

ProxyDet: Synthesizing Proxy Novel Classes via Classwise Mixup for Open-Vocabulary Object Detection

Joonhyun Jeong^{1,2}, Geondo Park², Jayeon Yoo³, Hyungsik Jung¹, Heesu Kim^{1*}

¹NAVER Cloud, ImageVision

²Korea Advanced Institute of Science and Technology (KAIST)

³Seoul National University

{joonhyun.jeong, hyungsik.jung, heesu.kim89}@navercorp.com, geondopark@kaist.ac.kr, jayeon.yoo@snu.ac.kr

Abstract

Open-vocabulary object detection (OVOD) aims to recognize novel objects whose categories are not included in the training set. In order to classify these unseen classes during training, many OVOD frameworks leverage the zero-shot capability of largely pretrained vision and language models, such as CLIP. To further improve generalization on the unseen novel classes, several approaches proposed to additionally train with pseudo region labeling on the external data sources that contain a substantial number of novel category labels beyond the existing training data. Albeit its simplicity, these pseudo-labeling methods still exhibit limited improvement with regard to the truly unseen novel classes that were not pseudo-labeled. In this paper, we present a novel, yet simple technique that helps generalization on the overall distribution of novel classes. Inspired by our observation that numerous novel classes reside within the convex hull constructed by the base (seen) classes in the CLIP embedding space, we propose to synthesize proxy-novel classes approximating novel classes via linear mixup between a pair of base classes. By training our detector with these synthetic proxy-novel classes, we effectively explore the embedding space of novel classes. The experimental results on various OVOD benchmarks such as LVIS and COCO demonstrate superior performance on novel classes compared to the other state-of-the-art methods. Code is available at <https://github.com/clovaai/ProxyDet>.

1 Introduction

Object detection is a task in which the model simultaneously predicts where all objects are located within an image and to which class each object belongs. Conventional closed-set object detection frameworks (Lin et al. 2017b; He et al. 2017; Lin et al. 2017a) have limited the range of discernible objects to a fixed set of base classes, which are usually the classes annotated in training sets. Recently, the performance of those frameworks on the fixed set of classes has largely evolved, building upon the success of vision transformer architectures (Dosovitskiy et al. 2020; Liu et al. 2022; Carion et al. 2020; Zong, Song, and Liu 2022; Chen et al. 2022) and large-scale pretraining techniques (Su et al. 2022; Chen et al. 2022; Wang et al. 2022). However, beyond improving

the performance on the fixed class set, *open-vocabulary object detection* (OVOD) frameworks (Gu et al. 2021; Zareian et al. 2021) have emerged to tackle the limited range of discernible objects. The objective is to detect a set of novel classes which are not seen during training, without newly labeling the bounding boxes and class labels of these novel classes which is labor-intensive and non-scalable.

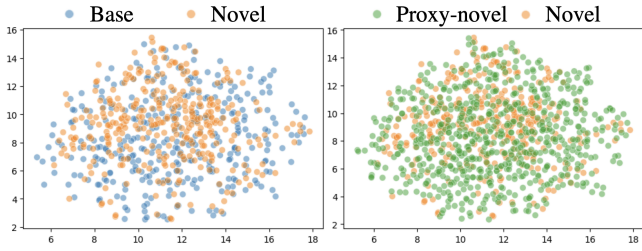
In order to achieve open-vocabulary property, OVOD frameworks should expand their classifier for novel classes. Therefore, previous OVOD frameworks (Gu et al. 2021; Zhou et al. 2022; Kuo et al. 2022) adopt pretrained vision-language model (VLM), CLIP (Radford et al. 2021), to obtain the weights of their classifier by prompting the names of base and novel classes. These VLM models provide an interface (i.e., encoder) that maps textual information (e.g., class name) to an embedding space where embeddings of semantically matched images are closely located. In line with previous work, our study also adopts this approach to enable open-vocabulary detection.

