Plain-Det: A Plain Multi-Dataset Object Detector

Cheng Shi[∗], Yuchen Zhu[∗], and Sibei Yang[⊠]

School of Information Science and Technology ShanghaiTech University {shicheng2022,yangsb}@shanghaitech.edu.cn

Abstract. Recent advancements in large-scale foundational models have sparked widespread interest in training highly proficient large vision models. A common consensus revolves around the necessity of aggregating extensive, high-quality annotated data. However, given the inherent challenges in annotating dense tasks in computer vision, such as object detection and segmentation, a practical strategy is to combine and leverage all available data for training purposes. In this work, we propose Plain-Det, which offers flexibility to accommodate new datasets, robustness in performance across diverse datasets, training efficiency, and compatibility with various detection architectures. We utilize Def-DETR, with the assistance of Plain-Det, to achieve a mAP of 51.9 on COCO, matching the current state-of-the-art detectors. We conduct extensive experiments on 13 downstream datasets and Plain-Det demonstrates strong generalization capability. Code is release at https://github. com/ChengShiest/Plain-Det.

Keywords: Object detection · Multiple datasets · Proposal generation

1 Introduction

Large-scale datasets have fostered significant advances in computer vision, ranging from ImageNet [6] for image classification to more recent datasets [9, 19] like SA-1B [13] for image segmentation. Object detection [1,24,25], as one of the fundamental tasks in computer vision, inherently demands large-scale annotated data. However, annotating such extensive and densely annotated objects is both costly and challenging. Another straightforward and practical approach is unifying multiple existing object detection datasets [9,19,29] to train a unified object detector [4,22,43]. Nonetheless, inconsistency between datasets, such as differing taxonomies and data distributions as illustrated in Fig 1(a), poses challenges to multi-dataset training.

In this paper, we aim to address the challenges to train an effective and unified detector using multiple object detection datasets, with the expectation that it should: (1) *Flexibility to new datasets* in a seamless and scalable

^{*} Equal contribution. \boxtimes Corresponding author.

Fig. 1: The benefits and challenges of multi-dataset object detection. (a) Various datasets span diverse taxonomies and data distributions. (b) Semantic space calibration. (c) Our approach leverages the advantages of training across multiple datasets to achieve performance enhancements through scaling up data volume.



manner without requiring manual adjustments, complex designs, and training from scratch. (2) **Robustness in performance** when incrementally incorporating new datasets, consistently leading to improved performance or, at the very least, maintaining stable performance. (3) **Training efficiency**. The number of training iterations required for multi-dataset training is no greater than that for a single dataset. (4) **Compatibility with detection families**, such as Faster-RCNN series [7, 25, 31] and DETR-based detection architecture [1, 36, 44].

To start, we introduce a simple yet flexible multi-dataset object detection baseline, which boldly challenges some recent design principles while keeping other advances. Recent works [4,22,43] explicitly unify taxonomies across different datasets into a single, unified one. However, despite their automatic methods, they still require carefully hand-designed components and lack flexibility in scaling to more datasets. This is primarily because 1) the mapping from datasetspecific label spaces to a unified one, learned automatically, becomes noisier as the label space size grows, and 2) incorporating new datasets necessitates reconstructing the unified taxonomy. Therefore, we introduce a shared detector with entirely dataset-specific classification heads to naturally prevent conflicts between different taxonomies and ensure flexibility. Furthermore, following [4, 8, 41], we utilize text embeddings of category labels to build a shared semantic space of all labels. Notably, the semantic space implicitly establishes connections between labels from different classifiers, enabling full use of all the training data despite the dataset-specific classification head. Although our multi-dataset baseline model demonstrates flexibility, its performance notably lags behind that of the single-dataset object detector, with a -3.2% mAP decrease observed when comparing our baseline "o C+L+O+D" and "o single" model in Table 1.

Therefore, we probe the pivotal factors impacting the success of the baseline and offer three insights to empower it to be not only super flexible but also highly effective:

1) Semantic space calibration is inspired by questioning whether the classifier with frozen text embeddings is perfect for object detection. Fig 1(b)-"origin"

Fig. 2: The insights for sparse proposal generation and emergent property. (a) Difference between dense proposal generation and sparse proposal generation. (b) Analysis of two types of proposal generation under multi-dataset object detection training. (c) The emergent property in multi-dataset training. The detector trained on COCO+O365+LVIS shows unstable performance on LVIS.



shows the similarity matrix of text embeddings between categories, which is noticeably different from the one generated by learnable classification weights (Fig 1b-"learnable"). The bias originates from CLIP's training data distribution; for instance, the text-image pairs in CLIP typically exhibit a long-tail distribution in the frequency of nouns. This results in a high similarity between the text embeddings of frequently occurring nouns (such as 'person' in Fig 1b) and other words (including NULL). In turn, we discover that an infrequently occurring NULL exhibits high similarity with frequently occurring words and low similarity with infrequently occurring ones. Therefore, we can treat the empty string NULL as a meaningless basis to extract the basis driven by frequency, resulting in the calibrated similarity matrix shown in Fig 1(b)-"modified".

