

Grounding DINO: Marrying DINO with Grounded Pre-Training for Open-Set Object Detection

Shilong Liu^{1,2*}, Zhaoyang Zeng², Tianhe Ren², Feng Li^{2, 3}, Hao Zhang^{2, 3},
 Jie Yang^{2, 4}, Qing Jiang^{2, 6}, Chunyuan Li⁵, Jianwei Yang⁵,
 Hang Su¹, Jun Zhu^{1**}, Lei Zhang^{2**}.

¹ Dept. of Comp. Sci. and Tech., BNRist Center, State Key Lab for Intell. Tech. & Sys.,
 Institute for AI, Tsinghua-Bosch Joint Center for ML, Tsinghua University

² International Digital Economy Academy (IDEA)

³ The Hong Kong University of Science and Technology

⁴ The Chinese University of Hong Kong (Shenzhen) ⁵ Microsoft Research, Redmond

⁶ South China University of Technology

liusl20@mails.tsinghua.edu.cn, leizhang@idea.edu.cn

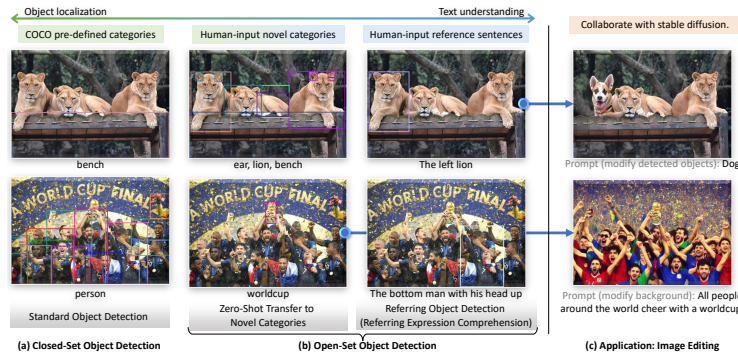


Fig. 1: (a) Closed-set object detection requires models to detect objects of pre-defined categories. (b) We evaluate models on novel objects and standard Referring expression comprehension (REC) benchmarks for model generalizations on novel objects with attributes. (c) We present an image editing application by combining Grounding DINO and Stable Diffusion [40]. Best viewed in colors.

Abstract. In this paper, we develop an open-set object detector, called Grounding DINO, by marrying Transformer-based detector DINO with grounded pre-training, which can detect arbitrary objects with human inputs such as category names or referring expressions. The key solution of open-set object detection is introducing language to a closed-set detector for open-set concept generalization. To effectively fuse language and vision modalities, we conceptually divide a closed-set detector into three phases and propose a tight fusion solution, which includes a feature enhancer, a language-guided query selection, and a cross-modality decoder for modalities fusion. We first pre-train Grounding DINO on large-scale datasets,

* This work was done when Shilong Liu, Feng Li, Hao Zhang, Jie Yang, and Qing Jiang were interns at IDEA.

** Corresponding authors.

including object detection data, grounding data, and caption data, and evaluate the model on both open-set object detection and referring object detection benchmarks. Grounding DINO performs remarkably well on all three settings, including benchmarks on COCO, LVIS, ODinW, and RefCOCO+/g. Grounding DINO achieves a 52.5 AP on the COCO zero-shot¹ detection benchmark. It sets a new record on the ODinW zero-shot benchmark with a mean 26.1 AP. We release some checkpoints and inference codes at <https://github.com/IDEA-Research/GroundingDINO>.

Keywords: Object Detection · Image Grounding · Multi-modal learning

1 Introduction

A key indicator of an Artificial General Intelligence (AGI) system’s capability is its proficiency in handling open-world scenarios. In this paper, we aim to develop a strong system to detect arbitrary objects specified by human language inputs, a task commonly referred to as *open-set object detection*². The task has wide applications for its great potential as a generic object detector. For example, we can cooperate with generative models for image editing (as shown in Fig. 1 (b)).

In pursuit of this goal, we design the strong open-set object detector **Grounding DINO** by following the two principles: tight modality fusion based on DINO [56] and large-scale grounded pre-train for concept generalization.

Tight modality fusion based on DINO. The key to open-set detection is introducing language for unseen object generalization [1, 6, 24]. Most existing open-set detectors are developed by extending closed-set detectors to open-set scenarios with language information. As shown in Fig. 2, a closed-set detector typically has three important modules, a backbone for feature extraction, a neck for feature enhancement, and a head for region refinement (or box prediction). A closed-set detector can be generalized to detect novel objects by learning language-aware region embeddings so that each region can be classified into novel categories in a language-aware semantic space. The key to achieving this goal is using contrastive loss between region outputs and language features at the neck and/or head outputs.

To help a model align cross-modality information some work tried to fuse features before the final loss stage. We summarize the modularized design




Fig. 2: Extending closed-set detectors to open-set scenarios.

¹ In this paper, ‘zero-shot’ refers to scenarios where the training split of the test dataset is not utilized in the training process.

² We view the terms *open-set object detection*, *open-world object detection*, and *open-vocabulary object detection* the same task in this paper. To avoid confusion, we always use *open-set object detection* in our paper.

of object detectors in Fig. 2. Feature fusion can be performed in three phases: neck (phase A), query initialization (phase B), and head (phase C). For example, GLIP [24] performs early fusion in the neck module (phase A), and OV-DETR [54] uses language-aware queries as head inputs (phase B). We argue that introducing more feature fusion into the pipeline can facilitate better alignment between different modality features, thereby achieving better performance.

Although conceptually simple, it is hard for previous work to perform feature fusion in all three phases. The design of classical detectors like Faster RCNN makes it hard to interact with language information in most blocks. Unlike classical detectors, the Transformer-based detector method such as DINO has a consistent structure with language blocks. The layer-by-layer design enables it to interact with language information easily. Under this principle, we design three feature fusion approaches in the neck, query initialization, and head phases. More specifically, we design a feature enhancer by stacking self-attention, text-to-image cross-attention, and image-to-text cross-attention as the neck module. We then develop a language-guided query selection method to initialize queries for the detection head. We also design a cross-modality decoder for the head phase with image and text cross-attention layers to boost query representations.

