

V3Det: Vast Vocabulary Visual Detection Dataset

Jiaqi Wang^{*1}, Pan Zhang^{*1}, Tao Chu^{*1}, Yuhang Cao^{*2},
Yujie Zhou¹, Tong Wu², Bin Wang¹, Conghui He¹, Dahua Lin^{1,2,3}
¹Shanghai AI Laboratory, ²The Chinese University of Hong Kong,
³Centre of Perceptual and Interactive Intelligence

{wangjiaqi, zhangpan}@pjlab.org.cn, dhlin@ie.cuhk.edu.hk

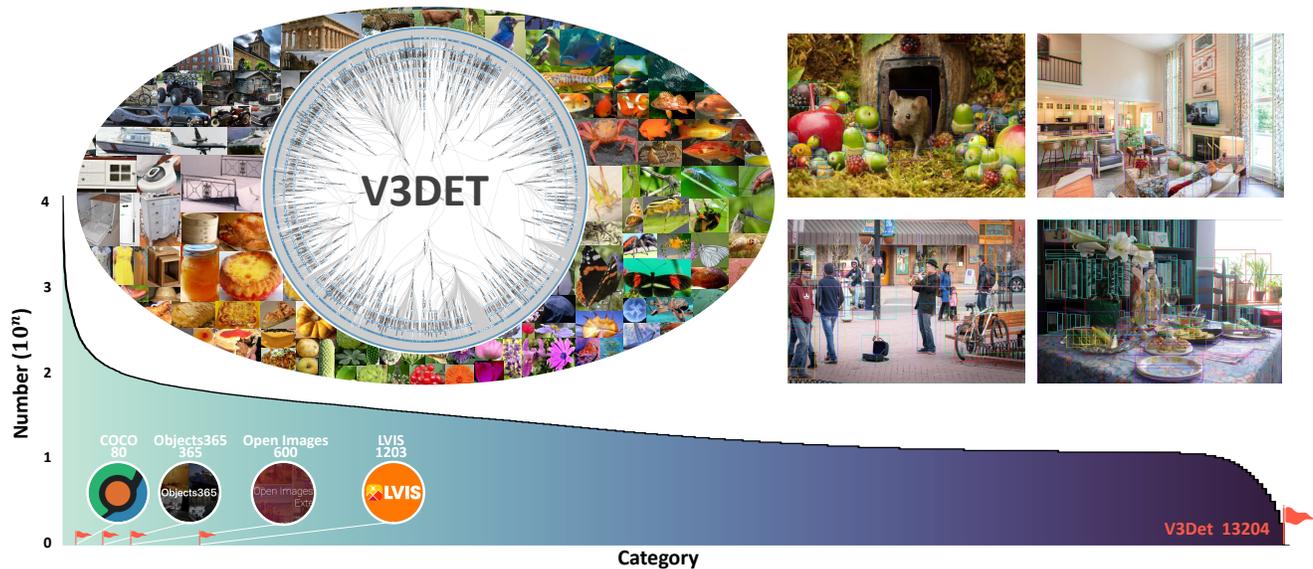


Figure 1: **V3Det** is a **vast vocabulary visual detection** dataset with precisely annotated bounding boxes of 13,204 categories, significantly more diverse than existing visual detection datasets, e.g., COCO [34], Objects365 [45], OpenImages [29], and LVIS [23]. The hierarchical category tree, per-category object distributions, and annotated image samples of V3Det are shown in this figure.

Abstract

Recent advances in detecting arbitrary objects in the real world are trained and evaluated on object detection datasets with a relatively restricted vocabulary. To facilitate the development of more general visual object detection, we propose **V3Det**, a vast vocabulary visual detection dataset with precisely annotated bounding boxes on massive images. V3Det has several appealing properties: **1) Vast Vocabulary:** It contains bounding boxes of objects from 13,204 categories on real-world images, which is 10 times larger than the existing large vocabulary object detection dataset, e.g., LVIS. **2) Hierarchical Category Organization:** The vast vocabulary of V3Det is organized by a hierarchical category tree which annotates the inclusion relationship among categories, encouraging the exploration of category relationships in vast and open vocabulary ob-

ject detection. **3) Rich Annotations:** V3Det comprises precisely annotated objects in 243k images and professional descriptions of each category written by human experts and a powerful chatbot. By offering a vast exploration space, V3Det enables extensive benchmarks on both vast and open vocabulary object detection, leading to new observations, practices, and insights for future research. It has the potential to serve as a cornerstone dataset for developing more general visual perception systems. V3Det is available at <https://v3det.openxlab.org.cn/>.

1. Introduction

Object detection [21, 20, 43, 2, 35, 33, 4, 47, 52, 50, 53, 51, 62] is the cornerstone of various real-world applications, e.g., autonomous driving, robotics, and augmented reality. Taking images as inputs, it localizes and classifies objects within a given vocabulary, where each detected object is

* equal contribution.

denoted as a bounding box with a class label.

Detecting arbitrary objects has been a long-standing goal in computer vision. Since the real world contains a vast variability of objects in countless classes, an ideal visual detection system should be able to detect objects from an extremely large set of classes and be readily applied to open-vocabulary categories. Therefore, it is imperative to equip the community with an object detection dataset with a vast vocabulary, as this can accelerate the exploration toward more general visual detection systems.

While there have been extensive efforts to create object detection datasets, they have only partially fulfilled the requirements for a comprehensive dataset with a vast vocabulary of categories. For example, as shown in Table 1, COCO [34] is one of the most widely adopted datasets comprising 123k images with 80 categories. Objects365 [45] and Open Images [29] are large-scale datasets aimed at pre-training, where Objects365 annotates 638k images with 365 categories, and Open Images contains 1,515k images with 600 categories. While these two datasets provide significant amounts of annotated objects, their limited category numbers still fall short in training class generalizable detectors. LVIS [23] is a large vocabulary object detection and instance segmentation [26, 5, 55, 48, 9, 59, 60, 5] dataset containing 120k images with 1,203 categories. Despite its success in recent advances in large / open vocabulary object detection, the vocabulary of LVIS (1,203 categories) is still insufficient to represent the vast diversity of classes in the real world. Therefore, a new object detection dataset with a significantly larger vocabulary would benefit the training and evaluation of visual detectors.

