. MAL architecture

MAL can be divided into two symmetric networks: the task network and the teacher network. The task network




Figure 4. (a) The fully connected decoder (b) The fully convolutional Decoder (c) The attention-based decoder (used in MAL) (d) The query-based Decoder.

consists of an image encoder denoted as  $E$ , and a mask decoder denoted as  $D$ , demonstrated in Fig. 3. The architecture of the teacher network is identical to the task network. We denote the segmentation output of the task network and the teacher network as  $\mathbf{m}, \mathbf{m}^t \in \{0, 1\}^N$ , respectively.

**Image encoder.** We use Standard ViTs [11] as the image encoder and drop the classification head of Standard ViTs. We compare different image encoders in Sec. 4.4. We also try feature pyramid networks on top of Standard ViTs, e.g., FPN [53], but it causes a performance drop. Similar conclusions were also found in ViTDet [43].

**Mask decoder.** For the mask decoder  $D$ , we use a simple attention-based network inspired by YOLACT [30], which includes an instance-aware head  $K$  and a pixel-wise head  $V$ , where  $D(E(I)) = K(E(I)) \cdot V(E(I))$ , and “ $\cdot$ ” represents the inner-product operator.

For the instance-aware head  $K$ , we use a max-pooling layer followed by a fully connected layer. The input channel dimension of  $K$  is equivalent to the output channel dimension of  $E$ . The output channel dimension of  $K$  is 256.

For the pixel-wise head  $V$ , we use four sequential convolutional layers. Each is followed by a ReLU layer. Between the second and the third convolutional layer, we insert a bilinear interpolation layer to increase the feature resolution by 2. The input channel dimension is equivalent to the output channel dimension of  $E$ . We use 256 dimensions for hidden channels and output channels. We also compare different design choices of mask decoders in Sec. 4.5.

**Exponential moving average (EMA) teacher.** Instead of training the teacher network directly, we leverage exponential moving averages (EMA) to update the parameters in the teacher network using the parameters in the task network similar to MOCO [54]. The goal of using EMA Teacher is to eliminate the loss-explosion issues in training since optimizing Standard Vision Transformers is usually non-trivial [13, 14, 16]. We do not observe any significant performance drop or improvement on DeiT-small-based MAL after removing the teacher network. However, it makes the training more stable when we use larger-scale image encoders in MAL, e.g. ViT-MAE-Base [13].

### 3.3. Losses

We use Multiple Instance Learning Loss  $\mathcal{L}_{\text{mil}}$  and Conditional Random Field Loss  $\mathcal{L}_{\text{crf}}$  as the box-supervised loss:

$$\mathcal{L} = \alpha_{\text{mil}} \mathcal{L}_{\text{mil}} + \alpha_{\text{crf}} \mathcal{L}_{\text{crf}} \quad (1)$$

**Multiple Instance Learning Loss.** The motivation of the Multiple Instance Segmentation is to exploit the priors of tight-bounding box annotations.

After the student network produces the output  $\mathbf{m}$ , we apply the Multiple Instance Learning (MIL) Loss on the output mask  $\mathbf{m}$ . We demonstrate the process in Fig. 3.

We denote  $\mathbf{m}_{i,j}$  as the mask score at the location  $i, j$  in the image  $I^c$ . We define each pixel as an instance in the MIL loss. Inspired by BBTP [4], we treat each row or column of pixels as a bag. We determine whether a bag is positive or negative based on whether it passes a ground-truth box. We define the bags as  $\mathbf{B}$ , and each bag  $\mathbf{B}_i$  contains a row or column of pixels. Additionally, we define the label for each bag  $\mathbf{g}$ , and each label  $\mathbf{g}_i$  corresponds to a bag  $\mathbf{B}_i$ .