Large-scale grounded pre-train for zero-shot transfer. Most existing open-set models [13, 20] rely on pre-trained CLIP models for concept generalization. Nevertheless, the efficacy of CLIP, specifically pre-trained on image-text pairs, is limited for region-text pair detection tasks, as identified in the RegionCLIP study by RegionCLIP [60]. In contrast, GLIP [24] presents a different way by reformulating object detection as a phrase grounding task and introducing contrastive training between object regions and language phrases on large-scale data. It shows great flexibility for heterogeneous datasets and remarkable performance on closed-set and open-set detection.

We have adopted and refined the grounded training methodology. GLIP’s approach involves concatenating all categories into a sentence in a random order. However, the direct category names concatenation does not consider the potential influence of unrelated categories on each other when extracting features. To mitigate this issue and improve model performance during grounded training, we introduce a technique that utilizes sub-sentence level text features. It removes the attention between unrelated categories during word feature extractions. Further elaboration on this technique can be found in Section 3.4.

We pre-train the Grounding DINO on a large-scale dataset and evaluate the performance on mainstream object detection benchmarks like COCO [28]. While some studies have examined open-set detection models under a "partial label" framework—training on a subset of data (e.g., base categories) and testing on additional categories—we advocate for a fully zero-shot approach to enhance practical applicability. Moreover, we extend the model to another important scenario Referring Expression Comprehension (REC) [29, 33]³, where objects are described with attributes.

³ We use the term *Referring Expression Comprehension (REC)* and *Referring (Object) Detection* exchangeable in this paper.

We conduct experiments on all three settings, including closed-set detection, open-set detection, and referring object detection, as shown in Fig. 1, to comprehensively evaluate open-set detection performance. Grounding DINO outperforms competitors by a large margin. For example, Grounding DINO reaches a 52.5 AP on COCO minival without any COCO training data. It also establishes a new state of the art on the ODinW [22] zero-shot benchmark with a 26.1 mean AP.

Model	Model Design			Text Prompt Represent. Level (Sec. 3.4)	Closed-Set Settings COCO	Zero-Shot Transfer			Referring Detection RefCOCO/ \pm /g
	Base Detector	Fusion (Fig. 2)	CLIP			COCO	LVIS	ODinW	
ViLD [13]	Mask R-CNN	-	✓	sentence	✓	partial label	partial label		
RegionCLIP [60]	Faster RCNN	-	✓	sentence	✓	partial label	partial label		
FindIt [20]	Faster RCNN	A		sentence	✓	partial label			
MDETR [17]	DETR	A,C		word			fine-tune	zero-shot	fine-tune
DQ-DETR [44]	DETR	A,C		word	✓		zero-shot		fine-tune
GLIP [24]	DyHead	A		word	✓	zero-shot	zero-shot	zero-shot	
GLIPv2 [57]	DyHead	A		word	✓	zero-shot	zero-shot	zero-shot	
OV-DETR [54]	Deformable DETR	B	✓	sentence	✓	partial label	partial label		
OWL-ViT [34]	-	-	✓	sentence	✓	partial label	partial label	zero-shot	
DetCLIP [51]	ATSS	-	✓	sentence			zero-shot	zero-shot	
OmDet [59]	Sparse R-CNN	C	✓	sentence	✓			zero-shot	
Grounding DINO (Ours)	DINO	A,B,C		sub-sentence	✓	zero-shot	zero-shot	zero-shot	zero-shot

Table 1: A comparison of previous open-set object detectors. Our summarization is based on the experiments in their paper, but not the ability to extend their models to other tasks. It is worth noting that some related works may not (only) be designed for the open-set object detection initially, like MDETR [17] and GLIPv2 [57], but we list them here for a comprehensive comparison with existing work. We use the term “partial label” for the settings, where models are trained on partial data (e.g. base categories) and evaluated on other cases. [55]

2 Related Work

Detection Transformers. Grounding DINO is built upon the DETR-like model DINO [56], which is an end-to-end Transformer-based detector. DETR was first proposed in [2] and then has been improved from many directions [4, 5, 12, 16, 32, 47, 63] in the past few years. DAB-DETR [30] introduces anchor boxes as DETR queries for more accurate box prediction. DN-DETR [23] proposes a query-denoising approach to stabilizing the bipartite matching. DINO [56] further develops several techniques including contrastive de-noising and sets a new record on the COCO object detection benchmark. However, such detectors mainly focus on closed-set detection and are difficult to generalize to novel classes because of the limited pre-defined categories.

Open-Set Object Detection. Open-set object detection is trained using existing bounding box annotations and aims at detecting arbitrary classes with the help of language generalization. OV-DETR [55] uses image and text embedding encoded by a CLIP model as queries to decode the category-specified boxes in the DETR framework [2]. ViLD [13] distills knowledge from a CLIP teacher model into a R-CNN-like detector so that the learned region embeddings contain the semantics of language. GLIP [11] formulates object detection as a grounding problem and leverages additional grounding data to help learn aligned semantics

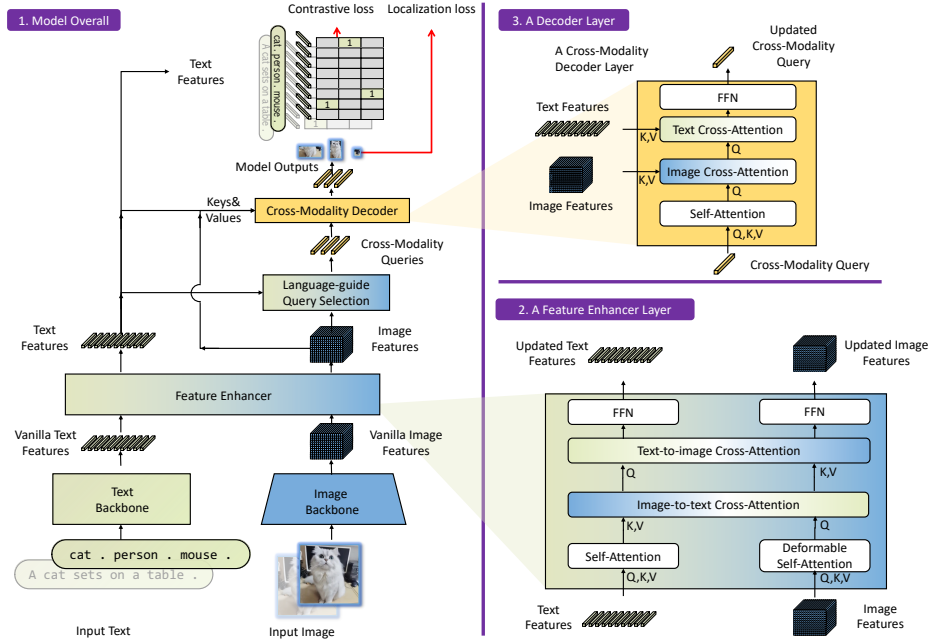


Fig. 3: The framework of Grounding DINO. We present the overall framework, a feature enhancer layer, and a decoder layer in block 1, block 2, and block 3, respectively.

