

CP-DETR: Concept Prompt Guide DETR Toward Stronger Universal Object Detection

Qibo Chen*, Weizhong Jin, Jianyue Ge, Mengdi Liu, Yuchao Yan, Jian Jiang, Li Yu, Xuanjiang Guo, Shuchang Li, Jianzhong Chen

China Mobile(Zhejiang) Research & Innovation Institute
chenqibo@zj.chinamobile.com

Abstract

Recent research on universal object detection aims to introduce language in a SoTA closed-set detector and then generalize the open-set concepts by constructing large-scale (text-region) datasets for training. However, these methods face two main challenges: (i) how to efficiently use the prior information in the prompts to genericise objects and (ii) how to reduce alignment bias in the downstream tasks, both leading to sub-optimal performance in some scenarios beyond pre-training. To address these challenges, we propose a strong universal detection foundation model called CP-DETR, which is competitive in almost all scenarios, with only one pre-training weight. Specifically, we design an efficient prompt visual hybrid encoder that enhances the information interaction between prompt and visual through scale-by-scale and multi-scale fusion modules. Then, the hybrid encoder is facilitated to fully utilize the prompted information by prompt multi-label loss and auxiliary detection head. In addition to text prompts, we have designed two practical concept prompt generation methods, visual prompt and optimized prompt, to extract abstract concepts through concrete visual examples and stably reduce alignment bias in downstream tasks. With these effective designs, CP-DETR demonstrates superior universal detection performance in a broad spectrum of scenarios. For example, our Swin-T backbone model achieves 47.6 zero-shot AP on LVIS, and the Swin-L backbone model achieves 32.2 zero-shot AP on ODinW35. Furthermore, our visual prompt generation method achieves 68.4 AP on COCO val by interactive detection, and the optimized prompt achieves 73.1 fully-shot AP on ODinW13.

Introduction

Universal object detection aims to detect objects of any category in any scene with one model weight. The trend in research is to incorporate language modality, where textual descriptions of objects are encoded as text prompt vectors through language model (Devlin et al. 2018; Radford et al. 2021), and the classification results are represented by the similarity between the vectors and the image regions. This flexible conceptual representation allows different object detection data to be trained jointly, aligning textual descriptions with visual representations. Ultimately, in downstream

tasks, universal object detection with zero-shot is achieved by modifying the textual descriptions of objects.

While using text prompts has been primarily favored in universal detection, they suffer from sub-optimal performance in downstream applications, where universal detectors fail to compete with specialist models in many scenarios and categories outside of pre-training. A significant factor is the matching deficiency, where the detector produces mismatched results with the text description. This deficiency arises from alignment mistakes between language and visual representations in pre-training, and there are both objective and subjective aspects to this bias. Objectively, text descriptions follow a long-tailed pattern and different descriptions can refer to the same image region, so it is impractical to align all the texts and image regions accurately during pre-training. Subjectively, it is difficult for users to accurately describe complex objects, such as specific mechanical devices, through language. Most works (Kamath et al. 2021; Minderer et al. 2022; Minderer, Gritsenko, and Houlsby 2023; Yao et al. 2024; Wu et al. 2024) have been devoted to constructing larger pre-train datasets to address the alignment problems, but this requires significant costs.

Another factor is the paradigm of utilizing prompt information. The work (Li et al. 2022b) has shown that the early fusion paradigm performs significantly better than the late fusion paradigm after eliminating alignment bias through prompt tuning in the downstream tasks. Late fusion paradigms (Li et al. 2019) only use prompt vectors in the classification part, the location dependent on pre-training data distributions, which is poor in utilizing prompt information. In contrast, the early fusion paradigm (Liu et al. 2023) has an additional cross-modal fusion phase. It is easy to observe that the success of the early fusion paradigm lies in the cross-modal information interaction through fusion, where visual features are updated based on prompt information, and both classification and localization can be generalized in downstream scenes through prompt information. Therefore, we believe that a key to improving the performance of universal detection lies in achieving effective cross-modal interaction between prompt and visual.

In this paper, our research is interested in constructing a strong universal detector that not only has superior zero-shot capability but also competes with specific models in all downstream tasks through a model weight. For this, we pro-

*Corresponding Author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

pose CP-DETR, a model based on the early fusion paradigm that not only supports text prompts but also introduces visual prompts and optimized prompts to address alignment biases beyond pre-training. Visual prompts avoid misalignment arising from subjective user description errors by providing visual examples to represent objects, e.g., by marking specific objects with boxes. An optimized prompt provides a more direct solution by prompt tuning through downstream data annotation to align regions without changing the pre-training weights. Interestingly, we note text prompts, visual prompts, and optimized prompts represent object concepts through high-dimensional vectors, so we use concept prompts to represent these vectors in a unified way and divide the whole model into two parts: detector and concept prompt generation.

The detector part determines the universal detection capability of the model, so we build the detector based on the SoTA DETR (Zhang et al. 2023) framework and exploit the prompting information through effective cross-modal interactions. For effective cross-modal interaction, we design an efficient prompt visual hybrid encoder that updates visual and concept prompts via progressive single-scale fusion (PSF) and multi-scale fusion gating (MFG), avoiding confusion due to semantic gaps between different levels of visual features. Due to DETR being a sparse detector framework, we added an auxiliary detection head and a prompt multi-label loss to facilitate the hybrid encoder to fully utilize different modal information in the interaction.

For the concept prompt generation part, CP-DETR supports both text prompts, visual prompts, and optimized prompts. With text prompts, we use sentence-level representation to reduce computational overhead and encode them via CLIP (Radford et al. 2021) encoder because of its better discriminability using larger-scale contrast learning. For visual prompt, we design a visual prompt encoder that encodes the bbox as a query and adaptively aggregates concept representations from multi-scale features output by the visual backbone. For optimized prompt, we design a super-class representation prompt tuning method to further improve the performance in downstream tasks by representing single categories through multiple vectors.

