

OvarNet: Towards Open-vocabulary Object Attribute Recognition

Keyan Chen^{1*}, Xiaolong Jiang^{2*}, Yao Hu², Xu Tang², Yan Gao², Jianqi Chen¹, Weidi Xie^{3,4†}
Beihang University¹, Xiaohongshu Inc²,
CMIC, Shanghai Jiao Tong University³, Shanghai AI Laboratory⁴

<https://kyanchen.github.io/OvarNet>

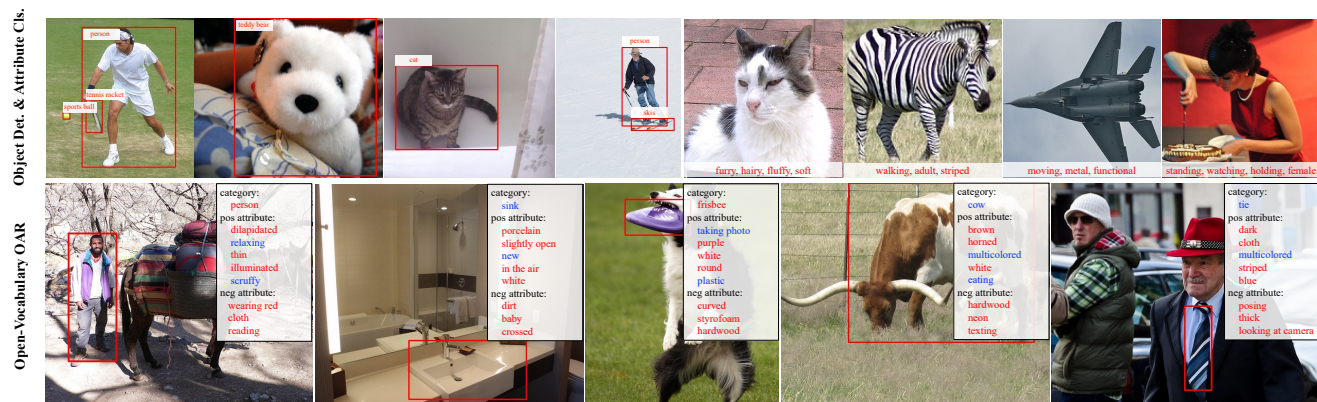


Figure 1. The first row depicts the tasks of object detection and attribute classification in a close-set setting, *i.e.*, train and test on the same vocabulary set. The second row gives qualitative results from our proposed OvarNet, which simultaneously localizes, categorizes, and characterizes arbitrary objects in an open-vocabulary scenario. We only show one object per image for ease of visualization, **red** denotes the base category/attribute *i.e.*, seen in the training set, while **blue** represents the novel category/attribute unseen in the training set.

Abstract

In this paper, we consider the problem of simultaneously detecting objects and inferring their visual attributes in an image, even for those with no manual annotations provided at the training stage, resembling an open-vocabulary scenario. To achieve this goal, we make the following contributions: (i) we start with a naïve two-stage approach for open-vocabulary object detection and attribute classification, termed CLIP-Attr. The candidate objects are first proposed with an offline RPN and later classified for semantic category and attributes; (ii) we combine all available datasets and train with a federated strategy to fine-tune the CLIP model, aligning the visual representation with attributes, additionally, we investigate the efficacy of leveraging freely available online image-caption pairs under weakly supervised learning; (iii) in pursuit of efficiency, we train a Faster-RCNN type model end-to-end with knowledge distillation, that performs class-agnostic object proposals and classification on semantic categories and attributes with classifiers generated from a text encoder; Finally, (iv) we conduct extensive experiments on VAW, MS-COCO, LSA, and OVAD datasets, and show that recog-

* Equal contribution. † Corresponding author.

nition of semantic category and attributes is complementary for visual scene understanding, *i.e.*, jointly training object detection and attributes prediction largely outperform existing approaches that treat the two tasks independently, demonstrating strong generalization ability to novel attributes and categories.

1. Introduction

Understanding the visual scene in terms of objects has been the main driving force for development in computer vision [46], for example, in object detection, the goal is to localise objects in an image and assign one of the predefined semantic labels to them, such as a ‘car’, ‘person’ or ‘bus’, despite tremendous success has been made by the community, such task definition has largely over-simplified our understanding of the visual world, as a visual object can often be characterised from many aspects other than semantic category, for example, a bus can be ‘yellow’ or ‘black’, a shirt can be ‘striped’ or ‘unpatterned’, learning attributes can thus complement category-level recognition, acquiring more comprehensive visual perception.

In the literature, numerous work has shown that understanding the objects’ attributes can greatly facilitate object recognition and detection, even with few or no examples

of visual objects [5, 16, 21, 36, 45], for example, Farhadi *et al.* proposed to shift the goal of object recognition from ‘naming’ to ‘description’, which allows naming familiar objects with attributes, but also to say something about unfamiliar objects (“hairy and four-legged”, not just “unknown”) [5]; Lampert *et al.* considered the open-set object recognition, that aims to recognise objects by human-specified high-level description, *e.g.*, arbitrary semantic attributes, like shape, color, or even geographic information, instead of training images [16]. However, the problem considered in these seminal work tends to be a simplification from today’s standard, for example, attribute classification are often trained and evaluated on object-centric images under the close-set scenario, *i.e.*, assuming the bounding boxes/segmentation masks are given [11, 25, 31], or sometimes even the object category are known as a prior [22, 25].

In this paper, we consider the task of simultaneously detecting objects and classifying the attributes in an open-vocabulary scenario, *i.e.*, the model is only trained on a set of base object categories and attributes, while it is required to generalise towards ones that are unseen at training time, as shown in Fig. 1. Generally speaking, we observe three major challenges: *First*, in the existing foundation models, *e.g.*, CLIP [27] and ALIGN [13], the representation learned from image-caption pairs tends to bias towards object category, rather than attributes, which makes it suffer from feature misalignment when used directly for attribute recognition. We experimentally validate this conjecture by showing a significant performance drop in attribute recognition, compared to category classification; *Second*, there is no ideal training dataset with three types of annotations, object bounding boxes, semantic categories, and attributes; as far as we know, only the COCO Attributes dataset [24] provides such a degree of annotations, but with a relatively limited vocabulary size (196 attributes, 29 categories); *Third*, training all three tasks under a unified framework is challenging and yet remains unexplored, *i.e.*, simultaneously localising (‘where’), classifying objects’ semantic categories and attributes (‘what’) under the open-vocabulary scenario.