2) Sparse proposal generation. In object detection, object proposal generation is crucial, especially in multi-dataset scenarios. This is because the same object proposals used as anchors to predict different object sets for different datasets. For example, while COCO and LVIS share the same image set, there are significant differences in annotated categories. This necessitates that the same object proposals within the same image can anchor different objects from both COCO's 80 categories and LVIS's 1203 categories. Currently, object proposal generation methods can be broadly categorized into two types [31]: 1) dense or dense-to-sparse proposal generation [24,25], which generates proposals across all image grids or selects a small subset from dense proposals, and 2) sparse proposal generation [1,31,44], which typically directly generates a set of learnable proposals (see Fig 2a). Therefore, we conduct preliminary experiments and comparisons of these two types of proposal generation methods in multi-dataset object detection on the COCO and LVIS datasets. The results suggest that sparse proposal generation methods consistently outperform the dense approach across both object detector families, as shown in Fig 2b. One possible reason is that compared to dense proposal generation, sparse proposals (*i.e.*, sparse queries) [1,17,36,44]have been demonstrated to capture the distribution of the dataset, making it easier to learn the joint distribution from multiple datasets. However, the perfor-

mance of multi-dataset training still falls below that of single-dataset training, due to the need for the same queries to capture the priors of different datasets. Therefore, we improve sparse queries to class-aware queries based on the unified semantic space and image prior, which mitigates the challenge of a set of queries having to accommodate multiple datasets.

3) Dynamic sampling strategy inspired by the emergent property. Despite the two insights above unlocking the potential of training a unified detector on multiple datasets like COCO [19] and LVIS [9], the inclusion of dataset Objects 365 [29] leads to large fluctuations in the detection performance during training (see Fig 2b-"static sampler"), primarily due to noticeable imbalances in dataset sizes (see Fig 2c). Surprisingly, we observe that even when the detector in a given iteration has low precision on a dataset, it can substantially enhance its precision by undergoing a few additional training iterations on that specific dataset (see Fig 2b-"emergent"). We attribute this phenomenon to an emergent property of multi-dataset detection training: a detector trained on multiple datasets inherently possesses a more general detection capability than training on a single dataset, and the ability can be activated and adapted to the particular dataset by a few dataset-specific iterations. Inspired by the property, we propose a dynamic sampling strategy to achieve better balance among different datasets, which dynamically adapts the multi-dataset sampling strategy in subsequent iterations based on the dataset-specific loss observed previously.

Finally, we introduce Plain-Det, a simple yet effective multi-dataset object detector that can be easily implemented by directly applying the three proposed insights to the baseline, thanks to the baseline's flexibility. In summary, our contributions are:

- We offer three key insights to unlock the challenges of multi-dataset object detection training, including the calibration of the label space, the application and improvement of sparse queries, and the emergent property with few-iteration dataset-specific training.
- Building upon these three insights, we introduce a simple yet flexible multidataset detection framework, denoted as Plain-Det, which satisfies the following criteria: flexibility to accommodate new datasets, robustness in performance across diverse datasets, training efficiency, and compatibility with various detection architectures.
- We integrate Plain-Det into the Def-DETR model and conduct joint training on common public datasets, comprising 2,249 categories and 4 million images. This integration boosts the performance of the Def-DETR model from 46.9% mAP on COCO to 51.9%, achieving performance on par with the current state-of-the-art object detectors. Furthermore, it establishes new state-of-the-art results on multiple downstream datasets.

2 Related work

Multi-Dataset Object Detection is gaining prominence as a crucial step for advancing large vision models. It aims to break down the barriers between datasets and empower detection models with extensive data integration. The fundamental challenges arise from the incongruities in dataset distributions, various taxonomies, and disparate data scales. Early works [16, 38, 39] leverage external priors or human knowledge to construct a unified label space during multi-dataset training. However, this approach becomes inflexible when scaling dataset categories to thousands. Recently, with the rise of joint image-text training, Many works [4,8,22,30,33] generate label spaces using text embeddings from pre-trained vision-language models. However, due to the imbalanced data distribution in pre-trained text corpora, the text embedding space exhibits biases towards specific datasets. We are the first to identify and propose a solution to address this issue during the multi-dataset training process.

Query-based Object Detection has rapidly evolved with the ascendancy of the query-based [1, 12, 17, 31, 36, 44] framework, deriving various query initialization strategies. The pioneering work DETR [1] proposes a sequence of learnable queries for capturing object-level information. Naive zero-initialization is adopted originally. Later works [3,20,31,44,45], including Deformable DETR [44] and Sparse RCNN [31], follow such design, accompanied by pixel-initialized queries or sparse proposals. In the context of multi-dataset training, Det-Hub [22] introduces dataset-aware queries to adapt to various data distributions. To effectively integrate the spatial and dataset priors into our model, we meticulously devised the Class-Aware Query Compositor, initialing class-aware queries with weak priors.

Data Sampling Strategy holds critical importance in multi-dataset training [4, 22, 43] due to substantial variations in the number of images and classes across different datasets. For instance, LVIS [9] encompasses nearly four times the number of categories present in Objects365 [29], while the number of images is 17 times fewer. UniDet [43] has demonstrated that incorporating a uniform sampling strategy for various datasets yields benefits. However, such "balance" between different datasets solely refers to the sampled number of images. Therefore, we propose a hardness-indicated sampling strategy to dynamically adjust sampler weights based on the training hardness of different datasets.