Through effective design, the CP-DETR demonstrates amazing universal detection capabilities, e.g., Using text prompt, it achieved a significant 32.2 zero-shot AP on the ODinW35 (Li et al. 2022a). In the visual prompt interactive evaluation, it achieved 68.4 AP on the COCO (Lin et al. 2014) val. Furthermore, using the optimized prompt method, it outperforms the previous SoTA model (Zhang et al. 2022) 5.1 average AP on ODinW13 (Li et al. 2022a) and can compete with full-model fine-tuned specialist models.

Related Work

Text Prompted Universal Detection

The recent work can be divided into early fusion and late fusion, depending on the degree of exploitation of the prompt. The late fusion-based method only utilizes the prompt information in the classification. ViLD (Gu et al. 2022), RegionCLIP (Zhong et al. 2022) focuses on transferring

knowledge from CLIP to detection. The OWL-ViT (Minderer et al. 2022; Minderer, Gritsenko, and Houlsby 2023) and DetCLIP series (Yao et al. 2022, 2023, 2024) tend to directly align language and image regions through pre-training, therefore scaling up the data to 10B and 50M levels by pseudo-labeling, respectively. The early fusion-based method considers the effect of the prompt on both classification and localization, using the prompt as a condition for image feature encoding. GLIP (Li et al. 2022b) fuses word-level text prompts with multi-scale image features through cross-attention and leverages grounding data to help learn aligned semantics. Grounding DINO (Liu et al. 2023) further proposes language-guided query selection and cross-modality decoder to achieve denser fusion. Then, APE (Shen et al. 2024) and GLEE (Wu et al. 2024) reduce the number of text prompts using sentence-level text encoding methods, significantly reducing the computational overhead of the fusion layer, and thus allowing more negative categories to be used during pre-training. However, previous work uses all visual features to interact with prompts, ignoring the semantic gap of features at different levels in the backbone. For this reason, we design a hybrid encoder to achieve efficient cross-modal interaction through progressive fusion from single to global scales.

Visual Prompt

Unlike text prompts, visual prompts use image information directly to refer to objects, avoiding misalignment due to incorrect descriptions. Since late fused detectors have a double-tower structure, work (Minderer et al. 2022; Zang et al. 2022) adopts raw image as a visual prompt and leverages image-text-aligned representation to transfer the concept to a visual prompt. MQ-Det (Xu et al. 2023) uses a mixed representation of visual prompts and text prompts. T-Rex2 (Jiang et al. 2024) uses visual instructions to achieve interactive detection, with input boxes and dots generating visual prompts to avoid context loss in cropped images.

Optimized Prompt

The optimized prompt is generated by prompt tuning, which has proved effective for alignment in the classification (Zhou et al. 2022). PromptDet (Feng et al. 2022) uses this prompt as the context of the text prompt to guide the classification foundation model to achieve text and region alignment. GLIP (Li et al. 2022b) aligns concepts in downstream tasks by using optimized prompts as offsets to text prompts, noting that deep cross-modal fusion is critical to improving the effectiveness of prompt tuning. Recent work (Chen et al. 2024) directly learning prompts avoids the dependence on text prompts and further improves performance. The specificity of prompt tuning is that the optimization object is the activation value, which only reduces the alignment bias in the downstream task without changing the model. Therefore, we believe that the evaluation metrics of the optimized prompt can better reflect detector universality.

Method

The overall architecture of the proposed CP-DETR is illustrated in figure 1, which consists of two parts: concept

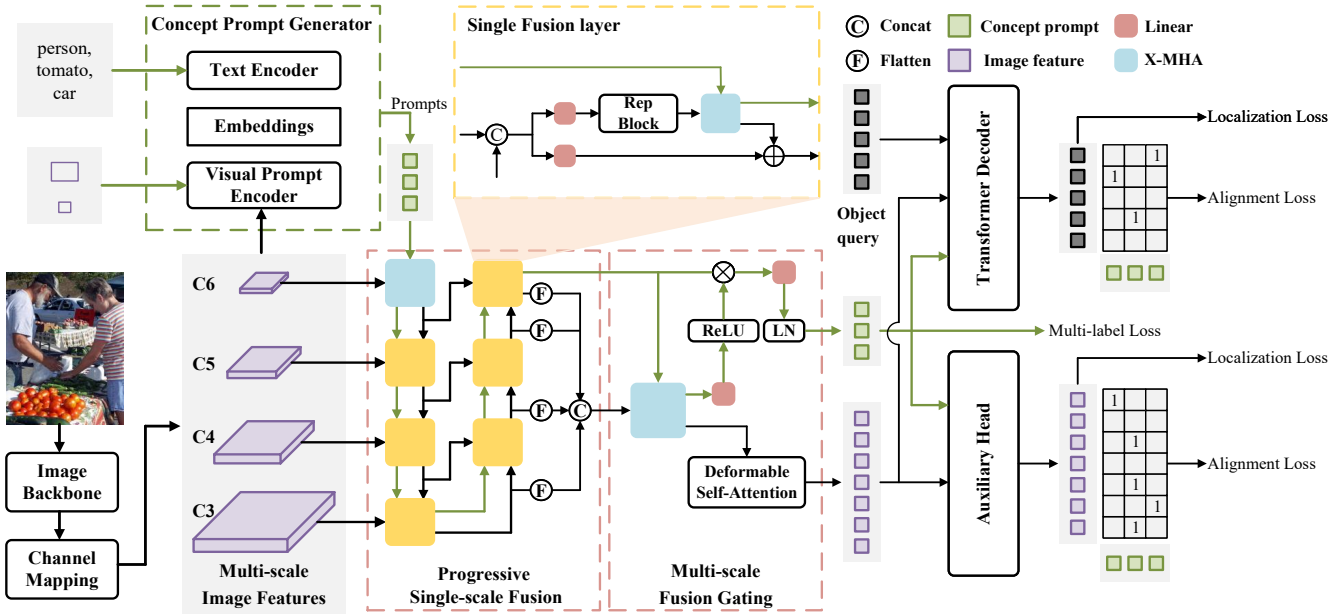


Figure 1: Overall architecture of CP-DETR. First, the concept prompt generator (shown in green dashed box) encodes textual descriptions, referring boxes, or annotations as concept prompts. Then, the detector encodes the image as multi-scale feature maps and performs a cross-modal fusion of concepts and images using the proposed hybrid encoder (shown as the red dashed box). Finally, the transformer decoder predicts results.

prompt generation and detection conditional on concept prompts. We use concept prompt generators to encode different object references (e.g., text, box coordinates, etc.) into uniform vector space, which represent the object concepts and serve as conditional input detectors. With different concept prompt generators, our model enables different workflows to handle alignment bias efficiently.