To facilitate future research on more general object detection, we introduce **V3Det: Vast Vocabulary Visual Detection Dataset**, which comprises a vast vocabulary of precisely annotated objects in real-world images. The proposed dataset has several appealing properties: **1) Vast Vocabulary:** It contains bounding boxes of objects from 13,204 categories on real-world images, which is 10 times larger than the existing large vocabulary object detection dataset, such as 1,203 categories in LVIS. **2) Hierarchical Category Organization:** The vast vocabulary of V3Det is organized by a hierarchical category tree that annotates the inclusion relationship among categories. This tree begins with ‘entity’ category, which is the ancestor of all classes, containing both the categories of annotated objects and their parent categories as shown in Figure 1. This hierarchical tree enriches the vocabulary of V3Det, enabling the exploration of category relationships in open-world recognition. **3) Rich Annotations:** V3Det comprises precisely annotated objects in 243k images. The large image collections ensure the diversity of objects and scenes. Besides category names, V3Det provides rich descriptions of each category written by human experts and a powerful chatbot,

Dataset	Categories	Images	Boxes	Class Tree	Descriptions
Pascal VOC [16]	20	11.5k	27k	No	No
COCO [7]	80	123k	896k	No	No
Objects365 [45]	365	638k	10,101k	No	No
Open Images [29]	600	1,515k	14,815k	Yes	No
LVIS [23]	1,203	120k	1,525k	No	No
ELEVATER [31]	314	132k	938k	No	Yes
BigDetection [1]	600	3,480k	35,960k	No	No
V3Det	13,204	213k	1,514k	Yes	Yes
V3Det (All)	13,204	243k	1,727k	Yes	Yes

Table 1: Dataset statistics comparison between the V3Det and existing object detection benchmarks on the released annotations, including training and validation data. V3Det (All) includes all of training, validation, and testing data of V3Det.

i.e., chatgpt, which can serve as language prompts for open-vocabulary object detection.

Extensive benchmarks are conducted on V3Det to reveal insights into this vast vocabulary object detection dataset. Specifically, we perform extensive experiments on both classical close-set object detection and open-vocabulary object detection. 1) In close-set object detection, we evaluate various wide-adopted object detection frameworks. We introduce the best practices of data samplers, optimizers, and classifier designs to achieve higher performance on V3Det. 2) In the open-vocabulary setting, we benchmark recent representative methods on V3Det, shedding light on the challenges of open-vocabulary object detections on a vast vocabulary of categories. Moreover, we show that pretraining on V3Det significantly boosts the class generalizability of detectors, demonstrating the solid benefits of a vast vocabulary dataset for open vocabulary algorithms.

In summary, we introduce **V3Det: Vast Vocabulary Visual Detection Dataset** with 13,204 categories, hierarchically organized with a category tree. To the best of our knowledge, V3Det is the first object detection dataset with more than ten thousand categories. Rich annotations of detection bounding boxes in 243k images, along with professional category descriptions, ensure the diversity and versatility of the dataset. Taking advantage of the vast exploration space offered by V3Det, extensive benchmarks reveal new observations, practices, and insights for future research on vast and open vocabulary object detection. It has the potential to serve as a cornerstone dataset for developing more general visual perception systems.

2. Related Work

Visual Detection Datasets. A large-scale dataset of broad categories is the key to building a robust visual detection system, and the community has made many efforts. The early PASCAL VOC [16] dataset contains 11.5K images of 20 categories. The MS COCO [34] dataset introduces 123K images of 80 categories to depict complex everyday scenes

of common objects. Objects365 [45], Open Images [29], and BigDetection [1] are large-scale datasets for pretraining, where the largest BigDetection contains 3,480k images of 600 categories. ELEVATER [31] is a multi-domain benchmark that comprises 35 datasets. While these datasets provide significant amounts of annotated objects, their limited category numbers still fall short in training class generalizable detectors. LVIS [23] is the detection dataset of the largest vocabulary size yet, which has 1,203 categories and exhibits a long-tailed distribution. Our V3Det dataset pursues a broader categories distribution that covers 13,204 classes, enabling training a more general detector.

Object Detection. Object Detection aims to localize and classify objects, which has witnessed remarkable progress in recent years. The classical CNN-based detectors can be divided into two sets. Two-stage detectors [21, 20, 43, 2, 67] first predict a set of proposals and then refine them in the second stage. One-stage detectors [33, 42, 35], which directly classify and regress the predefined anchor boxes, or densely searching geometric cues like points [49], centers [15], corners [30]. Recent transformer-based detectors [4, 68, 58] remove hand-designed components like NMS and anchors and show superior performance than CNN-based ones. In this work, we test various methods of different paradigms on our V3Det dataset.

Open-Vocabulary Object Detection. Open-Vocabulary Object Detection (OVD) aims to detect novel classes not seen during training. To cope with the limitation of traditional detectors that detect only a predefined set of categories, OVD approaches usually adopt vision-language models [40, 12] or utilize external data with weak annotations [44, 7, 46]. Specifically, OV-RCNN [57] uses BERT [12] to pretrain Faster R-CNN on a corpus of image-caption pairs. ViLD [22] distills general knowledge from CLIP [40] into a two-stage detector. GLIP [32] and MDETR [27] convert detection as a grounding task with captions. Region-CLIP [65] enhances CLIP to match image regions and textual concepts better. Detic [66] extends the classification capacity to twenty-thousand with image-level annotations from ImageNet-21K [44]. Though these works achieve excellent performance on two benchmarks, MS COCO [34], and LVIS [23], we found by experiments that they can hardly apply in a scenario of vast categories.

3. V3Det Dataset

This section presents a comprehensive exposition of the V3Det dataset, covering its data acquisition, annotation pipeline, and dataset analysis.

3.1. Data Acquisition

The construction of the V3Det dataset is based on the Bamboo [63] classification dataset, where most categories and images utilized in the V3Det dataset are drawn from.

Bamboo Classification Dataset. The extensive Bamboo [63] classification dataset, containing 69 million annotations for image classification across 119,000 visual categories, serves as a rich source of potential images and categories for the V3Det dataset. Although the Bamboo classification dataset is a mega-scale dataset, not all of its categories are suitable for detection purposes, such as ‘sunny day’ and ‘happiness’. Additionally, many classification images are less complex for the detection task. To address these issues, we developed a category construction and image selection process to obtain suitable categories and images for the V3Det dataset.

