

Detection Hub: Unifying Object Detection Datasets via Query Adaptation on Language Embedding

Lingchen Meng^{1,2} Xiyang Dai³ Yinpeng Chen³ Pengchuan Zhang³ Dongdong Chen³
 Mengchen Liu³ Jianfeng Wang³ Zuxuan Wu^{1,2†} Lu Yuan³ Yu-Gang Jiang^{1,2}

¹Shanghai Key Lab of Intell. Info. Processing, School of CS, Fudan University

²Shanghai Collaborative Innovation Center of Intelligent Visual Computing

³Microsoft

Abstract

Combining multiple datasets enables performance boost on many computer vision tasks. But similar trend has not been witnessed in object detection when combining multiple datasets due to two inconsistencies among detection datasets: taxonomy difference and domain gap. In this paper, we address these challenges by a new design (named Detection Hub) that is dataset-aware and category-aligned. It not only mitigates the dataset inconsistency but also provides coherent guidance for the detector to learn across multiple datasets. In particular, the dataset-aware design is achieved by learning a dataset embedding that is used to adapt object queries as well as convolutional kernels in detection heads. The categories across datasets are semantically aligned into a unified space by replacing one-hot category representations with word embedding and leveraging the semantic coherence of language embedding. Detection Hub fulfills the benefits of large data on object detection. Experiments demonstrate that joint training on multiple datasets achieves significant performance gains over training on each dataset alone. Detection Hub further achieves SoTA performance on UODB benchmark with wide variety of datasets.

1. Introduction

Recent computer vision development has demonstrated significant benefits of leveraging large-scale data for computer vision tasks, such as image retrieval [31], image recognition [47] and video recognition [39]. However, limited effort has been explored for the object detection task due to the lack of a unified large-scale data. A naive attempt is to combine all annotated data from different sources.

[†] Corresponding author.

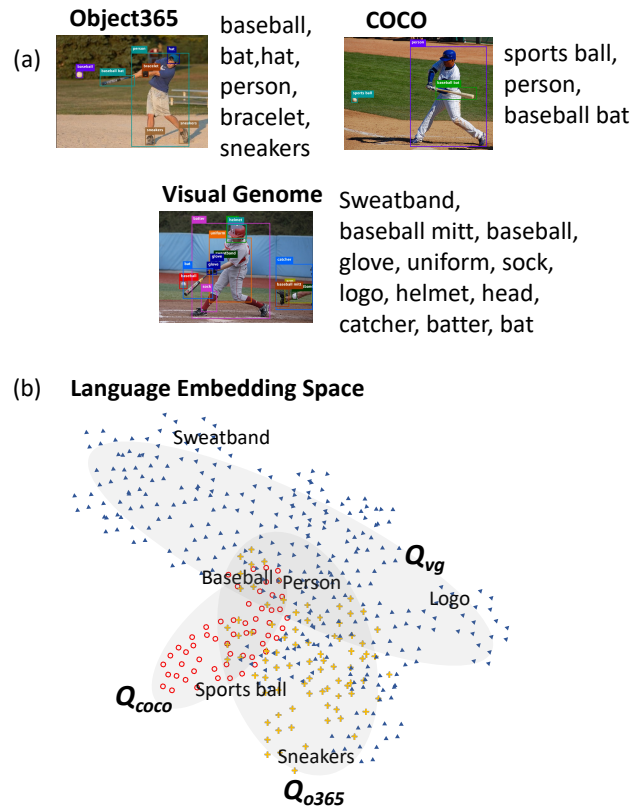


Figure 1. An illustration of the challenges of combining multiple datasets. (a) There are significant domain gaps between different datasets: the taxonomy are different and the annotations are inconsistent; (b) Even under a unified language embedding, the distribution of categories among datasets are significant different and requires special handling

However, due to the diversity of objects and the cost of annotating bounding boxes in images, traditional object detection datasets are collected in a domain-specific way, where a limited number of interested categories are regarded as foregrounds and the rest of objects as backgrounds. This results

in a non-trivial domain shift between different datasets and limits detectors to be trained and tested on a single dataset in order to achieve the best performance on such dataset.

In this paper, we attempt to answer the question: “*How can we unify multiple object detection datasets training with a single general object detector?*”. Towards this goal, we observe two challenges: taxonomy difference and annotation inconsistency, both of which introduce the domain gap issue, shown in Figure 1 (a). More specifically, for the taxonomy difference issue, the semantic names of similar concepts in different datasets may be very different. And for the annotation inconsistency issue, given similar images, foreground objects in one dataset may be labeled as background in another dataset. The existence of these two challenges may be the key underlying reason why most current detectors only focus on a specific training set, rather than deriving a universal object detector for multiple datasets.

Compared to traditional methods that regard semantic categories as class indices, converting them into language embedding can naturally unify the taxonomy and eases the challenge in recent works [19, 31]. However, it won’t solve the problem, as shown in Figure 1 (b), the distribution of categories among datasets are significant different, which contributes most to the domain gap.

