

High Quality Entity Segmentation

Lu Qi^{1*}, Jason Kuen^{2*}, Tiancheng Shen^{4*}, Jiuxiang Gu², Wenbo Li⁴,
 Weidong Guo^{3†}, Jiaya Jia⁴, Zhe Lin², Ming-Hsuan Yang¹,

¹The University of California, Merced ²Adobe Research

³QQ Brower Lab, Tencent, ⁴The Chinese University of Hong Kong




Figure 1: Segmentation results on in-the-wild test images: (1) Mask2Former [7] trained for panoptic segmentation [20] on COCO dataset [31]; (2) Mask2Former trained for entity segmentation [41] on COCO dataset [31]; (3) our CropFormer trained on the proposed EntitySeg Dataset. For a fair comparison, all three models use Swin-Large backbone [33] and are trained until full convergence. Our approach provides far more desirable and useful results for many real-world applications.

Abstract

Dense image segmentation tasks (e.g., semantic, panoptic) are useful for image editing, but existing methods can hardly generalize well in an in-the-wild setting where there are unrestricted image domains, classes, and image resolution & quality variations. Motivated by these observations, we construct a new entity segmentation dataset, with a strong focus on high-quality dense segmentation in the wild. The dataset contains images spanning diverse image domains and entities, along with plentiful high-resolution images and high-quality mask annotations for training and testing. Given the high-quality and -resolution nature of the dataset, we propose CropFormer which is designed

to tackle the intractability of instance-level segmentation on high-resolution images. It improves mask prediction by fusing high-res image crops that provides more fine-grained image details and the full image. CropFormer is the first query-based Transformer architecture that can effectively fuse mask predictions from multiple image views, by learning queries that effectively associate the same entities across the full image and its crop. With CropFormer, we achieve a significant AP gain of 1.9 on the challenging entity segmentation task. Furthermore, CropFormer consistently improves the accuracy of traditional segmentation tasks and datasets. The dataset and code are released at <http://luqi.info/entityv2.github.io/>.

*Equal contribution. † indicates corresponding author.

1. Introduction

Entity segmentation [41] is an emerging segmentation task that focuses on open-world and class-agnostic dense image segmentation, designed to have superior generalization capabilities for unseen-category segmentation. As such, it has great potential in image editing applications, where users can interact with visual entities that exist “in the wild”. Recent works have leveraged entity segmentation for image manipulation [51, 54] and image retouching [34], where users usually specify in-the-wild text queries (*e.g.*, *Oculus*, *platypus*) for local editing. In comparison to traditional closed-set category-based segmentation tasks such as panoptic/instance/semantic segmentation, entity segmentation has the advantage of being all-inclusive and unrestricted, making it easier to find masks that semantically match well with user-specified queries.

Previous work on entity segmentation [41] has shown promising results, yet it only employed existing panoptic segmentation datasets to verify the feasibility and benefits of entity segmentation. However, existing panoptic segmentation datasets (*e.g.*, COCO [31], Cityscapes [9], ADE20K [65]) suffer from the same issues of closed-set categories that entity segmentation [41] aims to avoid. These datasets are collected and annotated for their original panoptic segmentation task requirements. In view of the lack of a proper entity segmentation dataset, it is difficult to further study and explore on this important task. To address this issue, we present EntitySeg dataset, a new large-scale and high-quality dataset for entity segmentation.

Apart from in-the-wild generalization, we design EntitySeg to prioritize high-quality mask annotations and high-resolution images, which are essential for user-driven image editing in the current era of high-resolution imagery. Image editing workflows usually require pixel-precise segmentation masks to ensure a smooth user experience. However, existing panoptic segmentation datasets generally contain low-resolution images (*e.g.*, COCO[31] has zero images that are at least 1K) and coarsely-annotated masks, making it difficult for models trained on such datasets to produce accurate mask predictions for higher resolution images. This is why we focus on high-resolution images and high-quality masks in the EntitySeg dataset, which clearly distinguishes it from existing datasets.

The EntitySeg dataset poses new challenges to existing segmentation methods due to its high-resolution images and high-quality mask annotations. Standard practices of resizing images and mask annotations to 800px used in existing methods [43, 15, 7], result in a heavy loss of fine-grained details. On the other hand, training or evaluating existing segmentation models directly on our high-resolution/quality dataset is intractable due to excessive computation and memory overheads. A naive workaround is to divide the high-resolution image into smaller patches and apply seg-

mentation model to them separately, but it is unable to exploit global context to resolve ambiguity. Moreover, it is unclear how to combine mask predictions from multiple patches using state-of-the-art Transformer-based segmentation methods which rely on set-prediction [7, 6, 50, 61], as shown in Fig 4 and our supplementary material.

To this end, we introduce CropFormer, a novel approach for tackling the problem of high-resolution entity segmentation. CropFormer learns to dynamically generate tokens to associate the full image with its high-resolution crops and then fuses them to generate the final mask prediction. During training, one of the predefined four corner crops is randomly selected and fused with the full image. For inference, CropFormer is applied to all four corner crops and the full image to achieve complete fusion. The strength of CropFormer lies in its ability to dynamically generate queries to link the same entities in the full image and its higher-resolution crops, allowing us to benefit from both global image context and high-quality image/annotation details to obtain improved results.

The main contributions of this work are as follows:

- We present a large-scale, high-quality entity segmentation dataset, comprising of images collected from various domains with diverse resolutions and high-quality masks. Our dataset allows for the evaluation of the generalization and robustness of segmentation models in a unified and open-world manner, due to its domain-diverse and high-resolution setting.
- We benchmark existing segmentation models on our newly collected dataset. We find that the traditional training setting cannot accommodate our dataset, which comprises over 70% images with very high resolution.
- We propose a novel model, CropFormer, with an innovative association module and batch-level decoder to enable the fusion of the full image and its higher-resolution crops.
- Extensive experiments and thorough analysis demonstrate the efficacy of the proposed model and dataset on various segmentation tasks and settings. In Fig. 1, we compare the results from our best CropFormer (trained on EntitySeg) with the best Mask2Formers [7] trained on COCO-Panoptic [20] and COCO-Entity [41].