To address the aforementioned issues, we start with a naïve architecture, termed as CLIP-Attr, which first proposes object candidates with an offline RPN [30], and then performs open-vocabulary object attribute recognition by comparing the similarity between the attribute word embedding and the visual embedding of the proposal. To better align the feature between attribute words and proposals, we introduce learnable prompt vectors with parent attributes on the textual encoder side and finetune the original CLIP model on a large corpus of the freely available image-caption datasets. To further improve the model efficiency, we present OvarNet, a unified framework that performs detection and attributes recognition at once, which is trained by leveraging datasets from both object detec-

tion and attribute prediction, as well as absorbing knowledge from CLIP-Attr to improve the performance and robustness of unseen attributes. As a result, our proposed OvarNet, being the first scalable pipeline, can simultaneously localize objects and infer their categories with visual attributes in an open-vocabulary scenario. Experimental results demonstrate that despite only employing weakly supervised image-caption pairs for distillation, OvarNet outperforms previous the state-of-the-art on VAW [25], MSCOCO [18], LSA [26] and OVAD [3] datasets, exhibiting strong generalization ability on novel attributes and categories.

2. Related Work

Attribute Prediction. Visual attribute aims to describe one object/scene from various aspects, for example, color, texture, shape, material, state, etc, allowing to represent object categories in a combinatorial manner. However, annotating attributes can be very time-consuming, early efforts only focus on specific domains such as fashion [39, 40], face [9, 41], animals [1, 34], posing severe limitations for real-world deployment. With the release of large-scale datasets including COCO Attributes [24], Visual Genome [15], and VAW [25], recent work considers building models for large-vocabulary attributes classification [25, 37]. Nonetheless, these methods only perform multi-class classification on pre-computed image patches, which not only fail to acquire object localization ability but also endure extra computation overhead due to redundant feature extraction passes. Additionally, other methods such as SCoNE [25] require object category as input to perform attribute prediction, leading to extra complexity in practice. In this work, we aim to build a unified framework that can jointly settle object localization, category prediction, and attribute prediction in an open-vocabulary scenario, relieving the aforementioned practical limitations.

Open-vocabulary Object Detection. Open-vocabulary object detection strives to detect all objects, including those that are unseen at the training stage. Existing approaches [2, 6, 8, 44] achieve open-vocabulary capability by replacing the detector’s classifier with object category word embedding from the pre-trained visual-language model, *e.g.*, CLIP, and perform category classification via embedding matching. In specific, OVR-CNN [38] proposes an efficient training approach with image-caption pairs that can be easily obtained from the website. ViLD [8] adopts distillation to infuse open-vocabulary knowledge into a two-stage detector, Detic [44] increases the size of detector’s vocabulary to twenty-thousand by exploiting the large dataset (ImageNet-21K) with image-level annotations. PromptDet [6] leverages the pre-trained CLIP [27] and aligns the detector’s visual embedding with text embedding

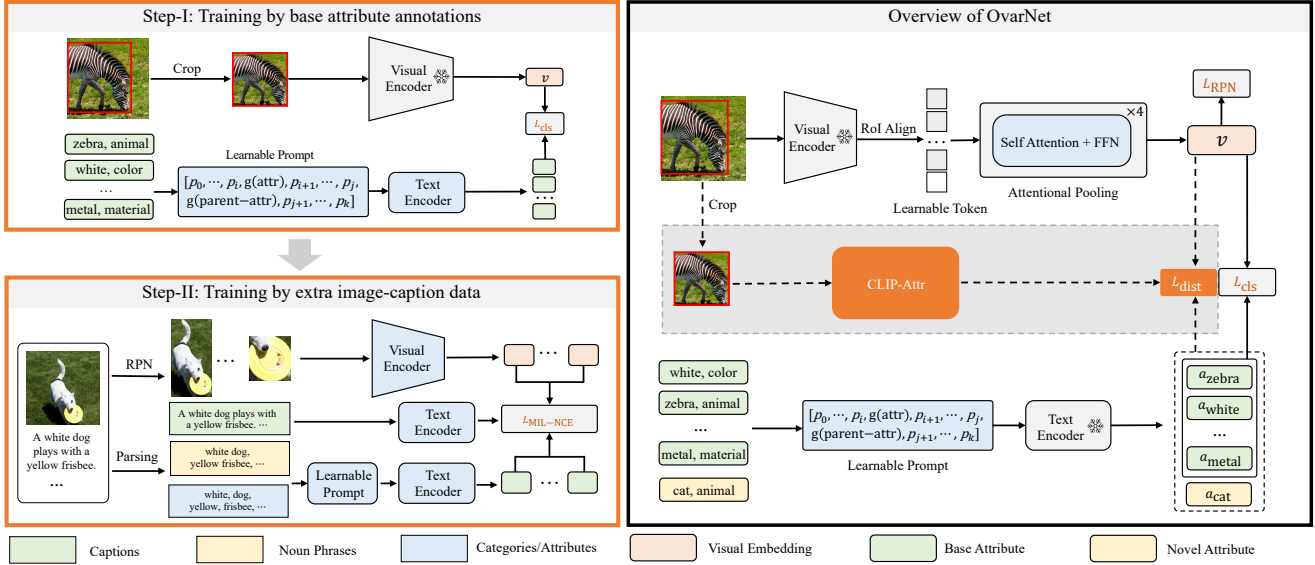


Figure 2. An overview of the proposed method. **Left:** the two-step training procedure for finetuning the pre-trained CLIP to get CLIP-Attr that better aligns the regional visual feature to attributes. **Step-I:** naïve federate training by base attribute annotations. **Step-II:** training by image-caption pairs. We first conduct RPN on the whole image to get box-level crops, parse the caption to get noun phrases, categories, and attributes, and then match these fine-grained concepts for weakly supervised training. **Right:** the proposed one-stage framework OvarNet. We inherit the CLIP-Attr for open-vocabulary object attribute recognition. Regional visual feature is learned from the attentional pooling of proposals; while attribute concept embedding is extracted from the text encoder. Solid lines declare the standard federated training regime. Dashed lines denote training by knowledge distillation with CLIP-Attr.

with learnable prompts. However, none of these models considers simultaneously inferring attributes for detected objects.

Zero-shot Learning. Zero-shot learning aims to extend the model’s capability towards recognising objects beyond those seen categories at training time [2, 8, 28, 29]. In the context of object detection, early zero-shot solutions rely on visual attributes to infer unseen categories [12, 17, 20, 36], aiming to represent category by attributes, such that it can generalize from seen to unseen category. Recent methods adopt vision-language feature alignment to achieve zero-shot learning, based on similarity computation between the visual feature and text concepts.