Inspired by recent success of visual-language models [19, 31], we propose a simple yet effective method “*Detection Hub*” that can enjoy the synergy of multiple datasets. It builds upon current end-to-end detectors which use learnable object queries to produce final detection results [37]. To overcome the aforementioned problems, we have two key ideas accordingly: 1) map the semantic category names of different datasets into a category-aligned embedding, and 2) more importantly, use the embedding to dynamically adapt object queries so that each dataset has its own specific query set. We further change the one-hot based classification branch into vision-language alignment, which can well align categories of different datasets. Accordingly, we adopt a region-to-word alignment loss instead of classical cross entropy loss, which makes our “*Detection Hub*” not limited by a fixed category vocabulary.

Our method leverages the linguistic property of pre-trained language encoders (such that categories with similar semantic meanings across datasets will be automatically embedded together without using an expert-designed label mappers). In addition, categories specific to different datasets will be preserved by dedicated embedding. Moreover, as each dataset has its own adapted object query set to generate detection results, the detector can learn how to adapt its behavior to each dataset based on its specific query set. Therefore, the potential competition or disturbance caused by the annotation inconsistency among different datasets can be avoided.

To demonstrate the effectiveness of our method, we

train our “*Detection Hub*” on three standard object detection datasets jointly: COCO [21], Object365 [35] and Visual-Genome [14]. These large-scale datasets have different properties of taxonomy, vocabulary size and annotation quality. *Detection Hub* achieves 45.3, 23.2 and 5.7 AP on each dataset, with significant performance gain of **+2.3, +1.0, +0.9** compared with each independently model. To further verify the effectiveness on datasets with larger variance, we conduct experiments on UODB [40], a combination of 11 extremely varied datasets. *Detection Hub* achieves an average score of 71 and outperforms the previous SoTA UniversalDA [40] by a large margin of **6.8** point.

2. Related Work

Object Detection. Object detection [9, 10, 18, 20, 34, 36, 38] has been studied for a long time and become a predominant direction. Beginning with the initial success of R-CNN [10], a two-stage paradigm is formulated by combining a region proposal detector and a region-wise CNN classifier. Fast R-CNN [9] formalizes the idea of region-wise feature extraction using ROI pooling to extract fixed-size region features and then predicts their classes and boxes with fully connected layers. Faster R-CNN [34] further introduces a Region Proposal Network (RPN) to generate region proposals and merges the two-stage into a single network by sharing their convolutional features. Since then, this two-stage architecture has become a leading object detection framework. Further works continue to improve this framework by introducing more stages [1, 16]. Recently, Vision Transformers [5] and end-to-end object detection methods have attracted such attention. DETR [2] first turns the object detection problem into a query-based set prediction problem and formulates object detection as an end-to-end framework. Many follow-ups [17, 24, 25, 51] further developed these methods into other tasks like acceleration and pseudo labeling. Beside, multiple recent methods were proposed to improve the problems in different directions. Deformable DETR [51] first introduces multi-scale deformable attention to replace transformers in DETR and reduces the key sampling points around the reference point to significantly accelerate the training speed. Dynamic DETR [3] further models dynamic attention among multiple dimensions with both a convolution-based encoder and a decoder to further reduce learning difficulty and improves the performance. Sparse R-CNN [37] proposes a sparse set of learnable proposal boxes and a dynamic head to perform end-to-end detection upon the R-CNN framework. Our proposed method is also a query-based object detector. Unlike those DETR query designs, our linguistic adapted query is irrelevant to position and plays a role like a semantic-guided filter.

Multiple Datasets Training. There are a few early attempts on training with multiple datasets to recognize a wide vari-

ety of objects. YOLO9000 [33] first proposed a method to jointly train on object detection and image classification, which expanded the detection vocabulary via classification labels. Later, [45] expanded such an approach to two large scale datasets, ImageNet and OpenImages with only a small fraction of fully annotated classes. [40] attempted to build a universal object detection system that is capable of working on various image domains, from diverse datasets by introducing a series of adaptation layers based on the principles of squeeze and excitation. [48] further proposed to improve loss functions that integrated partial but correct annotations with complementary but noisy pseudo labels among different datasets. Most recently, [50] proposed to train a shared detector with dataset-specific heads and then learned to projected the outputs into a unified label space. Thanks to the advantages of leveraging language embedding, our method does not need to learn a joint label space to merge different datasets. This naturally addresses the scaling up of categories and joint of different datasets.

Language Embedding. Generative and efficient language representation is an attractive topic in the past few decades. As an integral role in the modern NLP system, word embedding can encode rich semantic in a structured way. Most widely used word embeddings methods, Word2Vec [26], GloVe [28], leverage the co-occurrence statistics in a large corpus. As a result, words with similar meanings are close in the embedding space, which can encode rich semantic information. Furthermore, many follow-ups [4, 29, 32] extract context-sensitive features rather than static features. They perform sequence modeling with RNNs or Transformers on the top of traditional static word embeddings. Most recently, learning visual representation with language embedding from language supervision is a promising trend due to rich semantic information and flexible transferable ability of natural language. As a milestone, CLIP [31] designs an image-text alignment pre-train task and performs contrastive learning on a large amount of image-text pairs. Motivated by CLIP, many following works [11, 19, 49] try to improve the training strategy. GLIP [19] reformulate the object detection as phrase grounding, which makes detection benefit from large grounding datasets. RegionCLIP [49] and ViLD [11] leverage the semantics space of CLIP for open-vocabulary object detection. Unlike previous methods attempting to learn a language embedding specifically, we freeze a pre-trained language embedding and design an adaption mechanism to dynamically adapt queries on categories based on the different distributions of datasets.