The detection part takes (prompts, image) pairs as input and outputs object boxes for the prompt’s corresponding concepts. For the image, the detector first obtains multi-scale image feature maps in 256 dimensions by image backbone and channel mapping. In this paper, we only use four scales: 1/8, 1/16, 1/32, and 1/64. Then, a prompt visual hybrid encoder, which contains progressive single-scale fusion and multi-scale fusion gating, will be used for the mutual fusion of prompt and image features. Following the previous work (Liu et al. 2023), after obtaining fused features, 900 object queries are initialized language-guided query selection and updated by the 6-layer cross-modality decoder. The training objectives for the transformer decoder are as follows:

$$L_{decoder} = L_{localization} + L_{alignment} \quad (1)$$

where $L_{localization}$ contains GIoU (Rezatofighi et al. 2019) loss and L1 loss, and $L_{alignment}$ is focal (Li et al. 2020) loss.

Due to the sparsity of object query, which could cause hybrid encoder sub-optimization, we introduce prompt multi-label classification loss and anchor-based auxiliary detection head in training as auxiliary supervision to facilitate cross-modal and cross-scale feature fusion. The auxiliary supervi-

sion part will be removed during inference.

Prompt Visual Hybrid Encoder

Previous early fusion-based work (Shen et al. 2024; Wu et al. 2024; Liu et al. 2023; Li et al. 2022b; Zhang et al. 2022) fused full-scale image feature maps and prompts simultaneously, which ignores the semantic gaps that exist between features at different scales. However, due to the lack of semantic concepts and feature duplication, it is inefficient to perform cross-modal interaction on low-level feature maps in the early stages of fusion. Therefore, we use a progressive single-scale fusion module that performs fusion scale-by-scale from high-level feature maps. In order to avoid multi-scale information loss during scale-by-scale fusion, we also designed multi-scale fusion gating to enhance the fusion of critical information.

Progressive Single-scale Fusion. The structure is illustrated in the left red dashed box of figure 1, which follows the top-down and bottom-up flow paths in (Zhao et al. 2024b; Liu et al. 2018). The deepest $C6 \in \mathcal{R}^{H/64 \times W/64 \times D}$ feature map has richer semantic concepts that help initially establish the connection between prompt and visual. Therefore, we first use a cross-modality multi-head attention (Li et al. 2022b)(X-MHA) to fuse $C6$ and prompt P by:

$$C6^{t=1}, P^{l+1} = X\text{-MHA}(C6^{t=0}, P^l) \quad (2)$$

Where l denotes the number of prompt fusions, $t \in (0, 1, 2)$ denotes the stage, and 0,1,2 denotes no fusion, top-down fusion, and bottom-up fusion, respectively.

Then, during top-down and bottom-up, we design a single fusion layer, as shown in the yellow dashed box of figure 1,

with two neighboring scales of image features and prompts as inputs. Specifically, neighboring feature maps are concatenated in the channel to obtain the hybrid feature C_{ij} , and the channels are adjusted through the linear layer and block (Ding et al. 2021) to achieve cross-scale and implicit cross-modal information fusion simultaneously. Then, using X-MHA to direct cross-modal fusion, obtains the updated prompt P^{l+1} and image features ΔC . Finally, the image features C_j^t of j scale at stage t are output by element-wise summation, which fuses ΔC with C_{ij} after a linear layer. The formula is as follows:

$$\begin{aligned} C_{ij} &= \text{concat}(\text{resize}(C_i^t), C_j^{t-1}) \\ P^{l+1}, \Delta C &= X\text{-MHA}(\text{Block}(\text{Linear}(C_{ij})), P^l) \quad (3) \\ C_j^t &= \Delta C + \text{Linear}(C_{ij}) \end{aligned}$$

Multi-scales Fusion Gating. To avoid information loss due to scale-by-scale fusion processes, we propose to interact simultaneously at multi-scale feature maps. The four-scale feature maps are flattened and then concatenated in the spatial dimension to form the full-scale feature C_{all} . The fusion process of C_{all} and prompt P^l from PSF is as follows:

$$\begin{aligned} P^{l+1}, C'_{all} &= X\text{-MHA}(C_{all}, P^l) \\ P_{end} &= \text{LN}(\text{Linear}(\text{ReLU}(\text{Linear}(P^{l+1}) * P^l))) \quad (4) \\ C''_{all} &= \text{DeformAttn}(C'_{all}) \end{aligned}$$

Where DeformAttn is deformable self-attention (Zhu et al. 2021), LN is Layernorm, P_{end} denotes final concept prompts after full-scales information gating through dot product, C''_{all} denotes the image partial output of the hybrid encoder after full-scale image feature interaction by deformable self-attention.

Auxiliary Supervision

In the DETR architecture of detector training, both classification and location losses are implemented on the object queries. However, due to the number of object query much smaller than the image features and using a one-to-one set matching scheme of label assignment, encoder output features get sparse supervision signals from the transformer decoder. We argue that these sparse supervision signals will reduce the learning efficiency of cross-scale and cross-modal interactions in the hybrid encoder, leading to sub-optimal results. Therefore, we introduce the auxiliary detection head and prompt multi-label loss to apply additional supervision to image features and conceptual prompts, respectively, which will facilitate fusion learning in the hybrid encoder.