at phrase and region levels. It shows that such a formulation can even achieve stronger performance on fully-supervised detection benchmarks. DetCLIP [51] involves large-scale image captioning datasets and uses the generated pseudo labels to expand the knowledge database. The generated pseudo labels effectively help extend the generalization ability.

However, previous works only fuse multi-modal information in partial phases, which may lead to sub-optimal language generalization ability. For example, GLIP only considers fusion in the feature enhancement (phase A) and OV-DETR only injects language information at the decoder inputs (phase B). Moreover, the REC task is normally overlooked in evaluation, which is an important scenario for open-set detection. We compare our model with other open-set methods in Table 1.

3 Grounding DINO

Grounding DINO outputs multiple pairs of object boxes and noun phrases for a given (Image, Text) pair. For example, as shown in Fig. 3, the model locates a cat and a table from the input image and extracts word `cat` and `table` from the input text as corresponding labels. Both object detection and REC tasks can be aligned with the pipeline. Following GLIP [24], we concatenate all category names as input texts for object detection tasks. REC requires a bounding box for

each text input. We use the output object with the largest scores as the output for the REC task.

Grounding DINO is a dual-encoder-single-decoder architecture. It contains an image backbone for image feature extraction, a text backbone for text feature extraction, a feature enhancer for image and text feature fusion (Sec. 3.1), a language-guided query selection module for query initialization (Sec. 3.2), and a cross-modality decoder for box refinement (Sec. 3.3).

For each (Image, Text) pair, we first extract vanilla image features and vanilla text features using an image backbone and a text backbone, respectively. The two vanilla features are fed into a feature enhancer module for cross-modality feature fusion. After obtaining cross-modality text and image features, we use a language-guided query selection module to select cross-modality queries from image features. Like the object queries in most DETR-like models, these cross-modality queries will be fed into a cross-modality decoder to probe desired features from the two modal features and update themselves. The output queries of the last decoder layer will be used to predict object boxes and extract corresponding phrases.




Fig. 4: Comparisons of text representations.

3.1 Feature Extraction and Enhancer

Given an (Image, Text) pair, we extract multi-scale image features with an image backbone like Swin Transformer [31], and text features with a text backbone like BERT [7]. Following previous DETR-like detectors [56,63], multi-scale features are extracted from the outputs of different blocks. After extracting vanilla image and text features, we fed them into a feature enhancer for cross-modality feature fusion. The feature enhancer includes multiple feature enhancer layers. We illustrate a feature enhancer layer in Fig. 3 block 2. We leverage the Deformable self-attention to enhance image features and the vanilla self-attention for text feature enhancers. Inspired by GLIP [24], we add an image-to-text and a text-to-image cross-attention modules for feature fusion. These modules help align features of different modalities.

3.2 Language-Guided Query Selection

Grounding DINO aims to detect objects from an image specified by an input text. To effectively leverage the input text to guide object detection, we design a

language-guided query selection module to select features that are more relevant to the input text as decoder queries.

Let’s denote the image feature as $\mathbf{X}_I \in \mathbb{R}^{N_I \times d}$ and the text features as $\mathbf{X}_T \in \mathbb{R}^{N_T \times d}$. Here, N_I represents the number of image tokens, N_T indicates the number of text tokens, and d corresponds to the feature dimension. In our experiments, we specifically utilize a feature dimension of $d = 256$. Typically, in our models, the value of N_I exceeds 10,000, while N_T remains below 256. Our objective is to extract N_q queries from the encoder’s image features to be used as inputs for the decoder. In alignment with the DINO method, we set N_q to be 900. The top N_q query indices for the image feature, denoted as \mathbf{I}_{N_q} , are selected using the following expression:

$$\mathbf{I}_{N_q} = \text{Top}_{N_q}(\text{Max}^{(-1)}(\mathbf{X}_I \mathbf{X}_T^\top)). \quad (1)$$

In this expression, Top_{N_q} represents the operation to pick the top N_q indices. The function $\text{Max}^{(-1)}$ executes the \max operation along the -1 dimension, and the symbol \top denotes matrix transposition. We present the query selection process in Algorithm 1 in PyTorch style. The language-guided query selection module outputs N_q indices. We can extract features based on the selected indices to initialize queries. Following DINO [56], we use mixed query selection to initialize decoder queries. Each decoder query contains two parts: content part and positional part [32], respectively. We formulate the positional part as dynamic anchor boxes [30], which are initialized with encoder outputs. The other part, the content queries, are set to be learnable during training.

3.3 Cross-Modality Decoder

We develop a cross-modality decoder to combine image and text modality features, as shown in Fig. 3 block 3. Each cross-modality query is fed into a self-attention layer, an image cross-attention layer to combine image features, a text cross-attention layer to combine text features, and an FFN layer in each cross-modality decoder layer. Each decoder layer has an extra text cross-attention layer compared with the DINO decoder layer, as we need to inject text information into queries for better modality alignment.

3.4 Sub-Sentence Level Text Feature

Two kinds of text prompts are explored in previous works, which we named as sentence level representation and word level representation, as shown in Fig. 4. Sentence level representation [34, 51] encodes a whole sentence to one feature. If some sentences in phrase grounding data have multiple phrases, it extracts these phrases and discards other words. In this way, it removes the influence between words while losing fine-grained information in sentences. Word level representation [11, 17] enables encoding multiple category names with one forward but introduces unnecessary dependencies among categories, especially when the input text is a concatenation of multiple category names in an arbitrary order.