Category Construction. By carefully checking the available 119,000 visual categories, three researchers with ample expertise in object detection manually select visual concepts suitable for object detection. Each concept is equipped with a category name, a set of descriptions, and classification images. The description is sourced from human experts (e.g., WordNet, Wikidata) or generated by ChatGPT. The concepts unanimously chosen by all researchers will be incorporated into the V3Det dataset, resulting in 11,922 categories. In addition to selecting from the Bamboo classification dataset, we meticulously collect 1,282 object detection categories from web data. As a result, the V3Det dataset comprises a total of 13,204 categories.

Image Selection. For categories obtained from the Bamboo classification dataset, we adopt their corresponding images. For categories collected from web data, we crawl 120 images per category from Flickr ¹ and manually remove any mismatched images. To filter the simple classification images, we employ a region proposal network (RPN) [43] jointly trained on four detection datasets, namely COCO [34], LVIS [23], Objects365 [45], and Open-Images [29], to extract region proposals for each image. Images with proposal counts exceeding 4,000 are removed to avoid overcrowding in the dataset. The remaining 20 images per category with the largest number of proposals are kept in the V3Det dataset.

3.2. Annotation

The overall annotation pipeline of V3Det follows a coarse-to-fine strategy, as shown in Figure 2. In the coarse-grained annotation and joint verification stages, objects belonging to coarse-grained categories and fine-grained image classification labels are annotated with bounding boxes. In the fine-grained annotation and intra-category verification stages, the coarse-grained category of each bounding box is further refined to the corresponding fine-grained category.

Annotation Team. To improve the accuracy of the annotation, we establish four annotation teams. One large Annotator Team performs all annotation steps in the entire annotation pipeline. Two Inspector Teams are assigned to review

¹<https://www.flickr.com>

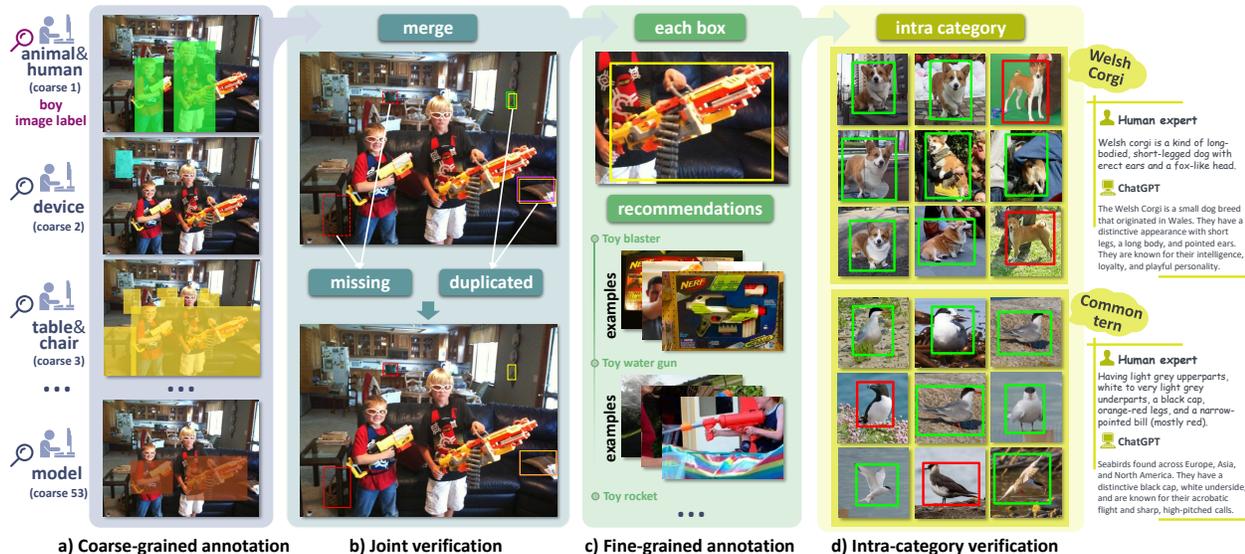


Figure 2: **The illustration of our annotation pipeline.** a) Each group is assigned to annotate one coarse-grained category and image-level categories belonging to the coarse-grained category. b) Annotations from all groups on the same image are merged and presented together for joint verification. c) For each box with coarse-grained annotation, some recommendations are given to the annotators for the fine-grained annotation. d) Objects annotated with the same fine-grained category are presented together for intra-category verification.

all the annotated images. Inspector Team I is responsible for reviewing all the annotations made by the Annotator Team. Once an annotation error is found in a data package (containing 100 images), the package will be returned for re-annotation. The data package approved by Inspector Team I will be further reviewed by Inspector Team II. Upon successful approval by the two Inspection Teams, one Examiner Team will finally examine the data packages. Any annotation errors will result in the data package being returned for re-annotation. The described entire process will be conducted in every stage of the annotation pipeline.

Coarse-grained Annotation. It is not feasible for each annotator to memorize and annotate all 13,204 categories. Therefore, we first divide the 13,204 categories into 53 coarse-grained categories, e.g., ‘animal and human’, ‘vehicle’, and ‘device’. The annotators are then trained to comprehend a single coarse-grained category and memorize the corresponding subcategories. Due to the similar appearance in each subcategory set, it is easier for annotators to complete the above task. To avoid misannotations, the annotators are divided into 53 groups, and each group annotates bounding boxes corresponding to a single coarse category across the entire dataset. For each group, numerous images may not contain any objects corresponding to the assigned coarse-grained category, leading to a large number of skipped images with no annotations. But extensive annotations are necessary to ensure the annotation accuracy.

When annotating coarse-grained bounding boxes, annotators also identify and annotate objects of image-level labels if it belongs to their assigned coarse category. As shown in Figure 2 a), the annotator of coarse group 1 (‘an-

imal and person’) also annotates the bounding box of the image label (‘boy’). Image-level classification labels are either from the Bamboo classification dataset or assigned after crawling images from the internet. Upon completion of the coarse-grained annotation stage, the resulting dataset comprises images annotated with coarse-grained bounding boxes and image-level fine-grained bounding boxes.