3 Our Method

As shown in Fig 3, we propose a general framework and training strategy for multi-dataset object detection, free from pursuing a particular detection architecture, and compatible with various detection families. In Sec 3.1, we first introduce the preliminaries on query-based object detectors and then abstract three primary components of multi-dataset object detection: 1) the object detector architecture with dataset-specific classification heads and frozen classifiers (see Section 3.2), 2) object queries generation via the class-aware query compositor (see Section 3.3) and 3) the hardness-indicated sampler for multi-dataset training (see Section 3.4).

Fig. 3: Method overview. Our multi-dataset detector Plain-Det is compatible with various query-based detection families. (a) Our multi-dataset joint training framework for object detection. (b) Overview of query compositor: it takes images and the label embeddings of datasets as inputs and outputs class-aware query.



3.1 Preliminaries

Query-based object detector. By reformulating object detection as a set prediction problem [1], recent query-based object detectors, such as DETR-based methods [1, 17, 36, 44] or the Sparse-RCNN [31], leverage learnable or dynamically selected object queries to directly generate predictions for the final object set. This approach eliminates the requirement for hand-crafted components like anchor presets [7, 25] and post-processing NMS [7, 25]. A query-based detector consists of three components: a set of object queries, an image encoder (*e.g.*, Transformer encoder [32] in DETR or CNN [11, 14] in Sparse-RCNN), and a decoder (*e.g.*, Transformer decoder [32] in DETR or dynamic head [5] in Sparse-RCNN). For a given image I, the image encoder $Enc(\cdot)$ extracts the image features, which are subsequently input to the decoder $Dec(\cdot)$ along with the object queries Q to predict the category C and bounding box B for each query. Typically, the classification head $\mathcal{H}_c(\cdot)$ and the box regression head $\mathcal{H}_b(\cdot)$ consist of several layers of multi-layer perceptrons (MLPs). The overall detection pipeline can be viewed as follows:

$$\hat{\mathcal{Q}} = Dec(Enc(I), \mathcal{Q}), C = \mathcal{H}_c(\hat{\mathcal{Q}}), \quad B = \mathcal{H}_b(\hat{\mathcal{Q}}),$$
(1)

where $\hat{\mathcal{Q}}$ is the query feature after query refinement by decoder layer. To simplify the demonstration, we use $f(\cdot)$ to represent the object detector, where the encoder $Enc(\cdot)$, decoder $Dec(\cdot)$, learnable or selected queries \mathcal{Q} , classification head $\mathcal{H}_c(\cdot)$, and box regression head $\mathcal{H}_b(\cdot)$ are components of it.

Single-dataset object detection training. For training on a single dataset D, the optimization objective of a query-based object detector can be formulated

as follows,

$$\underset{\Theta = \{Enc, Dec, \mathcal{Q}, \mathcal{H}_b, \mathcal{H}_c\}}{\operatorname{argmin}} \quad \mathbb{E}_{(I, \hat{B}) \sim D}[\ell(f(I; \Theta), \hat{B})], \tag{2}$$

where (I, \hat{B}) represents a pair of image and annotations from the dataset D. The loss function ℓ is typically the cross-entropy loss for the class predition and the generalized intersection over union loss for the box regression [27].

3.2Dataset-specific Head with Frozen Classifier

In this section, we introduce our multi-dataset object detection framework, which is compatible with any query-based object detection architecture. To support multiple datasets, our framework features a distinct dataset-specific classification head for each dataset. Within these heads, the classifiers are pre-extracted and frozen during training.

Object detector with dataset-specific classification head. Multiple datasets $D_1, D_2, ..., D_M$ with respective label spaces $L_1, L_2, ..., L_M$, may have inconsistent taxonomies. For example, the "dolphin" class in the Obj365 dataset [29] is labeled as background in the COCO dataset [19]. Recent works thus either manually or automatically create a unified label space for the M datasets by concatenating every dataset-specific label space [4], learning mappings from each label space to the unified one [43], or assigning soft labels to the sub-word set [22] of class names. However, a unified label space lacks flexibility when scaling to more datasets and tends to become noisier as the label space size grows. Therefore, we propose to keep each label space separate to directly and naturally address the issue of inconsistent taxonomies. Specifically, we augment the query-based object detector by adding M dataset-specific classification heads $\mathcal{H}^1_c(\cdot), \mathcal{H}^2_c(\cdot)$

..., $\mathcal{H}_{c}^{M}(\cdot)$, each specializing in classifying objects within its corresponding label space:

$$\hat{\mathcal{Q}} = Dec(Enc(I), \mathcal{Q}),$$

$$C^m = \mathcal{H}_c^m(\hat{\mathcal{Q}}), \quad B = \mathcal{H}_b(\hat{\mathcal{Q}}),$$
(3)

where $\mathcal{H}_{c}^{m}(\cdot)$ is the dataset D_{m} 's classification head on the label space L_{m} . The encoder $Enc(\cdot)$, decoder $Dec(\cdot)$, object queries \mathcal{Q} , and class-agnostic box regression head $\mathcal{H}_b(\hat{\mathcal{Q}})$ are shared across datasets. Notably, while our detector is formally similar to the partitioned detector [43], our classification heads are independently optimized with their respective objectives. In contrast, partitioned detector [43] subsequently optimizes the outputs of the partitioned detector with the objective of unified taxonomy.