3. Our Method

3.1. Revisiting End-to-end Object Detection

End-to-end detectors [2, 3, 37] utilize object queries to encode the content and position statistics over the train-

ing dataset and drive the detector to predict desired objects localization without non-differential components (such as pre-defined anchors and non-maximum suppression).

Given a set of N learnable queries $v \in \mathbb{R}^{N \times d}$, an end-to-end detector utilizes a transformer \mathcal{T} to generate N corresponding predictions. Then the optimization process can be abstracted as solving a matching problem L_{match} among images in the training dataset D :

$$\sigma = \operatorname{argmin}_{\sigma \in Q_N} \sum_i^D L_{match}(y_i, \mathcal{T}(Q_N, i)) \quad (1)$$

Ideally, these N learnable queries Q_N can cover an arbitrary number of boxes in each image. However, considering Q_N was optimized on the whole dataset, it will cause each object query to be in charge of multiple boxes and categories, meanwhile competing with each other to overfit to the training data [24, 41]. This limitation will be further amplified when learning across multiple datasets.

In this paper, we build upon Sparse-RCNN [37] as our default instantiation due to its intuitive explanation and training efficiency. Although achieving great performance on object detection, Sparse-RCNN has limitations on multi-dataset training: the learnable queries and features are dataset-dependant, which suffer from inconsistencies when the object statistics of datasets vary. Moreover, the classification branch is dataset-specific and hard to be jointly learned across different datasets. To enjoy the synergy and avoid such limitations, we propose category-aligned embedding and use it to adapt the object queries for each dataset to effectively formulate dataset-aware queries. In this way, our “*Detection Hub*” can learn how to adapt its behavior for each dataset based on its own query set.

3.2. Category-aligned Embedding

Given two datasets D_1, D_2 with n_1, n_2 categories respectively, we desire to have a embedding function \mathcal{E} that can map them together in a common space. Traditional one-hot embedding cannot serve this purpose as it treats category solely as an index. Later works tried to manually or automatically learn a mapping function across datasets with limited success [15].

It is natural to think about using category names directly, *e.g.*, [31, 48] use language models to generate unified embedding. In our method, we further leverage the semantic properties of the category set, through using a pretrained language model to encode the language embedding. Specifically, we concatenate the category names from a dataset C together and then use a language model to obtain the language embedding E of the dataset:

$$\begin{aligned} \text{CategorySet} &= \{c_1, c_2, c_3, \dots, c_n''\} \\ E &= \text{Embed}(\text{Tokenizer}(\text{CategorySet})) \end{aligned} \quad (2)$$

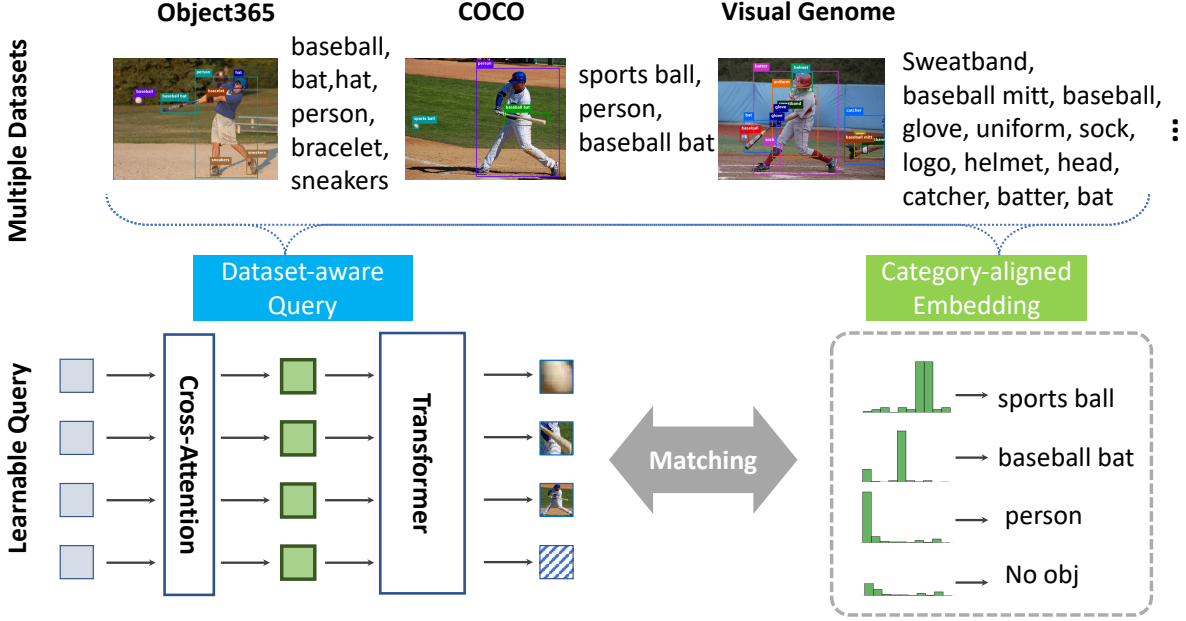


Figure 2. An illustration of our Detection Hub design. Based on a category-aligned embedding, our dataset-aware queries are learned to dynamically adapt on category based on the different distributions of datasets.