3. Methodology

In this section, we start by introducing the problem scenario (Sec. 3.1), followed by describing a naïve architecture for open-vocabulary attribute classification by steering a pre-trained CLIP model, dubbed CLIP-Attr (Sec. 3.2), and finally, we further distill the knowledge from CLIP-Attr into a more efficient two-stage detection architecture called OvarNet, which can perform detection and attribute prediction in a unified framework (Sec. 3.3).

3.1. Problem Scenario

Assuming we are given a training dataset, *i.e.*, $\mathcal{D}_{\text{train}} = \{(\mathcal{I}_1, y_1), \dots, (\mathcal{I}_N, y_N)\}$, where $\mathcal{I}_i \in \mathbb{R}^{H \times W \times 3}$ refers to

an image, and $y_i = \{b_i, c_i, a_i\}$ denotes its corresponding ground-truth annotations, with the coordinates for n object bounding boxes ($b_i \in \mathbb{R}^{n_i \times 4}$), their corresponding semantic categories ($c_i \in \mathbb{R}^{n_i \times \mathcal{C}_{\text{base}}}$), and a set of binary attributes for each object ($a_i \in \{0, 1\}^{n_i \times \mathcal{A}_{\text{base}}}$). Our goal is to train a model that can process any image from a test set ($\mathcal{I}_k \sim \mathcal{D}_{\text{test}}$), simultaneously localising the objects and inferring their semantic categories, and visual attributes:

$$\{\hat{b}_k, \hat{c}_k, \hat{a}_k\} = \Phi_{\text{CLS}} \circ \Phi_{\text{LOC}}(\mathcal{I}_k)$$

where the image is progressively processed by a class-agnostic object localization, and open-vocabulary attributes classification, to produce the $\hat{b}_k \in \mathbb{R}^{n_k \times 4}$, $\hat{c}_k \in \mathbb{R}^{n_k \times \mathcal{C}_{\text{test}}}$ and $\hat{a}_k \in \{0, 1\}^{n_k \times \mathcal{A}_{\text{test}}}$. Note that, at inference time, the objects may be of unseen/novel semantic categories or attributes, *i.e.*, $\mathcal{C}_{\text{test}} = \mathcal{C}_{\text{base}} \cup \mathcal{C}_{\text{novel}}$, $\mathcal{A}_{\text{test}} = \mathcal{A}_{\text{base}} \cup \mathcal{A}_{\text{novel}}$, thus the considered problem falls into open-vocabulary object attributes recognition. For simplicity, we will omit the subscript k while describing the proposed models. To avoid redundancies, we treat the category as a super-attribute for modeling our pipeline unless otherwise specified.

3.2. Two-stage Object Attribute Recognition

In this section, we describe a two-stage open-vocabulary attribute classification method, termed CLIP-Attr, that first uses a class-agnostic region proposal network (RPN) to generate object candidates, then verifies the candidates with

category and attributes using a finetuned CLIP:

$$\begin{aligned} \{\hat{b}_k\} &= \Phi_{\text{LOC}} = \Phi_{\text{crpn}}(\mathcal{I}) \\ \{\hat{c}_k, \hat{a}_k\} &= \Phi_{\text{CLS}} = \Phi_{\text{cls}} \circ \Phi_{\text{clip-v}} \circ \Phi_{\text{crop}}(\mathcal{I}, \{\hat{b}_k\}) \end{aligned}$$

where $\Phi_{\text{crpn}}(\mathcal{I})$ is a class-agnostic RPN, $\Phi_{\text{cls}}(\cdot)$ represents attributes classification, $\Phi_{\text{clip-v}}(\cdot)$ denotes the CLIP visual encoder, and $\Phi_{\text{crop}}(\cdot)$ is an operation that crops \hat{b}_k the box region from input image.

3.2.1 Object-centric Visual Encoding

Class-agnostic Region Proposal. To propose the candidate regions that potentially have objects situated, we employ a Faster-RCNN [30] based region proposal network that parametrises the anchor classification and bounding box regression in a class-agnostic manner, *i.e.*, $\Phi_{\text{crpn}}(\cdot)$ shares parameters for all categories. Inspired by the observation in [6, 14, 44], we train the proposal network only on base categories offline, and it shows sufficient generalization ability towards unseen categories.

RoI Visual Pooling. Given the pre-defined object boxes, we acquired the image crops ($\Phi_{\text{crop}}(\cdot)$) and feed them into the CLIP image encoder ($\Phi_{\text{clip-v}}(\cdot)$) to compute regional visual embeddings $\hat{v}_i \in \mathbb{R}^{1 \times D}$, i denotes the i th region.

3.2.2 Open-vocabulary Attributes Classification

Generating Attribute Embedding. To compute attribute embeddings, we employ the pre-trained text encoder from CLIP ($\Phi_{\text{clip-t}}(\cdot)$), and use two variants of prompts for better aligning the attribute with the visual region features: (i) for each attribute, we employ prior knowledge of ontologies, and encode their parent-class words along with the attribute, for example, the embedding for the ‘wet’ attribute can be expanded as: $\Phi_{\text{clip-t}}(\text{wet, state})$ to better distinguish from $\Phi_{\text{clip-t}}(\text{water, material})$, or $\Phi_{\text{clip-t}}(\text{in water, place})$; (ii) we augment it with multiple learnable prompt vectors, as a consequence, the attribute embeddings can be computed as:

$$\hat{t}_j = \Phi_{\text{clip-t}}([p_0, \dots, p_i, \mathbf{g}(\text{attribute}), p_{i+1}, \dots, p_j, \mathbf{g}(\text{parent-attribute}), p_{j+1}, \dots, p_k]) \quad (1)$$

where $\mathbf{g}(\cdot)$ denotes the tokenisation procedure, and p_i ($i \in 0, 1, \dots, k$) has the same dimension with the attribute word embeddings, denoting the learnable prompt vectors, that are shared across all attributes, and can generalize towards unseen attributes at inference time.

Attribute Classification. Attribute prediction can be obtained by computing the similarity between visual region feature and attribute concept embedding as:

$$\hat{s}_{ij} = \Phi_{\text{cls}}(\hat{v}_i, \hat{t}_j) = \sigma(\langle \hat{v}_i^T, \hat{t}_j \rangle / \tau), \quad (2)$$

where both v_i and t_j are L2 normalised, and \hat{s}_{ij} denotes the likelihood that the i th region contains the j th attribute. τ is a temperature parameter and σ denotes sigmoid function.