Frozen classifiers with a shared semantic space. While dataset-specific heads address conflicts arising from inconsistent taxonomies, they do not fully leverage similar semantic classes, such as the common class "person", from different datasets for comprehensive learning. To address this problem and transfer the common knowledge across diverse datasets, we follow [4] to leverage the pretrained CLIP [23] model's feature space as the shared semantic space for class labels. Specifically, for each dataset D_m with label space L_m , we utilize its labels

Fig. 4: Comparison of different proposal generation methods and ours. (a) Proposal generation from sparse queries. (b) Proposal generation from top-K dataset-specific pixel features in dense image feature map. (c) Our class-aware query generation relies on weak priors associated with the dataset and the image. Dataset-specific head shows we use the different frozen classification heads to calculate the loss.



CLIP text embeddings as the classifier W^m within its classification head $\mathcal{H}^m_c(\cdot)$:

$$W^m = Enc_{\text{text}}(\text{Prompt}(L_m)),\tag{4}$$

where $\operatorname{Prompt}(L_m)$ generates text prompts "the photo is [class name]" for every class in label space L_m , and $Enc_{\text{text}}(\cdot)$ is the frozen text encoder of CLIP. To correct the bias comes from CLIP's training data distribution, We calibrate text embeddings by removing the basis bias as follows:

$$\hat{W}^m = \operatorname{Norm}(W^m - \operatorname{Enc}_{\text{text}}(\text{NULL})), \tag{5}$$

where $Enc_{text}(NULL)$ is the text embedding of the empty string, and the Norm is the L2 normalization.

3.3 Class-Aware Query Compositor

Object query generation, as the essential component of query-based object detectors [1,3,17,20,31,36,44,45], has been extensively explored in single-dataset training, yielding diverse types based on their independence from images [36]. In multi-dataset object detection, initializing object queries becomes even more crucial due to the diversity of multiple datasets involved. This extends beyond the scope of query initialization in single-dataset object detection, where queries are initialized randomly or generated from input image feature map based on dataset-specific top-K score (see Fig 4a and b). In our preliminary experiments on multi-dataset object detection, selecting the top-K pixel features [1] from the encoder (Def-DETR ++ [44] in Fig 2) led to a significant drop in performance, compared to single-dataset training. This is because the top-K candidate objects within an image depend significantly on the dataset taxonomy and are strongly correlated with the dataset. An excessively strong dataset prior conditions the detector towards dataset-specific decoding, which in turn hinders the decoder from fully exploiting multiple datasets for comprehensive learning. Conversely, dataset-agnostic query initialization (Def-DETR [44] in Fig 2) shares the same learnable object queries across all datasets. Based on these observations and insights, we propose a novel query initialization method for multi-dataset object detection (see Fig 4c). Neither dataset-agnostic nor strongly dataset-dependent, our class-aware query initialization relies on weak priors associated with the dataset and the image. Given an image I and its corresponding dataset D_m 's classifier \hat{W}^m , we first construct a dataset-specific weak query embedding Q^b based on the classifier as follows:

$$\mathcal{Q}^b = \mathrm{MLP}(\hat{W}^m). \tag{6}$$

Notably, despite being dataset-specific, unlike strong priors that directly select dataset-specific image content, we obtain a weak prior by employing its dataset-specific classifier that shares the same semantic space across different datasets. Through this weak prior, similar semantic labels across different datasets can be shared. Subsequently, we opt to extract the global image feature rather than the top-K local content features as the weak image prior. We then combine this with dataset-specific queries as follows:

$$\mathcal{W} = \mathrm{MLP}(\mathrm{Max}\operatorname{-Pool}(Enc(I))),$$

$$\mathcal{Q}^{c} = \mathcal{W}\mathcal{Q}^{b},$$
(7)

where Max-Pool performs max pooling over the entire image, and \mathcal{W} can be regarded as a weak image prior. And \mathcal{Q}^c represents the final query feature fed to the decoder. Importantly, our modifications are solely focused on the classification head and query initialization, making them easily applied to and compatible with all query-based object detectors.

3.4 Training with Hardness-indicated Sampling

Apart from detector architecture adjustments discussed in Sec 3.2 and Sec 3.3 to adapt to multiple datasets, training a multi-dataset detector brings forth additional challenges due to notable disparities in dataset distribution, image quantities, label space sizes, and more. In this section, we first formulate the objective of multi-dataset training and then improve the training strategy based on our observation of the emergent property introduced in Fig 2.