In order to further explicitly expand the range of detection vocabulary on the novel classes, recent OVOD frameworks (Zhou et al. 2022; Feng et al. 2022; Zhao et al. 2022; Gao et al. 2022; Lin et al. 2022) rely on pseudo-labeling a subset of novel classes given in the inexpensive supplementary datasets (Ridnik et al. 2021) consisting of pairs of image and its class label (or caption) without box annotations. The missing box annotations are predicted using pretrained region proposal network (RPN) and their class labels are assigned using image-level labels (Zhou et al. 2022; Rasheed et al. 2022) and the words within the captions (Lin et al. 2022), which contain a substantial number of novel classes. These pseudo-labeling approaches are essentially identical to the conventional closed-set object detection, the only difference being that the human-labeled box annotations are not used. Consequently, it has limitations in fundamentally solving generalization to the truly unseen novel classes that were not pseudo-labeled, only to improve performance for the pseudo-labeled novel classes as shown in Table 1.

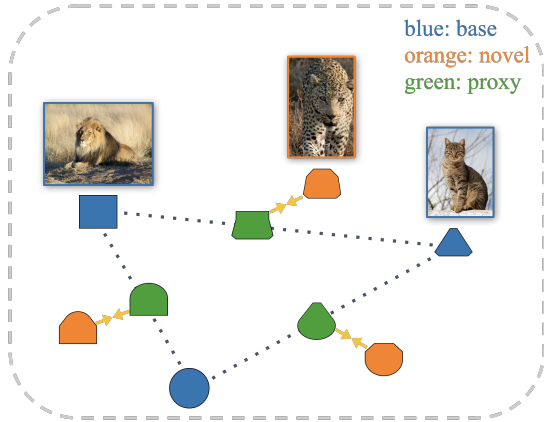
To overcome this limitation, we present ProxyDet, a new OVOD framework that helps better generalization on the overall distribution of novel classes including these truly unseen novel ones that have never been presented during training. Empirically, we observe that a wide range of novel classes resides within the convex hull constructed by the

*Corresponding Author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



(a) UMAP (McInnes, Healy, and Melville 2018) visualization of CLIP text embeddings for base/novel classes (left), proxy-novel/novel classes (right) in LVIS dataset.



(b) Proxy-novel class synthesis via convex combination of the base classes in CLIP embedding space.

Figure 1: Visualization of embedding space for all class groups (a) and our proposed proxy-novel synthesis (b).

base classes in the VLM embedding space as shown in the left side of Figure 1a. Motivated by this observation, we propose a synthesis of *proxy-novel* classes, which approximates novel classes via convex combination between the base classes (Figure 1b). Notably, we observe that the overall distribution of novel classes can be better approximated through our proposed proxy-novel synthesis as shown in Figure 2. Learning on these proxy-novel classes provides an extensive proximal exploration of numerous novel classes during training (the right side of Figure 1a), thus effectively expanding the detection vocabulary. Consequently, as shown in Table 1, our method improves Average Precision (AP) on the non-pseudo-labeled novel classes up to 3.4, while the previous pseudo-labeling approach (Zhou et al. 2022) rather degrades by 0.3 (See the Appendix¹ for the performance of the other pseudo-labeling methods).

Our extensive experiments on various OVOD benchmarks show that this straightforward technique significantly improves classification capability for novel classes, outperforming other state-of-the-art OVOD methods. On LVIS (Gupta, Dollar, and Girshick 2019) benchmark, our method shows largely improved performance with regard to novel class, +1.6 AP_r , compared to the baseline method (Zhou et al. 2022), while enhancing by +2.6 AP_{novel} on COCO (Lin et al. 2014). Furthermore, our proposed method shows

¹The full extended paper including the appendix is available at <https://arxiv.org/abs/2312.07266>.

| Method | AP_r | | |
|------------------------------------|----------------|-----------------|----------------|
| | Non-pseudo | Pseudo | Overall |
| Pseudo-labeling (Zhou et al. 2022) | 19.4 (-0.3) | 25.9 (+10.0) | 24.6 (+7.9) |
| Ours | 23.1 (+3.4) | 27.0 (+11.1) | 26.2 (+9.5) |

Table 1: Effectiveness of pseudo-labeling and our method on the novel classes. The evaluation is conducted on novel classes that are separated based on whether they are pseudo-labeled or not, abbreviated as "Non-pseudo" and "Pseudo". "Overall" denotes performance on the overall novel classes. The quantities in the parentheses indicate the amount of performance changes relative to the vanilla case without applying pseudo-labeling or our method.