3.2.3 Training Procedure

In this section, we describe the training procedure for open-vocabulary attributes classification, which strives to better align the regional visual feature to the attribute description.

Step-I: Federated Training. In order to align the regional visual feature to attributes, an ideal training dataset should contain three types of annotations, namely, object bounding boxes, semantic categories, and attributes, as far as we know, the COCO Attributes dataset [24] is the only one that provides such a level of annotations, but with a very limited vocabulary size (196 attributes, 29 categories).

To fully exploit the annotations in existing datasets, we combine the detection dataset, *e.g.*, COCO [18], and attribute prediction dataset, *e.g.*, VAW [25]. Specifically, we follow standard procedure for training the class-agnostic region proposal network with images from COCO, *i.e.*, SmoothL1 loss and Binary Cross Entropy (BCE) are applied for box coordinates regression and objectness prediction; while for training attribute/category classification, as illustrated in the top-left part of Fig. 2, we employ ground-truth bounding boxes to crop the objects, and compute their visual embeddings with a pre-trained visual encoder from CLIP, we finetune CLIP’s **text encoder** by optimising BCE loss with multi-label attribute classification, as follows,

$$\mathcal{L}_{\text{cls}} = \frac{1}{N} \sum_{i=1}^N w_i \cdot \text{BCE}(\hat{s}_i, s_i) \quad (3)$$

where $N = |\mathcal{C}_{\text{base}}| + |\mathcal{A}_{\text{base}}|$, denotes the class number of categories and attributes, i denotes the i -th category/attribute, \hat{s}_i is the predicted probability, and $s_i \in \{0, 1, \text{unk}\}$ denotes an attribute label being negative, positive or missing. By default, for the missing attributes we treat them as negative with a re-weight factor, *i.e.*, $s_i = 0$ during training. $w_i \propto 1/f_i^\gamma$, $\sum_{i=1}^N w_i = N$, where f_i indicates the occurrence frequency of the i -th attribute in the training set, $\gamma = 0.25$ is a smoothing factor. As a result, this step ends up with a finetuned CLIP text encoder that better aligns the regional visual feature to attributes, referred to $\Phi_{\text{CLIP-Attr}}(\cdot)$.

Step-II: Training with Image-caption Dataset. To further improve the alignment, especially for novel attributes, we also consider using freely available image-caption datasets, *e.g.*, $\mathcal{D}_{\text{img-cap}} = \{\{\mathcal{I}_1, s_1\}, \dots, \{\mathcal{I}_N, s_N\}\}$, where the \mathcal{I}_i, s_i refer to image and caption sentence respectively. We detect all the objects in each image with a class-agnostic object proposal as described in Sec. 3.2.1. We keep the largest box proposal (b^*) and those with top-K objectness scores (b^k), and crop original images with the inferred bounding boxes. We pass these crops through $\Phi_{\text{CLIP-Attr}}(\cdot)$,

to get the predictions for semantic categories and attributes, and keep those with confidence scores higher than 0.7 as pseudo-positive labels. In addition, for caption sentences, we use TextBlob [19] to parse all captions into ‘semantic category’, ‘attribute’, and ‘noun phrases’ based on COCO and VAW dictionaries. For example, the sentence “A striped zebra is eating green grass” is processed and converted to {category: ‘zebra’}, {attribute: ‘green’, ‘striped’}, {noun phrase: ‘striped zebra’, ‘green grass’}.

To this end, we continue finetuning the alignment model ($\Phi_{\text{CLIP-Attr}}(\cdot)$) with the pseudo groundtruths obtained from the pre-processing stage. In detail, we compute the visual and textual embeddings as in **Step-I**, however, as the labels obtained from captions or the model’s prediction are not guaranteed to be correct, that requires special actions. We adopt multi-instance contrastive learning (MIL-NCE) [23], that maximizes the accumulated similarity score of positive matches between the visual and textual embeddings as follows:

$$\mathcal{L}_{\text{MIL-NCE}} = -\log \frac{\sum_{(v,t) \in \mathcal{P}} \exp(\frac{\langle v^T, t \rangle}{\tau})}{\sum_{(v,t) \in \mathcal{P}} \exp(\frac{\langle v^T, t \rangle}{\tau}) + \sum_{(v',t') \sim \mathcal{N}} \exp(\frac{\langle v'^T, t' \rangle}{\tau})} \quad (4)$$

where \mathcal{P} is a set of *positive* pairs with image crop feature and textual concept embeddings, \mathcal{N} conversely refers to an associated set of *negative* pairs. Here, we pair the largest box (b^*) with the given caption, *i.e.*, noun phrases, attributes, and semantic categories. While for the other top-K boxes (b^k), we treat the **model inferred** categories and attributes as positives. Here, we continue training both **visual and text encoders** in $\Phi_{\text{CLIP-Attr}}$ by optimising the following loss:

$$\mathcal{L}_{\text{cls}} = 1/K \cdot \sum_{k=0}^K \mathcal{L}_{\text{MIL-NCE}}^k \quad (5)$$

where $\mathcal{L}_{\text{MIL-NCE}}^k$ denotes MIL-NCE loss over the k th box and the corresponding textual concepts (here, we treat the largest box b^* as the 0th). An overview is shown in the bottom-left of Fig. 2.

3.3. Distilled Object Attribute Recognition

Although open-vocabulary object attribute prediction can be realised by the above proposed $\Phi_{\text{CLIP-Attr}}$ with the pre-computed proposals, the inference procedure is time-consuming, because every cropped region is fed into the visual encoder. In this section, we aim to address the slow inference speed, and train a Faster-RCNN type model end-to-end for object detection and attribute prediction, termed as OvarNet (Open-vocabulary attribute recognition):

$$\{\hat{b}_k, \hat{c}_k, \hat{a}_k\} = \Phi_{\text{Ovar}} = \Phi_{\text{cls}} \circ \Phi_{\text{crpn}} \circ \Phi_{\text{v-enc}}(\mathcal{I})$$

where the image is sequentially processed by a visual encoder, class-agnostic region proposal, and open-vocabulary attributes classification, as illustrated in the right of Fig. 2.

Visual Encoder. To start with, the input image is fed into a visual backbone, obtaining multi-scale feature maps:

$$\mathcal{F} = \{f^1, \dots, f^l\} = \Phi_{\text{v-enc}}(\mathcal{I}) \quad (6)$$

where f^i refers to the feature map at i -th level, we adopt the visual encoder from $\Phi_{\text{CLIP-Attr}}$.