In general, the optimization objective for training our multi-dataset object detector on M datasets $D_1, D_2, ..., D_M$ can be formulated as follows,

$$\underset{\Theta \in \{Enc, Dec, \mathcal{Q}_c, \mathcal{H}_b\}, \{\mathcal{H}_c^m\}_{m=1}^M}{\operatorname{argmin}} \mathbb{E}_{D_m} \big[\mathbb{E}_{(I, \hat{B}) \sim D_m} \big[\ell_m(f(I, \Theta; \mathcal{H}_c^m), \hat{B}) \big] \big], \qquad (8)$$

Where, except for the task-specific classification head \mathcal{H}_c^m , the remaining components of the object detector $f(\cdot)$, including the encoder $Enc(\cdot)$, decoder $Dec(\cdot)$, lightweight MLPs for generating object queries \mathcal{Q}^c (Equ 7), and class-agnostic box regression head $\mathcal{H}_b(\cdot)$, are shared across different datasets. Thanks to our dataset-specific classification head $\mathcal{H}_c^m(\cdot)$, our loss ℓ_m can be naturally tailored to the specific dataset, ensuring the preservation of the original training loss and

sampling strategy for each dataset individually. For instance, we apply RFS [9] to the long-tailed LVIS [9] dataset but not to the COCO [19] dataset.

Although dataset-specific losses can adapt to each dataset's internal characteristics, significant differences between datasets, such as differences in dataset sizes, present training challenges that must be addressed. Therefore, we propose a hardness-indicated sampling strategy to balance the number of images across datasets and dynamically assess dataset difficulty during online training. We first periodically record the box loss, L_1, \ldots, L_m , for different datasets. Then compute the online sampling weight w_m as follows:

$$w_m = \frac{L_m}{\min(\{L_i\}_{i=1}^m)} \left[\frac{\max(\{S_i\}_{i=1}^m)}{S_m}\right]^{\frac{1}{2}},\tag{9}$$

where S_i means the number of images in the *i*-th dataset, and w_m will be involved in controlling the weight of each dataset in data sampling. The online sampler will sample data from each dataset according to the proportion of its corresponding weight w_m .

4 Experiment

We conduct the following experiments to demonstrate the effectiveness of our Plain-Det. We first introduce our training setups in Section 4.1. Then we analyze Plain-Det with an increasing number of datasets in Sec 4.2, report Plain-Det on various detection benchmarks in Sec 4.3 and Sec 4.4 and conduct ablation study in Sec 4.6.

4.1 Training setups

Implementation details. If not mentioned, we use a partitioned (with dataset-specific head) Deformable-DETR [44] as our default object detector and set the number of queries to 300. Profiting from our model's compatibility with the detection family, we also integrate Plain-Det with other query-based detectors, Sparse R-CNN [31], for experiments. Our implementation is based on the official implementation in Detectron2 [34] and Detrex [26].

4.2 Performance with growing number of datasets

Strict multi-dataset training. To demonstrate the effectiveness of our Plain-Det in multi-dataset training, showcasing how different datasets can mutually benefit from our approach, we test our method under the strict multi-dataset training setting. Under the strict multi-dataset training setting in Table 1, when the number of datasets increases, the total training iteration remains the same as the largest dataset among them. For example, we train O365 [29] for 450k iterations separately, then we train 'L+C+O' for 450k iterations too. Note that generally, increasing training iterations improves the performance of detectors.

Table 1: Performance with growing number of datasets. We train our Plain- D^{N} et under a strict multi-dataset setting: when the number of datasets increases, the total training iteration remains the same. For example, we train O365 [29] for 450k iterations separately, then we train 'C+L+O' for 450k iterations too. \circ : method w/o our Plain- D^{N} et, \bullet : method w/ our Plain- D^{N} et. Single: single-dataset training, mAP is the mean AP across datasets.

Method	COCO [19]	LVIS $[9]$	O365 [29]	OID [35]	mAP
\circ Single [44]	45.6	33.6	32.2	61.0	43.2
◦ C+L+O+D	44.2	29.1	27.4	59.4	40.0
• L	37.2	33.3	13.4	35.3	29.8
$\bullet L+C$	46.0	33.2	14.2	35.7	32.3
\bullet L+C+O	51.8	39.9	33.2	41.7	41.7
\bullet L+C+O+D	51.9	40.9	33.3	63.4	47.4

However, in strict settings, the performance improvement of multi-dataset training can only come from the mutual assistance of different datasets, which is precisely our focus. An exception arises with the 'L+C+O+D' dataset, where severe underfitting is observed when training only on O365 [29] iterations. Consequently, we double the training iterations.

Increasing performance when increasing number of datasets. We provide the single-dataset baseline by training official Deformable-DETR [44] which has a similar single-dataset performance (33.6% vs 33.3% on LVIS [9]). We also provide a simple baseline for merging multi-dataset training to demonstrate that improving performance across multiple datasets is not a trivial task. Starting from training on a single dataset, LVIS, the performance on COCO [19] improved from 37.2% to 46.0%, 51.8%, and 51.9%. The results of multi-dataset training also surpassed the performance of single-dataset training at 45.6%. The mAP of multi-dataset training also increased from 29.8% to 32.3%, 41.7%, 47.4%. Moreover, compared to the multi-dataset detector baseline, Plain-Det gives much better performance, even though overall training iteration does not extend which means that the utilization of Plain-Det facilitated mutual assistance among different datasets, leading to the learning of a unified detector.