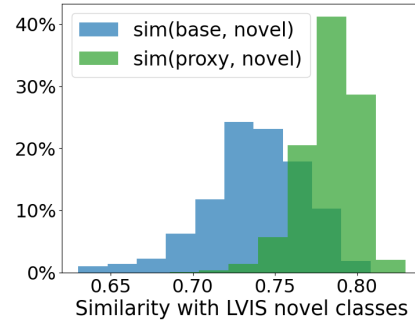


Figure 2: Histogram of similarity between training class groups (base and proxy-novel) and novel classes. We compute pairwise cosine similarity between all the classes in each of training class group and novel classes in LVIS dataset, using text representations obtained from CLIP.

the state-of-the-art performance on cross-dataset evaluation (57.0 AP on COCO and 19.1 AP on Objects365), showing the superiority of our method on expanding detection vocabulary. In summary, our contributions are outlined as follows:

- We empirically discover that mixing the representations of base classes generates proxy-novel classes that can approximate novel classes.
- Building upon the empirical observation, we propose a novel approach, introducing proxy-novel classes in training objective. By training with the proxy-novel classes, our proposed method encourages a model to explore the proximal representation space of novel classes.
- Extensive experiments on various benchmarks demonstrate the superiority of our method over the other state-of-the-art OVOD methods, with a large margin of performance improvement in detecting novel classes.

2 Related Work

Mixup Regularizing neural networks with data augmentation during training is essential to alleviate overfitting on fixed training dataset. Beyond the fundamental image transformations (cropping, flipping, and color jittering), image-mixing augmentations (Zhang et al. 2017; Yun et al. 2019; Kim, Choo, and Song 2020; Kim et al. 2021; Jeong et al. 2023) have successfully resolved the overfitting problem and achieve better generalization performance on training

classes, becoming a widely-accepted standard for training neural network based classifiers. Especially, mixup (Zhang et al. 2017) regularized the overfitting on training images by presenting synthesized images made by convex combinations between existing training instances (images). Furthermore, Verma et al. (2019) and So et al. (2022) expanded the application of mixup to the hidden representation (embedding) space of the uni-modal and multi-modal network, respectively. While these mixup variants (Zhang et al. 2017; Yun et al. 2019; Kim et al. 2021; Verma et al. 2019; Jeong et al. 2023) have primarily focused on instance-wise mixup in image space for better generalization on training classes, this paper concentrates on the class-wise mixup that explores the augmented multi-modal representation space spanned by the base training classes, for better understanding the novel classes which are not presented for training.

Open-vocabulary object detection Based on an extensive source of object detection data with high-quality annotations, object detection frameworks (Lin et al. 2017b; Redmon et al. 2016; Jeong et al. 2022; Zhu et al. 2020) have flourished on various benchmarks (Lin et al. 2014; Everingham et al. 2009; Gupta, Dollar, and Girshick 2019) where the range of classes to be detected is predefined and limited. In order to expand the predefined class set without the laborious annotation procedures, open-vocabulary object detection frameworks (Zareian et al. 2021; Gu et al. 2021; Zhong et al. 2022; Du et al. 2022; Kuo et al. 2022; Minderer et al. 2022) have been proposed. The pioneering work (Zareian et al. 2021) employed the backbone pretrained with a corpus of image-caption pairs. With the advent of massively pretrained VLMs (Radford et al. 2021; Jia et al. 2021), ViLD (Gu et al. 2021) proposed to distill knowledge from these VLMs to the detector. One step further, F-VLM (Kuo et al. 2022) utilized a visual encoder of VLM (Radford et al. 2021) as its representation generator. Several works (Du et al. 2022; Feng et al. 2022) proposed prompt engineering for VLM tailored for object detection. Although the above methods achieved state-of-the-art results, they still exhibit limited generalization to the novel classes since their detectors are supervised with the base classes only. Conversely, our proposed method additionally leverages the supervision with respect to the proxy-novel classes via class-wise mixup between a pair of base classes, improving generalization on the overall distribution of novel classes.

Open-vocabulary object detection with pseudo-labeling Another line of works (Zhou et al. 2022; Feng et al. 2022; Zhao et al. 2022; Gao et al. 2022; Rasheed et al. 2022; Lin et al. 2022) proposed to pseudo-label on inexpensive datasets, which do not contain bounding box annotations, but cover lots of novel classes, to improve the performance on novel classes. Detic (Zhou et al. 2022) proposed to use image-level labeled dataset (Ridnik et al. 2021) and pseudo-label the bounding box using the max-size proposal presented from RPN (Ren et al. 2015) where its class label is given by image-level labels. A similar approach (Rasheed et al. 2022) further refined the pseudo-labeling method via the class-agnostic object detector (Maaz et al. 2022). On the other hand, several works (Zhong et al. 2022; Lin

et al. 2022) exploited pseudo-labeling on image-caption pair dataset (Sharma et al. 2018) to manipulate VLM to be aware of the relationship between object and novel vocabularies within the caption. However, these pseudo-labeling methods exhibit limited generalization performance on the truly unseen novel classes that were not pseudo-labeled, while our method largely improves performance on these genuine novel ones, as shown in Table 1.