Therefore, we use the max pooling as the reduction function and dice loss [55]:

$$\mathcal{L}_{mil} = 1 - \frac{2 \sum_i \mathbf{g}_i \cdot \max\{\mathbf{B}_i\}^2}{\sum_i \max\{\mathbf{B}_i\}^2 + \sum_i \mathbf{g}_i^2} \quad (2)$$

**Conditional Random Field Loss.** The goal of CRF loss is to refine the mask prediction by imposing the smoothness priors via energy minimization. Then, we leverage this refined mask as pseudo-labels to self-train the mask prediction in an online-teacher manner. We use the average mask prediction  $\mathbf{m}^a = \frac{1}{2}(\mathbf{m} + \mathbf{m}^t)$  as the mask prediction to be refined for more stable training.

Next, we define a random field  $\mathbf{X} = \{\mathbf{X}_1, \dots, \mathbf{X}_N\}$ , where  $N = H^c \times W^c$  is the size of cropped image and each  $\mathbf{X}_i$  represents the label that corresponds to a pixel in  $I^c$ , therefore we have  $\mathbf{X} \in \{0, 1\}^N$ , meaning the background or the foreground. We use  $\mathbf{l} \in \{0, 1\}^N$  to represent a labeling of  $\mathbf{X}$  minimizing the following CRF energy:

$$E(\mathbf{l}|\mathbf{m}^a, \mathbf{X}^c) = \mu(\mathbf{X}|\mathbf{m}^a, I^c) + \psi(\mathbf{X}|I^c), \quad (3)$$

where  $\mu(\mathbf{X}|\mathbf{m}^a, I^c)$  represents the unary potentials, which is used to align  $\mathbf{X}_i$  and  $\mathbf{m}_i^a$  since we assume that most of the mask predictions are correct. Meanwhile,  $\psi(\mathbf{X}|I^c)$  represents the pairwise potential, which sharpens the refined mask. Specifically, we define the pairwise potentials as:

$$\psi(\mathbf{X}|I^c) = \sum_{\substack{i \in \{0..N-1\}, \\ j \in \mathcal{N}(i)}} \omega \exp\left(\frac{-|I_i^c - I_j^c|^2}{2\zeta^2}\right) [\mathbf{X}_i \neq \mathbf{X}_j], \quad (4)$$

where  $\mathcal{N}(i)$  represents the set of 8 immediate neighbors to  $\mathbf{X}_i$  as shown in Fig. 3. Then, we use the MeanField algorithm [7, 56] to efficiently approximate the optimal solution, denoted as  $\mathbf{l} = \text{MeanField}(I^c, \mathbf{m}^a)$ . We attach the derivation and PyTorch code in the supplementary. At last, we apply Dice Loss to leverage the refined masks  $\mathbf{l}$  to self-train the models as:

$$\mathcal{L}_{crf} = 1 - \frac{2 \sum_i \mathbf{l}_i \mathbf{m}_i}{\sum_i \mathbf{l}_i^2 + \mathbf{m}_i^2} \quad (5)$$

## 4. Experiments

We evaluate MAL on COCO dataset [1], and LVIS [57]. The main results on COCO and LVIS are shown in Tab. 1 and 2. The qualitative results are shown in Fig. 1 and Fig. 5.

### 4.1. Datasets

**COCO dataset.** contains 80 semantic categories. We follow the standard partition, which includes train2017 (115K images), val2017 (5K images), and test-dev (20k images).

**LVIS dataset.** contains 1200+ categories and 164K images. We follow the standard partition of training and validation.

### 4.2. Implementation Details

We use 8 NVIDIA Tesla V100s to run the experiments. **Phase 1 (mask auto-labeling).** We use AdamW [58] as the network optimizer and set the two momentums as 0.9, 0.9. We use the cosine and annealing scheduler to adjust the learning rate, which is set to  $1.5 \cdot 10^{-6}$  per image. The MIL loss weight  $\alpha_{mil}$ , CRF loss weight  $\alpha_{crf}$ ,  $\zeta$ , and  $\omega$  in CRF pairwise potentials are set to 4, 0.5, 0.5, 2, respectively. We analyze the sensitivity of the loss weights and CRF hyperparameters in Fig. 8. We use the input resolution of  $512 \times 512$ , and a batch size of 32 (4 per GPU). For EMA, we use a momentum of 0.996. For the task and teacher network, we apply random flip data augmentation. On top of that, we apply extra random color jittering, random grey-scale conversion, and random Gaussian blur for the task network. We train MAL for 10 epochs. It takes around 23 hours and 35 hours to train MAL with Standard ViT-Base [11] on the COCO and LVIS datasets, respectively.