Model	Backbone	Pre-Training Data	Zero-Shot 2017val	Fine-Tuning 2017val/test-dev
Faster R-CNN	RN50-FPN	-	-	40.2 / -
Faster R-CNN	RN101-FPN	-	-	42.0 / -
DyHead-T [5]	Swin-T	-	-	49.7 / -
DyHead-L [5]	Swin-L	-	-	58.4 / 58.7
DyHead-L [5]	Swin-L	O365,ImageNet21K	-	60.3 / 60.6
SoftTeacher [49]	Swin-L	O365,SS-COCO	-	60.7 / 61.3
DINO(Swin-L) [56]	Swin-L	O365	-	62.5 / -
DyHead-T† [5]	Swin-T	O365	43.6	53.3 / -
GLIP-T (B) [24]	Swin-T	O365	44.9	53.8 / -
GLIP-T (C) [24]	Swin-T	O365,GoldG	46.7	55.1 / -
GLIP-L [24]	Swin-L	FourODs,GoldG,Cap24M	49.8	60.8 / 61.0
DINO(Swin-T)† [56]	Swin-T	O365	46.2	56.9 / -
Grounding DINO T (Ours)	Swin-T	O365	46.7	56.9 / -
Grounding DINO T (Ours)	Swin-T	O365,GoldG	48.1	57.1 / -
Grounding DINO T (Ours)	Swin-T	O365,GoldG,Cap4M	48.4	57.2 / -
Grounding DINO L (Ours)	Swin-L	O365,OI [18],GoldG	52.5	62.6 / 62.7 (63.0 / 63.0)*
Grounding DINO L (Ours)	Swin-L	O365,OI,GoldG,Cap4M,COCO,RefC	60.7†	62.6 / -

Table 2: Zero-shot domain transfer and fine-tuning on COCO. * The results in brackets are trained with $1.5\times$ image sizes, i.e., with a maximum image size of 2000. †The models map a subset of O365 categories to COCO for zero-shot evaluations. ‡This is not a real zero-shot performance as we add COCO data during model training. We list the result here for a reference.

As shown in Fig. 4 (b), some unrelated words interact during attention. To avoid unwanted word interactions, we introduce attention masks to block attentions among unrelated category names, named “sub-sentence” level representation. It eliminates the influence between different category names while keeping per-word features for fine-grained understanding.

3.5 Loss Function

Following previous DETR-like works [2, 23, 30, 32, 56, 63], we use the L1 loss and the GIOU [39] loss for bounding box regressions. We follow GLIP [24] and use contrastive loss between predicted objects and language tokens for classification. Specifically, we dot product each query with text features to predict logits for each text token and then compute focal loss [27] for each logit. Box regression and classification costs are first used for bipartite matching between predictions and ground truths. We then calculate final losses between ground truths and matched predictions with the same loss components. Following DETR-like models, we add auxiliary loss after each decoder layer and after the encoder outputs.

4 Experiments

We conduct extensive experiments on three settings: a closed-set setting on the COCO detection benchmark (Sec. C.2), an open-set setting on zero-shot COCO, LVIS, and ODinW (Sec. 4.2), and a referring detection setting on RefCOCO/+g (Sec. 4.3). Ablations are then conducted to show the effectiveness of our model design (Sec. 4.5). We also explore a way to transfer a well-trained DINO to the

open-set scenario by training a few plug-in modules in Sec. C.1. The test of our model efficiency is presented in Sec. C.4.

4.1 Implementation Details

We trained two model variants, Grounding DINO T with Swin-T [31], and Grounding DINO L with Swin-L [31] as an image backbone, respectively. We leveraged BERT-base [7] from Hugging Face [48] as text backbones. As we focus more on the model performance on novel classes, we list zero-shot transfer and referring detection results in the main text.

By default, we use 900 queries in our model following DINO. We set the maximum text token number as 256. Using BERT as our text encoder, we follow BERT to tokenize texts with a BPE scheme [41]. We use six feature enhancer layers in the feature enhancer module. The cross-modality decoder is composed of six decoder layers as well. We leverage deformable attention [63] in image cross-attention layers.

Both matching costs and final losses include classification losses (or contrastive losses), box L1 losses, and GIOU [39] losses. Following DINO, we set the weight of classification costs, box L1 costs, and GIOU costs as 2.0, 5.0, and 2.0, respectively, during Hungarian matching. The corresponding loss weights are 1.0, 5.0, and 2.0 in the final loss calculation.

Our Swin Transformer Tiny models are trained on 16 Nvidia V100 GPUs with a total batch size of 32. We extract three image feature scales, from $8\times$ to $32\times$. It is named “4scale” in DINO since we downsample the $32\times$ feature map to $64\times$ as an extra feature scale. For the model with Swin Transformer Large, we extract four image feature scales from backbones, from $4\times$ to $32\times$. The model is trained on 64 Nvidia A100 GPUs with a total batch size of 64.

4.2 Zero-Shot Transfer of Grounding DINO

In this setting, we pre-train models on large-scale datasets and directly evaluate models on new datasets. We also list some fine-tuned results for a more thorough comparison of our model with prior works.

COCO Benchmark We compare Grounding DINO with GLIP and DINO in Table 2. We pre-train models on large-scale datasets and directly evaluate our model on the COCO benchmark. As the O365 dataset [42] has (nearly⁴) covered all categories in COCO, we evaluate an O365 pre-trained DINO on COCO as a zero-shot baseline. The result shows that DINO performs better on the COCO zero-shot transfer than DyHead. Grounding DINO outperforms all previous models on the zero-shot transfer setting, with +0.5AP and +1.8AP compared with DINO and GLIP under the same setting. Grounding data is still helpful for Grounding DINO, introducing more than 1AP (48.1 vs. 46.7) on the zero-shot

⁴ It is not an exact mapping between O365 and COCO categories. We made some approximations during evaluation.

transfer setting. With stronger backbones and larger data, Grounding DINO sets a new record of 52.5 AP on the COCO object detection benchmark without seeing any COCO images during training. Grounding DINO obtains a 62.6 AP on COCO minival, outperforming DINO’s 62.5 AP. When enlarging the input images by $1.5\times$, the benefits reduce. We suspect that the text branch enlarges the gap between models with different input images. Even though the performance plateaus with larger input size, Grounding DINO gets an impressive 63.0 AP on COCO test-dev with fine-tuning on the COCO dataset(See the number in brackets of Table 2).