3 Method

We briefly review the overall framework of OVOD methods with notation clarification in Section 3.1. In Section 3.2, we present our method, ProxyDet, which ventures into an unexplored embedding space via proxy-novel classes, enhancing the classification ability on the novel classes.

3.1 Notations and Preliminaries

An object detector is trained on a dataset $D = \{(I_i, g_i)\}_{i=1}^{|D|}$ where I_i refers to an i -th image and $g_i = \{(b_k, c_k)\}_{k=1}^{|g_i|}$ denotes a set of ground-truth object labels of the i -th image, consisting of a bounding box (b_k) and the corresponding class label ($c_k \in \mathcal{C}^B$) where \mathcal{C}^B is the entire set of base classes. At the inference phase, OVOD aims to detect not only the base classes \mathcal{C}^B but also novel classes \mathcal{C}^N which satisfy $\mathcal{C}^B \cap \mathcal{C}^N = \emptyset$.

In OVOD training, RPN (Zhou, Koltun, and Krähenbühl 2021) generates a set of region proposals \mathcal{P} for the images in the batch and positive proposals, $\mathcal{p} \subset \mathcal{P}$, with sufficient IoUs with ground-truth bounding boxes of base classes are selected and labeled according to the class of the nearest ground-truth box. Then, each positive proposal $p \in \mathcal{p}$ is vectorized into a region embedding $r \in \mathbb{R}^M$ through detector head with RoI-Align (He et al. 2017), forming a set of region embeddings \mathcal{R} . For classification training of base class objects, we adopt binary cross entropy (BCE) loss (Zhou et al. 2022) using cosine similarity between r and text embeddings of base classes $\mathcal{T}(\mathcal{C}^B) \in \mathbb{R}^{|\mathcal{C}^B| \times M}$, where $\mathcal{T}(\cdot)$ denotes text encoder of CLIP (Radford et al. 2021).

3.2 ProxyDet: Learning Proxy Novel Classes via Class-wise Mixup

In order to facilitate the detector to recognize novel classes not seen during the training phase, it is crucial for the detector to explore the embedding space that is proximate to the space of novel classes. Inspired by our observation that a number of novel classes reside within the convex hull delineated by the base classes (Figure 1a), we synthesize the proxy-novel classes via class-wise mixup between the representations of the base classes. After obtaining the visual and textual embedding of these proxy-novel classes, we regulate them to be close to each other. The overall framework of our method is illustrated in Figure 3.

Synthesizing proxy-novel classes Given the visual and textual prototypes of base classes in a batch, \bar{r}^B and $\mathcal{T}(\mathcal{C}^B)$, we randomly sample a pair of prototypes ($\bar{r}^{B_i}, \mathcal{T}(c^{B_i})$) and ($\bar{r}^{B_j}, \mathcal{T}(c^{B_j})$) that contain class-specific features of base