Class-agnostic Region Encoding. To extract regional visual embeddings for candidate objects, we make the class-agnostic region encoding as follows,

$$\{\hat{v}_1, \dots, \hat{v}_n\} = \Phi_{\text{crpn}} = \Phi_{\text{attn-pool}} \circ \Phi_{\text{roi-align}} \circ \Phi_{\text{rpn}}(\mathcal{F}) \quad (7)$$

specifically, the feature pyramid is used in the region proposal network to fuse multi-scale features. The ROI-align’s output ($\mathbb{R}^{14 \times 14 \times 256}$) is firstly down-sampled with a convolutional layer (stride 2 and kernel size 2×2), and then passed into a block with 4 Transformer encoder layers with a learnable token, acting as attentional pooling. As a result, $\hat{v}_i \in \mathbb{R}^{1 \times D}$ refers to the feature embedding of the i -th candidate object. We train the proposal network only on base categories as described in Sec. 3.2.1

Open-vocabulary Attributes Classification. We extracted attribute concept embeddings as in Sec. 3.2.2. After obtaining the embeddings for each of the proposed objects, we can classify them into arbitrary attributes or categories by measuring the similarity between visual and attribute embeddings (Eq. 2).

Federated Training. We combine both COCO and VAW, and adopt a similar federated training strategy as in CLIP-Attr, with the key difference being that we jointly supervise localization for class-agnostic region proposal and classification for attribute prediction. The overall loss function can be formulated as: $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{cls}} + \lambda_{\text{RPN}} \cdot \mathcal{L}_{\text{RPN}}$, λ_{RPN} is a re-weighted parameter.

Intuitively, if the embedding spaces for visual and textual can be well-aligned by training on a limited number of base categories/attributes, the model should enable open-vocabulary object attribute recognition with the aforementioned training procedure, however, in practice, we observe unsatisfactory performance on the novel categories and attributes. We further incorporate additional knowledge distillation from the CLIP-Attr model described in Sec. 3.2.3 to improve the model’s ability for handling unseen categories and attributes.

Training via Knowledge Distillation. In addition to the federated training loss $\mathcal{L}_{\text{total}}$, we introduce an extra distillation item $\mathcal{L}_{\text{dist}}$, that encourages similar prediction between $\Phi_{\text{CLIP-Attr}}(\cdot)$ and $\Phi_{\text{Ovar}}(\cdot)$:

$$\mathcal{L}_{\text{dist}}(\hat{s}, s) = \frac{1}{N} \sum_{i=1}^N \text{KL}(\hat{s}_i, s_i), \quad (8)$$

where \hat{s} is prediction probabilities over all attributes from OvarNet and s is the prediction by using image crops from the aligned $\Phi_{\text{CLIP-Attr}}$. KL denotes the Kullback-Leibler divergence loss.

4. Experimental Setup

4.1. Datasets

Here, we introduce the datasets for training and evaluation of our proposed models for open-vocabulary object attribute recognition. Note that, while training the model, we have to consider two aspects of the openness evaluation, one is on semantic category, and the other is on the attributes.

MS-COCO [18]. We follow the setup for generalized zero-shot detection as proposed in ZSD [2]: 48 classes are selected as base classes (C_{base}), and 17 classes are used as unseen/novel classes (C_{novel}). The train and minival sets are the same as standard MS-COCO 2017. At the training stage, only the images with base category objects are used.

VAW [25]. For attributes recognition, VAW is constructed with VGPhraseCut [35] and GQA [10], containing a large vocabulary of 620 attributes, for example, color, material, shape, size, texture, action, *etc.* Each instance is annotated with positive, negative, and missing attributes. In our experiments, we sample half of the ‘tail’ attributes and 15% of the ‘medium’ attributes as the novel set (A_{novel} , 79 attributes) and the remaining as the base (A_{base} , 541 attributes). More details are included in the supplementary material.

Image-Caption Datasets. Conceptual Captions 3M (CC-3M) [32] contains 3 million image-text pairs harvested from the web with wide diversities, and COCO Caption (COCO-Cap) [4] comprises roughly 120k images and 5-way image-caption curated style annotations. We only keep images whose pairing captions have overlapped attributes or categories in the COCO and VAW dictionaries. We refer to the two subsets as CC-3M-sub and COCO-Cap-sub.

Summary. We have constructed the COCO-base and VAW-base datasets for training, and COCO-novel and VAW-novel for evaluation purposes, with the former for *object category classification*, and the latter for *object attributes classification*. To align regional visual features with attributes in CLIP-Attr, we use COCO-base and VAW-base for **Step-I** training, and then use CC-3M-sub and COCO-Cap-sub for **Step-II** finetuning. Later, COCO-base and VAW-base are employed in distilling knowledge from CLIP-Attr to OvarNet for efficiency. We refer the reader to a more detailed table with dataset statistics in the supplementary material.

4.2. Evaluation Protocol and Metrics

Our considered open-vocabulary object attribute recognition involves two sub-tasks: open-vocabulary object detection and classifying the attributes for all detected objects.

We evaluate the two sub-tasks in both **box-given** and **box-free** settings on COCO for category detection and VAW, LSA, and OVAD for attribute prediction (evaluations of the VAW and LSA are in the supplemental material). Specifically, the box-given setting is widely used in attribute prediction and object recognition communities [7, 25, 31, 33], where the ground-truth bounding box annotations are assumed to be available for all objects, and the protocol only evaluates object category classification and multi-label attribute classification with mAP metric; In contrast, the box-free setting favors a more challenging problem, as the model is also required to simultaneously localise the objects, and classify the semantic category and attributes.

Note that, the annotations on existing attribute datasets, such as VAW, LSA, are **not exhaustive or object-centric**, (i) not all the objects are labeled in an image, (ii) some annotations are on stuffs, that represents uncountable amorphous regions, such as sky and grass. We have to strike a balance in the box-free setting for attributes by matching the predicted boxes to the ground-truth box with the largest IoU, and then evaluate the attribute predictions using mAP. We consider the aforementioned metrics over base set classes, novel set classes, and all classes.

4.3. Implementation details

CLIP-Attr Training. We use the pre-trained R50-CLIP as the visual backbone to get object-centric visual features; all cropped regions are resized to 224×224 based on the short side with the original aspect ratio kept. Similar to Detic [44], we use sigmoid activation and multi-label binary cross-entropy loss for classification. We adopt the Stochastic Gradient Descent (SGD) optimizer with a learning rate of 0.001, a weight decay of 0.0001, and a momentum of 0.9. In **Step-I** training, we train the prompt vectors and text encoder for 50 epochs. In **Step-II** training with the image-caption dataset, we further finetune the entire model (both visual and textual encoders) for another 40 epochs. We select 30 top-K proposals while pre-possessing COCO-Cap-sub, and 15 for CC-3M-sub, as the images are often object-centric in the latter case.