LVIS Benchmark LVIS [14] is a dataset for long-tail objects. It contains more than 1000 categories for evaluation. We use LVIS as a downstream task to test the zero-shot abilities of our model. We use GLIP and DetCLIPv2 as baselines for our models. The results are shown in Table 3.

Model	Backbone	Pre-Training Data	MiniVal [17]		
			AP	APr/APc/APf	
<i>Zero-Shot Setting</i>					
GLIP-T (C)	Swin-T	O365,GoldG	24.9	17.7/19.5/31.0	
GLIP-T	Swin-T	O365,GoldG,Cap4M	26.0	20.8/21.4/31.0	
DetCLIPv2	Swin-T	O365,GoldG,CC15M	40.4	36.0/41.7/40.0	
Grounding DINO T	Swin-T	O365,GoldG	25.6	14.4/19.6/32.2	
Grounding DINO T	Swin-T	O365,GoldG,Cap4M	27.4	18.1/23.3/32.7	
Grounding DINO L	Swin-L	O365,OI,GoldG,Cap4M, COCO,RefC	33.9	22.2/30.7/38.8	
<i>Fine-Tune Setting</i>					
MDETR	RN101	GoldG,RefC	24.2	20.9/24.9/24.3	
Mask R-CNN	RN101	-	33.3	26.3/34.0/33.9	
DetCLIPv2 [50]	Swin-T	O365,GoldG,CC15M	50.7	44.3/52.4/50.3	
Grounding DINO T	Swin-T	O365,GoldG	52.1	35.4/51.3/55.7	

Table 3: Model results on LVIS.

We found two interesting phenomena in the results. First, Grounding DINO works better than common objects than GLIP, but worse on rare categories. We reviewed DETR-like models on LVIS and noted these models often exhibit lower rare category AP despite similar overall AP, like Table 2 of [8] and Table 6 of [17]. To our knowledge, no existing DETR-like models effectively address the rarity challenge in LVIS without extra training data, which may be a characteristic limitation of the architecture.

The other phenomenon is that Grounding DINO has larger gains with more data than GLIP. For example, Grounding DINO introduces +1.8 AP gains with the caption data Cap4M, whereas GLIP has only +1.1 AP. We believe that Grounding DINO has better scalability compared with GLIP. A larger-scale training will be left as our future work.

Although achieving better results than GLIP, we found that Grounding DINO is inferior to DetCLIPv2, which is trained on a larger scale data. This performance difference might be attributed to the disparity in data distribution between the training dataset and the LVIS dataset.

To unveil the full potential of Grounding DINO, we fine-tuned it on the LVIS dataset. Table 3 highlights the commendable capability of our model. Remarkably, despite being pre-trained only on the O365 and GoldG datasets, Grounding DINO outperforms DetCLIPv2-T by a margin of 1.5 AP. This result shows that Grounding DINO might have learned a better object-level representation which helps yield a better performance after fine-tuning (aligning with the target dataset). In our future work, we will perform more studies, including varying the semantic concept coverage of the training data and increasing the scale of the training data, to further improve the zero-shot generalization performance.

Model	Language Input	Backbone	Model Size	Pre-Training Data	Test	
					$AP_{average}$	AP_{median}
<i>Zero-Shot Setting</i>						
MDETR [17]	✓	ENB5 [45]	169M	GoldG,RefC	10.7	3.0
OWL-ViT [34]	✓	ViT L/14 (CLIP)	>1243M	O365, VG	18.8	9.8
GLIP-T [24]	✓	Swin-T	232M	O365,GoldG,Cap4M	19.6	5.1
OmDet [59]	✓	ConvNeXt-B	230M	COCO,O365,LVIS,PhraseCut	19.7	10.8
GLIPv2-T [58]	✓	Swin-T	232M	O365,GoldG,Cap4M	22.3	8.9
DetCLIP [51]	✓	Swin-L	267M	O365,GoldG,YFCC1M	24.9	18.3
Florence [53]	✓	CoSwinH	≈841M	FLD900M,O365,GoldG	25.8	14.3
Grounding DINO T(Ours)	✓	Swin-T	172M	O365,GoldG	20.0	9.5
Grounding DINO T(Ours)	✓	Swin-T	172M	O365,GoldG,Cap4M	22.3	11.9
Grounding DINO L(Ours)	✓	Swin-L	341M	O365,OI,GoldG,Cap4M,COCO,RefC	26.1	18.4
<i>Few-Shot Setting</i>						
DyHead-T [5]	✗	Swin-T	≈100M	O365	37.5	36.7
GLIP-T [24]	✓	Swin-T	232M	O365,GoldG,Cap4M	38.9	33.7
DINO-Swin-T [56]	✗	Swin-T	49M	O365	41.2	41.1
OmDet [59]	✓	ConvNeXt-B	230M	COCO,O365,LVIS,PhraseCut	42.4	41.7
Grounding DINO T(Ours)	✓	Swin-T	172M	O365,GoldG	46.4	51.1
<i>Full-Shot Setting</i>						
GLIP-T [24]	✓	Swin-T	232M	O365,GoldG,Cap4M	62.6	62.1
DyHead-T [5]	✗	Swin-T	≈100M	O365	63.2	64.9
DINO-Swin-T [56]	✗	Swin-T	49M	O365	66.7	68.5
OmDet [59]	✓	ConvNeXt-B	230M	COCO,O365,LVIS,PhraseCut	67.1	71.2
DINO-Swin-L [56]	✗	Swin-L	218M	O365	68.8	70.7
Grounding DINO T(Ours)	✓	Swin-T	172M	O365,GoldG	70.7	76.2

Table 4: Model results on the ODinW benchmark.