Joint Verification. Although the ‘one annotator one coarse-grained category’ strategy significantly reduces the difficulty of the annotations, there are still some annotation problems that need to be addressed. One is duplicated annotation. The other is missing annotation. Due to the mega-scale categories in the dataset, it is hard to ensure that coarse-grained categories are mutually exclusive. For instance, the category ‘truck’ belongs to both the coarse-grained categories ‘vehicles’ and ‘heavy machinery’. As a result, one annotator may annotate a bounding box with ‘vehicles’ for a truck, while another annotator may annotate a box with ‘heavy machinery’ for the same truck. Meanwhile, some fine-grained categories may not fit within a coarse-grained category well, which may lead to missed annotations from a small number of annotators.

To solve the aforementioned problems, a joint verification stage is proposed. As shown in Figure 2 b), the annotations from 53 annotation groups on each image are merged and forwarded to proficient annotators who have demonstrated exceptional skill in the coarse-grained annotation stage. The annotators are tasked with correcting any duplicated or missed annotations. The inspector and examiner teams review and examine this process to ensure the high quality of the joint verification stage.

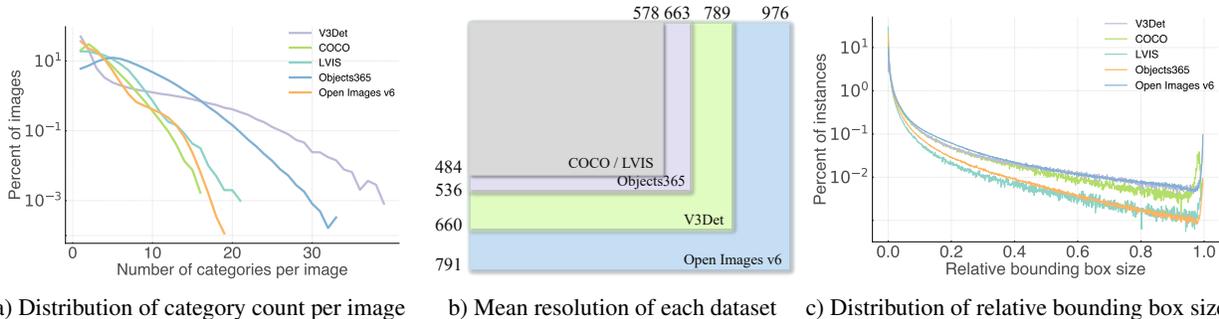


Figure 3: **Dataset analysis.** a) The percentage of images containing more than 16 categories in our V3Det is greater than that in other dataset. b) The mean image resolution of our V3Det is higher than COCO, LVIS and Objects365. c) The object scale in our V3Det is comparable to that of the COCO and Open Images v6 datasets.

Fine-grained Annotation. After completing the previous stages, all objects within the dataset are annotated by bounding boxes with coarse-grained categories or fine-grained image-level categories. This stage aims to refine the coarse-grained category of each bounding box to its corresponding fine-grained category.

To alleviate the complexity of the coarse-to-fine annotation task, fine-grained categories are ranked by leveraging bounding boxes with image-level labels to assist the annotators. Specifically, the image patches are firstly cropped based on annotated bounding boxes, with a fixed padding size of 25 pixels. Then an image encoder of a large pre-trained visual language model, *i.e.*, CLIP (ViT-L) [40] is utilized to extract the discriminative features of the cropped patches. To obtain the representation for a fine-grained category i , we cluster all patch features of category i with k-means [24] to get k prototypes $c_i^1, c_i^2, \dots, c_i^k$. Following CLUE [8], the similarity s_i^x between a coarse-grained image patch x with a fine-grained category i is defined as:

$$s_i^x = \max_{j=1,2,\dots,k} \frac{e(x) \cdot c_i^j}{\|e(x)\| \|c_i^j\|} \quad (1)$$

where $e(x)$ is the extracted feature of patch x , and k is set to 3. Due to mega-scale categories, ranking all fine-grained categories solely based on the similarity score s_i^x gives sub-optimal recommendations. Here we leverage the coarse-grained annotations. Specifically, for a coarse-grained category I , we construct a set of coarse-grained categories \mathcal{I}_m that are more likely to be misannotated as category I . For a given cropped patch annotated with I , a traversal list is constructed with top-20 fine-grained categories belonging to coarse-grained category I , along with top-10 fine-grained categories whose coarse-grained categories is in \mathcal{I}_m .

During the fine-grained annotation stage, annotators are required to match a given cropped patch with all fine-grained categories in the traversal list. If none of the categories in the traversal list is correct, annotators are required to search in all fine-grained categories based on their own knowledge. The search platform provides relevant cat-

egories based on the search keyword as suggestions. As shown in Figure 2 c), the annotators choose the corresponding fine-grained category according to example images and class descriptions.

After the above steps, there is still a small number of bounding boxes that are not successfully matched. To ensure the categories of these bounding boxes are not in the set of 13,204 categories, a small subset of the bounding boxes is sampled from each annotator, and a thorough check is conducted to confirm the results. An error rate greater than 2% results in re-annotation.

Intra-category Verification. After the fine-grained annotation stage, most annotation work of 13,204 categories has been completed. To enhance the accuracy of the annotations, a fine-grained intra-category verification stage is designed, where all the image patches cropped by bounding boxes with the same fine-grained category are displayed together, as in Figure 2 d). Annotators are tasked to verify if there are any misclassified objects or inconsistent annotations within the given fine-grained category with the help of example images and descriptions. The misclassified objects will be re-annotated in the fine-grained annotations stage, and the inconsistent annotations will be directly corrected by the annotator, following the same “largest bounding box” rule as the Objects365 [45] dataset. This stage guarantees the accuracy of our V3Det dataset with 13,204 categories.

3.3. Dataset Analysis

With the proposed annotation pipeline, approximately 243k images have been annotated with a total of 13,204 categories and 1.7M instances. The V3Det dataset exhibits the following main characteristics:

Vast Vocabulary. As indicated in Table 1. The number of the categories in the V3Det dataset is ten times greater than that of previous detection datasets. The extensive categories facilitate the development of the vast-vocabulary object detection. Moreover, since a real open-vocabulary object detector is also a vast-vocabulary object detector, the V3Det dataset, with its vast vocabulary, could provide a more ro-

bust benchmark to evaluate and support open-vocabulary object detection. Figure 3 a) illustrates the full categories-per-image distribution, where the percentage of images containing more than 16 categories in V3Det is greater than that in other datasets, indicating the complexity and difficulty of the proposed V3Det dataset.