4.3 Comparison to state-of-the-art multi-dataset detectors

In Table 2, we compare Plain-Det against other state-of-the-art multi-dataset detectors [4, 41, 43]. In Table 2a, We provide the performance of different methods [4, 41] on combinations of different datasets [9, 15, 19, 28, 29] for LVIS [9] (L) and COCO [19] (C) to demonstrate the robustness of different methods to different datasets. In Table 2b, We present the performance gap between joint training and individual training of different methods [4, 43] on the same multi-dataset COCO [19] (C), Object365 [29] (O365) and OID [15] (D) to demonstrate the benefits of different methods on the same multi-dataset.

Performance

(a)

two datasets. Model Dataset(s) \mathbf{C} L mAP Detic [41] L,C43.9 33.0 38.4Detic [41] L,C,IN21K 42.4 35.4 38.9 ScaleDet [4] L,C 44.9 33.3 39.1 ScaleDet [4] L,C,O365 47.0 36.5 41.7 ScaleDet [4] L,C,O365,D 47.1 36.8 41.9 Ours L,C 46.0 33.3 40.0Ours L,C,O365 $51.8 \ 39.9$ 45.9L,C,O365,D 51.9 40.9 46.4 Ours

of

ods [4, 41] on combinations of different

datasets [9, 15, 19, 28, 29]. We report the AP

on COCO (C), LVIS (L), and mean AP across

different

 Table 2: Compare to state-of-the-art multi-dataset detectors.

meth-

(b) Performance gap between joint training and individual training of different methods [4,43] on the multi-dataset training. We report the AP on COCO (C), Object365 (O365), and mean AP across two datasets.

Model	Dataset(s)	\mathbf{C}	O365	mAP
	single	42.5	24.9	33.7
UniDet [43]	multiple	45.5	24.6	35.0
	Δ	+3.0	+0.3	+1.3
	single	46.8	28.8	37.8
ScaleDet [4]	multiple	47.1	30.6	38.9
	Δ	+0.3	+1.8	+1.1
	single	45.6	30.0	37.8
Ours	multiple	51.9	33.3	42.6
	Δ	+6.3	+3.3	+4.8

Performance on different multi-datasets. Table 2a shows the comparison between Plain-Det and previous multi-dataset detectors [4,41] on combinations of different datasets. Compared with ScaleDet [4], our Plain-Det consistently outperforms ScaleDet [4] across different dataset combinations (in an apple-to-apple comparison). The improvement increases from 0.9% for 'L+C' to 4.2% for 'L+C+O', and further to 4.5% for 'L+C+O+D', demonstrating that our method scales better to larger datasets compared to ScaleDet. This also indicates that the three proposed improvements are highly effective.

Performance increment. Table 2b presents the comparison between Plain-Det and previous multi-dataset detectors [4, 43], showcasing the performance gap between joint training and individual training. Compared with Unidet [43], which learns a unified label space, Plain-Det achieves an absolute increase of 7.6% (42.6% vs. 35.0%) and a relative growth of over 3.5% (4.8% vs. 1.3%), considering the average AP of two datasets. Compared with ScaleDet [4], which also utilizes CLIP [23] text embeddings but overlooks the issues within text embeddings, Plain-Det achieves an absolute increase of 5.9% (42.6% vs. 36.7%) and a relative growth of over 2.4% (4.8% vs. 2.4%), considering the average AP.

4.4 Comparisons of performance and efficiency across different detection families

Comparison under non-RCNN family. Table 3a shows the comparison between Plain-Det and previous multi-dataset detectors [4,41] based on non-RCNN detectors [42, 44]. We report the number of epochs where COCO [19] appears during multi-dataset training to represent the efficiency of the multi-dataset object detector. A lower frequency of COCO appearances indicates higher training efficiency. Compared with the previous SOTA method ScaleDet [4], Plain-Det

Table 3: Comparisons of performance and efficiency across different detection families. We present the results of various multi-dataset training methods on the COCO validation set, along with the total number of iterations where COCO data appeared during multi-dataset training.

(a) Comparison under non-RCNN family

(b) Comparison under RCNN family

non-RCNN family	multi datasets	epoch	$\begin{array}{c} \mathrm{COCO} \\ \mathrm{AP^{box}} \end{array}$	RCNN family	multi datasets	epoch	C A
CenterNet2 [42]	X	12	42.9	Mask RCNN [10]	×	36	
+Detic [41]	1	-	42.4	+RegionCLIP [40]	1	-	4
+ScaleDet [4]	1	192	47.1	Sparse RCNN [31]	×	12	4
Def-DETR [44]	X	50	46.9	+ Det-Hub [22]	1	12	4
+ Ours	1	36	51.9	+ Ours	1	12	4

Table 4: Results of zero-shot transfer on 5 individual datasets on ODinW. R, T: ResNet50 [11], Swin-Tiny [21]. GoldG: 0.8M grounding data curated by MDETR [12], cap4M: 4M image-text pairs [2].