ODinW Benchmark ODinW (Object Detection in the Wild) [22] is a more challenging benchmark to test model performance under real-world scenarios. It collects more than 35 datasets for evaluation. We report three settings, zero-shot, few-shot, and full-shot results in Table 4. Grounding DINO performs well on this benchmark. With only O365 and GoldG for pre-train, Grounding DINO T outperforms DINO on few-shot and full-shot settings. Impressively, Grounding DINO with a Swin-T backbone outperforms DINO with Swin-L on the full-shot setting.

Grounding DINO outperforms GLIP under the same backbone for the zero-shot setting. Grounding DINO and GLIPv2-T show similar $AP_{average}$. However, a key distinction lies in the AP_{median} , where Grounding DINO significantly outperforms GLIPv2-T (11.9 vs 8.9). This suggests that while GLIPv2 may exhibit larger performance variance across different datasets, Grounding DINO maintains a more consistent performance level. GLIPv2 incorporates advanced techniques like masked text training and cross-instance contrastive learning, making it

more complex than our Grounding DINO model. Moreover, our model is more compact (172M parameters) compared to GLIPv2 (232M parameters). These factors combined—performance consistency, model complexity, and size—should address concerns about our model’s capability in true open-set scenarios.

Grounding DINO L set a new record on ODinW zero-shot with a 26.1 AP, even outperforming the giant Florence models [53]. The results show the generalization and scalability of Grounding DINO.

4.3 Referring Object Detection Settings

We further explore our models’ performances on the REC task. We leverage GLIP [24] as our baseline. We evaluate the model performance on RefCOCO/+g directly.⁵ The results are shown in Table 5. Grounding DINO outperforms GLIP under the same setting. Nevertheless, both GLIP and Grounding DINO perform not well without REC data. More training data like caption data or larger models help the final performance, but quite minor. After injecting RefCOCO/+g data into training, Grounding DINO obtains significant gains. The results reveal that most nowadays open-set object detectors need to pay more attention for a more fine-grained detection.

Method	Backbone	Pre-Training Data	Fine-tuning	RefCOCO			RefCOCO+			RefCOCOg	
				val	testA	testB	val	testA	testB	val	test
MAttNet [52]	R101	None	✓	76.65	81.14	69.99	65.33	71.62	56.02	66.58	67.27
VGTR [9]	R101	None	✓	79.20	82.32	73.78	63.91	70.09	56.51	65.73	67.23
TransVG [6]	R101	None	✓	81.02	82.72	78.35	64.82	70.70	56.94	68.67	67.73
VILLA L* [10]	R101	CC, SBU, COCO, VG	✓	82.39	87.48	74.84	76.17	81.54	66.84	76.18	76.71
RefTR [25]	R101	VG	✓	85.65	88.73	81.16	77.55	82.26	68.99	79.25	80.01
MDETR [17]	R101	GoldG,RefC	✓	86.75	89.58	81.41	79.52	84.09	70.62	81.64	80.89
DQ-DETR [44]	R101	GoldG,RefC	✓	88.63	91.04	83.51	81.66	86.15	73.21	82.76	83.44
GLIP-T(B)	Swin-T	O365,GoldG		49.96	54.69	43.06	49.01	53.44	43.42	65.58	66.08
GLIP-T	Swin-T	O365,GoldG,Cap4M		50.42	54.30	43.83	49.50	52.78	44.59	66.09	66.89
Grounding DINO T (Ours)	Swin-T	O365,GoldG		50.41	57.24	43.21	51.40	57.59	45.81	67.46	67.13
Grounding DINO T (Ours)	Swin-T	O365,GoldG,RefC		73.98	74.88	59.29	66.81	69.91	56.09	71.06	72.07
Grounding DINO T (Ours)	Swin-T	O365,GoldG,RefC	✓	89.19	91.86	85.99	81.09	87.40	74.71	84.15	84.94
Grounding DINO L (Ours)*	Swin-L	O365,OI,GoldG,Cap4M,COCO,RefC	✓	90.56	93.19	88.24	82.75	88.95	75.92	86.13	87.02

Table 5: Top-1 accuracy comparison on the referring expression comprehension task. We mark the best results in bold. All models are trained with a ResNet-101 backbone. We use the notations “CC”, “SBU”, “VG”, “OI”, “O365”, and “YFCC” for Conceptual Captions [43], SBU Captions [35], Visual Genome [19], OpenImage [21], Objects365 [62], YFCC100M [46] respectively. The term “RefC” is used for RefCOCO, RefCOCO+, and RefCOCOg three datasets. * There might be a data leak since COCO includes validation images in RefC. But the annotations of the two datasets are different.

4.4 Effects of RefC and COCO Data

We add the RefCOCO/+g (we note it as “RefC” in tables) and COCO into training in some settings. We explore the influence of these data in Table 6.

⁵ We used the official released code and checkpoints in <https://github.com/microsoft/GLIP>.

The results show that RefC helps improve the COCO zero-shot and fine-tuning performance but hurts the LVIS and ODinW results. With COCO introduced, the COCO results are greatly improved. It shows that COCO brings marginal improvements in LVIS and slight decreases on ODinW.

Model	Pre-Train	COCO minival		LVIS minival	ODinW
		Zero-Shot	Fine-Tune	Zero-Shot	Zero-Shot
Grounding DINO T	O365,GoldG	48.1	57.1	25.6	20.0
Grounding DINO T	O365,GoldG,RefC	48.5	57.3	21.9	17.7
Grounding DINO T	O365,GoldG,RefC,COCO	56.1	57.5	22.3	17.4

Table 6: Impacts of RefC and COCO data for open-set settings. All models are trained with a Swin Transformer Tiny backbone.

4.5 Ablations

We conduct ablation studies in this section. We propose a tight fusion grounding model for open-set object detection and a sub-sentence level text prompt. To verify the effectiveness of the model design, we remove some fusion blocks for different variants. Results are shown in Table 7. All models are pre-trained on O365 with a Swin-T backbone.