High Resolution. As shown in Figure 3 b), the higher image resolution of the V3Det, as compared to COCO, LVIS and Objects365, provides greater visual details that can aid in object detection. Figure 3 c) shows the relative size distribution of bounding boxes. The object scale of V3Det is comparable to that of the COCO and Open Images v6.

High Quality. To validate the annotation quality of the V3Det dataset, three researchers with ample expertise in object detection are asked to annotate 100 images randomly sampled from V3Det dataset and 100 images randomly sampled from LVIS dataset. The results demonstrate that V3Det performs competitively with LVIS in terms of class-specific recall, achieving 90.9% versus 91.2%. In addition, V3Det outperforms LVIS in terms of mean average precision, achieving a higher rate of 86.5% compared to 81.2%.

4. Benchmarks

4.1. Experimental Settings

Split Setup. V3Det is divided into three splits, *i.e.*, a *train* split, a *val* split, and a *test* split. The *train* split has 1,299,765 objects in 183,354 images, the *val* split has 214,416 objects in 29,821 images, and the *test* split has 212,437 objects in 29,863 images. We manually ensure at least one object of each category in the *val* and *test* split. The images and annotations of the *train* and *val* split will be publicly available. We will release the images of the *test* split and set up a public test server to evaluate the uploaded results for fair comparisons. If not further specified, the reported performance is trained on *train* split and evaluated on *val* split. We perform benchmarks on both close-set vast vocabulary object detection and open vocabulary object detection. 1) In the vast vocabulary setting, we take all the 13,204 categories for training and evaluation. 2) In the open vocabulary setting, we randomly sample 6,709 categories as the base classes C_{base} and the remaining 6,495 categories as the novel classes C_{novel} .

Evaluation Protocol. We follow the evaluation metrics of the COCO [7] dataset to report the mean average precise (AP) on different IoU thresholds (*i.e.*, 0.5 \sim 0.95).

Implementation Details. 1) In the vast vocabulary setting, we evaluate two-stage / cascaded detectors [43, 2, 67], single-stage detectors [33, 49, 61], and DETR-style detectors [68, 58]. We take their standard implementations in mmdetection [6], Detectron2 [56], and detrex [11]. All models are trained with AdamW [37] optimizer, multi-scale data augmentation, and repeat factor sampler [23] as the de-

fault setting. 2) In the open vocabulary object detection, we follow the officially released code and settings for training and inference. Please refer to supplementary materials for more implementation details.

4.2. Experiments on Vast Vocabulary Settings

Compared to previous object detection datasets, V3Det contains a vast vocabulary of 13,204 categories. We first perform benchmarks on close-set vast vocabulary object detection to reveal insights and best practices on V3Det.

Main Results. We evaluate detectors with two-stage, cascaded, single-stage, and DETR-style frameworks. Due to the limited resources, we can only test some representative detectors, but we believe the selected frameworks are sufficient to reveal important insights. As shown in Table 2, we report the performance of Faster R-CNN [43], CenterNet2 [67], and Cascade R-CNN [3] for two-stage / cascaded frameworks with ResNet-50 [25] and Swin-B [36] backbones. Experimental results show that cascaded frameworks (e.g., Cascade R-CNN, CenterNet2) significantly outperform two-stage detectors (e.g., Faster R-CNN) with around 8% on ResNet-50 and around 10% on Swin-B. Single-stage detectors, e.g., FCOS [49] and ATSS [61] perform much inferior to Faster R-CNN and Cascade R-CNN on V3Det. The experimental results show that the cascaded refinement of regression and classification is crucial in V3Det. The current leading paradigm DETR-style detectors, e.g., Deformable DETR [68] and DINO [58], achieve strong performance on V3Det. For example, DINO [58] achieve higher performance than Cascade R-CNN [3] with a large margin on ResNet-50 (33.5% v.s. 28.6%) and a clear margin on Swin-B (42.0% v.s. 40.2%).

Moreover, we evaluate the Norm Linear Layer [54], which is effective on large-vocabulary object detection dataset, LVIS. The Norm Linear Layer significantly improves the performance of Faster R-CNN, Cascade R-CNN, and FCOS with different backbones on the V3Det dataset.

Best Practices on V3Det Dataset. We explore the effectiveness of different data samplers, optimizers, and classifiers on V3Det. We take cascade R-CNN w/ ResNet-50 backbone with 12 epochs training as our baseline. In Table 3, we compare random sampler, class balanced sampler [39], and repeat factor sampler [23]. We observe that the repeat factor sampler and the class balanced sampler perform comparably, slightly better than the random sampler (\sim 0.2%). The results are consistent with the fact that V3Det dataset has a relatively balanced category distribution. In Table 4, we perform an ablation study of AdamW optimizer and Norm Linear Layer. Both AdamW optimizer and Norm Linear Layer bring significant gains on V3Det. Therefore, in Table 2, we take the AdamW optimizer and repeat factor sampler as the default setting.

Evaluate Large Foundation Model on V3Det Dataset.

Framework	Method	Epochs	Backbone	AP	AP_{50}	AP_{75}
Two Stage / Cascade	Faster R-CNN [43]	24	ResNet-50	20.1	28.2	22.9
			Swin-B	30.7	40.0	34.7
	Faster R-CNN w/ Norm Linear Layer [54]	24	ResNet-50	25.4	32.9	28.1
			Swin-B	37.6	46.0	41.1
	CenterNet2 [67]	24	ResNet-50	28.3	33.7	30.4
		Swin-B	39.8	46.1	42.4	
Cascade R-CNN [3]	24	ResNet-50	28.6	34.8	31.1	
		Swin-B	40.2	47.3	43.0	
Cascade R-CNN w/ Norm Linear Layer [54]	24	ResNet-50	31.6	37.3	33.7	
		Swin-B	42.5	49.1	44.9	
Single Stage	ATSS [5]	24	ResNet-50	4.4	5.3	4.6
			Swin-B	7.6	8.9	8.0
	FCOS [49]	24	ResNet-50	6.5	8.1	6.9
		Swin-B	15.0	17.9	16.0	
FCOS w/ Norm Linear Layer [54]	24	ResNet-50	9.4	11.7	10.1	
		Swin-B	21.0	24.8	22.3	
DETR Style	Deformable DETR [68]	50	ResNet-50	34.4	39.9	36.4
			Swin-B	42.5	48.3	44.7
DINO [58]	24	ResNet-50	33.5	37.7	35.0	
		Swin-B	42.0	46.8	43.9	

Table 2: We report benchmark results of close-set vast vocabulary object detection setting on V3Det. AdamW [37] optimizer and repeat factor sampler [23] are adopted for better performance in this table. We report AP , AP_{50} , and AP_{75} for comparisons.