**Phase 2 (Training instance segmentation models).** We select a couple of high-performance fully supervised instance segmentation models, which are ConvNeXts [44] with Cascade R-CNN [28], Swin Transformers [12] with Mask2Former [41], ResNets [59] and ResNeXts [60] with SOLOv2 [31]. MAL works extremely well with these architectures, which demonstrates the great power of Mask Auto-Labeler from the perspective of accuracy and generalization. We leverage the codebase in MMDetection [61] for phase 2. Again, we only replace the GT masks with MAL-generated mask pseudo-labels to adjust all these fully supervised models to box-supervised learning.

### 4.3. Instance segmentation results

**Retention Rate.** We argue that the sole mAP of instance segmentation is not fair enough to evaluate box-supervised instance segmentation since the performance gain can be achieved by improving box quality unrelated to segmentation quality. However, the retention rate can better reflect the real mask quality because the fully supervised counterparts also get boosted by the better box results.

**Results on COCO.** In table 1, we show that various modern instance segmentation models can achieve up to 94.5%

Method	Labeler Backbone	InstSeg Backbone	InstSeg Model	Sup	(%)Mask AP <sub>val</sub>	(%)Mask AP <sub>test</sub>	(%)Ret <sub>val</sub>	(%)Ret <sub>test</sub>
Mask R-CNN* [24]	-	ResNet-101	Mask R-CNN	Mask	38.6	38.8	-	-
Mask R-CNN* [24]	-	ResNeXt-101	Mask R-CNN	Mask	39.5	39.9	-	-
CondInst [33]	-	ResNet-101	CondInst	Mask	38.6	39.1	-	-
SOLOv2 [31]	-	ResNet-50	SOLOv2	Mask	37.5	38.4	-	-
SOLOv2 [31]	-	ResNet-101-DCN	SOLOv2	Mask	41.7	41.8	-	-
SOLOv2 [31]	-	ResNeXt-101-DCN	SOLOv2	Mask	42.4	42.7	-	-
ConvNeXt [44]	-	ConvNeXt-Small [44]	Cascade R-CNN	Mask	44.8	45.5	-	-
ConvNeXt [44]	-	ConvNeXt-Base [44]	Cascade R-CNN	Mask	45.4	46.1	-	-
Mask2Former [41]	-	Swin-Small	Mask2Former	Mask	<b>46.1</b>	<b>47.0</b>	-	-
BBTP† [4]	-	ResNet-101	Mask R-CNN	Box	-	21.1	-	59.1
BoxInst [5]	-	ResNet-101	CondInst	Box	33.0	33.2	85.5	84.9
BoxLevelSet [6]	-	ResNet-101-DCN	SOLOv2	Box	35.0	35.4	83.9	83.5
DiscoBox [7]	-	ResNet-50	SOLOv2	Box	30.7	32.0	81.9	83.3
DiscoBox [7]	-	ResNet-101-DCN	SOLOv2	Box	35.3	35.8	84.7	85.9
DiscoBox [7]	-	ResNeXt-101-DCN	SOLOv2	Box	37.3	37.9	88.0	88.8
BoxTeacher [8]	-	Swin-Base	CondInst	Box	-	40.0	-	-
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNet-50	SOLOv2	Box	35.0	35.7	93.3	93.0
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNet-101-DCN	SOLOv2	Box	38.2	38.7	91.6	92.6
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNeXt-101-DCN	SOLOv2	Box	38.9	39.1	91.7	91.6
Mask Auto-Labeler	ViT-MAE-Base [13]	ConvNeXt-Small [44]	Cascade R-CNN	Box	42.3	43.0	94.4	<b>94.5</b>
Mask Auto-Labeler	ViT-MAE-Base [13]	ConvNeXt-Base [44]	Cascade R-CNN	Box	42.9	43.3	<b>94.5</b>	93.9
Mask Auto-Labeler	ViT-MAE-Base [13]	Swin-Small [12]	Mask2Former [41]	Box	<b>43.3</b>	<b>44.1</b>	93.9	93.8