Auxiliary Detection Head. We choose an anchor-based detector head (Zhang et al. 2020) to facilitate training, which was shown effective in closed-set detection (Zong, Song, and Liu 2023). The auxiliary head employs one-to-many label assignment and computes losses by anchors whose number is normal to image features, thus applying denser and direct supervision signals to image features. We use a contrastive layer to replace the classification layer in the closed-set detector header and represent the category scores by the similarity s_{mn} of prompt and image features as follows:

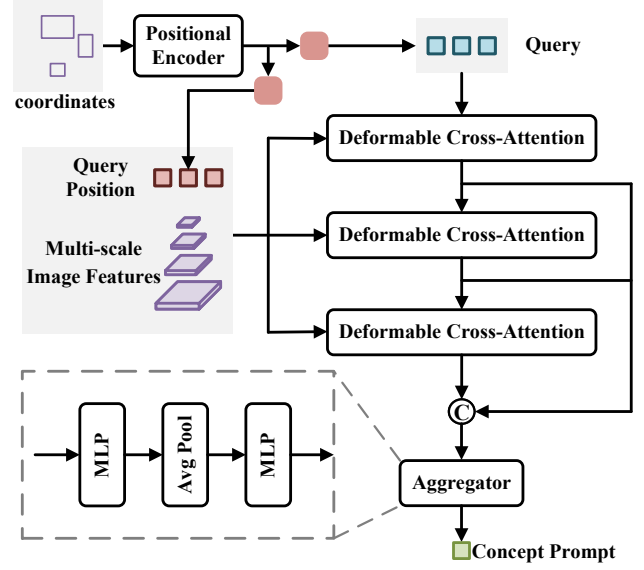


Figure 2: The overall architecture of the visual prompt encoder. Coordinates of 2D boxes are encoded as query and query position vectors, and the concept prompt is aggregated from image features via three layers of deformable cross-attention.

$$s_{mn} = \frac{a^m \times \text{Linear}(P_{end}^n)}{\sqrt{d}} + \text{bias} \quad (5)$$

Where d is the number of feature channels, bias is a learnable constant, a^m denotes the image feature corresponding to the m -th anchor, and P_{end}^n denotes the n -th concept vector. With this simple modification, the closed-set detector head is converted to open-set form and thus can be used for auxiliary supervision in pre-training with class uncertainty. The training objectives are as follows:

$$L_{aux.head} = L_{class} + L_{centerness} + L_{iou} \quad (6)$$

where L_{class} is focal loss, $L_{centerness}$ is binary cross entropy loss, and L_{iou} is GIoU loss.

Prompt Multi-label Loss. In open-set pre-training, there are both positive prompts and a large number of negative prompts in each (image, prompts) pair, and the negative prompts don't have a corresponding object in the image. Therefore, we could count the positivity and negativity of the prompts during the training process and automatically generate a multi-label annotation of g . The concept prompts output from the hybrid encoder are mapped to 1-dimensional through a MLP layer, and the loss is computed as follows:

$$L_{prompt} = \text{binary_cross_entropy}(MLP(P_{end}), g) \quad (7)$$

By applying a multi-label classification loss on a single modality, the concept prompts need learning to leverage the image information during the fusion process, thus rejecting the negative concept and retaining the positive prompts, making the fused concept prompts more discriminative.

Concept Prompt Generator

Text prompts successfully unify the training of different datasets and achieve zero-shot detection through a unified semantic space. However, due to alignment bias, the detector is prone to associate with wrong objects when meeting long-tailed or inaccurately described text in downstream tasks. For ODinW (Li et al. 2022a), we observe that the performance of all existing universal models with zero-shot significantly lags behind the closed-set trained models. Therefore, in order to reduce the impact of alignment bias on models in downstream tasks, CP-DETR also introduces two prompt generation methods, namely visual prompts and optimized prompts, to fully stimulate the universal detection capability of pre-trained models.

Text Prompt. We select the pre-trained CLIP text encoder to extract text features and use average pooling to aggregate token-level text features into sentence-level concept prompt. Only text prompts are used in CP-DETR pre-training, as this strategy was demonstrated efficiently in previous work (Li et al. 2022b). In order to reduce the detection hallucination, which is predicting objects that are not present in the input image, we randomly sampled 80 categories or descriptions from the text dictionary as negative samples in the pre-training. Unlike object detection datasets, Grounding and Referring Expression Comprehension (REC) datasets lack a unified category dictionary, so we construct a text dictionary online via a memory bank during training. Then, in pre-training, the overall training objective is a linear combination of $L_{decoder}$, $L_{aux.head}$, and L_{prompt} .

Visual Prompt. Figure 2 shows encoder structure, where N normalized box coordinates are encoded by sine-cosine position encoder to obtain vector $r \in \mathcal{R}^{N \times 128}$, which respectively through two linear layers generates query embedding q and query position embedding q_{pos} . Then, concept information is extracted from the image features by cross-attention, with q_{pos} limiting the extraction range to ensure information is relevant to box content. We use three layers of attention and concatenate the output query of each layer by channel, and finally generate the concept prompt for the corresponding category through an aggregator, which is shown in the dashed box in figure 2. When training the encoder, we freeze the pre-training weights and use the box sampling method of work (Jiang et al. 2024). Since we train the encoder after pre-training, where text prompts can be considered as concept prompt ground truths on the pre-training data, in addition to $L_{decoder}$, we also use MSE loss for direct supervision, with the overall training objective as follows:

$$L_{visual.prompt} = \frac{1}{K} \sum_{i=0}^K (P_v^i - P_t^i)^2 + L_{decoder} \quad (8)$$

where K is the number of positive categories, P_v denotes the concept prompt obtained by a visual prompt encoder, and P_t denotes the concept prompt obtained by the text encoder.