2. Related Work

Image Segmentation Dataset. Numerous datasets have been proposed for semantic, instance, and panoptic segmentation, *e.g.*, Microsoft COCO [31], ADE20K [65], KITTI [12], LVIS [14], PASCAL-VOC [11], Open-Images [23], Cityscapes [9], and so on. In addition, there are some datasets specially designed for specific scenarios, such as amodal segmentation (COCO-Amodal [66], KINS [39]), human segmentation (CelebAMask-HQ [25],

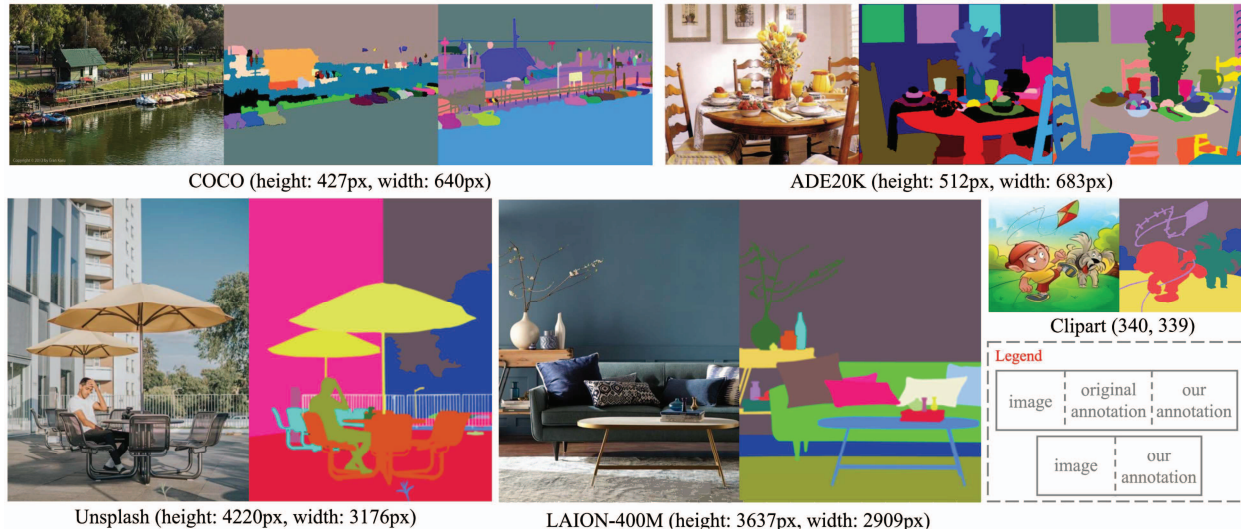


Figure 2: High-quality mask annotations for low- and high-resolution images collected from existing datasets such as COCO [31], ADE20K [65] and Cityscapes [9] as well as from the internet are presented. For the images collected from the aforementioned datasets, a visual comparison between the original and our annotations is provided in the middle and rightmost sub-figures, where the unannotated regions are shaded in black. It is worth noting that the RGB and mask images shown here have been downsampled, resulting in some quality degradation compared to the original datasets.

LIP [13], MHP [63]) and domain adaptation (Synscapes [57], GTA5 [44]). Despite the significant contributions from these datasets, there is still a need to fulfill the requirements of real-world applications with high-quality images of large diversity. For example, a segmentation model should be robust to high-resolution images from different domains and able to segment unseen objects in the open-world setting. Most relevant to our work is the ADE20K dataset [65], which has large-scale open-vocabulary categories. However, the collected images are of low resolution and from narrowed domains. In contrast, our dataset collects images from multiple domains, including indoor, outdoor, street scenes, cartoon, and remote images, with over 80% of the resolution of the images falling within the high-resolution range of 2000px to 8000px. Differently from the ADE20K [64] and LVIS [14] dataset with a predefined category list, our annotation process first conducts class-agnostic mask annotation on each entity and then labels the category information in an open-vocabulary manner.

Scene Parsing Methods. Scene parsing methods mainly rely on convolution-based dense prediction or transformer-based query prediction. Explicit localization information, such as bounding box proposals or pixel position, is often used for pixel-level mask prediction in these approaches. Examples include FCN [35], DeepLab [5], PSPNet [62], Mask R-CNN [15], PANet [32], SOLO [53], CondInst [48], PanopticFPN [19], and PanopticFCN [28]. Inspired by the success of transformer-based detection tasks such as DETR [4] [29, 30, 10, 40], a number of methods, such as Max-Deeplab [50] and Mask2Former [7], have been pro-

posed to directly predict segmentation masks without relying on bounding boxes. The key to these methods is to use learnable queries trained by a set-prediction loss to model the implicit location of the potential instance area. However, query-based set-prediction methods cannot effectively perform multi-view fusion, as the query-to-object assignment can drastically change with the change of image view. This is because the set prediction loss typically adopted by such methods does not enforce any assignment constraints across multiple image views.

3. EntitySeg Dataset

The EntitySeg dataset contains 33,227 images with high-quality mask annotations, as illustrated in Fig. 2. Compared with existing datasets, it has three distinctive features. Firstly, 71.25% and 86.23% of the images are of high quality, with at least 2000px × 2000px and 1000px × 1000px, respectively, staying in line with the current digital imaging trends. Secondly, our dataset is open-world and not restricted to predefined classes. We consider each semantically-coherent region in the images as an entity, even if it is blurred or hard to recognize semantically. Thirdly, the mask annotation along the boundaries is more precise than those in other datasets. In the following, we will provide an in-depth description of our EntitySeg dataset, as well as an extensive analysis.