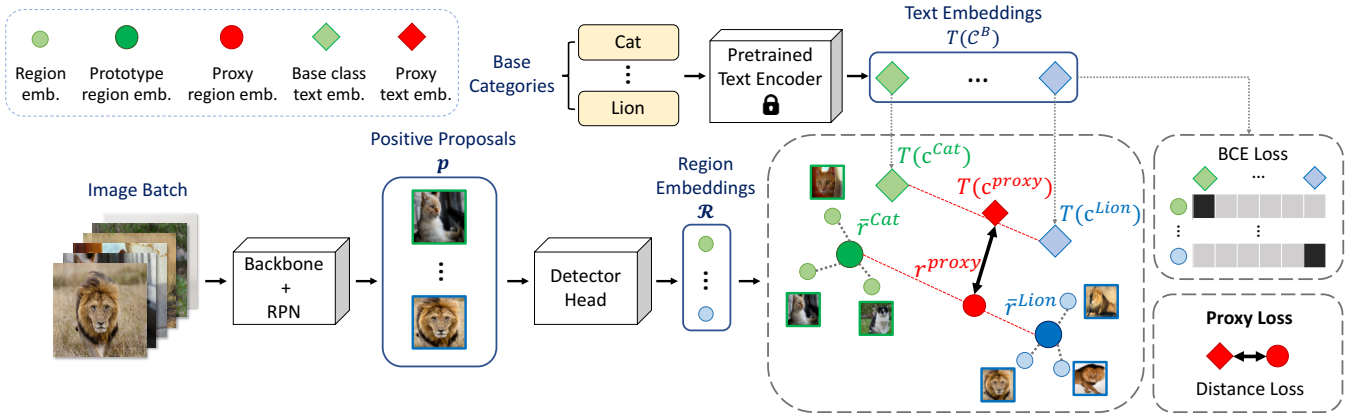


Figure 3: Overview of training pipeline for ProxyDet. Our proposed proxy loss supervises the detector to learn image-text modality with regard to proxy-novel classes via class-wise mixup between a pair of base classes.

classes c^{B_i} and c^{B_j} , respectively. The procedure for constructing the prototype will be elaborated on afterwards.

In order to obtain the visual and textual representations of proxy-novel class $c^{proxy} \in \mathcal{C}$, where \mathcal{C} is the set of all classes beyond the base and novel classes, we propose a class-wise mixup approach between these prototypes of base classes:

$$r^{proxy} = \mathcal{V}(\lambda \bar{r}^{B_i} + (1 - \lambda) \bar{r}^{B_j})$$

$$\mathcal{T}(c^{proxy}) = \mathcal{V}(\lambda \mathcal{T}(c^{B_i}) + (1 - \lambda) \mathcal{T}(c^{B_j})) \quad (1)$$

where r^{proxy} and $\mathcal{T}(c^{proxy})$ refer to the region (i.e., visual) and textual prototype embedding of c^{proxy} , respectively. $\mathcal{V}(\mathbf{x}) = \frac{\mathbf{x}}{\|\mathbf{x}\|_2}$ refers to the L2 normalization of a vector \mathbf{x} and λ follows Beta(γ, γ).

Constructing prototype The way of constructing visual and textual prototypes is a crucial factor in obtaining sophisticated representations of proxy-novel classes. Textual prototypes can be obtained by prompting the category name c^{B_i} into the text encoder of a pretrained VLM (Radford et al. 2021). On the other hand, obtaining class-specific region embedding \bar{r}^{B_i} is more complicated as there are multiple region embeddings from a batch of images that have the potential to represent a class c^{B_i} . A naive solution is to calculate the centroid of all these region embeddings which correspond to c^{B_i} :

$$\bar{r}^{B_i} = \mathcal{V}\left(\sum_{r \in \mathcal{R}^{B_i}} F(r) * r\right),$$

$$F(r) = \frac{1}{|\mathcal{R}^{B_i}|} \quad (2)$$

where \mathcal{R}^{B_i} refers to the set of region embeddings for positive proposals labeled as c^{B_i} . $F(r)$ denotes the weighting function that determines the contribution of each positive region embedding r towards constructing the prototype. In Eq. 2, all the region embeddings are given equal weight in the construction process.

Robust visual prototype However, simply averaging all region embeddings to construct a centroid prototype embedding could be sub-optimal since the positive proposals

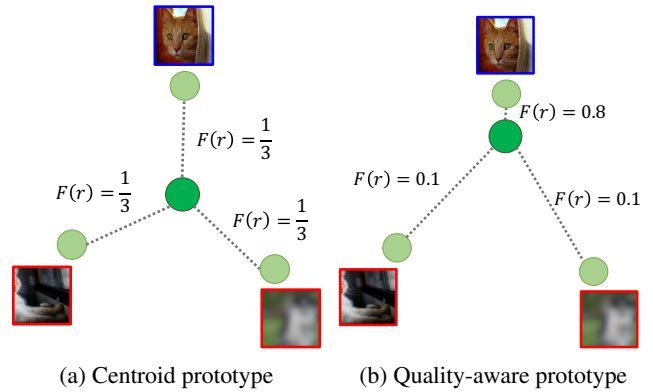


Figure 4: Construction of visual prototype embedding. Red boxes denote low-quality proposals (e.g., partially localized, blurry) while blue boxes denote high-quality proposals. Quality-aware prototype reduces the influence of low-quality proposals, resulting in robust prototype.