Consider that there are N_d datasets $D_{multi} = \{D_1, D_2, \dots, D_{N_d}\}$, and thus we can obtain a set of dataset-specific language embedding $E_{multi} = \{E_1, E_2, \dots, E_{N_d}\}$ accordingly. E_{multi} can be considered as an automatic “label mapper” by its very nature: all semantic category names of different datasets are mapped into a unified language space, where categories of similar meaning will be close while categories specific to each dataset will be represented by dedicated embeddings.

With this great property, we further convert the original index-based classification branch (dataset-specific) used in SparseRCNN into vision-language alignment based classification branch (shared). Accordingly, we adopt a region-word alignment loss rather than the traditional cross-entropy loss for optimization. In particular, we first use a BERT model to take dataset embedding E , and estimate the inner relationships among different categories in each of the dataset-specific language embedding and obtain an enhanced embedding E' . Then we project E' and object features x^{box*} into a same vision-language space via two MLP branches FC_E and FC_V respectively. Following [19, 31], the alignment scores S can be calculated by the dot product between E' and x^{obj} to reflect the similarities between visual objects and semantic categories:

$$E' = \text{BERT}(E) \quad (3)$$

$$S = \text{Sigmoid}(FC_E(E') \cdot FC_V(x^{obj*})) \quad (4)$$

Since we treat the classification as vision-language alignment, the classification target is converted into sub-word

grounding [19] instead of one-hot label [34, 37]. For each prediction and ground-truth pair, we mark the sub-word belonging to the target category as a positive pair, i.e., $\hat{T}_{i,j} = 1$, otherwise $\hat{T}_{i,j} = 0$. Finally, the alignment scores can be optimized with the vision-language alignment loss S , same as the binary cross-entropy loss in traditional detector:

$$\mathcal{L}_{align}(S, \hat{T}) = \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^L S_{i,j} \cdot \hat{T}_{i,j} + (1 - S_{i,j}) \cdot (1 - \hat{T}_{i,j}) \quad (5)$$

$$\hat{T}_{i,j} = \begin{cases} 1, & \text{if token}_{i,j} \text{ belongs to } T_i, \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

where N is the number of queries, and L is the length of the dataset embedding E .

3.3. Dataset-aware Query

As described above, to fully unleash the power of large amount of data in different datasets, our method proposes to use the dataset-specific language embedding to adapt the object queries so that the model can learn to adapt its behavior for each dataset. In detail, given a dataset D and its language embedding E , we can achieve the query adaptation by simply performing a cross-attention between the learnable queries Q and E . Then the following multi-head-self-attention is used to enhance the adapted queries.

$$Q_D = \text{Cross-Attn}(Q, E), \quad Q_D^* = \text{MHSA}(Q) \quad (7)$$

where Q_D is the adapted query for dataset D , and Q_D^* is the enhanced query. Conceptually, through cross-attention, we encode the dataset-specific language embedding into the adapted queries Q_D , which is used as object query as Fig 2 and makes our detector dataset-aware.

Following [37], enhanced queries Q_D^* and object features x^{obj} are interacted using dynamic convolutions. In contrast to [37], we use enhanced queries to generate kernels for dynamic convolutions with a few linear layers in the dynamic instance interactive head.

$$\begin{aligned} x^{obj*} &= \text{DyConv}(x^{obj}, Q_D^*) \\ &= \text{Conv}(K_2(Q_D^*), \text{Conv}(K_1(Q_D^*), x^{obj})), \\ K_i(Q_D^*) &= \text{Linear}_i(Q_D^*), i \in \{1, 2\} \end{aligned} \quad (8)$$

where x^{obj*} is the interacted object feature between enhanced adapted query and object feature; $K_i(Q^*) \in \mathbb{R}^{k \times k \times c_{in} \times c_{out}}$ is the dynamic kernel generated by the enhanced adapted query Q_D^* .

To further address the challenge of position variance, besides the head, we also use a lightweight query-based RPN with one layer of convolution and the adapted queries to generate the dataset-specific convolution kernels for the RPN dynamically. Then we use the proposals of the highest N scores generated by the query-based RPN as initial proposals boxes for the dynamic instance interactive head. The query-based RPN and decoder in our detector are both optimized in an end-to-end manner similar to [3, 37, 51].