Table 1. Main results on COCO. Ret means the retention rate of  $\frac{\text{box-supervised mask AP}}{\text{supervised mask AP}}$ . MAL with SOLOv2/ResNeXt-101 outperforms DiscoBox with SOLOv2/ResNeXt-101 by 1.6% on val2017 and 1.3% on test-dev. Our best model (Mask2former/Swin-Small) achieves 43.3% AP on val and 44.1% AP on test-dev.

performance with the pseudo-labels of the fully supervised oracles. Our best results are 43.3% mAP on COCO test-dev and 44.1% mAP on COCO val, achieved by using MAL (Standard ViT-Base [11] pretrained with MAE) for phase 1, and using Mask2Former (Swin-Small) [12, 41] for phase 2. There is no significant retention drop when we use the mask pseudo-labels to train more powerful instance segmentation models. On the contrary, the higher retention rates on COCO are achieved by the heavier instance segmentation models, e.g., Cascade R-CNN with ConvNeXts and Mask2Former with Swin-Small. However, other methods have significantly lower retention rates compared with MAL. The experiment results quantitatively imply that the mask quality outperforms other methods by a large margin. **Results on LVIS.** In table 2, we also observe that all instance segmentation models work very well with the mask pseudo-labels generated by MAL (Ret. = 93% ~ 98%). We visualize part of the results in figure 5. We also evaluate the open-vocabulary ability of MAL by training MAL on COCO dataset but generating mask pseudo-labels on LVIS, and thus training instance segmentation models using these mask pseudo-labels.

#### 4.4. Image encoder variation

To support our claim that Vision Transformers are good auto-labelers, we compare three popular networks as the image encoders of MAL: Standard Vision Transformers [11, 13, 16], Swin Transformer [12], ConvNeXts [44] in Tab. 4.

First, we compare the fully supervised pretrained weights of these three models. We choose the official fully supervised pre-trained weights of ConvNeXts and Swin Transformers. For Standard Vision Transformers, we adopt a popular fully supervised approach, DeiT [16]. We ob-

serve that fully supervised Standard Vision Transformers (DeiT) as image encoders of Mask Auto-Labeler are better than Swin Transformers and ConvNeXts even though the imaganet-1k performance of Swin Transformers and ConvNeXts is higher than that of DeiT. We argue that the success of Standard Vision Transformers might be owed to the self-emerging properties of Standard ViTs [9, 11] (visualized in Fig. 6), and the larger-receptive field brought by global multi-head self-attention layers.

Second, the mask pseudo-labels can be further improved by Mask AutoEncoder (MAE) pretraining [13]. The potential reason might be that MAE pretraining enhances Standard ViTs via learning pixel-level information, which is very important for dense-prediction tasks like segmentation.

#### 4.5. Mask decoder variation

We compare four different modern designs of mask decoders: the fully connected Decoder [62], the fully convolutional decoder [24, 63], the attention-based decoder [30, 31], and the query-based decoder [41] in Tab. 3. We visualize different designs of mask decoders in Figure 4. For the fully connected Decoder, we use two fully connected layers with a hidden dimension of 2048 and then output a confidence map for each pixel. We reshape this output vector as the 2D confidence map. We introduce the attention-based decoder in Sec 3.2. For the fully convolutional Decoder, We adopt the pixel-wise head  $V$  in the attention-based Decoder. For the query-based decoder, we follow the design in Mask2Former [41]. We spend much effort exploring the query-based Decoder on MAL since it performs extremely well on fully supervised instance segmentation. However, the results are surprisingly unsatisfactory. We suspect the slightly heavier layers might cause optimization issues un-



Figure 5. Qualitative results of mask pseudo-labels generated by Mask Auto-Labeler on LVIS v1.