The results show that encoder fusion significantly improves model performance on both COCO and LVIS datasets. The results from comparing model #1 with the baseline model #0 validate this observation. Other techniques, such as language-guided query selection, text cross-attention, and sub-sentence text prompt, also contribute positively to the LVIS performance, yielding significant gains of +3.0 AP, +1.8 AP, and +0.5 AP, respectively. Additionally, these methods enhance the COCO zero-shot performance, further underscoring their effectiveness. However, we observed that language-guided query selection and sub-sentence text prompt had minimal impact on the COCO fine-tune performance. This outcome is reasonable, given that these methods do not alter model parameters or add computational burdens. Text cross-attention, while introducing fewer parameters than encoder fusion, showed less performance improvement compared to encoder fusion (+0.6 vs. +0.8). This finding suggests that fine-tuning performance is predominantly influenced by the model’s parameters, indicating that scaling models is a promising direction for enhancing performance.

5 Conclusion

We have presented a Grounding DINO model in this paper. Grounding DINO extends DINO to open-set object detection, enabling it to detect arbitrary objects given texts as queries. We review open-set object detector designs and propose a tight fusion approach to better fusing cross-modality information. We propose a sub-sentence level representation to use detection data for text prompts in a more reasonable way. The results show the effectiveness of our model design and

#ID	Model	COCO minival		LVIS minival
		Zero-Shot	Fine-Tune	Zero-Shot
0	Grounding DINO (Full Model)	46.7	56.9	16.1
1	w/o encoder fusion	45.8	56.1	13.1
2	static query selection	46.3	56.6	13.6
3	w/o text cross-attention	46.1	56.3	14.3
4	word-level text prompt	46.4	56.6	15.6

Table 7: Ablations for our model. All models are trained on the O365 dataset with a Swin Transformer Tiny backbone.

fusion approach. Moreover, we extend open-set object detection to REC tasks and perform evaluation accordingly. We show that existing open-set detectors do not work well for REC data without fine-tuning. Hence we call extra attention to REC zero-shot performance in future studies.

Limitations: Despite the great performance on open-set object detection settings, Grounding DINO cannot be used for segmentation tasks like GLIPv2. Our training data is less than the largest GLIP model, which may limit our final performance. Moreover, we find that our model will produce false positive results in some cases, which may need more techniques or data to reduce the hallucination.

Social Impacts: The use of deep learning models, such as this one, exposes them to vulnerabilities through adversarial attacks. Additionally, the accuracy and correctness of the model’s outputs cannot be guaranteed. There is also the risk that the open-set detection capabilities of the model could be exploited for unlawful purposes.

Acknowledgement

We thank the authors of GLIP [24]: Liunian Harold Li, Pengchuan Zhang, and Haotian Zhang for their helpful discussions and instructions. We also thank Tiancheng Zhao, the author of OmDet [59], and Jianhua Han, the author of DetCLIP [51], for their response to their model details. We thank He Cao of The Hong Kong University of Science and Technology for his help on diffusion models.

References

1. Anderson, P., He, X., Buehler, C., Teney, D., Johnson, M., Gould, S., Zhang, L.: Bottom-up and top-down attention for image captioning and visual question answering. *computer vision and pattern recognition* (2017)
2. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *European Conference on Computer Vision*. pp. 213–229. Springer (2020)
3. Chen, K., Pang, J., Wang, J., Xiong, Y., Li, X., Sun, S., Feng, W., Liu, Z., Shi, J., Ouyang, W., et al.: Hybrid task cascade for instance segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4974–4983 (2019)

4. Chen, Q., Chen, X., Wang, J., Feng, H., Han, J., Ding, E., Zeng, G., Wang, J.: Group DETR: Fast detr training with group-wise one-to-many assignment (2022)
5. Dai, X., Chen, Y., Xiao, B., Chen, D., Liu, M., Yuan, L., Zhang, L.: Dynamic head: Unifying object detection heads with attentions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7373–7382 (2021)
6. Deng, J., Yang, Z., Chen, T., Zhou, W., Li, H.: Transvg: End-to-end visual grounding with transformers. arXiv: Computer Vision and Pattern Recognition (2021)
7. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
8. Dong, N., Zhang, Y., Ding, M., Lee, G.H.: Boosting long-tailed object detection via step-wise learning on smooth-tail data (2023), <https://arxiv.org/abs/2305.12833>
9. Du, Y., Fu, Z., Liu, Q., Wang, Y.: Visual grounding with transformers. (2021)
10. Gan, Z., Chen, Y.C., Li, L., Zhu, C., Cheng, Y., Liu, J.: Large-scale adversarial training for vision-and-language representation learning. neural information processing systems (2020)
11. Gao, P., Geng, S., Zhang, R., Ma, T., Fang, R., Zhang, Y., Li, H., Qiao, Y.: Clip-adapter: Better vision-language models with feature adapters. arXiv preprint arXiv:2110.04544 (2021)
12. Gao, P., Zheng, M., Wang, X., Dai, J., Li, H.: Fast convergence of DETR with spatially modulated co-attention. arXiv preprint arXiv:2101.07448 (2021)
13. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. Learning (2021)
14. Gupta, A., Dollar, P., Girshick, R.: Lvis: A dataset for large vocabulary instance segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5356–5364 (2019)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
16. Jia, D., Yuan, Y., He, H., Wu, X., Yu, H., Lin, W., Sun, L., Zhang, C., Hu, H.: DETRs with hybrid matching (2022)
17. Kamath, A., Singh, M., LeCun, Y., Synnaeve, G., Misra, I., Carion, N.: MDETR-modulated detection for end-to-end multi-modal understanding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1780–1790 (2021)
18. Krasin, I., Duerig, T., Alldrin, N., Ferrari, V., Abu-El-Haija, S., Kuznetsova, A., Rom, H., Uijlings, J., Popov, S., Veit, A., et al.: Openimages: A public dataset for large-scale multi-label and multi-class image classification. Dataset available from <https://github.com/openimages> **2**(3), 18 (2017)
19. Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., Bernstein, M.S., Fei-Fei, L.: Visual genome: Connecting language and vision using crowdsourced dense image annotations. International Journal of Computer Vision (2017)
20. Kuo, W., Bertsch, F., Li, W., Piergiovanni, A., Saffar, M., Angelova, A.: Findit: Generalized localization with natural language queries (2022)
21. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., Duerig, T., Ferrari, V.: The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. arXiv: Computer Vision and Pattern Recognition (2018)