3.1. Image Collection

Inspired by the collection criterion of COCO [31], we collected most non-iconic images with plentiful contextual information and non-canonical viewpoints. To en-

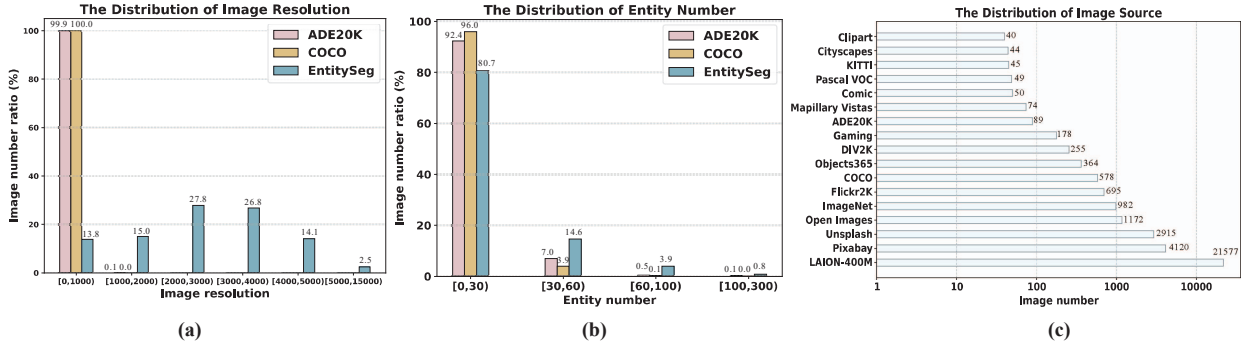


Figure 3: Sub-figures (a) and (b) present the distributions of image resolutions and average number of entities among ADE20K, COCO, and EntitySeg, respectively. Moreover, sub-figure (c) displays the distribution of image sources from which the EntitySeg images were collected.

sure substantial domain diversity, we have sourced the images from various public datasets and the Internet, with permission for academic research. Specifically, part of the images are from COCO [31], ADE20K [65], Pascal VOC [11], Cityscapes [9], Mapillary Vistas [38], ImageNet [22], Open Images [23], DIV2K [3], Flickr2K [55], Clipart [17], Comic [17], DOTA [59], Synscapes [57] and GTA5 [44]. Moreover, we have crawled high-resolution images from Pixabay [1], Unsplash [2], and LAION-400M [46]¹. As shown in Fig. 3(c), most of the images are from high-resolution sources like Pixabay, Unsplash, and LAION-400M. Moreover, Fig. 3(a) reveals that our dataset has more high-resolution images, with resolutions normally distributed within the range of 0px to 15,000px, in comparison to ADE20K and COCO with images all under 1000px.

3.2. Image Annotation

For our dataset, the annotation process consists of the following steps. Firstly, entities in the image are annotated with class-agnostic pixel-level masks, ensuring that no overlaps exist between masks. Subsequently, each mask is annotated with a class label in an open vocabulary manner. Two special considerations should be taken into account here: 1) we annotate the entity as ‘unknown’ or ‘blurred’ if it cannot be easily named or is severely blurred; 2) the entity is annotated as a ‘supercategory_other’ if we only know it is a ‘supercategory’ but unable to determine its fine-grained class. Our annotation process is the opposite of those of the popular segmentation datasets like COCO [31] and ADE20K [64], which first predefine their class sets before annotating masks based on the predefined classes.

Table 2 indicates that our dataset has a more considerable proportion of covered image areas and more masks on average compared to COCO- and ADE20K-Panoptic. This is attributed to our novel annotation process that enables us to take all semantically-coherent regions into account and assign class labels accordingly. Besides, our annotation procedure is more similar to the human visual system. As

¹We removed not-safe-for-work images via strict CLIP filtering.

	ann ₁	ann ₂	ann ₃
ann ₁	-	-	-
ann ₂	90.6	-	-
ann ₃	91.2	92.1	-

(a)

	ann ₁	ann ₂	ann ₃
ann ₁	-	-	-
ann ₂	95.4	-	-
ann ₃	94.8	95.2	-

(b)

Table 1: Annotation consistency of class-agnostic localization and class-aware categories among annotators.

Dataset	ImageRes (avg)↑	EntityNum (avg)↑	EntityNum (max)↑	Entity Complexity↓	Entity Simplicity↓	Valid Area↑
COCO	522.5	11.2	95	0.758	0.581	0.891
ADE20K	461.3	13.6	255	0.802	0.606	0.914
EntitySeg	2700.7	18.1	236	0.719	0.538	0.999

Table 2: The statistical comparisons between COCO [31], ADE20K [65] and EntitySeg. ‘ImageRes (avg)’, ‘EntityNum (avg)’, ‘Valid Area’, ‘Entity Complexity’, and ‘Entity Simplicity’ refer to the average value of image resolution size, entity numbers per image, valid area ratio per image, average entity complexity and simplicity respectively. ‘EntityNum (max)’ means the maximum per-image number of entities within each dataset.

evidenced in [36], the human vision system is intrinsically class-agnostic and can recognize entities without comprehending their usage and purpose.

Annotation consistency is crucial to any human-labeled dataset since it implies whether the annotation task is well-defined and proper. To study this, we randomly selected 500 images from EntitySeg (1.5% of the entire dataset), and asked another two annotators to annotate them after four months since the first round. Table 1(a) shows that the class-agnostic mask mAP in the first step of the annotation process compared to those of the other two annotators. We use one as ground truth and another one as prediction results. In Table 1(b), we also evaluate the category consistency by accuracy (ACC) under the same mask annotations. These two tables demonstrate that our EntitySeg has high annotation consistency in the mask and category labeling stages.

3.3. Dataset Statistics

In our EntitySeg dataset, we annotate all images at the entity level, regardless of whether they belong to thing or stuff in previous datasets. As shown in Fig. 1, we show

some low- and high-resolution photos and their mask annotations to highlight our high-quality pixel-level annotation quality. For example, one can zoom in on the Unsplash image to see how high-quality the railing’s mask is.

Table 2 shows the quantitative comparison among our dataset, COCO-Panoptic [31], and ADE20K-Panoptic [65]. In the EntitySeg dataset, each image has 18.1 entities on average, which is more than 11.2 and 13.6 entities in COCO and ADE20K. For a fair comparison, we treat each uncountable stuff region as one segment although it is disconnected. The detailed comparison among those three datasets on the distribution of entity number is shown in Fig. 3(b). Furthermore, the shapes of entities in our dataset are more complex than those COCO [31] and ADE20K [65] as represented by the columns ‘Entity Complexity’ and ‘Entity Simplicity’ of Table 2 where an entity with a large convexity and simplicity value means it is a simple shape (and both metrics achieve their maximum value of 1.0 for a circle [66]). More details about how complexity and simplicity are calculated can be found in the supplementary material.