Method	Autolabeler Backbone	InstSeg Backbone	InstSeg Model	Training Data	Sup	(%)Mask AP <sub>val</sub>	(%)Ret <sub>val</sub>
Mask R-CNN [24]	-	ResNet-50-DCN	Mask R-CNN [24]	-	Mask	21.7	-
Mask R-CNN [24]	-	ResNet-101-DCN	Mask R-CNN [24]	-	Mask	23.6	-
Mask R-CNN [24]	-	ResNeXt-101-32x4d-FPN	Mask R-CNN [24]	-	Mask	25.5	-
Mask R-CNN [24]	-	ResNeXt-101-64x4d-FPN	Mask R-CNN [24]	-	Mask	25.8	-
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNet-50-DCN	Mask R-CNN [24]	LVIS v1	Box	20.7	95.4
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNet-101-DCN	Mask R-CNN [24]	LVIS v1	Box	23.0	<b>97.4</b>
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNeXt-101-32x4d-FPN	Mask R-CNN [24]	LVIS v1	Box	23.7	92.9
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNeXt-101-64x4d-FPN	Mask R-CNN [24]	LVIS v1	Box	<b>24.5</b>	95.0
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNeXt-101-32x4d-FPN	Mask R-CNN [24]	COCO	Box	23.3	91.8
Mask Auto-Labeler	ViT-MAE-Base [13]	ResNeXt-101-64x4d-FPN	Mask R-CNN [24]	COCO	Box	<b>24.2</b>	<b>93.8</b>

Table 2. Main results on LVIS v1. Training data means the dataset we use for training MAL. We also finetune it on COCO and then generate pseudo-labels of LVIS v1. Compared with trained on LVIS v1 directly, MAL finetuned on COCO only caused around 0.35% mAP drop on the final results, which indicates the great potential of the open-set ability of MAL. Ret means the retention rate of  $\frac{\text{box-supervised mask AP}}{\text{supervised mask AP}}$ .

Mask decoder	(%)Mask AP <sub>val</sub>	(%)Ret <sub>val</sub>
Fully connected decoder	35.5	79.2
Fully convolutional decoder	36.1	80.5
Attention-based decoder	<b>42.3</b>	<b>94.4</b>
Query-based decoder	-	-

Table 3. Ablation study of box expansion. We use Standard ViT-MAE-Base as the image encoder of MAL in phase 1 and Cascade RCNN with ConvNext-Small as the instance segmentation models in phase 2. The numbers are reported in % Mask mAP. Among different designs, the attention-based decoder performs the best. We can not obtain reasonable results with Query-based Decoder.

Backbone	IN-1k Acc@1	Mask AP <sub>val</sub>	Ret <sub>val</sub>
ConvNeXt-Base [44]	<b>83.8</b>	39.6	88.4
Swin-Base [12]	83.5	40.2	89.7
ViT-DeiT-Small [64]	79.9	40.8	91.0
ViT-DeiT-Base [64]	81.8	<b>41.1</b>	<b>91.7</b>
ViT-MAE-Base [13]	83.6	<b>42.3</b>	<b>94.4</b>
ViT-MAE-Large [13]	85.9	<b>42.3</b>	<b>94.4</b>

Table 4. Ablation study of different backbones. All models are pre-trained on ImageNet-1k. ConvNeXt and Swin Transformer outperform DeiT on image classification, but standard ViT-Small [16] (ViT-DeiT-Small) outperforms ConvNeXt-base and Swin-Base on mask Auto-labeling. Standard ViT-Base (ViT-MAE-Base) and Standard ViT-Large (ViT-MAE-Large) pretrained via MAE achieve the best performance on mask Auto-labeling.

der the box-supervised losses.

Experiments show that box-supervised instance segmentation favors the attention-based decoder. However, state-of-the-art instance segmentation and object detection methods often adopt the fully convolutional decoder [15, 43] or the query-based decoder [41]. Our proposed two-phase framework resolves this dilemma and allows the networks to enjoy the merits of both the attention-based Decoder and the non-attention-based Decoders.