Model	Datasets	#Data $ $ #1	#2	#3	#4	#5	mAP
GLIP-T [18]	O+GoldG+Cap4M	5.5M 18.	450.0	49.6	57.8	44.1	44.0
ScaleDet-R [4]	L+C	0.2M 9.1	2.1	12.4	41.2	25.5	18.1
ScaleDet-R [4]	L+C+O	1.9M 8.7	1.4	25.0	47.3	20.8	20.6
ScaleDet-R [4]	L+C+O+D	3.6M 23.	$2\ 45.1$	38.9	48.3	40.6	39.2
Ours-R	L+C	0.2M 16.	0 7.0	15.5	45.9	26.6	22.2
Ours-R	L+C+O	1.9M 16.	$5\ 13.3$	31.4	49.4	52.8	32.7
Ours-R	L+C+O+D	3.6M 27.	9 43.3	47.3	58.1	54.1	46.1

achieves superior performance (51.9% vs 47.1%) while utilizing fewer instances of COCO data (36 vs 192). Compared to our underlying object detector Def-DETR [44], the advantages of multi-dataset training are also significant. We achieved a notable performance boost of 5.0% in COCO validation AP while using only 72% of the data.

Comparison under RCNN family. Table 3b presents the comparison between Plain-Det and previous multi-dataset detectors [22, 40] based on RCNN detectors [10, 31]. Sparse RCNN, as the first work to introduce queries into the RCNN series, has shown significant improvements in both performance and efficiency. By integrating Plain-Det into Sparse RCNN and utilizing the same quantity of COCO data, we achieved a performance increase of 3.1%, reaching a COCO validation performance of 46.1%, matching the best multi-dataset object detector under the RCNN series, Det-Hub [22].

4.5 ODinW benchmark

Zero-shot transfer on ODinW. Following ScaleDet [4], table 4 shows the zero-shot transfer performance on 5 individual datasets from Object Detection in Wild (ODinW [37]). For an apple-to-apple comparison, our method significantly

Table 5: Ablation on partition head, class-aware query and sampler.

) -L	we a , and	sam	the pa	C+L+C	nead, spars	e query	(b) Analysis on ratio w.
	Par	Que	Tex	C [19]	L [9]	mAP	.º 3.5
1	X	X	1	39.3	23.8	31.6	
2	1	X	1	38.1	24.0	31.1	guilo
3	1	1	×	44.2	28.7	36.5	
4	1	1	1	45.3	30.2	37.8	
	Samj	pler	C [19]	L [9]	O365 [29]	mAP	\overline{C} 0.5 \rightarrow L \rightarrow C \rightarrow O \rightarrow Base
4	2	x	43.5	26.8	25.5	31.9	
5	۷	/	47.1	32.4	25.6	35.0	Epochs

surpasses the previous state-of-the-art ScaleDet [4] in various training dimensions (+4.1% in L+C, +12.1% in L+C+O, +6.9% in L+C+O+D). In comparison to GLIP [37], which was trained on a much larger combination of multi-datasets, our method still outperforms GLIP with the use of fewer data (3.6M vs 5.5M), achieving a higher performance of 46.1\% compared to 44.0\%.

4.6 Ablation Study

Ablation on partition head, sparse query, and label corr. As we gradually introduce the components proposed, including Que (class-aware query) and Tex (label space correction), the mAP steadily increases from 31.1% to 36.5% and then to 37.8%. One exception is the partition head. With the introduction of the partition head, there is a slight performance decrease (-0.5%). However, we want to emphasize that a partition head is essential for the flexibility of scaling up datasets. Therefore, a minor performance loss is deemed acceptable.

Ablation on sampler. We conducted experiments on samplers using three datasets with highly uneven sizes (C+L+O), where L has the most categories and O has the most annotations and our online sampler results in a 4.1% increase in mAP. As shown in Fig 5b, the online sampler oversampled from L and undersampled from O, which aligns with our prediction of the difficulty levels of the L and O datasets.

5 Conclusion and Limitations

We introduce Plain-Det, an efficient multi-dataset detector, that meets the following requirements: adaptability to incorporate new datasets, consistency in performance across a wide range of datasets, efficiency in training, and compatibility with different detection architectures. We validated our conclusions using over 17 datasets and achieved consistent improvements across multiple datasets. **Limitations:** Given that our semantic label space is derived from the CLIP [23] models, we acknowledge that biases and controversies inherent in the training data for these models may be introduced into our model. **Acknowledgment:** This work was supported by the National Natural Science Foundation of China (No.62206174) and MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration (ShanghaiTech University).