4. Experiment

4.1. Implementation Detail

We mainly evaluate our model on three popular object detection benchmarks: COCO [21], Object365 [35] and Visual-Genome [14] following the common practice. COCO dataset contains 118,000 images collected from web images on 80 common classes. Object365 dataset contains around 740,000 images with 365 classes. Each image is densely annotated by human labelers to ensure quality. Visual-Genome dataset contains around 108,077 images with about 1600 classes. The images are very diverse and often contain complex categories with a long tail distribution. Since boxes are generated by algorithms, they are more noise compared to COCO and Object365. Due to different category numbers, image diversity, and annotation density, these three datasets provide a good test bed on the performance of joint training.

We implement our method in PyTorch and train our models using V100 GPUs. We use AdamW [23] optimizer and step down the learning rate by a rate of 0.1 at 78% and 93% of epochs. For ablation studies, we use standard ImageNet-1k pre-trained ResNet-50 [12] as the backbone with $5e^{-5}$ learning rate and $1e^{-4}$ weight decay and train it with stan-

dard $1 \times$ schedule. We demonstrate the effectiveness of each component and also compare it with other methods under this standard setup. For multi-dataset joint training, our detector can treat any datasets in a unified way. We only need to sample a batch from all datasets, and sample their corresponding categories to build the dataset specific language embedding. We apply dataset re-sampling to make sure that each dataset was trained at only $1 \times$ schedule. We also apply category re-balancing to handle the long-tailed distributions of Object365. Besides, we also evaluate our method on a variety of backbones to demonstrate robustness. For experiments comparing with state-of-the-art methods, we train our method with Swin Large [22] backbone at the $2x$ schedule with multi-scale training. We set the learning rate and weight decay the same as ResNet-50. During evaluation, we evaluate our best model with multi-scale testing to compare with state-of-the-art methods reported without using test time augmentation. There is no other augmentation or optimization tricks used during training.

4.2. Effectiveness of Multi-dataset Training

For the baseline, we build a Sparse R-CNN model with a ResNet-50 and train on three single-dataset separately. As for baseline for multi-dataset joint training, we follow [40], which simply concatenates all datasets and maps their taxonomy together. Then we can train it with a bigger label mapping layer with a common loss. For our proposed method, we evaluate both joint and separate training ways.

As shown in Table 1, under standard $1 \times$ schedule, the jointly trained baseline model is observed a clear performance drop on COCO and VG compared to models trained separately. It demonstrated that simply mapping different labels together doesn't work well due to the domain gaps and taxonomy differences between the datasets. On the contrary, our query-adapted joint training obtains higher performance on each dataset, especially offering a significant gain at 2.3, 1.0, and 0.9 mAP gain on COCO, Object365, and Visual-Genome, respectively. The experiment well demonstrated that our single detector can leverage benefits from multiple datasets training simultaneously.

We also train our method with multiple backbones, such as ResNet-50, ResNext-101, Swin-Tiny, Swin-Large to verify its robustness. As shown in Table 2, it is clear to see that our method obtains significant performance gains on all datasets when jointly trained compared to separate trained.

Furthermore, we also evaluate our method's multi-dataset training stability under different training lengths. As shown in Fig 3, the performance gain preserves even when prolonging the training beyond full convergence. These experiments well demonstrate the effectiveness and robustness of our method under multiple datasets joint training.

	Separate Training			Joint Training		
	COCO	O365	VG	COCO	O365	VG
Sparse-RCNN [37]	40.6	16.1	4.9	–	–	–
Sparse-RCNN [37] + Simple merge [15]	40.6	16.1	4.9	40.4↓ 0.2	17.4↑ 1.3	4.3↓ 0.6
Ours	43.0	22.2	4.8	45.3↑ 2.3	23.2↑ 1.0	5.7↑ 0.9

Table 1. Compared with baseline methods under single and multiple dataset training.




Figure 3. Comparison with different training length under single and multiple dataset training.

	Separate Training			Joint Training		
	COCO	O365	VG	COCO	O365	VG
R-50 [12]	43.0	22.2	4.8	45.3↑ 2.3	23.2↑ 1.0	5.7↑ 0.9
Swin-T [22]	46.3	25.1	5.6	49.2↑ 2.9	27.2↑ 2.1	6.4↑ 0.8
RX-101-DCN [43]	49.2	28.3	6.1	52.0↑ 2.8	29.9↑ 1.6	7.4↑ 1.3
Swin-L [22]	52.5	35.1	8.0	55.4↑ 2.9	36.1↑ 1.0	8.9↑ 0.9

Table 2. Comparison with different backbones under separate and joint training.

Method	Separate Training			Joint Training		
	COCO	O365	VG	COCO	O365	VG
Instance Sampling	15.5	2.3	0.1	14.5	2.2	0.1
Global Sampling	34.5	19.2	4.8	39.8	20.2	5.0
Dataset-aware Sampling	43.0	22.2	4.8	45.3	23.2	5.7

Table 3. Effectiveness of query adaption under separate and joint training. All models are trained with the R-50 backbone.

Query num	Separate Training			Joint Training		
	COCO	Object365	VG	COCO	Object365	VG
100	41.5	20.2	3.8	42.3	20.3	5.0
300	43.0	22.2	4.8	45.3	23.2	5.7

Table 4. Ablation study on the number of adapted query under separate and joint training. All models are trained with the R-50 backbone under $1\times$ schedule.