Data Sampler	AP	AP_{50}	AP_{75}
Random Sampler	25.7	31.7	28.1
Class Balanced Sampler [39]	25.9	31.8	28.3
Repeat Factor Sampler [23]	25.9	32.0	28.4

Table 3: Study on different data samplers using Cascade R-CNN ResNet-50 with 12 epochs training as the baseline. AdamW optimizer is adopted in this table.

Recent advances in the foundation model dramatically boost the performance on existing object detection benchmarks, e.g., COCO and LVIS. Taking a recent foundation model EVA [17] as an example, its ViT-G [13] version is first pre-trained on Objects365 [45] dataset and then fine-tuned on the target dataset, achieving 64.2% AP on COCO and 62.2% AP on LVIS. Notably, LVIS is always considered a much more challenging dataset than COCO due to its large vocabulary and long-tailed distributions. However, EVA achieves comparable performance on both datasets. We explore whether such a powerful model can still achieve such a strong performance on V3Det. As in Table 5, we report the performance of EVA w/ ViT-G on COCO, LVIS, and V3Det. To have a fair comparison, all training settings and model architecture are the same in experiments. Compared to 60+% AP on COCO and LVIS, there exists still wide space to improve the performance on V3Det.

Performance Differences Between *val* and *test* Split. We evaluate Cascade R-CNN [3] and DINO [58] with 24 epochs and Swin-B backbone on *val* and *test* split of V3Det. The performance on *test* is just slightly higher than *val* split

AdamW	Norm Linear Layer	AP	AP_{50}	AP_{75}
		20.9	25.3	22.6
✓		25.9	31.8	28.3
	✓	27.3	33.8	29.8
✓	✓	28.3	34.5	30.8

Table 4: Study on AdamW optimizer and Norm Linear Layer using Cascade R-CNN ResNet-50 with 12 epochs training baseline.

Dataset	AP	AP_{50}	AP_{75}
COCO [34]	64.2	81.9	70.6
LVIS [23]	62.2	76.2	65.4
V3Det	49.4	54.8	51.4

Table 5: Comparison of the detection performance of a strong foundation model, *i.e.*, EVA [17] on COCO, LVIS, and V3Det.

with 0.6% AP and 0.4% AP on the two methods, respectively, showing the consistency between the two splits.

4.3. Experiments on Open Vocabulary Settings.

Detecting objects of arbitrary categories is a long-standing goal of object detection. In recent years, open vocabulary object detection has attracted more and more attention that aims to detect objects of any category after training on a close set of categories. In this section, we evaluate existing open vocabulary object detection (OVD) works on V3Det for more comprehensively measuring their category generalization ability.

Evaluate Pre-trained OVD models. We first test various

Method	Framework	Training Data	AP	bAP	nAP
PBL [19]	MRCNN-R50	COCO,CCCap,VG,SBU	0.24	0.22	0.26
VL-PLM [64]	FRCNN-R50	COCO	0.10	0.08	0.19
PromptDet [18]	MRCNN-R50	LVIS,LAION	0.61	0.58	0.63
ViLD [14]	MRCNN-R50	LVIS	0.52	0.49	0.55
DetPro [14]	MRCNN-R50	LVIS	0.53	0.52	0.55
GLIP [32]	DyHead[10]-SWT	O365,GoldG,CC3M,SBU	-	-	0.91
OC-OVD [41]	FRCNN-R50	LVIS,ImageNet	0.97	0.92	1.03
RegionClip [65]	FRCNN-R50	LVIS	0.72	0.71	0.72
Detic [66]	CenterNet2-R50	LVIS,ImageNet	1.00	0.95	1.04

Table 6: Evaluation results of directly testing pretrained open-vocabulary object detectors on V3Det. The framework and training data are also reported. AP , bAP , and nAP indicates the performance on all classes, base classes C_{base} , and novel classes C_{novel} of V3Det. GLIP is evaluated only on nAP due to its high inference cost with a large class number. For framework, MRCNN, FRCNN, R-50, and SWT indicate Mask R-CNN, Faster R-CNN, ResNet-50, and Swin Transformer Tiny Model. For training data, CCCap, VG, SBU, O365, and CC3M indicate COCO Caption [7], Visual Genome [28], SBU Caption [38], Objects365 [45], and Conceptual Caption 3M [46], respectively. GoldG indicates gold grounding data curated by MDETR [27].

Method	Framework	Training Data	AP	bAP	nAP
Detic [66]	CenterNet2-R50	V3Det- C_{base}	17.7	30.2	6.7
RegionClip [65]	FR-CNN-R50	V3Det- C_{base}	12.6	22.1	3.1

Table 7: Evaluation of two representative open-vocabulary object detection works with training on C_{base} in *train* split of V3Det.

OVD methods via direct inference on V3Det *val* split. As shown in Table 6, we report the results on AP of all categories, bAP of base classes C_{base} , nAP of novel classes on C_{novel} . Experimental results show that all OVD methods trained on existing object detection datasets, e.g., COCO, LVIS, and objects365, fail to perform well on V3Det. Among them, Detic achieves the highest performance of only 1% AP . According to Table 2, fully supervised detectors can obtain much superior performance on V3Det. This observation shows that OVD methods still have a large exploration space to improve themselves.