Class-Aware Annotation. We select a subset of 11,580 images from the dataset for class labeling, for investigating the diversity of semantic classes in our dataset. In our class annotation process, the class is labeled in an open-vocabulary and free-form manner, then the class set is continually expanded. After the class annotation process, we organize and merge the annotated classes according to Word-Tree [45]. This results in 535 things and 109 stuff classes for those 11,580 images that we use to construct the class-aware EntityClass dataset. Please refer to our supplementary file for more analysis on EntityClass dataset.

4. CropFormer

We propose CropFormer which is the first Transformer-based set-prediction segmentation method that is capable of fusing the mask predictions from multiple image views. Inspired by existing methods [8, 7, 50, 61], we design CropFormer to learn N number of K -dimensional queries $\mathbf{Q} \in \mathbb{R}^{N \times K}$ to generate mask embeddings $\mathbf{E} \in \mathbb{R}^{N \times 1 \times 1 \times 1 \times K}$. \mathbf{E} then acts as convolution filters that are applied to pixel-level mask features $\mathbf{P}_2 \in \mathbb{R}^{T \times H \times W}$ to generate N segmentation masks $\mathbf{U}^m \in \mathbb{R}^{N \times T \times H \times W}$, where T , H and W correspond to the image view, height, and width dimensions. However, such a rudimentary design can only work well with single-view inputs and cannot effectively fuse the results from multiple views. Here, we introduce a novel association module and batch-level decoder to achieve the goal of CropFormer, which is to exploit global context from full image and fine-grained local details from crops for high-quality segmentation. CropFormer is illustrated in Fig. 4.

Crop Dataloader. In CropFormer framework, the crop dataloader is designed to generate a batch of images that si-

multaneously include the full images and their corresponding crops. There are two steps in our dataloader: *crop* and *resize*. First, we augment the full image I^o with a crop that is part of the full image. The crop is randomly extracted from one of the fixed image corners: upper-left, upper-right, bottom-left, bottom-right. The crop size is controlled by a fixed ratio hyperparameter $\delta \in \mathbb{R}$ relative to the full image size. Then, we resize both the full image and crop to the same size. In this way, we construct an input tensor $\mathbf{I} \in \mathbb{R}^{1 \times 2 \times H' \times W' \times 3}$ with 2 elements (full image & 1 crop) along the view dimension. Compared to the full image, our corner crops preserve more image details (*e.g.*, local information) that are useful for high-quality segmentation. Specifically, given an dataloader \mathcal{D} , we define the data pre-processing as:

$$\mathbf{I} = \{I^o, I^c\}, I^o \sim \mathcal{D} \quad (1)$$

where I^o is sampled from \mathcal{D} , and I^c is the cropped image controlled by fixed ratio δ . Unlike the random cropping in [58, 16], we adopt a fixed cropping solution to keep the training and inference consistent. Such a rigid cropping strategy provides better inference efficiency since, rather than bunches of random crops, evaluating a fixed number of crops during inference is sufficient enough.

Image Encoder and Decoder. Based on learnable image-level queries $\mathbf{Q}_i \in \mathbb{R}^{N \times K}$, we use image encoder Θ and image-level decoder Φ_i to generate image-level embeddings $\mathbf{E}_i \in \mathbb{R}^{N \times 2 \times 1 \times 1 \times K}$ for the full input image and its crop. Given the input tensor (\mathbf{I}) and queries (\mathbf{Q}_i), we define \mathbf{E}_i as follows:

$$\mathbf{E}_i = \Phi_i(\mathbf{Q}_i, \Theta(\mathbf{I})), \quad (2)$$

where $\Phi_i(\cdot)$ is a Transformer-based image-level decoder. \mathbf{E}_i is used for image-level entityness prediction and pixel-level prediction using the low-level image feature \mathbf{P}_2 (derived from the image encoder). The prediction process can be formulated as:

$$\mathbf{U}_i^e, \mathbf{U}_i^m = \text{PredHead}_i(\mathbf{E}_i, \mathbf{P}_2), \quad (3)$$

where \mathbf{U}_i^e and \mathbf{U}_i^m denote the entityness prediction and pixel-level mask outputs. Here, we employ i subscript to differentiate image-level embeddings and mask outputs from those of the association module.

Association Module. Following [6], our association module aims to generate batch queries \mathbf{Q}_b that are fully shared by the full image and its crop to represent the same entities consistently. However, instead of treating \mathbf{Q}_b as learnable network parameters, considering that \mathbf{E}_i already contains strong segmentation features, we generate \mathbf{Q}_b directly from $\mathbf{E}_i = \{\mathbf{E}_{I^o}, \mathbf{E}_{I^c}\}$. Particularly, we use a Transformer architecture with cross-attention (f_{XAtt}) and self-

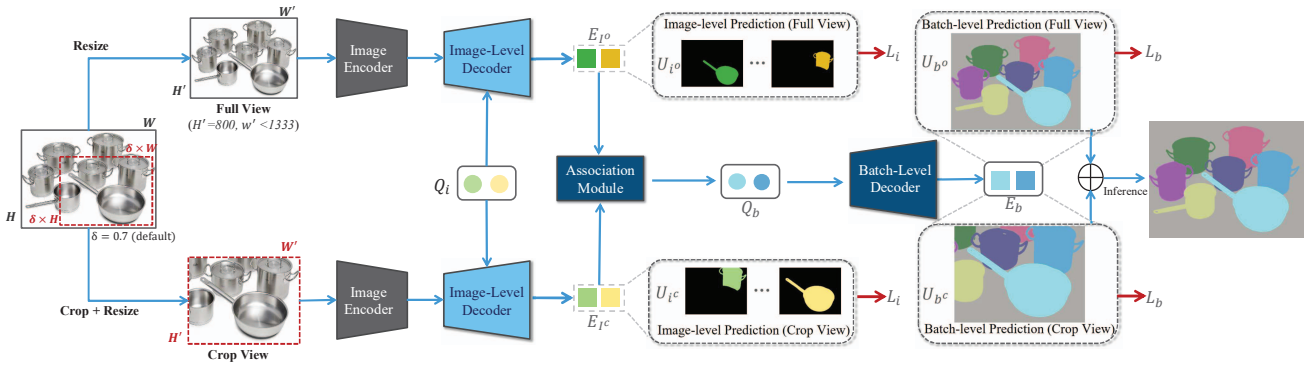


Figure 4: Framework of the proposed CropFormer. The red box indicates cropped region randomly sampled from four fixed image corners. In image-level prediction, the same entity across different image views may be assigned to different queries. Our association module and batch decoder can effectively associate the same entity across different views with a single query.

attention (f_{SAAtt}) to obtain \mathbf{Q}_b :

$$\mathbf{Q}_b = \text{FFN}(f_{\text{SAAtt}}(f_{\text{XAtt}}(f_q(\mathbf{E}_{I^o}), f_k(\mathbf{E}_i), f_v(\mathbf{E}_i))))), \quad (4)$$

where FFN is a feed-forward network, $f_{\{q,k,v\}}(\cdot)$ are linear transformations. Considering the importance of full image for entity segmentation, inspired by [49, 26] we take the image-level embeddings of full image \mathbf{E}_{I^o} as *query*, while treating all image-level embeddings \mathbf{E}_i as *key* and *value*.