4.3. Effectiveness of Query Adaption

We then evaluate the effectiveness of our proposed query adaption method. To compare, we design two baselines: “Instance Sampling” means only considering categories belonging to each training image. This limits the usage of category embedding to image level, where categories from different images are not interacted with each other;

Training set	COCO	VOC	Kitti	Clipart	WaterColor
O365	34.6	45.3	10.2	15.7	16.9
VG	22.6	32.0	12.9	12.7	17.1
O365 + VG	36.9	49.4	16.0	18.5	17.7

Table 5. Generalization capability of cross dataset evaluation. We perform OOD evaluation to verify the generalization capability. The jointly trained model are marked in the gray rows.

method	Separate	Joint		
	COCO	COCO	O365	VG
Sparse-RCNN [37] + Merge [15]	40.6	40.4	17.4	4.3
+ Cat-aligned Embed	40.1↓ 0.5	43.1↑ 2.7	21.3↑ 3.9	5.2↑ 0.9
+ D-Q-adapt Decoder	41.1↑ 1.0	44.0↑ 0.9	21.9↑ 0.6	5.4↑ 0.2
+ D-Q-adapt RPN	43.0↑ 1.9	45.6↑ 1.6	23.5↑ 1.6	5.7↑ 0.3

Table 6. Baseline evolution under separate and joint training.

“Global Sampling” means we consider the categories from all datasets. This allows all categories in different datasets to interact with each other. As shown in Table 3, query adaption is the key to effective multi-dataset training. With “Global Sampling”, separate and joint training drops significantly due to the negative interfering between different datasets. Finally, our dataset-aware query can support “Dataset-aware Sampling”, which samples category names according to the source dataset of each training image specifically to improve performance.

In addition, we also conduct an ablation on utilizing different numbers of adapted queries. As shown in Table 4, it is obvious that more adapted queries further enlarge the performance gains from joint training compared to separate training. To further verify the generalization capability of our adapted queries, we conduct the cross dataset evaluation by evaluating on five datasets out of the training set as

Method	Backbone	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
Deformable-DETR [30]	R-50	37.2	55.5	40.5	21.1	40.7	50.5
Dynamic-DETR [3]	R-50	42.9	61.0	46.3	24.6	44.9	54.4
Sparse-RCNN [37]	R-50	39.5	57.7	42.8	21.8	42.3	54.4
Ours (separate training)	R-50	43.0	60.8	47.1	27.5	45.4	55.6
Ours (joint training)	R-50	45.3	63.5	49.7	30.1	48.3	57.3

Table 7. **Comparison with SoTA query-based models** under standard $1\times$ setup using the ResNet-50 backbone on COCO val.

Method	Clipart	COCO	Comic	Deeplesion	Dota	Kitchen	KITTI	LISA	VOC	WaterColor	Widerface	Avg
Single [40]	32.1	47.3	45.8	51.3	57.5	87.7	64.3	88.3	78.5	52.4	48.9	59.4
UniDA [40]	55.8	47.0	53.4	53.4	56.3	90.0	68.0	87.6	82.4	60.6	51.3	64.2
Ours	63.5	56.8	49.7	62.5	65.2	94.6	89.3	97.0	81.3	57.1	63.7	71.0

Table 8. **Ablation on UODB**, where Avg means the average score of 11 datasets. Single means the baseline trained on each single dataset; we report the UniDA [40] of best Avg score in the original paper. Ours means our method jointly trained on UODB. All models take ResNet-50 as backbone.

shown in Table 5. Our models, especially the joint-training model, obtain non-trivial performance. Overall, above ablation studies well justify the effectiveness of our query adaptation.

4.4. Effectiveness of Detector

We also evaluate our query adapted detector separately to compare it with popular detectors. To conduct a fair comparison and ensure reproduction, we first train our method with a ResNet-50 on COCO data only using a standard $1\times$ schedule without augmentation and multi-dataset training. We compare our result with state-of-the-art query-based methods reported in a similar setting, such as [3, 37, 51]. Shown in Table 7, our method achieves state-of-the-art performance at 43.0 AP. In addition, our method really shines when trained under multi-dataset joint training and outperforms previous methods by clear margins.

4.5. Effectiveness of Each Component

To further clarify the baseline evolution from Sparse-RCNN to our model, we study the effectiveness of each component under separate and joint training. The two main differences can be summarized as follows: 1) *Category-aligned Embedding* replaces the traditional classifier through visual-word alignment and 2) *Dataset-aware Query* is incorporated for both the decoder and the extra RPN. As shown in Tab 6 Such a combination offers a 2.4 AP improvement on COCO under separate training and best boost the detector under joint training.