References

- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: Endto-end object detection with transformers. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16. pp. 213–229. Springer (2020)
- Changpinyo, S., Sharma, P., Ding, N., Soricut, R.: Conceptual 12m: Pushing webscale image-text pre-training to recognize long-tail visual concepts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3558–3568 (2021)
- 3. Chen, Q., Chen, X., Zeng, G., Wang, J.: Group detr: Fast training convergence with decoupled one-to-many label assignment. arXiv preprint arXiv:2207.13085 (2022)
- Chen, Y., Wang, M., Mittal, A., Xu, Z., Favaro, P., Tighe, J., Modolo, D.: Scaledet: A scalable multi-dataset object detector. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7288–7297 (2023)
- Chen, Y., Dai, X., Liu, M., Chen, D., Yuan, L., Liu, Z.: Dynamic convolution: Attention over convolution kernels. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11030–11039 (2020)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A largescale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Girshick, R.: Fast r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 1440–1448 (2015)
- Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. arXiv preprint arXiv:2104.13921 (2021)
- Gupta, A., Dollar, P., Girshick, R.: LVIS: A dataset for large vocabulary instance segmentation. In: CVPR (2019)
- 10. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: CVPR (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: CVPR (2016)
- Kamath, A., Singh, M., LeCun, Y., Synnaeve, G., Misra, I., Carion, N.: Mdetrmodulated detection for end-to-end multi-modal understanding. In: CVPR (2021)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
- Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems 25 (2012)
- Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Malloci, M., Kolesnikov, A., et al.: The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. International Journal of Computer Vision 128(7), 1956–1981 (2020)
- Lambert, J., Liu, Z., Sener, O., Hays, J., Koltun, V.: Mseg: A composite dataset for multi-domain semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2879–2888 (2020)
- Li, F., Zhang, H., Liu, S., Guo, J., Ni, L.M., Zhang, L.: Dn-detr: Accelerate detr training by introducing query denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13619–13627 (2022)

- 16 Cheng et al.
- Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., Chang, K.W., Gao, J.: Grounded language-image pretraining. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10965–10975 (June 2022)
- Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: ECCV (2014)
- Liu, S., Li, F., Zhang, H., Yang, X., Qi, X., Su, H., Zhu, J., Zhang, L.: Dab-detr: Dynamic anchor boxes are better queries for detr. arXiv preprint arXiv:2201.12329 (2022)
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
- 22. Meng, L., Dai, X., Chen, Y., Zhang, P., Chen, D., Liu, M., Wang, J., Wu, Z., Yuan, L., Jiang, Y.G.: Detection hub: Unifying object detection datasets via query adaptation on language embedding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11402–11411 (2023)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
- 24. Redmon, J., Farhadi, A.: Yolo9000: better, faster, stronger. In: CVPR (2017)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. NeurIPS 28 (2015)
- Ren, T., Liu, S., Li, F., Zhang, H., Zeng, A., Yang, J., Liao, X., Jia, D., Li, H., Cao, H., Wang, J., Zeng, Z., Qi, X., Yuan, Y., Yang, J., Zhang, L.: detrex: Benchmarking detection transformers (2023)
- Rezatofighi, H., Tsoi, N., Gwak, J., Sadeghian, A., Reid, I., Savarese, S.: Generalized intersection over union: A metric and a loss for bounding box regression. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 658–666 (2019)
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. IJCV 115(3), 211–252 (2015)
- Shao, S., Li, Z., Zhang, T., Peng, C., Yu, G., Zhang, X., Li, J., Sun, J.: Objects365: A large-scale, high-quality dataset for object detection. In: ICCV (2019)
- Shi, C., Yang, S.: Edadet: Open-vocabulary object detection using early dense alignment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15724–15734 (2023)
- 31. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., et al.: Sparse r-cnn: End-to-end object detection with learnable proposals. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 14454–14463 (2021)
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. Advances in neural information processing systems **30** (2017)
- Wang, Z., Li, Y., Chen, X., Lim, S.N., Torralba, A., Zhao, H., Wang, S.: Detecting everything in the open world: Towards universal object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11433–11443 (2023)

- Wu, Y., Kirillov, A., Massa, F., Lo, W.Y., Girshick, R.: Detectron2. https://github.com/facebookresearch/detectron2 (2019)
- Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: CVPR (2021)
- Zhang, H., Li, F., Liu, S., Zhang, L., Su, H., Zhu, J., Ni, L.M., Shum, H.Y.: Dino: Detr with improved denoising anchor boxes for end-to-end object detection. arXiv preprint arXiv:2203.03605 (2022)
- Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: Glipv2: Unifying localization and vision-language understanding. Advances in Neural Information Processing Systems 35, 36067–36080 (2022)
- 38. Zhao, J., Ou, M., Xue, L., Cui, Y., Wu, S., Chen, G.: Joining datasets via data augmentation in the label space for neural networks (2021)
- Zhao, X., Schulter, S., Sharma, G., Tsai, Y.H., Chandraker, M., Wu, Y.: Object detection with a unified label space from multiple datasets (2020)
- 40. Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., et al.: Regionclip: Region-based language-image pretraining. arXiv preprint arXiv:2112.09106 (2021)
- Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twentythousand classes using image-level supervision. arXiv preprint arXiv:2201.02605 (2022)
- Zhou, X., Koltun, V., Krähenbühl, P.: Probabilistic two-stage detection. arXiv preprint arXiv:2103.07461 (2021)
- Zhou, X., Koltun, V., Krähenbühl, P.: Simple multi-dataset detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7571–7580 (2022)
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. arXiv preprint arXiv:2010.04159 (2020)
- Zong, Z., Song, G., Liu, Y.: Detrs with collaborative hybrid assignments training. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6748–6758 (2023)