Batch-level Decoder. Given $\mathbf{Q}_b \in \mathbb{R}^{N \times 1 \times 1 \times 1 \times K}$, we obtain batch embeddings $\mathbf{E}_b \in \mathbb{R}^{N \times 1 \times 1 \times 1 \times K}$:

$$\mathbf{E}_b = \Phi_b(\mathbf{Q}_b, \Theta(\mathbf{I})), \quad (5)$$

where $\Phi_b(\cdot)$ denotes the batch-level decoder. Thus, we broadcast \mathbf{E}_b to the shape of $N \times 2 \times 1 \times 1 \times K$ to share it between the full image and its crop. Finally, the broadcasted \mathbf{E}_b and low-level batch features \mathbf{P}_2 are used for batch-level mask classification and pixel-wise prediction:

$$\mathbf{U}_b^e, \mathbf{U}_b^m = \text{PredHead}_b(\mathbf{E}_b, \mathbf{P}_2), \quad (6)$$

For the image-level decoder, the queries are separate for each image/crop. Whereas, for the batch-level decoder, the queries are shared between the image and its crop. We illustrate the differences of the image- and batch-level decoders in Fig. 5, where we regard two image views as an entirety.

Training. During training, we employ two separate losses \mathcal{L}_i and \mathcal{L}_b for image- and batch-level predictions, with respect to ground truth \mathbf{G} . The main difference between the two losses lies in whether the same entities in full image and crop are tied to the same queries or not.

The overall training loss for our model is defined as:

$$\mathcal{L} = \sum_{k \in \{i,b\}} \mathcal{L}_k^{\text{ce}}(\mathbf{U}_k^e, \mathbf{G}_k^e) + \sum_{k \in \{i,b\}} \mathcal{L}_k^{\text{bce}}(\mathbf{U}_k^m, \mathbf{G}_k^m) + \sum_{k \in \{i,b\}} \mathcal{L}_k^{\text{dice}}(\mathbf{U}_k^m, \mathbf{G}_k^m), \quad (7)$$

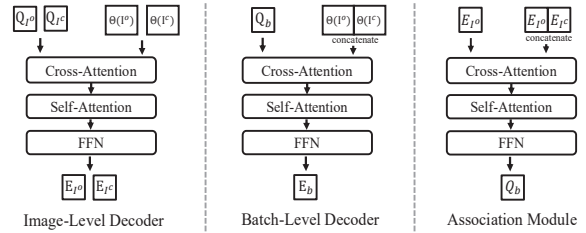


Figure 5: The illustration of image- and batch-level decoder and association module.

where $\mathcal{L}_i^{\text{ce}}$ and $\mathcal{L}_b^{\text{ce}}$ denote binary cross-entropy loss for image- and batch-level entityness prediction. Similarly, $\mathcal{L}_i^{\text{bce}}$, $\mathcal{L}_b^{\text{bce}}$, $\mathcal{L}_i^{\text{dice}}$ and $\mathcal{L}_b^{\text{dice}}$ denote the binary cross-entropy and dice loss for image- and batch-level mask prediction.

Inference. In CropFormer, the same entities between the full image and its crop are represented by the same queries by virtue of the association module. For the final segmentation output, we fuse the per-pixel mask predictions obtained from the full image and every crop of the 4 corners by using an average operation. The confidence score of each entity comes from the batch-level entityness prediction score.

5. Experiments

We conduct our main experiments on the entity segmentation with proposed EntitySeg Dataset in this work. First, we benchmark our dataset in several segmentation tasks, including entity, instance, semantic and panoptic segmentation. Then, we ablate our proposed CropFormer to demonstrate its effectiveness. Note that, in Sec. 4, we omit the conventional *data batch dimension* to improve clarity, but in practice, we train CropFormer with a batch of multiple dataset images and their crops, following existing methods.

Training Details. We leave out some details like the training schedule and some ablation studies due to paper length limitations (see supplementary for details).

Task	Metrics	Method	Backbone	Pretrain	Performance
SEM	mIoU	Deeplabv3	R-50	ImageNet	27.9
		Mask2Former	R-50	ImageNet	37.8
			R-50	COCO	43.3
			Swin-T	COCO	45.0
			Swin-L	COCO	50.5
INS	AP	Mask-RCNN	R-50	ImageNet	5.0
		Mask2Former	R-50	ImageNet	13.0
			R-50	COCO	20.3
			Swin-T	COCO	22.7
			Swin-L	COCO	30.3
PAN	PQ	PanopticFPN	R-50	ImageNet	3.6
		Mask2Former	R-50	ImageNet	5.5
			R-50	COCO	9.6
			Swin-T	COCO	9.8
			Swin-L	COCO	13.4
ENT	AP ^e	Mask-RCNN	R-50	ImageNet	24.9
		Mask2Former	R-50	ImageNet	26.0
			R-50	COCO	35.2
			Swin-T	COCO	42.9
			Swin-L	COCO	46.2

Table 3: Benchmark on several segmentation tasks on the EntitySeg Dataset. The ‘SEM’, ‘INS’, ‘PAN’, and ‘ENT’ indicate the task of semantic, instance, panoptic and entity segmentation. We note that we leave out all details and more experiments in the supplementary file due to paper length limitations.