4.6. Effectiveness on Extreme Varied Datasets

We further evaluate on UODB [40], a combination of COCO [21], KITTI [7], WiderFace [46], VOC [6],

LISA [27], DOTA [42], Watercolor [13], Clipart [13], Comic [13], Kitchen [8], and DeepLesion [44]. These datasets cover wide variations in category, camera view, image style, etc. Following UODB [40], we train our model with ResNet-50 as backbone and report the AP_{50} for each dataset separately. As shown in Table 8, our Detection Hub can outperform UDA [40] by a large margin on avg score. It provides a convincing results, highlighting the effectiveness on the combinations of extremely varied datasets.

4.7. Comparison with SoTA

Finally, we compare our method with state-of-the-art object detectors using a large backbone. We report our best performance with Swin Large jointly trained on COCO, Object365, and Visual-Genome datasets. Unlike previous methods that require separate pre-training and fine-tuning steps, our method directly reports the performance after joint training without domain-specific fine-tuning. As shown in Table 9, we compare with recent methods that leverage multiple datasets. Compared to previous best results, our method archives new state-of-the-art results on both Object365 and Visual-Genome. On COCO, our method is competitive with recent work HTC++ [22]. However, HTC++ is first pre-trained on Object365 and COCO, then finetuned on COCO for a better result. Besides, unlike HTC++ combining the segmentation task and detection task, our method achieves this performance without the help of instance segmentation information. Meanwhile, our method demonstrates significant improvements on both COCO and Object365 over concurrent work [50]. In conclusion, thanks to our effective query adaption under multi-dataset joint training, our method can take advantage of different data to further advance the state-of-the-art results.

Method	Training Data	Data Size	COCO AP	O365 AP	VG	
					AP	AP50
HTC++ [22]	COCO+O365	0.7M	57.7	–	–	–
VinVL [47]	COCO+O365+VG+OID	2.6M	–	–	–	13.7
UniDet [50]	COCO+O365+OID+Mapillary	2.5M	52.9	33.7	–	–
Ours	COCO+O365+VG	0.8M	57.0	37.2	8.9	14.4

Table 9. Compared to SoTA on all three datasets. “-” indicates the numbers are not available for us.




Figure 4. Visualization of different predictions results from our method. The first odd rows show the predicted category names. The even rows show which dataset contributes to our final prediction.

4.8. Visualization

Finally, we visualize the predictions in Fig 4, where the odd rows show the predicted category names and the even rows show which dataset contributes to the prediction. Our predictions effectively combine the taxonomies from multiple datasets and yield better predictions. This well justifies our motivation for proposing a universal object detector.

5. Conclusion

Unifying multiple object detection datasets with a single object detector is a non-trivial problem. We propose

“*Detection Hub*” to address the inherent taxonomy differences and annotation inconsistency challenges. It can enjoy the synergy of multiple datasets and achieve substantial improvements over the independently trained detector baseline. Meanwhile, we also find that the performance on the dataset with a large vocabulary may be constrained by the maximum length of language embedding. In the future, we will try to expand our work to open-world object detection by combining more datasets from different domains to cover a wide variety of concepts.

Acknowledgement This project was supported by National Key R&D Program of China (No. 2021ZD0112805).

References

- [1] Z. Cai and N. Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 2018. 2
- [2] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 2, 3
- [3] X. Dai, Y. Chen, J. Yang, P. Zhang, L. Yuan, and L. Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *ICCV*, 2021. 2, 3, 5, 7
- [4] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2018. 3
- [5] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 2
- [6] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 7
- [7] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 2013. 7
- [8] G. Georgakis, M. A. Reza, A. Mousavian, P.-H. Le, and J. Košecká. Multiview rgb-d dataset for object instance detection. In *3DV*, 2016. 7
- [9] R. Girshick. Fast r-cnn. In *CVPR*, 2015. 2
- [10] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *CVPR*, 2014. 2
- [11] X. Gu, T.-Y. Lin, W. Kuo, and Y. Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *ICLR*, 2021. 3
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5, 6
- [13] N. Inoue, R. Furuta, T. Yamasaki, and K. Aizawa. Cross-domain weakly-supervised object detection through progressive domain adaptation. In *CVPR*, 2018. 7
- [14] R. Krishna, Y. Zhu, O. Groth, J. Johnson, K. Hata, J. Kravitz, S. Chen, Y. Kalantidis, L.-J. Li, D. A. Shamma, et al. Visual genome: Connecting language and vision using crowd-sourced dense image annotations. *IJCV*, 2017. 2, 5
- [15] J. Lambert, Z. Liu, O. Sener, J. Hays, and V. Koltun. Mseg: A composite dataset for multi-domain semantic segmentation. In *CVPR*, 2020. 3, 6
- [16] S. Lan, Z. Ren, Y. Wu, L. S. Davis, and G. Hua. Saccadenet: A fast and accurate object detector. In *CVPR*, 2020. 2
- [17] S. Lan, X. Yang, Z. Yu, Z. Wu, J. M. Alvarez, and A. Anandkumar. Vision transformers are good mask auto-labelers. *arXiv preprint arXiv:2301.03992*, 2023. 2
- [18] H. Li, Z. Wu, C. Zhu, C. Xiong, R. Socher, and L. S. Davis. Learning from noisy anchors for one-stage object detection. In *CVPR*, 2020. 2
- [19] L. H. Li, P. Zhang, H. Zhang, J. Yang, C. Li, Y. Zhong, L. Wang, L. Yuan, L. Zhang, J.-N. Hwang, et al. Grounded language-image pre-training. In *CVPR*, 2022. 2, 3, 4
- [20] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 2
- [21] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 2, 5, 7
- [22] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 5, 6, 7, 8
- [23] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. 5
- [24] D. Meng, X. Chen, Z. Fan, G. Zeng, H. Li, Y. Yuan, L. Sun, and J. Wang. Conditional detr for fast training convergence. In *ICCV*, 2021. 2, 3
- [25] L. Meng, H. Li, B.-C. Chen, S. Lan, Z. Wu, Y.-G. Jiang, and S.-N. Lim. Adavit: Adaptive vision transformers for efficient image recognition. In *CVPR*, 2022. 2
- [26] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. In *ICLR*, 2013. 3
- [27] A. Mogelmoose, M. M. Trivedi, and T. B. Moeslund. Vision-based traffic sign detection and analysis for intelligent driver assistance systems: Perspectives and survey. *TPAMI*, 2012. 7
- [28] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 3
- [29] M. E. Peters, M. Neumann, M. Iyyer, M. Gardner, C. Clark, K. Lee, and L. Zettlemoyer. Deep contextualized word representations. In *NAACL*, 2018. 3
- [30] H. Qiu, Y. Ma, Z. Li, S. Liu, and J. Sun. Borderdet: Border feature for dense object detection. In *ECCV*, 2020. 7
- [31] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1, 2, 3, 4
- [32] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al. Language models are unsupervised multi-task learners. *OpenAI blog*, 2019. 3
- [33] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 3
- [34] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *PAMI*, 2016. 2, 4
- [35] S. Shao, Z. Li, T. Zhang, C. Peng, G. Yu, X. Zhang, J. Li, and J. Sun. Objects365: A large-scale, high-quality dataset for object detection. In *ICCV*, 2019. 2, 5
- [36] B. Singh, H. Li, A. Sharma, and L. S. Davis. R-fcn-3000 at 30fps: Decoupling detection and classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018. 2