#### 4.6. Clustering analysis

As the results are shown in Tab. 4, we wonder why the Standard ViTs outperform other modern image encoders in auto-labeling. As the comparison of classification ability does not seem to reflect the actual ability of auto-labeling, we try to use the ability clustering to evaluate the image encoders because foreground(FG)/background(BG) segmentation is very similar to the binary clustering problem.

Specifically, we extract the feature map output by the last layers of Swin Transformers [12], ConvNeXts [44], Stan-

dard ViTs [11]. Then, we use the GT mask to divide the feature vectors into the FG and BG feature sets. By evaluating the average distance from the FG/BG feature vectors to their clustering centers, we can reveal the ability of the networks to distinguish FG and BG pixels empirically.

Formally, we define the feature vector of token  $i$  generated by backbone E as  $\mathbf{f}_i^E$ . We define the FG/BG clustering centers  $\mathbf{f}'_1, \mathbf{f}'_0$  as the mean of the FG/BG feature vectors. Then, we use the following metric as the clustering score:

$$S = \frac{1}{N} \sum_i \left( \frac{\mathbf{f}_i^E}{|\mathbf{f}_i^E|} - \frac{\mathbf{f}'_{\gamma(i)}}{|\mathbf{f}'_{\gamma(i)}|} \right)^2, \quad (6)$$

where if pixel  $i$  is FG,  $\gamma(i) = 1$ , otherwise  $\gamma(i) = 0$ .

We show the clustering evaluation on the COCO val 2017 in Tab. 6. The results align our conclusion that Standard Vision Transformers are better at mask auto-labeling.




Figure 6. Attention visualization of two RoI images produced by MAL. In each image group, the left-most image is the original image. We visualize the attention map output by the 4<sup>th</sup>, 8<sup>th</sup>, 12<sup>th</sup> MHSA layers of the Standard ViTs in MAL.




Figure 7. The lateral comparison between MAL-generated pseudo-labels (top) and GT masks (bottom) on COCO val2017. On the left, we observe that MAL-generated pseudo-labels are sharper and more boundary-sticky than GT masks in some cases. On the right, we observe that in highly occluded situations, human-annotated masks are still better.




Figure 8. Sensitivity analysis of loss weights and CRF hyper-parameters. We use ViT-Base [11] pretrained via MAE [13] as the image encoder for the first phase and SOLOv2 (ResNet-50) for the second phase. The x-axis and y-axis indicate the hyper-parameter values and the (%)mask AP, respectively.

$\theta$	Mask AP <sub>val</sub>	Ret. <sub>val</sub>	Backbone	Score (↓)
0.6	41.3	92.2	ConvNeXt-Base [44]	0.459
0.8	41.7	93.1	Swin-Base [12]	0.425
1.0	42.2	94.2	ViT-DeiT-Small [64]	0.431
1.2	<b>42.3</b>	<b>94.4</b>	ViT-DeiT-Base [64]	0.398
1.4	42.0	93.8	ViT-MAE-Base [13]	0.324
1.6	41.8	93.3	ViT-MAE-Large [13]	<b>0.301</b>

Table 5. Ablation on box expansion ratio. We use Standard ViT-Base pretrained via MAE (ViT-MAE-Base) and Cascade R-CNN (ConvNeXt-Small) for phase 1 and 2.

Table 6. Clustering scores for different image encoders. The smaller clustering scores imply a better ability to distinguish foreground and background features.

#### 4.7. MAL masks v.s. GT masks

We show the apple-to-apple qualitative comparison in Fig. 7. We find out two prominent patterns in this qualitative comparison. First, MAL-generated mask pseudo-labels are sharper and boundary-sticky than the human-annotated masks since many mask annotations in COCO are poorly annotated. Second, occlusion is also a challenging problem

in MAL-generated mask pseudo-labels.