Transfer Dataset	Type	SEM (mIoU)	INS (AP)	PAN (PQ)	ENT (AP ^e)
ADE20K	Image	57.2(+1.0)	36.0(+1.1)	49.0(+0.9)	40.6(+1.5)
YVIS2019 [60]	Video	-	61.3(+0.9)	-	-
VIPSeg [37]	Video	-	-	37.0(+1.1)	-
UVO [52]	Video	-	-	-	29.5(+2.2)

Table 4: The relative performance gains of EntitySeg pre-training over COCO-Panoptic pretraining when fine-tuned on various image/video-level segmentation datasets.

5.1. Dataset Benchmarks

Benchmark on EntitySeg Dataset. In Table 3, we provide results on semantic, instance, panoptic, and entity segmentation on EntitySeg dataset with some baseline methods. Similar to existing segmentation datasets like COCO and ADE20K, the use of COCO pre-trained weight and a stronger backbone can benefit each segmentation performance of EntitySeg Dataset. Compared to COCO and ADE20K, EntitySeg dataset is more challenging as evidenced by the smaller performance numbers (e.g., Swin-L Mask2Former’s PQ on COCO-PAN is 57.8%).

Transferring EntitySeg to other general segmentation tasks. Table 4 shows the performance comparison of transferring models pretrained on either COCO or EntitySeg dataset to other image/video-level segmentation tasks. Owing to its emphasis on in-the-wild generalization, EntitySeg provides greater improvements than COCO dataset.

Benefit of EntitySeg for high-quality image segmentation. In Table 5, we transfer Mask2Former pretrained on EntitySeg dataset for finetuning on Dichotomous Image Segmentation [42]. Compared to state-of-the-art ISNet [42]

Method	$F_{\beta}^{m,x} \uparrow$	$F_{\beta}^{w} \uparrow$	$M \downarrow$	$S_{\alpha} \uparrow$	$E_{\phi}^m \uparrow$	$HCE_{\gamma} \downarrow$
IS-Net [42]	.791	.717	.074	.813	.856	1116
Ours	.872	.789	.035	.896	.913	969

Table 5: Evaluation on Dichotomous Image Segmentation (DIS5K dataset [42]). ‘Ours’ indicates Mask2Former pre-trained on EntitySeg and then finetuned on DIS5K.

Train Dataset	Task	LVIS [14]	OCID [47]	CAMO [24]	FSS [27]	Avg. AR
COCO	PAN	24.38	39.94	43.80	64.96	43.27
	ENT	27.75	45.88	59.70	75.60	52.23
SAM-H [21]	Prompt	49.25	55.17	75.33	73.60	63.34
	EntitySeg	33.56	76.63	73.20	86.26	67.41

Table 6: In-the-wild segmentation performance (AR@100 with IoU threshold 0.5). We directly evaluate models trained on either EntitySeg or COCO on large-vocabulary segmentation LVIS [14], object clutter indoor segmentation for robot grab [47], camouflaged instance segmentation CAMO [24], few-shot segmentation FSS [27]. All the models are based on Mask2Former with Swin-Large backbone.

that is carefully designed for the task, Mask2Former pre-trained on EntitySeg outperforms ISNet by a large margin, despite not having any task-specific design. This is by virtue of the high-quality mask annotations in EntitySeg dataset.

Segmentation ‘in the wild’ Table 6 compares in-the-wild segmentation performance between Mask2Formers trained on either EntitySeg or COCO. Here, we directly use the trained models without finetuning on the evaluated datasets. The EntitySeg outperforms COCO-Panoptic [20] and COCO-Entity [41] by large margins.

5.2. A Comprehensive Analysis of CropFormer

We first ablate CropFormer on the entity segmentation task and then report results on other segmentation datasets.

Table 7 shows the benefits of CropFormer on segmentation performance and computation cost compared to the Mask2Former [7], its multi-scale extension, as well as related methods [6, 16] that supports multi-view fusion. The first four rows are our baseline Mask2Former with the single-scale inference (800, 1040, or 2700 for shortest side) and (1333, 1732 or 4500 for longest side) for the full images, where $1040=800 \times \delta$ and $1732=1333 \times \delta$ with $\delta = 0.7$ by default. For the results in the fifth row, we use test-time Hungarian matching to associate the same entities obtained with multi-scale images. We find minimal improvement with such an inference strategy. In the sixth and seventh rows, directly using video-level Mask2Former or VITA, the state-of-the-art video instance segmentation framework, merely brings marginal performance gains. Whereas, using the crop output from CropFormer’s batch decoder achieves a significant AP^e gain as indicated by the last row. By combining the full image and four crop outputs from CropFormer (final row), we obtain even stronger 1.5 AP^e gains compared to the baseline (first row). Compared to the

Method	Input Preprocessing	GPU Type	AP ^e	GPU Memory (Train)	GPU Memory (Inference)	Train Time (GPU day)	Inference Time (ms)
Mask2Former	high-res image $\xrightarrow{\text{resize}}$ (800, 1333)	A100-40G	39.5	5.8G	3.0G	2.4	637
Mask2Former	high-res image $\xrightarrow{\text{resize}}$ (2700, 4500)	A100-40G	OOM	OOM	-	OOM	-
Mask2Former	high-res image $\xrightarrow{\text{resize}}$ (2700, 4500)	A100-80G	40.9	58.8G	10.4G	60.8	3312
V-Mask2Former [6]	$\xrightarrow{\text{resize}}$ (800, 1333)	A100-40G	39.7	10.4G	6.7G	6.8	1503
VITA [16]	high-res image $\xrightarrow{\text{crop}(\delta)}$ high-res crop $\xrightarrow{\text{resize}}$ (800, 1333)	A100-40G	39.9	11.9G	6.8G	10.2	1512
CropFormer		A100-40G	41.0	11.4G	6.9G	9.8	1503

Table 7: Comparison with high-resolution methods on EntitySeg with multi-scale training. ‘(x, y)’ indicates number of pixels along the shortest (x) and longest (y) sides of image. Row 1, 2, & 3 are single-scale Mask2Formers trained and evaluated at various single-scale resolutions. (V-Mask2Former & VITA) mask predictions from full image and four crops are fused with either video-level Mask2Former or VITA.