Optimized Prompt. We freeze all model parameters and initiate concept prompts with learnable embedding layers,

which are fine-tuned to get aligned concept prompts. In addition, we propose the super-class representation considering the case where different classes may be labeled as the same class in downstream scenarios. Specifically, class I corresponds to M prompts, and the correspondence is saved through a mapping table. Finally, the maximum similarity value was extracted from the M prompts as the classification score. Since hybrid encoder optimization is not required, the training objective contains only $L_{decoder}$.

Experiments

Training Datasets

We use multiple datasets with region-text annotations from different sources for joint training. For the object level, we use publicly available detection datasets, which contain Objects365 (Shao et al. 2019) (O365), OpenImages(Kuznetsova et al. 2020) (OI), V3Det (Wang et al. 2023), LVIS (Gupta, Dollar, and Girshick 2019) and COCO (Lin et al. 2014) datasets. For grounding or REC data, we used the GoldG (Kamath et al. 2021), RefCOCO+/g (Yu et al. 2016; Mao et al. 2016), Visual Genome (Krishna et al. 2017) (VG) and PhraseCut (Wu et al. 2020) datasets, with a memory bank set length of 1000 in pre-training. where GoldG, RefCOCO+/g, we used the cleaned labels from GLIP (Li et al. 2022b) and we combined RefCOCO+/g into RefC by removing duplicate samples. For GoldG, PhraseCut, and VG, where object phrases are treated as categories. For RefC, we treat the entire description as a category. It is worth noting that the training labels we use all come from publicly available datasets and do not scale up the data by pseudo-labeling image-text pair data as most work (Yao et al. 2024; Wu et al. 2024) does.

Implementation Details

In our experiments, we developed two model variants, CP-DETR-T and CP-DETR-L, by using Swin-Tiny and Swin-Large (Liu et al. 2021) as image backbone, respectively. We used CLIP-L (Fang et al. 2024) as the text encoder in all variants and only fine-tuned it during pre-training. For CP-DETR-T, we use O365, V3Det, and GoldG for pre-training with a total training epoch of 30. For CP-DETR-L, we train $1M$ iterations using all training datasets. In all experiments, we use AdamW as the optimizer with weight decay set to $1e-4$ and set a minibatch to 32 on 8 A100 40GB GPUs. In pre-training, the learning rate was set to $1e-5$ for the text encoder and image backbone and $1e-4$ for the rest of the modules, and a decay of 0.1 was applied at 80% and 90% of the total training steps. In visual prompt training, the O365, V3Det, GoldG, and OI datasets are used, the learning rate of the visual prompt encoder is set to $1e-4$, and the training is performed for $0.5M$ iterations. In the optimized prompt, the learning rate of the embedding layer is set to $5e-2$, the total number of training epochs is 24, and a decay of 0.1 is applied at 80% of the total training steps.

Evaluation Benchmark

We evaluated the universal detection ability on the COCO, LVIS, ODinW (Li et al. 2022a) and RefCOCO+/g bench-

Method	Backbone	COCO		LVIS		RefC	ODinW35
		val	test-dev	minival	val	refcoco/+g	test
GLIP-T (Li et al. 2022b)	Swin-T	46.3	-	26.0	17.2	50.4/49.5/66.1	19.6
Grounding-DINO-T (Liu et al. 2023)	Swin-T	48.4	-	27.4	20.1	50.8/51.6/60.4	22.3
YOLO-World-L (Cheng et al. 2024)	YOLOv8-L	45.1	-	35.4	-	-	-
DetCLIPv3-T (Yao et al. 2024)	Swin-T	47.2	-	47.0	38.9	-	-
T-Rex2-T (Jiang et al. 2024)	Swin-T	45.8	-	42.8	34.8	-	18.0
CP-DETR-T	Swin-T	52.0	52.2	47.6	39.9	43.7/42.2/52.6	27.3
GLIPv2-H (Zhang et al. 2022)	Swin-H	-	60.6	59.8	-	-	-
Grounding-DINO-L (Liu et al. 2023)	Swin-L	60.7	-	33.9	-	90.6/82.8/86.1	26.1
OmDet-Turbo-B (Zhao et al. 2024a)	ConvNeXt-B	53.4	-	34.7	-	-	30.1
T-Rex2-L (Jiang et al. 2024)	Swin-L	52.2	-	54.9	45.8	-	22.0
OWL-ST (Minderer, Gritsenko, and Hounsby 2023)	CLIP L/14	-	-	40.9	35.2	-	-
UNINEXT-H (Yan et al. 2023)	ViT-H	60.6	-	18.3	14.0	92.6/85.2/88.7	-
DetCLIPv2-L (Yao et al. 2023)	Swin-L	-	-	44.7	36.6	-	-
DetCLIPv3-L (Yao et al. 2024)	Swin-L	48.5	-	48.8	41.4	-	-
GLEE-Pro (Wu et al. 2024)	ViT-L	62.0	62.3	-	55.7	91.0/82.6/86.4	-
APE(D) (Shen et al. 2024)	ViT-L	58.3	-	64.7	59.6	84.6/76.4/80.0	28.8
CP-DETR-L	Swin-L	62.8	62.7	65.9	60.3	90.7/81.4/85.6	32.2

Table 1: Comparison with state-of-the-art universal models on multiple datasets through text prompts. Black numbers indicate zero-shot. Gray numbers indicate that the model pre-training contains the training parts of this dataset.