OvarNet Training. In OvarNet, we initialise its visual backbone with the trained CLIP-Attr (Resnet50 without AttentionPool2d layer) and keep its text encoder frozen for efficiency. We adopt the AdamW optimizer with a learning rate of 0.0001. The models are trained for 30 epochs with the distillation term, and 60 epochs without distillation. We employ 640-800 scale jittering and horizontal flipping, and the temperature parameter τ is configured to be trainable. Following the observation and related prior, we empirically set: $\gamma = 0.25$, and $\lambda_{\text{RPN}} = 1$. All the experiments are conducted on 8 NVIDIA A100 GPUs.

Prompt Engineering. We have experimented with differ-

Attribute	Parent Attribute	M/L	VAW		COCO	
			AP _{novel}	AP _{all}	AP _{novel}	AP _{all}
✓	✗	none	52.15	59.16	40.53	49.84
✓	✗	M	53.64	62.22	41.65	52.35
✓	✓	M	53.78	62.76	41.97	52.81
✓	✗	L	55.73	64.54	42.77	53.80
✓	✓	L	57.39	66.92	45.82	55.21

Table 1. Ablation study on prompt engineering with CLIP-Attr model. **M/L** denotes whether manually designed prompts or learnable prompts are used.

b^* -cap.	MIL-NCE			VAW		COCO	
	b^* -phr.	b^* -attr.	b^k -attr.	AP _{novel}	AP _{all}	AP _{novel}	AP _{all}
				57.39	66.92	45.82	55.21
✓				57.45	66.94	45.87	55.36
✓	✓			57.42	67.87	48.29	57.92
✓	✓	✓		57.61	69.33	51.83	63.80
✓	✓	✓	✓	57.73	69.03	52.61	65.17

Table 3. The effect of different weakly supervised loss terms in Step-II training. We conduct ablation studies with COCO-Cap-sub dataset. b^* and b^k refer to the largest object proposal and top-K objectness proposals of an image respectively.

ent numbers of prompt vectors, empirically, we take 30 vectors and divide them into 10 inserting before, between, and after the attribute and the parent-class attributes. In terms of prompts used for encoding noun phrases, we use 16 learnable vectors, *i.e.*, 8 before and 8 after phrase embedding.

4.4. Ablation Study

We conduct ablation studies on VAW and COCO datasets, to thoroughly validate the effectiveness of proposed components, including prompt learning with the parent-class attribute, different losses for training CLIP-Attr, and the effect of **Step-I** and **Step-II** training. Finally, we validate the effectiveness of knowledge distillation.

Prompt Learning with Parent-class Attribute. In attribute embedding, we employ two variants of prompts for better aligning the attribute with the visual region features. We compare the learned prompt to the manual prompt while training $\Phi_{\text{CLIP-Attr}}$ with the pre-annotated object boxes. As shown in Tab. 1, comparing to the results from only using plain attribute words, using carefully designed prompt [8, 43], for example, “It is a photo of [category]” and “The attribute of the object is [attribute]” for the category and attribute words, indeed delivers improvements; While adding the parent-class word to the prompt template, *i.e.*, use the “The attribute of the object is [attribute], and it is a [parent-attribute]”, the lexical ambiguity can be alleviated, leading to a considerable improvement on novel categories and attributes. Finally, our proposed prompt learning with parent-class attribute words further brings a performance improvement by 3.61/3.85 AP on novel attributes/categories, compared to the manual prompt with parent-class words.

Method	Training Data	VAW			COCO		
		AP _{base}	AP _{novel}	AP _{all}	AP _{base}	AP _{novel}	AP _{all}
Plain CLIP	none	47.69	46.15	47.53	38.56	41.13	39.45
$\Phi_{\text{CLIP-Attr}}$	COCO-base	49.03	47.07	48.75	59.33	42.49	53.93
$\Phi_{\text{CLIP-Attr}}$	VAW-base	67.71	57.28	66.82	38.90	42.54	39.98
$\Phi_{\text{CLIP-Attr}}$	COCO-base + VAW-base	67.90	57.39	66.92	58.26	45.82	55.21
$\Phi_{\text{CLIP-Attr}}$	+ CC-3M-sub	69.79	59.16	68.87	65.79	48.90	61.36
$\Phi_{\text{CLIP-Attr}}$	+ COCO-Cap-sub	70.24	57.73	69.03	69.62	52.61	65.17

Table 2. Oracle test for Step-I and Step-II training with **objects’ boxes given**. ‘Plain CLIP’ directly classifies cropped images with a manual prompt.

Effect of Step-I and Step-II Training. We first compare the performance of attributes classification on regional visual feature (assuming ground-truth object boxes are given) before and after **Step-I** training. As illustrated in Tab. 2, the original plain CLIP model has certainly exhibited attribute classification on an elementary level. We can see a substantial improvement by further training on COCO-base and VAW-base, for example, from 46.15 to 57.39 AP for novel attribute classification, and 41.13 to 45.82 AP for novel category classification. Furthermore, by incorporating image-caption datasets in **Step-II** training, the performance has been improved to 69.03/65.17 AP on all attributes/categories. In the following experiments, we employ the model $\Phi_{\text{CLIP-Attr}}$ that exploits the COCO-Cap-Sub for Step-II training.

Effect of Different Losses for Step-II Training. We investigate the performance variance by adjusting the different supervisions in Step-II training (Eq. 5). As illustrated in Tab. 3, performance tends to grow monotonically with the increased supervision terms, from 66.92/55.21 to 69.03/65.17 AP on all attributes/categories, indicating that all supervision signals count.

Knowledge Distillation. We validate the necessity for knowledge distillation while training OvarNet. Specifically, we experiment by training the model with federated loss only, and two other knowledge distillation approaches, *i.e.*, the regional visual features (Feat.) from the visual encoder of $\Phi_{\text{CLIP-Attr}}$, like ViLD [8], and the prediction probability (Prob.) over all attributes from the matching scores of $\Phi_{\text{CLIP-Attr}}$. We achieve the distillation by constraining the regional visual feature or prediction probability of OvarNet to be the same as that of $\Phi_{\text{CLIP-Attr}}$, employing L2/L1 loss on features, and KL loss on probability. As shown in Tab. 4, we make two observations, *first*, knowledge distillation is essential; *second*, knowledge gained from attribute prediction probabilities is more beneficial to improving performance on novel sets, *e.g.*, from 51.87/33.17 to 56.43/54.10 AP when compared to L2 loss on visual features, in particular for semantic classification on COCO.