Iterations (Train)	Mask2Former	CropFormer	Decoder	AP ^c	AP ^e ₅₀	AP ^e ₇₅	δ	AP ^e	Train	Test	AP ^e	XAtt	SAtt	FFN	AP ^e
1×	39.5	41.0 (+1.5)	I-O	39.3	56.7	39.8	0.5	38.5	Random	Fixed (4)	39.7	✓	○	○	40.7
2×	40.2	42.1 (+1.9)	B-O	39.1	56.6	39.7	0.6	40.2	Fixed (4)	Fixed (4)	41.0	✓	○	✓	40.8
3×	40.9	42.8 (+1.9)	B-C	40.2	57.5	40.8	0.7	41.0	Fixed (4)	Fixed (8)	41.3	✓	✓	○	40.8
4×	40.9	42.7 (+1.8)	B-OC	41.0	58.4	41.9	0.8	40.9	Fixed (8)	Fixed (8)	41.0	✓	✓	✓	41.0

(a) metric: AP^c

(b)

(a)

(b)

(c)

Table 8: Ablation study on: (a) training schedules; (b) CropFormer’s decoder combinations. In sub-table(b), the ‘Decoder’ column indicates whether we use the predictions of the full image (‘O’), four cropped patches (‘C’), or both of them (‘OC’) from the ‘Image’ or ‘Batch’ decoder.

Mask2Former baselines with a naive increase in image resolution, CropFormer is far more efficient computationally.

Longer Training Schedules Table 8(a) indicates that CropFormer consistently improves over the baseline Mask2Former when we adopt 2×/3× training schedules. A similar trend is observed when using Swin-L [33] backbone, which is provided in our supplementary file.

Decoder Performance Table 8(b) shows the impact of using different decoder combinations in CropFormer for mask prediction during inference. Using just the crop outputs from CropFormer’s batch decoder (B-C) outperforms using just the full image view (I-O, B-O). Combining both the full image and crops (B-OC) gives the best result.

CropDataLoader Table 9(a) and (b) show the ablation study on the usage of crop ratio and crop type. As indicated by Table 9(a), CropFormer obtains the best performance with a crop ratio δ of 0.7. Smaller crop ratios perform worse and it might be a problem caused by queries that are not robust against drastic image transformation.

Association Module Table 9(c) shows the ablation study on the structure of the association module. There is a slight performance difference between using self-attention or a feedforward module following cross-attention.

Comparison to other high-quality segmentation methods. Table 10(a) shows that our CropFormer significantly outperforms recent high-quality instance-level segmentation methods like Mask Transfiner [18] and PatchDCT [56] on both traditional and boundary APs.

Table 9: Ablation study on the usage of crop ratio δ , crop type and association module in CropFormer. In sub-table (b), ‘Random’ indicates random crops and ‘Fixed (4/8)’ indicates 4 or 8 fixed corner crops. In sub-table(c), ✓ and ○ means whether we use the module or not.

Method	AP ^e	AP ^b	Task	Mask2Former	CropFormer
Mask Transfiner [18]	33.7	25.0	COCO-INS	50.1	51.9
PatchDCT [56]	35.4	27.4	COCO-PAN	57.8	59.1
Mask2Former	39.5	31.2	ADE20K-INS	34.9	36.4
Ours	41.0	33.7	ADE20K-PAN	48.1	49.7

(a)

(b) metrics: AP for INS; PQ for PAN

Table 10: (a) Comparison with high-quality segmentation methods on EntitySeg task/dataset using AP^e and boundary AP (AP^b). (b) evaluation on instance (INS) and panoptic (PAN) segmentation tasks on COCO and ADE20K datasets.

Improvement on class-aware segmentation tasks Table 10(b) shows the class-aware task performance of CropFormer with Swin-L [33] backbone and 3× training schedule. CropFormer can still achieve competitive gains on the instance and panoptic segmentation tasks on COCO and ADE20K datasets. This suggests that CropFormer may provide other benefits that are beyond high-resolution segmentation (*e.g.*, prediction robustness).

6. Conclusion

This paper proposes an EntitySeg dataset and a CropFormer framework for high-quality entity segmentation, with a strong focus on in-the-wild generalization and high-quality dense segmentation. The EntitySeg dataset contains about 33K images from various domains and resolutions, accompanied by high-quality mask annotations for training and testing. A novel CropFormer framework is also proposed to overcome the computation challenges introduced by EntitySeg dataset’s high-quality and high-resolution characteristics. Specifically, it enables fusing multi-view predictions from a low-res full image and its high-res crops by learning batch-level queries. We hope this work will serve as a catalyst for future research on open-world and high-quality image segmentation tasks.

References

- [1] Pixabay. <https://pixabay.com>. Accessed: 2022-04-17. 4
- [2] Unsplash. <https://unsplash.com>. Accessed: 2022-04-17. 4
- [3] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *CVPR Workshops*, 2017. 4
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 3
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *PAMI*, 2017. 3
- [6] 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, 5, 7, 8
- [7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *CVPR*, 2022. 1, 2, 3, 5, 7
- [8] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Pixel classification is not all you need for semantic segmentation. *NeurIPS*, 2021. 5
- [9] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *CVPR*, 2016. 2, 3, 4
- [10] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, 2017. 3
- [11] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 2, 4
- [12] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, 2012. 2
- [13] Ke Gong, Xiaodan Liang, Dongyu Zhang, Xiaohui Shen, and Liang Lin. Look into person: Self-supervised structure-sensitive learning and a new benchmark for human parsing. In *CVPR*, 2017. 3
- [14] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 2, 3, 7
- [15] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 2, 3
- [16] Miran Heo, Sukjun Hwang, Seoung Wug Oh, Joon-Young Lee, and Seon Joo Kim. Vita: Video instance segmentation via object token association. In *NeurIPS*, 2022. 5, 7, 8
- [17] Naoto Inoue, Ryosuke Furuta, Toshihiko Yamasaki, and Kiyoharu Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *CVPR*, 2018. 4
- [18] Lei Ke, Martin Danelljan, Xia Li, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Mask transfiner for high-quality instance segmentation. In *CVPR*, 2022. 8
- [19] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *CVPR*, 2019. 3
- [20] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *CVPR*, 2019. 1, 2, 7
- [21] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *arXiv*, 2023. 7
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012. 4
- [23] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4. *International Journal of Computer Vision*, 128(7):1956–1981, 2020. 2, 4
- [24] Trung-Nghia Le, Yubo Cao, Tan-Cong Nguyen, Minh-Quan Le, Khanh-Duy Nguyen, Thanh-Toan Do, Minh-Triet Tran, and Tam V Nguyen. Camouflaged instance segmentation in-the-wild: Dataset, method, and benchmark suite. *TIP*, 2021. 7
- [25] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *CVPR*, 2020. 2
- [26] Peizhao Li, Jiuxiang Gu, Jason Kuen, Vlad I Morariu, Handong Zhao, Rajiv Jain, Varun Manjunatha, and Hongfu Liu. Selfdoc: Self-supervised document representation learning. In *CVPR*, pages 5652–5660, 2021. 6
- [27] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *CVPR*, 2020. 7
- [28] Yanwei Li, Hengshuang Zhao, Xiaojuan Qi, Yukang Chen, Lu Qi, Liwei Wang, Zeming Li, Jian Sun, and Jiaya Jia. Fully convolutional networks for panoptic segmentation with point-based supervision. *TPAMI*, 2022. 3
- [29] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 3
- [30] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 3
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 1, 2, 3, 4, 5
- [32] Shu Liu, Lu Qi, Haifang Qin, Jianping Shi, and Jiaya Jia. Path aggregation network for instance segmentation. In *CVPR*, 2018. 3
- [33] 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*, 2021. 1, 8