have the potential to include low-quality localizing results as shown in Figure 4. To address this issue, we propose an approach to construct a robust prototype by incorporating the weighting function $F(r)$ with a measure that is aware of the quality of proposals, denoted as $\phi(\cdot)$.

$$F(r) = \frac{\exp(\phi(r))}{\sum_{\hat{r} \in \mathcal{R}^{B_i}} \exp(\phi(\hat{r}))} \quad (3)$$

For designing the proposal quality-aware measure function $\phi(\cdot)$, we leverage *objectness score* from RPN or *IoU* with the nearest ground-truth box. By doing so, our proxy-novel synthesis in Eq. 1 is conditioned on more robust prototype embeddings which can more precisely estimate the class-wise representations with respect to the base classes.

Proxy loss Building upon the visual and textual representations of synthesized proxy-novel classes with the robust formulation of prototype embedding, we derive our proposed training objective as follows:

$$\mathcal{L}_{proxy} = \|\mathcal{T}(c^{proxy}) - r^{proxy}\|_1 \quad (4)$$

Intuitively, we train the OVOD model by regularizing the region embeddings of proxy-novel classes to minimize its distance with the proxy-novel text embeddings, which effectively encourages the model to explore the proximal embedding space of the novel classes.

Inference As our proxy loss (Eq. 4) is designed to improve the classification performance of novel classes, the objective is different from the conventional BCE loss which prioritizes the classification performance of base classes. To prevent the potential contention between these two losses, we apply each loss separately on two separate detector heads with identical architecture. During the inference phase, we compute classification scores of the two region embeddings from the separated heads by calculating cosine similarity against the text embeddings of all the classes. Then, we merge the two classification scores by geometric mean, following Gu et al. (2021) and Kuo et al. (2022):

$$s_i = \begin{cases} (rw^\top)_i^\alpha * (r'w^\top)_i^{1-\alpha}, & \text{if } c^i \in \mathcal{C}^B \\ (rw^\top)_i^\beta * (r'w^\top)_i^{1-\beta}, & \text{if } c^i \in \mathcal{C}^N \end{cases} \quad (5)$$

where s_i is the final classification score for i -th class c^i , classifier weight w is L2 normalized text embeddings from $\mathcal{T}(\mathcal{C}^B \cup \mathcal{C}^N)$, and r, r' are L2 normalized region embeddings from the head trained with proxy loss and the other head with BCE loss, respectively. $\alpha, \beta \in [0, 1]$ determines the ratio of reflecting classification score from proxy loss head for base/novel classes, respectively.

4 Experiments

4.1 Datasets

We evaluate our method on the widely-used OVOD benchmark datasets, LVIS (Gupta, Dollar, and Girshick 2019), COCO (Lin et al. 2014), and Objects365 (Shao et al. 2019), to compare the performance against the previous works.

LVIS and ImageNet-LVIS. The LVIS dataset consists of 100k annotated images with bounding boxes and segmentation mask labels for 1203 classes. The object classes in LVIS are divided into three distinct groups, namely "frequent", "common", and "rare", based on the occurrence frequency within the dataset. We follow Open-Vocabulary LVIS (OV-LVIS) benchmark (Gu et al. 2021) that uses common and frequent class groups for base classes (866 categories) and the rare group for novel classes (337 categories). Optionally, we utilize ImageNet-LVIS (IN-L) as an additional training dataset with image-level labels (Zhou et al. 2022), which is a subset of ImageNet-21k (Ridnik et al. 2021) containing only the categories that intersect with those of LVIS.

COCO and COCO Caption. We adopt Open-Vocabulary COCO (OV-COCO) benchmark proposed in Bansal et al. (2018) and Zareian et al. (2021), which divides COCO classes into base classes (48 categories) and novel classes (17 categories). We additionally use captions in COCO dataset as image-level labels, following Zhou et al. (2022).