Method	Tune	PascalVOC	AerialDrone	Aquarium	Rabbits	EgoHands	Mushrooms	Packages	Raccoon	Shellfish	Vehicles	Pistols	Pothole	Thermal	Average
		GLEE-Pro (Wu et al. 2024)	full	72.6	36.5	58.1	80.5	74.1	92.0	67.0	76.5	66.4	70.5	66.4	55.7
GLIP-L (Li et al. 2022b)	full	69.6	32.6	56.6	76.4	79.4	88.1	67.1	69.4	65.8	71.6	75.7	60.3	83.1	68.9
GLIPv2-H (Zhang et al. 2022)	full	74.4	36.3	58.7	77.1	79.3	88.1	74.3	73.1	70.0	72.2	72.5	58.3	81.4	70.4
OmDet-B (Zhao et al. 2022)	full	71.2	27.5	52.7	76.5	77.4	93.6	73.7	74.3	57.7	64.5	74.2	56.9	83.3	68.0
DetCLIPv2-L (Yao et al. 2023)	full	74.4	44.1	54.7	80.9	79.9	90	74.1	69.4	61.2	68.1	80.3	57.1	81.1	70.4
DetCLIPv3-L (Yao et al. 2024)	full	76.4	51.2	57.5	79.9	80.2	90.4	75.1	70.9	63.6	69.8	82.7	56.2	83.8	72.1
GLIP-L (Li et al. 2022b)	prompt	72.9	23.0	51.8	72.0	75.8	88.1	75.2	69.5	73.6	72.1	73.7	53.5	81.4	67.9
GLIPv2-H (Zhang et al. 2022)	prompt	71.2	31.1	57.1	75.0	79.8	88.1	68.6	68.3	59.6	70.9	73.6	61.4	78.6	68.0
Grounding-DINO-T (Chen et al. 2024)	prompt	71.7	34.2	53.0	75.8	73.4	88.1	75.6	74.3	58.7	68.0	73.6	52.3	81.5	67.7
CP-DETR-T	prompt	74.2	37.7	54.4	78.4	75.5	88.1	72.0	72.8	61.0	72.9	75.9	54.4	82.2	69.2
CP-DETR-L	prompt	80.5	47.9	60.3	77.5	79.0	90.4	76.4	77.0	68.9	73.4	81.5	55.9	81.5	73.1

Table 2: Comparison with state-of-the-art universal models on multiple datasets through fine-tuning. A tune of type full indicates fine-tuning the full model. A tune of type prompt indicates optimizing prompt only.

Method	Backbone	COCO-val	LVIS-minival	ODinW35
T-Rex2-T	Swin-T	56.6	59.3	37.7
T-Rex2-L	Swin-L	58.5	62.5	39.7
CP-DETR-T	Swin-T	61.8	64.1	41.0
CP-DETR-L	Swin-L	68.4	71.6	50.6

Table 3: Comparison with universal models on multiple datasets through interactive object detection.

marks. ODinW contains 35 real-world scenarios that can reflect the model’s universality in downstream tasks. For COCO, LVIS, and ODinW, the AP is an evaluation metric. Following work (Liu et al. 2023), we also used RefCO-CO/+g to evaluate the ability of the model to understand complex textual descriptions with the $P@0.5$ metric.

Comparison with Universal Detectors

By switching among the three concept prompt generation methods, we demonstrate the universality and effectiveness of CP-DETR as an object detection model, both in the pre-training domain and downstream scenarios, while ensuring state-of-the-art performance. In all evaluations, CP-DETR only uses one weight.

Text Prompt Direct Evaluation. In this evaluation, we use all category names or description sentences of the benchmark as text prompt inputs, consistent with previous work (Shen et al. 2024) settings. Depending on whether the benchmark’s training set is used in pre-training, the text prompt-based evaluation can be categorized into zero-shot and full-shot. We primarily use CP-DETR-T to evaluate the effectiveness of our method on zero-shot. As shown in table 1, CP-DETR-T outperforms all similarly sized previous models in COCO and LVIS benchmarking, with $+3.6AP$

and +20.2*AP* compared to baseline Grounding DINO. The method closest to ours in terms of zero-shot performance is DetCLIPv3-T, which not only uses 1.61*M* of O365, V3Det, and GoldG as we do, but also an extra 50*M* of private data GranuCap50M, which indicates that our method is sufficiently effective in terms of concept generalization. CP-DETR-T has limitations on RefC, which we believe are due to the pre-training containing only object phrases and lacking the descriptive sentences required in the RefC evaluation. On the ODinW35 benchmark, we observed that several datasets showed significant quality issues in terms of annotated category names, so we followed the APE (Shen et al. 2024) evaluation setup, and our CP-DETR-L set a new zero-shot record with an average of 32.2*AP* across 35 datasets.

A universal model should have concept generalization capabilities and perform well in scenarios that have already been seen in pre-training. Due to CP-DETR-L’s pre-training data containing COCO, LVIS, and RefC, we use it for full-shot comparisons with other state-of-the-art universal models. As shown at the bottom of table 1, the CP-DETR-L simultaneously achieves state-of-the-art performance or competitive performance in all object detection benchmarks, with +2.1*AP* in COCO-val compared to baseline (Liu et al. 2023). On the RefC benchmark, CP-DETR achieved comparable results to Grounding DINO-L, showing that the sentence feature as a concept prompt is sufficient to represent complex textual descriptions. Notably for the COCO and LVIS parts of the evaluation, the state-of-the-art APE (Shen et al. 2024) and GLEE (Wu et al. 2024) performed well on only one of them, even though they used a larger backbone and stronger large-scale jittering data augmentation methods. And CP-DETR performs well on both benchmarks, proving that our method remembers and distinguishes all seen concepts well.

Fine-tuning Evaluation. Table 2 shows the comparison results with the state-of-the-art universal detection models on 13 subsets in ODinW, which are fine-tuned using prompt or full model. Optimized prompts reduce alignment bias by adjusting concept prompts and can truly reflect the universality of the detector. The significant performance advantage of our approach in this setting, along with the +5.1*AP* compared to GLIPv2-H (Zhang et al. 2022) in terms of average metrics, demonstrates the strong generalization of CP-DETR in downstream scenarios, and we believe that this advantage stems from our design, which better facilitates the use of prompt information. Even compared to the approach of applying full model fine-tuning, CP-DETR-L still achieved state-of-the-art or competitive performance in 13 subsets with only optimized prompts, and set a new record of 73.1*AP* on average. This phenomenon indicates that CP-DETR can achieve competitive performance with a specific model in downstream scenarios by using a pre-trained weight, greatly enhancing the application value of the universal model in the real world.