- [34] Zerun Liu, Fan Zhang, Jingxuan He, Jin Wang, Zhangye Wang, and Lechao Cheng. Text-guided mask-free local image retouching. *arXiv preprint arXiv:2212.07603*, 2022. [2](#)
- [35] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015. [3](#)
- [36] David Marr. *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman. Print., 1982. [4](#)
- [37] Jiayu Miao, Xiaohan Wang, Yu Wu, Wei Li, Xu Zhang, Yunchao Wei, and Yi Yang. Large-scale video panoptic segmentation in the wild: A benchmark. In *CVPR*, 2022. [7](#)
- [38] Gerhard Neuhof, Tobias Ollmann, Samuel Rota Bulo, and Peter Kotschieder. The mapillary vistas dataset for semantic understanding of street scenes. In *ICCV*, 2017. [4](#)
- [39] Lu Qi, Li Jiang, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Amodal instance segmentation with kins dataset. In *CVPR*, 2019. [2](#)
- [40] Lu Qi, Jason Kuen, Jiuxiang Gu, Zhe Lin, Yi Wang, Yukang Chen, Yanwei Li, and Jiaya Jia. Multi-scale aligned distillation for low-resolution detection. In *CVPR*, 2021. [3](#)
- [41] Lu Qi, Jason Kuen, Yi Wang, Jiuxiang Gu, Hengshuang Zhao, Zhe Lin, Philip Torr, and Jiaya Jia. Open-world entity segmentation. *TPAMI*, 2022. [1](#), [2](#), [7](#)
- [42] Xuebin Qin, Hang Dai, Xiaobin Hu, Deng-Ping Fan, Ling Shao, et al. Highly accurate dichotomous image segmentation. In *ECCV*, 2022. [7](#)
- [43] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. [2](#)
- [44] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *ECCV*, 2016. [3](#), [4](#)
- [45] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. [5](#)
- [46] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. In *NeurIPS Workshop*, 2021. [4](#)
- [47] Markus Suchi, Timothy Patten, David Fischinger, and Markus Vincze. Easylabel: A semi-automatic pixel-wise object annotation tool for creating robotic RGB-D datasets. In *ICRA*, 2019. [7](#)
- [48] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *ECCV*, 2020. [3](#)
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [6](#)
- [50] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *CVPR*, 2021. [2](#), [3](#), [5](#)
- [51] Jianan Wang, Guansong Lu, Hang Xu, Zhenguo Li, Chun-jing Xu, and Yanwei Fu. Manitrans: Entity-level text-guided image manipulation via token-wise semantic alignment and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10707–10717, 2022. [2](#)
- [52] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. *arXiv preprint arXiv:2104.04691*, 2021. [7](#)
- [53] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. Solo: Segmenting objects by locations. In *ECCV*, 2020. [3](#)
- [54] Yikai Wang, Jianan Wang, Guansong Lu, Hang Xu, Zhenguo Li, Wei Zhang, and Yanwei Fu. Entity-level text-guided image manipulation. *arXiv preprint arXiv:2302.11383*, 2023. [2](#)
- [55] Yingqian Wang, Longguang Wang, Jungang Yang, Wei An, and Yulan Guo. Flickr1024: A large-scale dataset for stereo image super-resolution. In *ICCV Workshops*, 2019. [4](#)
- [56] Qinrou Wen, Jirui Yang, Xue Yang, and Kewei Liang. Patchdct: Patch refinement for high quality instance segmentation. In *ICLR*, 2023. [8](#)
- [57] Magnus Wrenninge and Jonas Unger. Synscapes: A photorealistic synthetic dataset for street scene parsing. *arXiv preprint arXiv:1810.08705*, 2018. [3](#), [4](#)
- [58] Junfeng Wu, Yi Jiang, Song Bai, Wenqing Zhang, and Xiang Bai. Seqformer: Sequential transformer for video instance segmentation. In *ECCV*, 2022. [5](#)
- [59] Gui-Song Xia, Xiang Bai, Jian Ding, Zhen Zhu, Serge Belongie, Jiebo Luo, Mihai Datcu, Marcello Pelillo, and Liangpei Zhang. Dota: A large-scale dataset for object detection in aerial images. In *CVPR*, 2018. [4](#)
- [60] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *ICCV*, 2019. [7](#)
- [61] Qihang Yu, Huiyu Wang, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. k-means mask transformer. In *ECCV*, 2022. [2](#), [5](#)
- [62] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *CVPR*, 2017. [3](#)
- [63] Jian Zhao, Jianshu Li, Yu Cheng, Terence Sim, Shuicheng Yan, and Jiashi Feng. Understanding humans in crowded scenes: Deep nested adversarial learning and a new benchmark for multi-human parsing. In *ACMMM*, 2018. [3](#)
- [64] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6):1452–1464, 2017. [3](#), [4](#)
- [65] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. [2](#), [3](#), [4](#), [5](#)
- [66] Yan Zhu, Yuandong Tian, Dimitris Metaxas, and Piotr Dollár. Semantic amodal segmentation. In *CVPR*, 2017. [2](#), [5](#)