22. Li, C., Liu, H., Li, L.H., Zhang, P., Aneja, J., Yang, J., Jin, P., Lee, Y.J., Hu, H., Liu, Z., Gao, J.: Elevater: A benchmark and toolkit for evaluating language-augmented visual models (2022)
23. 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)
24. Li, L.H., Zhang, P., Zhang, H., Yang, J., Li, C., Zhong, Y., Wang, L., Yuan, L., Zhang, L., Hwang, J.N., et al.: Grounded language-image pre-training. arXiv preprint arXiv:2112.03857 (2021)
25. Li, M., Sigal, L.: Referring transformer: A one-step approach to multi-task visual grounding. arXiv: Computer Vision and Pattern Recognition (2021)
26. Li, Y., Liu, H., Wu, Q., Mu, F., Yang, J., Gao, J., Li, C., Lee, Y.J.: GLIGEN: Open-set grounded text-to-image generation. CVPR (2023)
27. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
28. 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: European conference on computer vision. pp. 740–755. Springer (2014)
29. Liu, J., Wang, L., Yang, M.H.: Referring expression generation and comprehension via attributes. international conference on computer vision (2017)
30. 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. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=oMI9Pj0b9J1>
31. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. arXiv preprint arXiv:2103.14030 (2021)
32. Meng, D., Chen, X., Fan, Z., Zeng, G., Li, H., Yuan, Y., Sun, L., Wang, J.: Conditional DETR for fast training convergence. arXiv preprint arXiv:2108.06152 (2021)
33. Miao, P., Su, W., Wang, L., Fu, Y., Li, X.: Referring expression comprehension via cross-level multi-modal fusion. ArXiv [abs/2204.09957](https://arxiv.org/abs/2204.09957) (2022)
34. Minderer, M., Gritsenko, A., Stone, A., Neumann, M., Weissenborn, D., Dosovitskiy, A., Mahendran, A., Arnab, A., Dehghani, M., Shen, Z., Wang, X., Zhai, X., Kipf, T., Houlsby, N.: Simple open-vocabulary object detection with vision transformers (2022)
35. Ordóñez, V., Kulkarni, G., Berg, T.L.: Im2text: Describing images using 1 million captioned photographs. neural information processing systems (2011)
36. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In: Proceedings of the IEEE international conference on computer vision. pp. 2641–2649 (2015)
37. Plummer, B.A., Wang, L., Cervantes, C.M., Caicedo, J.C., Hockenmaier, J., Lazebnik, S.: Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. International Journal of Computer Vision (2015)
38. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. IEEE Transactions on Pattern Analysis and Machine Intelligence **39**(6), 1137–1149 (2017)

39. 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)
40. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2021)
41. Sennrich, R., Haddow, B., Birch, A.: Neural machine translation of rare words with subword units. meeting of the association for computational linguistics (2015)
42. 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: Proceedings of the IEEE international conference on computer vision. pp. 8430–8439 (2019)
43. Sharma, P., Ding, N., Goodman, S., Soricut, R.: Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. meeting of the association for computational linguistics (2018)
44. Shilong, L., Yaoyuan, L., Shijia, H., Feng, L., Hao, Z., Hang, S., Jun, Z., Lei, Z.: DQ-DETR: Dual query detection transformer for phrase extraction and grounding. In: Proceedings of the AAAI Conference on Artificial Intelligence (2023)
45. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. international conference on machine learning (2019)
46. Thomee, B., Shamma, D.A., Friedland, G., Elizalde, B., Ni, K., Poland, D.N., Borth, D., Li, L.J.: Yfcc100m: the new data in multimedia research. Communications of The ACM (2016)
47. Wang, Y., Zhang, X., Yang, T., Sun, J.: Anchor DETR: Query design for transformer-based detector. national conference on artificial intelligence (2021)
48. Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al.: Huggingface’s transformers: State-of-the-art natural language processing. arXiv preprint arXiv:1910.03771 (2019)
49. Xu, M., Zhang, Z., Hu, H., Wang, J., Wang, L., Wei, F., Bai, X., Liu, Z.: End-to-end semi-supervised object detection with soft teacher. arXiv preprint arXiv:2106.09018 (2021)
50. Yao, L., Han, J., Liang, X., Xu, D., Zhang, W., Li, Z., Xu, H.: DetCLIPv2: Scalable Open-Vocabulary Object Detection Pre-training via Word-Region Alignment (2023)
51. Yao, L., Han, J., Wen, Y., Liang, X., Xu, D., Zhang, W., Li, Z., Xu, C., Xu, H.: Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection (2022)
52. Yu, L., Lin, Z., Shen, X., Yang, J., Lu, X., Bansal, M., Berg, T.L.: Mattnet: Modular attention network for referring expression comprehension. computer vision and pattern recognition (2018)
53. Yuan, L., Chen, D., Chen, Y.L., Codella, N., Dai, X., Gao, J., Hu, H., Huang, X., Li, B., Li, C., Liu, C., Liu, M., Liu, Z., Lu, Y., Shi, Y., Wang, L., Wang, J., Xiao, B., Xiao, Z., Yang, J., Zeng, M., Zhou, L., Zhang, P.: Florence: A new foundation model for computer vision (2022)
54. Zang, Y., Li, W., Zhou, K., Huang, C., Loy, C.C.: Open-vocabulary DETR with conditional matching (2022)
55. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14393–14402 (2021)
56. 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 (2022)

57. Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L.H., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: GLIPv2: Unifying Localization and Vision-Language Understanding (2022)
58. Zhang, H., Zhang, P., Hu, X., Chen, Y.C., Li, L.H., Dai, X., Wang, L., Yuan, L., Hwang, J.N., Gao, J.: Glipv2: Unifying localization and vision-language understanding (2022)
59. Zhao, T., Liu, P., Lu, X., Lee, K.: Omdet: Language-aware object detection with large-scale vision-language multi-dataset pre-training (2022)
60. Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., Gao, J.: Regionclip: Region-based language-image pretraining (2022)
61. Zhou, X., Girdhar, R., Joulin, A., Krähenbühl, P., Misra, I.: Detecting twenty-thousand classes using image-level supervision. In: ECCV (2022)
62. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. arXiv preprint arXiv:1904.07850 (2019)
63. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable DETR: Deformable transformers for end-to-end object detection. In: ICLR 2021: The Ninth International Conference on Learning Representations (2021)