- [37] P. Sun, R. Zhang, Y. Jiang, T. Kong, C. Xu, W. Zhan, M. Tomizuka, L. Li, Z. Yuan, C. Wang, et al. Sparse r-cnn: End-to-end object detection with learnable proposals. In *CVPR*, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [38] Z. Tian, C. Shen, H. Chen, and T. He. Fcos: Fully convolutional one-stage object detection. In *CVPR*, 2019. [2](#)
- [39] R. Wang, D. Chen, Z. Wu, Y. Chen, X. Dai, M. Liu, Y.-G. Jiang, L. Zhou, and L. Yuan. Bevt: Bert pretraining of video transformers. In *CVPR*, 2022. [1](#)
- [40] X. Wang, Z. Cai, D. Gao, and N. Vasconcelos. Towards universal object detection by domain attention. In *CVPR*, 2019. [2](#), [3](#), [5](#), [7](#)
- [41] Y. Wang, X. Zhang, T. Yang, and J. Sun. Anchor detr: Query design for transformer-based detector. In *AAAI*, 2022. [3](#)
- [42] G.-S. Xia, X. Bai, J. Ding, Z. Zhu, S. Belongie, J. Luo, M. Datcu, M. Pelillo, and L. Zhang. Dota: A large-scale dataset for object detection in aerial images. In *CVPR*, 2018. [7](#)
- [43] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017. [6](#)
- [44] K. Yan, X. Wang, L. Lu, and R. M. Summers. Deeplesion: automated mining of large-scale lesion annotations and universal lesion detection with deep learning. *Journal of medical imaging*, 2018. [7](#)
- [45] H. Yang, H.-Y. Wu, and H. Chen. Detecting 11k classes: Large scale object detection without fine-grained bounding boxes. In *ICCV*, 2019. [3](#)
- [46] S. Yang, P. Luo, C.-C. Loy, and X. Tang. Wider face: A face detection benchmark. In *CVPR*, 2016. [7](#)
- [47] P. Zhang, X. Li, X. Hu, J. Yang, L. Zhang, L. Wang, Y. Choi, and J. Gao. Vinvl: Revisiting visual representations in vision-language models. In *CVPR*, 2021. [1](#), [8](#)
- [48] X. Zhao, S. Schulter, G. Sharma, Y.-H. Tsai, M. Chandraker, and Y. Wu. Object detection with a unified label space from multiple datasets. In *ECCV*, 2020. [3](#)
- [49] Y. Zhong, J. Yang, P. Zhang, C. Li, N. Codella, L. H. Li, L. Zhou, X. Dai, L. Yuan, Y. Li, et al. Regionclip: Region-based language-image pretraining. In *CVPR*, 2022. [3](#)
- [50] X. Zhou, V. Koltun, and P. Krähenbühl. Simple multi-dataset detection. In *CVPR*, 2022. [3](#), [7](#), [8](#)
- [51] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *ICLR*, 2020. [2](#), [5](#), [7](#)