Evaluate OVD methods trained on V3Det. We further explore the effectiveness of current OVD methods after training on V3Det. Specifically, we train two well-known OVD methods, *i.e.*, Detic [66] and RegionCLIP [65] on base classes C_{base} in *train* split of V3Det, and test them on both C_{base} and C_{novel} in *val* split of V3Det. As shown in Table 7, after training on C_{base} of V3Det, both methods gain considerable improvements on both bAP and nAP . Notably, nAP of Detic and RegionCLIP are significantly boosted from 1.04% to 6.7% and 0.72% to 3.1%, respectively. Although the performance is still far from that of supervised models, these experimental results confirm that

Training Data	LVIS- nAP	Objects365- AP
LVIS* [23]	7.0	1.6
LVIS* + ImageNet-22k* [44]	13.2	3.4
V3Det*	14.5	4.4

Table 8: Comparisons of category generalization. Detic is adopted as the OVD framework. 1) LVIS*: LVIS *train* split without categories of LVIS-rare and objects365, 2) LVIS* + ImageNet-22k*: LVIS *train* split without categories of LVIS-rare and objects365 + ImageNet-22k images with class labels in LVIS and objects365, 3) V3Det*: V3Det *train* split without categories of LVIS-rare and objects365. LVIS- nAP and Objects365- AP indicate the performance of LVIS rare classes and all Objects365 classes.

a vast vocabulary dataset will lead to stronger category generalization ability.

Category Generalization of V3Det. To further verify the effectiveness of V3Det in training open vocabulary object detectors, we evaluate Detic trained with V3Det on LVIS-rare and objects365 categories. We form a new subset of categories $C_{v3det}^* = C_{v3det} \setminus (C_{lv\text{is-rare}} \cup C_{\text{objects365}})$ which remove LVIS rare classes and Objects365 classes from the vocabulary of V3Det, and $C_{lv\text{is}}^* = C_{lv\text{is}} \setminus (C_{lv\text{is-rare}} \cup C_{\text{objects365}})$ which remove LVIS rare classes and Objects365 classes from the vocabulary of LVIS. After learning from C_{v3det}^* and $C_{lv\text{is}}^*$, we report the performance of LVIS- nAP of LVIS rare classes and Objects365- AP of objects365 on the *val* split of the two datasets. As shown in Table 8, Detic trained on C_{v3det}^* significantly outperforms the one trained on $C_{lv\text{is}}^*$, even on rare classes of LVIS itself. Moreover, we also try to leverage ImageNet images with class labels in LVIS and objects365, which means ‘LVIS* + ImageNet-22k*’ setting contains target categories, *i.e.*, LVIS-rare and objects365. Surprisingly, V3Det still performs better, showing its capability to serve as a dataset for further research on open vocabulary object detection.

5. Conclusion

We introduce **V3Det**, a vast vocabulary visual detection dataset with precisely annotated bounding boxes on massive real-world images. V3Det comprises 243k images in 13,204 categories. The vast vocabulary is organized by a hierarchical category tree. Professional descriptions of each category written by human experts and a powerful chatbot are available. Extensive benchmarks are performed on both vast and open vocabulary object detection, leading to new observations, practices, and insights for future research.

Acknowledgement. This project is funded by Shanghai AI Laboratory (P23KS00010, 2022ZD0160201), the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK, and CUHK Interdisciplinary AI Research Institute.

References

- [1] Likun Cai, Zhi Zhang, Yi Zhu, Li Zhang, Mu Li, and Xiangyang Xue. Bigdetection: A large-scale benchmark for improved object detector pre-training. In *CVPR*, 2022. 2, 3
- [2] Zhaowei Cai and Nuno Vasconcelos. Cascade R-CNN: Delving into high quality object detection. In *CVPR*, 2018. 1, 3, 6
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: High quality object detection and instance segmentation. *CVPR*, 2018. 6, 7
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 1, 3
- [5] Kai Chen, Jiangmiao Pang, Jiaqi Wang, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. Hybrid task cascade for instance segmentation. In *CVPR*, 2019. 2, 7
- [6] 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. 6
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv*, 2015. 2, 3, 6, 8
- [8] Yixin Chen, James Z Wang, and Robert Krovetz. Content-based image retrieval by clustering. In *Proceedings of the 5th ACM SIGMM international workshop on Multimedia information retrieval*, pages 193–200, 2003. 5
- [9] Tao Chu, Pan Zhang, Qiong Liu, and Jiaqi Wang. Buol: A bottom-up framework with occupancy-aware lifting for panoptic 3d scene reconstruction from a single image. In *CVPR*, 2023. 2
- [10] Xiyang Dai, Yinpeng Chen, Bin Xiao, Dongdong Chen, Mengchen Liu, Lu Yuan, and Lei Zhang. Dynamic head: Unifying object detection heads with attentions. In *CVPR*, 2021. 8
- [11] detrex contributors. detrex: An research platform for transformer-based object detection algorithms. <https://github.com/IDEA-Research/detrex>, 2022. 6
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv*, 2018. 3
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. 7
- [14] Yu Du, Fangyun Wei, Zihe Zhang, Miaoqing Shi, Yue Gao, and Guoqi Li. Learning to prompt for open-vocabulary object detection with vision-language model. In *CVPR*, 2022. 8
- [15] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *ICCV*, 2019. 3
- [16] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 2
- [17] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. *arXiv*, 2022. 7
- [18] Chengjian Feng, Yujie Zhong, Zequn Jie, Xiangxiang Chu, Haibing Ren, Xiaolin Wei, Weidi Xie, and Lin Ma. Promptdet: Towards open-vocabulary detection using uncurated images. In *ECCV*, 2022. 8
- [19] Mingfei Gao, Chen Xing, Juan Carlos Niebles, Junnan Li, Ran Xu, Wenhao Liu, and Caiming Xiong. Open vocabulary object detection with pseudo bounding-box labels. *arXiv*, 2021. 8
- [20] Ross Girshick. Fast R-CNN. In *ICCV*, 2015. 1, 3
- [21] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 1, 3
- [22] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv*, 2021. 3
- [23] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 1, 2, 3, 6, 7, 8
- [24] John A Hartigan and Manchek A Wong. Algorithm as 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)*, 28(1):100–108, 1979. 5
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 6
- [26] Zhaojin Huang, Lichao Huang, Yongchao Gong, Chang Huang, and Xinggang Wang. Mask Scoring R-CNN. In *CVPR*, 2019. 2
- [27] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *ICCV*, 2021. 3, 8
- [28] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: connecting language and vision using crowdsourced dense image annotations. *corr abs/1602.07332* (2016). *arXiv*, 2016. 8
- [29] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, S Popov, M Mallocci, T Duerig, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *arXiv*, 2018. *arXiv*, 2018. 1, 2, 3
- [30] Hei Law and Jia Deng. CornerNet: Detecting objects as paired keypoints. In *ECCV*, 2018. 3