Different Architectures in $\Phi_{\text{CLIP-Attr}}$. We have evaluated the performance on different pre-trained CLIP architectures for attribute classification, such as R50 and ViT-

Distil.	VAW		COCO	
	AP _{novel}	AP _{all}	AP _{novel}	AP _{all}
none	50.53	61.74	30.43	59.83
Feat. L2	51.87	63.34	33.17	59.16
Feat. L1	52.57	64.65	32.92	59.62
Prob. KL	56.43	68.52	54.10	67.23

Table 4. Ablation study on knowledge distillation **with boxes given**.

Model	Arch.	VAW		COCO	
		AP _{novel}	AP _{all}	AP _{novel}	AP _{all}
$\Phi_{\text{CLIP-Attr}}$	R50	57.73	69.03	52.61	65.17
$\Phi_{\text{CLIP-Attr}}$	ViT-B/16	57.69	71.72	58.30	70.94
Φ_{Ovar}	R50	56.43	68.52	54.10	67.23
Φ_{Ovar}	ViT-B/16	56.41	68.79	55.70	68.02

Table 5. Experiments with different $\Phi_{\text{CLIP-Attr}}$ architectures in a **box-given setting**.

Initialisation	Freeze	VAW		COCO	
		AP _{novel}	AP _{all}	AP50 _{novel}	AP50 _{all}
ImageNet	✗	50.27	59.22	31.06	48.62
$\Phi_{\text{CLIP-Attr}}$	✗	54.85	68.02	35.25	54.31
$\Phi_{\text{CLIP-Attr}}$	✓	55.47	67.62	35.17	54.15

Table 6. Ablation study on finetuning the visual backbone of OvarNet with different initializations under **box-free setting**.

Method	Training Data	VAW			COCO		
		AP _{base}	AP _{novel}	AP _{all}	AP50 _{base}	AP50 _{novel}	AP50 _{all}
SConE [25]	fully supervised	-	-	68.30	-	-	-
TAP [26]	fully supervised	-	-	65.40	-	-	-
OVR-RCNN [38]	COCO Cap	-	-	-	46.00	22.80	39.90
OVR-RCNN [38]	CC 3M	-	-	-	-	-	34.30
ViLD [8]	CLIP400M	-	-	-	59.50	27.60	51.30
Region CLIP [42]	COCO Cap	-	-	-	54.80	26.80	47.50
Region CLIP [42]	CC 3M	-	-	-	57.10	31.40	50.40
PromptDet [6]	Web Images	-	-	-	-	26.60	50.60
Detic [44]	COCO Cap	-	-	-	47.10	27.80	45.00
OvarNet (box-given)	COCO-base + VAW-base	68.27	53.75	66.85	60.94	41.44	55.85
OvarNet (box-given)	+CC 3M-sub	69.30	55.44	67.96	68.35	52.34	64.18
OvarNet (box-given)	+COCO Cap-sub	69.80	56.43	68.52	71.88	54.10	67.23
OvarNet (box-free)	COCO-base + VAW-base	67.71	53.42	66.03	56.20	32.02	49.77
OvarNet (box-free)	+CC 3M-sub	67.32	54.26	66.75	59.50	33.68	52.40
OvarNet (box-free)	+COCO Cap-sub	68.93	55.47	67.62	60.35	35.17	54.15

Table 7. Comparison for open-vocabulary object detection and attribute prediction on the VAW test set and COCO validation.

B/16, and then conducted knowledge distillation from the different $\Phi_{\text{CLIP-Attr}}$ models. As seen in Tab. 5, both architectures perform competitively, with transformer-based architectures consistently outperforming the ConvNet ones.

Updating OvarNet’s Visual Backbone. We experiment by updating or freezing the OvarNet’s visual backbone from different initializations at training. As shown in Tab. 6, initialising the visual backbone from aligned $\Phi_{\text{CLIP-Attr}}$ is advantageous, whereas finetuning or freezing it makes little difference. For efficiency, we opt to freeze the visual backbone in other experiments.

4.5. Comparison with the State-of-the-Art

Benchmark on COCO and VAW. In Tab. 7, we compare OvarNet to other attribute prediction methods and open-vocabulary object detectors on the VAW test set and COCO validation set. As there is no open-vocabulary attribute prediction method developed on the VAW dataset, we re-train two models on the *full* VAW dataset as an oracle comparison, namely, SConE [25] and TAP [26]. Our best model achieves 68.52/67.62 AP across all attribute classes for the box-given and box-free settings respectively. On COCO open-vocabulary object detection, we compare with OVR-RCNN [38], ViLD [8], Region CLIP [42], PromptDet [6], and Detic [44], our best model obtains 54.10/35.17 AP for novel categories, surpassing the recent state-of-the-art ViLD-ens [8] and Detic [44] by a large margin, showing that attributes understanding is beneficial for open-vocabulary




Figure 3. Qualitative visualization from OvarNet. **Red**: base category/attributes. **Blue**: category/attributes.

object recognition. Fig. 3 shows some prediction results of OvarNet.

5. Conclusion

In the paper, we consider the problem of open-vocabulary object detection and attribute recognition, *i.e.*, simultaneously localising objects and inferring their semantic categories and visual attributes. We start with a naïve two-stage framework (CLIP-Attr) that uses a pre-trained CLIP to classify the object proposals, to better align the object-centric visual feature with attribute concepts, we use learnable prompt vectors on the textual encoder side. On the training side, we adopt a federated training strategy to exploit both object detection and attribute prediction datasets, and explore a weakly supervised training regime with external image-text pairs to increase the robustness for recognising novel attributes. Finally, for computational efficiency, we distill the knowledge of CLIP-Attr into a Faster-RCNN type model (termed as OvarNet), while evaluating on four different benchmarks, *e.g.*, VAW, MS-COCO, LSA, and OVAD, we show that jointly training object detection and attributes prediction is beneficial for visual scene understanding, largely outperforming the existing approaches that treat the two tasks independently, demonstrating strong generalization ability to novel attributes and categories.

Acknowledgement

This work is supported by National Key R&D Program of China (No. 2022ZD0161400).