## 5. Conclusions

In this work, we propose a novel two-phase framework for box-supervised instance segmentation and a novel Transformer-based architecture, Mask Auto-Labeler (MAL), to generate high-quality mask pseudo-labels in phase 1. We reveal that Standard Vision Transformers are good mask auto-labelers. Moreover, we find that random using box-expansion RoI inputs, the attention-based Decoder, and class-agnostic training are crucial to the strong mask auto-labeling performance. Moreover, thanks to the two-phase framework design and MAL, we can adjust almost all kinds of fully supervised instance segmentation models to box-supervised learning with little performance drop, which shows the great generalization of MAL.

**Limitations.** Although great improvement has been made by our approaches in mask auto-labeling, we still observe many failure cases in the occlusion situation, where human annotations are much better than MAL-generated masks. Additionally, we meet saturation problems when scaling the model from Standard ViT-Base to Standard ViT-Large. We leave those problems in the future work.

**Broader impacts.** Our proposed Transformer-based mask auto-labeler and the two-phase architecture serve as a standard paradigm for high-quality box-supervised instance segmentation. If follow-up work can find and fix the issues under our proposed paradigm, there is great potential that expansive human-annotated masks are no longer needed for instance segmentation in the future.



## References

- [1] 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. [1](#), [5](#)
- [2] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, Tom Duerig, and Vittorio Ferrari. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020. [1](#)
- [3] 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. [1](#)
- [4] Cheng-Chun Hsu, Kuang-Jui Hsu, Chung-Chi Tsai, Yen-Yu Lin, and Yung-Yu Chuang. Weakly supervised instance segmentation using the bounding box tightness prior. *Advances in Neural Information Processing Systems*, 32, 2019. [1](#), [3](#), [5](#), [6](#)
- [5] Zhi Tian, Chunhua Shen, Xinlong Wang, and Hao Chen. Boxinst: High-performance instance segmentation with box annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5443–5452, 2021. [1](#), [3](#), [6](#)
- [6] Wentong Li, Wenyu Liu, Jianke Zhu, Miaomiao Cui, Xi-ansheng Hua, and Lei Zhang. Box-supervised instance segmentation with level set evolution. *arXiv preprint arXiv:2207.09055*, 2022. [1](#), [3](#), [6](#)
- [7] Shiyi Lan, Zhiding Yu, Christopher Choy, Subhashree Radhakrishnan, Guilin Liu, Yuke Zhu, Larry S Davis, and Anima Anandkumar. Discobox: Weakly supervised instance segmentation and semantic correspondence from box supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3406–3416, 2021. [1](#), [3](#), [5](#), [6](#)
- [8] Tianheng Cheng, Xinggang Wang, Shaoyu Chen, Qian Zhang, and Wenyu Liu. Boxteacher: Exploring high-quality pseudo labels for weakly supervised instance segmentation. *arXiv preprint arXiv:2210.05174*, 2022. [1](#), [3](#), [6](#)
- [9] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. [2](#), [3](#), [6](#)
- [10] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [2](#)
- [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. [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [12] 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](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [2](#), [4](#), [6](#), [7](#), [8](#)
- [14] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. [2](#), [4](#)
- [15] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. [2](#), [7](#)
- [16] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. [2](#), [4](#), [6](#), [7](#)
- [17] Hugo Touvron, Matthieu Cord, and Herve Jegou. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022. [2](#)
- [18] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 32–42, October 2021. [2](#)
- [19] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Adavit: Adaptive vision transformers for efficient image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12309–12318, 2022. [2](#)
- [20] Wenhui Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 568–578, 2021. [2](#)
- [21] Wenhui Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3):415–424, 2022. [2](#)
- [22] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12009–12019, 2022. [2](#), [3](#)
- [23] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021. [2](#)
- [24] 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. [2](#), [3](#), [6](#), [7](#)
- [25] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, and Yichen Wei. Fully convolutional instance-aware semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2359–2367, 2017. [2](#)
- [26] Ke Xu, Kaiyu Guan, Jian Peng, Yunan Luo, and Sibao Wang. Deepmask: An algorithm for cloud and cloud shadow detection in optical satellite remote sensing images using deep residual network. *arXiv preprint arXiv:1911.03607*, 2019. [2](#)
- [27] Hexiang Hu, Shiyi Lan, Yuning Jiang, Zhimin Cao, and Fei Sha. Fastmask: Segment multi-scale object candidates in one shot. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 991–999, 2017. [2](#)
- [28] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6154–6162, 2018. [2](#), [3](#), [5](#)
- [29] 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. [2](#)
- [30] Daniel Bolya, Chong Zhou, Fanyi Xiao, and Yong Jae Lee. Yolact: Real-time instance segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9157–9166, 2019. [2](#), [4](#), [6](#)
- [31] Xinlong Wang, Rufeng Zhang, Tao Kong, Lei Li, and Chunhua Shen. Solov2: Dynamic and fast instance segmentation. *Advances in Neural information processing systems*, 33:17721–17732, 2020. [2](#), [5](#), [6](#)
- [32] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Solo: A simple framework for instance segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [2](#)
- [33] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *European conference on computer vision*, pages 282–298. Springer, 2020. [2](#), [6](#)
- [34] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask: Single shot instance segmentation with polar representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12193–12202, 2020. [2](#)
- [35] Enze Xie, Wenhai Wang, Mingyu Ding, Ruimao Zhang, and Ping Luo. Polarmask++: Enhanced polar representation for single-shot instance segmentation and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. [2](#)
- [36] 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. [2](#)
- [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 *International Conference on Learning Representations*, 2021. [2](#)
- [38] 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](#)
- [39] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. [2](#)
- [40] Zhiqi Li, Wenhai Wang, Enze Xie, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, Ping Luo, and Tong Lu. Panoptic segformer: Delving deeper into panoptic segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1280–1289, 2022. [2](#)
- [41] Bowen Cheng, Anwesa Choudhuri, Ishan Misra, Alexander Kirillov, Rohit Girdhar, and Alexander G Schwing. Mask2former for video instance segmentation. *arXiv preprint arXiv:2112.10764*, 2021. [2](#), [3](#), [5](#), [6](#), [7](#)
- [42] 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](#)
- [43] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527*, 2022. [3](#), [4](#), [7](#)
- [44] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. [3](#), [5](#), [6](#), [7](#), [8](#)
- [45] Thibaut Durand, Taylor Mordan, Nicolas Thome, and Matthieu Cord. Wildcat: Weakly supervised learning of deep convnets for image classification, pointwise localization and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 642–651, 2017. [3](#)
- [46] Bin Jin, Maria V Ortiz Segovia, and Sabine Susstrunk. Weakly supervised semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3626–3635, 2017. [3](#)