- [31] Chunyuan Li, Haotian Liu, Liunian Harold Li, Pengchuan Zhang, Jyoti Aneja, Jianwei Yang, Ping Jin, Yong Jae Lee, Houdong Hu, Zicheng Liu, et al. Elevater: A benchmark and toolkit for evaluating language-augmented visual models. *arXiv*, 2022. **2, 3**
- [32] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *CVPR*, 2022. **3, 8**
- [33] Tsung-Yi Lin, Priya Goyal, Ross B. Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. **1, 3, 6**
- [34] 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, 7**
- [35] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. In *ECCV*, 2016. **1, 3**
- [36] 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 *ICCV*, 2021. **6**
- [37] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. **6, 7**
- [38] Vicente Ordonez, Girish Kulkarni, and Tamara Berg. Im2text: Describing images using 1 million captioned photographs. *NIPS*, 2011. **8**
- [39] Wanli Ouyang, Xiaogang Wang, Cong Zhang, and Xiaokang Yang. Factors in finetuning deep model for object detection with long-tail distribution. In *CVPR*, 2016. **6, 7**
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*. PMLR, 2021. **3, 5**
- [41] Hanoona Rasheed, Muhammad Maaz, Muhammad Uzair Khattak, Salman Khan, and Fahad Shahbaz Khan. Bridging the gap between object and image-level representations for open-vocabulary detection. In *NIPS*, 2022. **8**
- [42] Joseph Redmon and Ali Farhadi. YOLO9000: Better, faster, stronger. In *CVPR*, 2017. **3**
- [43] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. **1, 3, 6, 7**
- [44] 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. **3, 8**
- [45] 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 *ICCV*, 2019. **1, 2, 3, 5, 7, 8**
- [46] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 2018. **3, 8**
- [47] Peize Sun, Rufeng Zhang, Yi Jiang, Tao Kong, Chenfeng Xu, Wei Zhan, Masayoshi Tomizuka, Lei Li, Zehuan Yuan, Changhu Wang, and Ping Luo. Sparse r-cnn: End-to-end object detection with learnable proposals. In *CVPR*, 2020. **1**
- [48] Zhi Tian, Chunhua Shen, and Hao Chen. Conditional convolutions for instance segmentation. In *ECCV*, 2020. **2**
- [49] Zhi Tian, Chunhua Shen, Hao Chen, and Tong He. FCOS: Fully convolutional one-stage object detection. *CoRR*, abs/1904.01355, 2019. **3, 6, 7**
- [50] Jiaqi Wang, Kai Chen, Rui Xu, Ziwei Liu, Chen Change Loy, and Dahua Lin. CARAFE: Content-Aware ReAssembly of FEatures. In *ICCV*, 2019. **1**
- [51] Jiaqi Wang, Kai Chen, Rui Xu, Ziwei Liu, Chen Change Loy, and Dahua Lin. Carafe++: Unified content-aware reassembly of features. *TPAMI*, 2021. **1**
- [52] Jiaqi Wang, Kai Chen, Shuo Yang, Chen Change Loy, and Dahua Lin. Region proposal by guided anchoring. In *CVPR*, 2019. **1**
- [53] Jiaqi Wang, Wenwei Zhang, Yuhang Cao, Kai Chen, Jiangmiao Pang, Tao Gong, Jianping Shi, Chen Change Loy, and Dahua Lin. Side-aware boundary localization for more precise object detection. In *ECCV*, 2020. **1**
- [54] Jiaqi Wang, Wenwei Zhang, Yuhang Zang, Yuhang Cao, Jiangmiao Pang, Tao Gong, Kai Chen, Ziwei Liu, Chen Change Loy, and Dahua Lin. Seesaw loss for long-tailed instance segmentation. In *CVPR*, 2021. **6, 7**
- [55] Xinlong Wang, Tao Kong, Chunhua Shen, Yuning Jiang, and Lei Li. SOLO: Segmenting objects by locations. In *ECCV*, 2020. **2**
- [56] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019. **6**
- [57] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *CVPR*, 2021. **3**
- [58] Hao Zhang, Feng Li, Shilong Liu, Lei Zhang, Hang Su, Jun Zhu, Lionel Ni, and Harry Shum. Dino: Detr with improved denoising anchor boxes for end-to-end object detection. In *ICLR*, 2022. **3, 6, 7**
- [59] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, Yong Wang, and Fang Wen. Prototypical pseudo label denoising and target structure learning for domain adaptive semantic segmentation. In *CVPR*, 2021. **2**
- [60] Pan Zhang, Bo Zhang, Ting Zhang, Dong Chen, and Fang Wen. Robust mutual learning for semi-supervised semantic segmentation. *arXiv*, 2021. **2**
- [61] Shifeng Zhang, Cheng Chi, Yongqiang Yao, Zhen Lei, and Stan Z Li. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, 2020. **6**
- [62] Shilong Zhang, Xinjiang Wang, Jiaqi Wang, Jiangmiao Pang, Chengqi Lyu, Wenwei Zhang, Ping Luo, and Kai Chen. Dense distinct query for end-to-end object detection. In *CVPR*, June 2023. **1**

- [63] Yuanhan Zhang, Qinghong Sun, Yichun Zhou, Zexin He, Zhenfei Yin, Kun Wang, Lu Sheng, Yu Qiao, Jing Shao, and Ziwei Liu. Bamboo: Building mega-scale vision dataset continually with human-machine synergy. *arXiv*, 2022. 3
- [64] Shiyu Zhao, Zhixing Zhang, Samuel Schulter, Long Zhao, BG Vijay Kumar, Anastasis Stathopoulos, Manmohan Chandraker, and Dimitris N Metaxas. Exploiting unlabeled data with vision and language models for object detection. In *ECCV*, 2022. 8
- [65] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *CVPR*, 2022. 3, 8
- [66] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Philipp Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. In *ECCV*, 2022. 3, 8
- [67] Xingyi Zhou, Vladlen Koltun, and Philipp Krähenbühl. Probabilistic two-stage detection. *arXiv*, 2021. 3, 6, 7
- [68] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv*, 2020. 3, 6, 7