References

- [1] Abrar H Abdalnabi, Gang Wang, Jiwen Lu, and Kui Jia. Multi-task cnn model for attribute prediction. *IEEE Transactions on Multimedia*, 17(11):1949–1959, 2015. 2
- [2] Ankan Bansal, Karan Sikka, Gaurav Sharma, Rama Chellappa, and Ajay Divakaran. Zero-shot object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 384–400, 2018. 2, 3, 6
- [3] María A Bravo, Sudhanshu Mittal, Simon Ging, and Thomas Brox. Open-vocabulary attribute detection. *arXiv preprint arXiv:2211.12914*, 2022. 2
- [4] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 6
- [5] Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *2009 IEEE conference on computer vision and pattern recognition*, pages 1778–1785. IEEE, 2009. 2
- [6] Chengjian Feng, Yujie Zhong, Zequn Jie, Xiangxiang Chu, Haibing Ren, Xiaolin Wei, Weidi Xie, and Lin Ma. Promptdet: Expand your detector vocabulary with uncurated images. *arXiv preprint arXiv:2203.16513*, 2022. 2, 4, 8
- [7] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 6
- [8] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*, 2021. 2, 3, 7, 8
- [9] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Deep imbalanced learning for face recognition and attribute prediction. *IEEE transactions on pattern analysis and machine intelligence*, 42(11):2781–2794, 2019. 2
- [10] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. 6
- [11] Phillip Isola, Joseph J Lim, and Edward H Adelson. Discovering states and transformations in image collections. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1383–1391, 2015. 2
- [12] Dinesh Jayaraman and Kristen Grauman. Zero-shot recognition with unreliable attributes. *Advances in neural information processing systems*, 27, 2014. 3
- [13] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 2
- [14] Prannay Kaul, Weidi Xie, and Andrew Zisserman. Label, verify, correct: A simple few shot object detection method. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14237–14247, 2022. 4
- [15] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017. 2
- [16] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *2009 IEEE conference on computer vision and pattern recognition*, pages 951–958. IEEE, 2009. 2
- [17] Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot learning of object categories. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (01):1–1, 2013. 3
- [18] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 4, 6
- [19] Steven Loria et al. Textblob. <https://textblob.readthedocs.io>. 5
- [20] Qiaomei Mao, Chong Wang, Shenghao Yu, Ye Zheng, and Yuqi Li. Zero-shot object detection with attributes-based category similarity. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 67(5):921–925, 2020. 3
- [21] Qier Meng and Satoh Shin’ichi. Adinet: Attribute driven incremental network for retinal image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. 2
- [22] Kareem Metwaly, Aerin Kim, Elliot Branson, and Vishal Monga. Glidenet: Global, local and intrinsic based dense embedding network for multi-category attributes prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4835–4846, 2022. 2
- [23] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9879–9889, 2020. 5
- [24] Genevieve Patterson and James Hays. Coco attributes: Attributes for people, animals, and objects. In *European Conference on Computer Vision*, pages 85–100. Springer, 2016. 2, 4
- [25] Khoi Pham, Kushal Kafle, Zhe Lin, Zhihong Ding, Scott Cohen, Quan Tran, and Abhinav Shrivastava. Learning to predict visual attributes in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13018–13028, 2021. 2, 4, 6, 8
- [26] Khoi Pham, Kushal Kafle, Zhe Lin, Zhihong Ding, Scott Cohen, Quan Tran, and Abhinav Shrivastava. Improving closed and open-vocabulary attribute prediction using transformers. In *European Conference on Computer Vision*, pages 201–219. Springer, 2022. 2, 8

- [27] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [2](#)
- [28] Shafin Rahman, Salman Khan, and Nick Barnes. Improved visual-semantic alignment for zero-shot object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11932–11939, 2020. [3](#)
- [29] Shafin Rahman, Salman Khan, and Fatih Porikli. Zero-shot object detection: Learning to simultaneously recognize and localize novel concepts. In *Asian Conference on Computer Vision*, pages 547–563. Springer, 2018. [3](#)
- [30] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [2](#), [4](#)
- [31] Nirat Saini, Khoi Pham, and Abhinav Shrivastava. Disentangling visual embeddings for attributes and objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13658–13667, 2022. [2](#), [6](#)
- [32] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernamed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. [6](#)
- [33] Jasper RR Uijlings, Koen EA Van De Sande, Theo Gevers, and Arnold WM Smeulders. Selective search for object recognition. *International journal of computer vision*, 104(2):154–171, 2013. [6](#)
- [34] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. [2](#)
- [35] Chenyun Wu, Zhe Lin, Scott Cohen, Trung Bui, and Subhransu Maji. Phrasecut: Language-based image segmentation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10216–10225, 2020. [6](#)
- [36] Wenjia Xu, Yongqin Xian, Jiuniu Wang, Bernt Schiele, and Zeynep Akata. Attribute prototype network for zero-shot learning. *Advances in Neural Information Processing Systems*, 33:21969–21980, 2020. [2](#), [3](#)
- [37] Yu Yun, Sen Wang, Mingzhen Hou, and Quanxue Gao. Attributes learning network for generalized zero-shot learning. *Neural Networks*, 150:112–118, 2022. [2](#)
- [38] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021. [2](#), [8](#)
- [39] Sanyi Zhang, Zhanjie Song, Xiaochun Cao, Hua Zhang, and Jie Zhou. Task-aware attention model for clothing attribute prediction. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(4):1051–1064, 2019. [2](#)
- [40] Yuwei Zhang, Peng Zhang, Chun Yuan, and Zhi Wang. Texture and shape biased two-stream networks for clothing classification and attribute recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13538–13547, 2020. [2](#)
- [41] Yang Zhong, Josephine Sullivan, and Haibo Li. Face attribute prediction using off-the-shelf cnn features. In *2016 International Conference on Biometrics (ICB)*, pages 1–7. IEEE, 2016. [2](#)
- [42] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022. [8](#)
- [43] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9):2337–2348, 2022. [7](#)
- [44] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Phillip Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. *arXiv preprint arXiv:2201.02605*, 2022. [2](#), [4](#), [6](#), [8](#)
- [45] Yaohui Zhu, Weiqing Min, and Shuqiang Jiang. Attribute-guided feature learning for few-shot image recognition. *IEEE Transactions on Multimedia*, 23:1200–1209, 2020. [2](#)
- [46] Zhengxia Zou, Keyan Chen, Zhenwei Shi, Yuhong Guo, and Jieping Ye. Object detection in 20 years: A survey. *Proceedings of the IEEE*, 111(3):257–276, 2023. [1](#)