Visual Prompt Interactive Evaluation. Since the concept prompt generation of visual prompt requires boxes as input, we use interactive evaluation, unlike the interactive process of T-Rex2 (Jiang et al. 2024), we avoid introduc-

Row	Model Set	LVIS minival	ODinW13
		Zero-shot	Full-shot
0	CP-DETR(base model)	44.3	64.0
1	replaced by DINO encoder	42.2	58.5
2	w/o MFG	42.8	62.3
3	add prompt multi-label loss	44.8	64.2
4	add row3 and auxiliary head	44.7	64.9
5	add row3, row4 and super-class	44.7	67.0

Table 4: Ablations for our model with a Swin-T backbone. The full shot is achieved by the optimized prompt.

ing category priors. For the test image with M total dataset categories and N positive categories, we randomly chose a GT box as the visual prompt input for the positive category and used text prompts for the remaining $M - N$ negative categories. As shown in table 3, our method significantly outperforms previous work (Jiang et al. 2024) in all benchmarks. CP-DETR is the first to implement interactive detection based on visual prompts in the early fusion paradigm, which is more effective in exploiting prompt information than the late fusion paradigm (Jiang et al. 2024). Comparing table 1 and 3, it can be observed that visual prompts outperform text prompts, with +18.4*AP* on ODinW35, indicating that visual prompts can reduce alignment bias and have a strong application in interactive scenarios such as auxiliary labelling.

Ablation

In this section, we conducted ablation experiments. The different variants models all use the Swin-T backbone and are trained on O365, V3Det, and GoldG for 12 epochs. In order to show the impact of various components on the universality of the detector, we not only performed a zero-shot evaluation on LVIS but also employed a full-shot optimized prompt on ODinW13.

Table 4 demonstrates the effectiveness of the different designs, where row 0 is the CP-DETR base model without the inclusion of auxiliary supervision and super-class representation. The prompt visual hybrid encoder was ablated in rows 1 and 2; the results show that the hybrid fusion approach of PSF and MSG reduces the difficulty of alignment and contributes to the zero-shot generalization of concepts, and the metric is improved by 0.6*AP* and 1.5*AP* on LVIS, respectively. The hybrid encoder is the most important improvement, with the ODinW13 full-shot metric upgraded from 58.5*AP* in row 1 to 64.0*AP*. This improvement reveals the importance of effective cross-modal fusion for universal location, encouraging the model to decode object boxes based on information in the prompt. Rows 3 and 4 experiments show that auxiliary supervision facilitates the hybrid encoder in learning cross-modal knowledge during the training phase and modestly improves the zero-shot and full-shot metrics performance. Row 5 further adds super-class in prompt fine-tuning, with the number set to 10 in the experiment, i.e., 10 prompt vectors represent a category. This design can effectively handle situations where different objects in downstream scenes are represented in the same category,

thus further improving the ODinW13 metric to 67.0AP. See the Technical Appendix for more results.

Conclusion

In this paper, we propose a universal detector CP-DETR that enables prompt-conditional detection through efficient prompt visual fusion. We focus on downstream applications and achieve SoTA zero-shot performance through text prompts. Furthermore, with visual prompt and optimized prompt, CP-DETR with only one weight can compete with the full model fine-tuned methods in downstream scenarios, demonstrating its superior universality.

References

Chen, Q.; Jin, W.; Li, S.; Liu, M.; Yu, L.; Jiang, J.; and Wang, X. 2024. Exploration of visual prompt in Grounded pre-trained open-set detection. In *ICASSP*, 6115–6119. IEEE.

Cheng, T.; Song, L.; Ge, Y.; Liu, W.; Wang, X.; and Shan, Y. 2024. Yolo-world: Real-time open-vocabulary object detection. In *CVPR*, 16901–16911.

Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv:1810.04805.

Ding, X.; Zhang, X.; Ma, N.; Han, J.; Ding, G.; and Sun, J. 2021. Repvgg: Making vgg-style convnets great again. In *CVPR*, 13733–13742.

Fang, Y.; Sun, Q.; Wang, X.; Huang, T.; Wang, X.; and Cao, Y. 2024. Eva-02: A visual representation for neon genesis. *Image and Vision Computing*, 105171.

Feng, C.; Zhong, Y.; Jie, Z.; Chu, X.; Ren, H.; Wei, X.; Xie, W.; and Ma, L. 2022. Promptdet: Towards open-vocabulary detection using uncurated images. In *ECCV*, 701–717. Springer.

Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2022. Open-vocabulary Object Detection via Vision and Language Knowledge Distillation. In *ICLR*.

Gupta, A.; Dollar, P.; and Girshick, R. 2019. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, 5356–5364.

Jiang, Q.; Li, F.; Zeng, Z.; Ren, T.; Liu, S.; and Zhang, L. 2024. T-Rex2: Towards Generic Object Detection via Text-Visual Prompt Synergy. arXiv:2403.14610.

Kamath, A.; Singh, M.; LeCun, Y.; Synnaeve, G.; Misra, I.; and Carion, N. 2021. Mdetr-modulated detection for end-to-end multi-modal understanding. In *ICCV*, 1780–1790.

Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123: 32–73.

Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Mallocci, M.; Kolesnikov, A.; et al. 2020. 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.

Li, C.; Liu, H.; Li, L.; Zhang, P.; Aneja, J.; Yang, J.; Jin, P.; Hu, H.; Liu, Z.; Lee, Y. J.; and Gao, J. 2022a. ELEVATER: A Benchmark and Toolkit for Evaluating Language-Augmented Visual Models. In *NeurIPS*, 9287–9301.

Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022b. Grounded language-image pre-training. In *CVPR*, 10965–10975.

Li, X.; Wang, W.; Wu, L.; Chen, S.; Hu, X.; Li, J.; Tang, J.; and Yang, J. 2020. Generalized Focal Loss: Learning Qualified and Distributed Bounding Boxes for Dense Object Detection. In *NeurIPS*, 21002–21012.