- [47] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4981–4990, 2018. 3
- [48] Ruochen Fan, Qibin Hou, Ming-Ming Cheng, Gang Yu, Ralph R Martin, and Shi-Min Hu. Associating inter-image salient instances for weakly supervised semantic segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 367–383, 2018. 3
- [49] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *European conference on computer vision*, pages 347–365. Springer, 2020. 3
- [50] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. 3
- [51] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. 3
- [52] Daquan Zhou, Zhiding Yu, Enze Xie, Chaowei Xiao, Animesh Anandkumar, Jiashi Feng, and Jose M Alvarez. Understanding the robustness in vision transformers. In *International Conference on Machine Learning*, pages 27378–27394. PMLR, 2022. 3
- [53] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 4
- [54] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 4
- [55] Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep learning in medical image analysis and multimodal learning for clinical decision support*, pages 240–248. Springer, 2017. 5
- [56] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. *Advances in neural information processing systems*, 24, 2011. 5
- [57] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. 5
- [58] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [59] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 5
- [60] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017. 5
- [61] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. 5
- [62] Jifeng Dai, Kaiming He, and Jian Sun. Instance-aware semantic segmentation via multi-task network cascades. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3150–3158, 2016. 6
- [63] 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. 6
- [64] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, volume 139, pages 10347–10357, July 2021. 7, 8