Li, Z.; Yao, L.; Zhang, X.; Wang, X.; Kanhere, S.; and Zhang, H. 2019. Zero-shot object detection with textual descriptions. In *AAAI*, 8690–8697.

Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*, 740–755. Springer.

Liu, S.; Qi, L.; Qin, H.; Shi, J.; and Jia, J. 2018. Path aggregation network for instance segmentation. In *CVPR*, 8759–8768.

Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Li, C.; Yang, J.; Su, H.; Zhu, J.; et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. arXiv:2303.05499.

Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 10012–10022.

Mao, J.; Huang, J.; Toshev, A.; Camburu, O.; Yuille, A. L.; and Murphy, K. 2016. Generation and comprehension of unambiguous object descriptions. In *CVPR*, 11–20.

Minderer, M.; Gritsenko, A.; and Houlsby, N. 2023. Scaling Open-Vocabulary Object Detection. In *NeurIPS*, 72983–73007.

Minderer, M.; Gritsenko, A.; Stone, A.; Neumann, M.; Weissenborn, D.; Dosovitskiy, A.; Mahendran, A.; Arnab, A.; Dehghani, M.; Shen, Z.; et al. 2022. Simple open-vocabulary object detection. In *ECCV*, 728–755. Springer.

Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, 8748–8763. PMLR.

Rezatofighi, H.; Tsoi, N.; Gwak, J.; Sadeghian, A.; Reid, I.; and Savarese, S. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *CVPR*, 658–666.

Shao, S.; Li, Z.; Zhang, T.; Peng, C.; Yu, G.; Zhang, X.; Li, J.; and Sun, J. 2019. Objects365: A large-scale, high-quality dataset for object detection. In *ICCV*, 8430–8439.

Shen, Y.; Fu, C.; Chen, P.; Zhang, M.; Li, K.; Sun, X.; Wu, Y.; Lin, S.; and Ji, R. 2024. Aligning and prompting everything all at once for universal visual perception. In *CVPR*, 13193–13203.

- Wang, J.; Zhang, P.; Chu, T.; Cao, Y.; Zhou, Y.; Wu, T.; Wang, B.; He, C.; and Lin, D. 2023. V3det: Vast vocabulary visual detection dataset. In *ICCV*, 19844–19854.
- Wu, C.; Lin, Z.; Cohen, S.; Bui, T.; and Maji, S. 2020. Phrasecut: Language-based image segmentation in the wild. In *CVPR*, 10216–10225.
- Wu, J.; Jiang, Y.; Liu, Q.; Yuan, Z.; Bai, X.; and Bai, S. 2024. General object foundation model for images and videos at scale. In *CVPR*, 3783–3795.
- Xu, Y.; Zhang, M.; Fu, C.; Chen, P.; Yang, X.; Li, K.; and Xu, C. 2023. Multi-modal queried object detection in the wild. In *NeurIPS*, 4452–4469.
- Yan, B.; Jiang, Y.; Wu, J.; Wang, D.; Luo, P.; Yuan, Z.; and Lu, H. 2023. Universal instance perception as object discovery and retrieval. In *CVPR*, 15325–15336.
- Yao, L.; Han, J.; Liang, X.; Xu, D.; Zhang, W.; Li, Z.; and Xu, H. 2023. Detclipv2: Scalable open-vocabulary object detection pre-training via word-region alignment. In *CVPR*, 23497–23506.
- Yao, L.; Han, J.; Wen, Y.; Liang, X.; Xu, D.; Zhang, W.; Li, Z.; XU, C.; and Xu, H. 2022. DetCLIP: Dictionary-Enriched Visual-Concept Paralleled Pre-training for Open-world Detection. In *NeurIPS*, 9125–9138.
- Yao, L.; Pi, R.; Han, J.; Liang, X.; Xu, H.; Zhang, W.; Li, Z.; and Xu, D. 2024. DetCLIPv3: Towards Versatile Generative Open-vocabulary Object Detection. In *CVPR*, 27391–27401.
- Yu, L.; Poirson, P.; Yang, S.; Berg, A. C.; and Berg, T. L. 2016. Modeling context in referring expressions. In *ECCV*, 69–85. Springer.
- Zang, Y.; Li, W.; Zhou, K.; Huang, C.; and Loy, C. C. 2022. Open-vocabulary detr with conditional matching. In *ECCV*, 106–122. Springer.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L.; and Shum, H.-Y. 2023. DINO: DETR with Improved DeNoising Anchor Boxes for End-to-End Object Detection. In *ICLR*.
- Zhang, H.; Zhang, P.; Hu, X.; Chen, Y.-C.; Li, L.; Dai, X.; Wang, L.; Yuan, L.; Hwang, J.-N.; and Gao, J. 2022. GLIPv2: Unifying Localization and Vision-Language Understanding. In *NeurIPS*, 36067–36080.
- Zhang, S.; Chi, C.; Yao, Y.; Lei, Z.; and Li, S. Z. 2020. Bridging the gap between anchor-based and anchor-free detection via adaptive training sample selection. In *CVPR*, 9759–9768.
- Zhao, T.; Liu, P.; He, X.; Zhang, L.; and Lee, K. 2024a. Real-time Transformer-based Open-Vocabulary Detection with Efficient Fusion Head. arXiv:2403.06892.
- Zhao, T.; Liu, P.; Lu, X.; and Lee, K. 2022. Omdet: Language-aware object detection with large-scale vision-language multi-dataset pre-training. arXiv:2209.05946.
- Zhao, Y.; Lv, W.; Xu, S.; Wei, J.; Wang, G.; Dang, Q.; Liu, Y.; and Chen, J. 2024b. Dets beat yolos on real-time object detection. In *CVPR*, 16965–16974.
- Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; et al. 2022. Region-clip: Region-based language-image pretraining. In *CVPR*, 16793–16803.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2021. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *ICLR*.
- Zong, Z.; Song, G.; and Liu, Y. 2023. Dets with collaborative hybrid assignments training. In *ICCV*, 6748–6758.