Objects365. Objects365 has 365 categories over 2M images. We use Objects365 (Shao et al. 2019) for cross-dataset evaluation to substantiate the generalization capability.

| Method | Backbone | Training source | AP _r | AP |
|---------------------------|----------|------------------|-----------------|-------------|
| ViLD (Gu et al.) | R50 | LVIS-base | 16.6 | 25.5 |
| RegionCLIP (Zhong et al.) | R50 | LVIS-base | 17.1 | 28.2 |
| F-VLM (Kuo et al.) | R50 | LVIS-base | 18.6 | 24.2 |
| ProxyDet (ours) | R50 | LVIS-base | 18.9 | 30.1 |
| Fully-supervised* | R50 | LVIS | 25.5 | 31.1 |
| Detic (Zhou et al.) | R50 | LVIS-base + IN-L | 24.6 | 32.4 |
| ProxyDet (ours) | R50 | LVIS-base + IN-L | 26.2 | 32.5 |
| Detic (Zhou et al.) | Swin-B | LVIS-base + IN-L | 33.8 | 40.7 |
| ProxyDet (ours) | Swin-B | LVIS-base + IN-L | 36.7 | 41.5 |

Table 2: Performance comparison on Open-Vocabulary LVIS benchmark. LVIS-base denotes training with base classes only, while * denotes fully supervised with the base and novel classes. R50 denotes ResNet50.

4.2 Implementation Details

In OV-LVIS setup, we follow the detector model architecture of Detic (Zhou et al. 2022) and its training recipe. Our object detection model is based on Cascade Mask-RCNN (Cai and Vasconcelos 2018) with ResNet50 (He et al. 2016) or Swin-B (Liu et al. 2022) backbone, employing CenterNet2 (Zhou, Koltun, and Krähenbühl 2021) as RPN. For obtaining the text embedding for the classifier, we use CLIP ViT-B/32 (Radford et al. 2021). The training recipe using ResNet50 backbone consists of federated loss (Zhou, Koltun, and Krähenbühl 2021) and binary cross-entropy loss. We trained the detector for 90,000 iterations (1x schedule) with 64 batch size and AdamW (Loshchilov and Hutter 2017) optimizer. For the learning rate schedule, we increased 0 to 2e-4 for the first 1000 warmup iterations and decayed by cosine annealing strategy (Loshchilov and Hutter 2016). In the case of using image-level supervision, we fine-tune the pretrained detector whose weight parameters were trained on LVIS base classes, and thus total training iteration is 2x schedule. For the image-level labeled dataset, IN-L, we do not supervise the detector with our proposed proxy loss. For evaluation, we report mask AP for all the classes and AP_r for the rare (novel) classes, respectively. AP_r is the main metric we primarily concentrate on.

In OV-COCO setup, we adopt Faster R-CNN (Ren et al. 2015) with ResNet50-C4 backbone. We trained for 1x schedule with batch size 16 and SGD optimizer. For evaluation, we report box AP₅₀ for the base (AP_{base}), novel (AP_{novel}), and all classes (AP), where AP_{novel} is the main metric. For the hyper-parameters of our method, we randomly sample the mixing coefficient λ (in Eq. 1), following Beta(1, 1) distribution. For the score fusion in Eq. 5, we use $\alpha = 0.45$ and $\beta = 0.65$.

4.3 Main Results

Open-vocabulary LVIS. In Table 2, we compare our proposed method, ProxyDet, with the other state-of-the-art OVOD methods on OV-LVIS benchmark. When only LVIS-base is used for training without any external dataset, our method shows the best performance on novel classes, 18.9 AP_r. This result supports the fact that ProxyDet flourishes open-vocabulary property without any other additional training source. However, the fully-supervised detector trained with all the categories including base and novel

| Method | AP _{novel} | AP _{base} | AP |
|---------------------------------|---------------------|--------------------|------|
| WSDDN (Bilen and Vedaldi 2016) | 19.7 | 19.6 | 19.6 |
| Cap2Det (Ye et al. 2019) | 20.3 | 20.1 | 20.1 |
| OVR-CNN (Zareian et al. 2021) | 22.8 | 46.0 | 39.9 |
| ViLD† (Gu et al. 2021) | 27.6 | 59.5 | 51.3 |
| RegionCLIP‡ (Zhong et al. 2022) | 26.8 | 54.8 | 47.5 |
| Detic (Zhou et al. 2022) | 27.8 | 47.1 | 45.0 |
| ProxyDet (ours) | 30.4 | 52.6 | 46.8 |

Table 3: Performance comparison on Open-Vocabulary COCO benchmark. † denotes using 8x training schedule. ‡ denotes leveraging an additional pretraining scheme for the visual encoder with extensive object concepts.

| Method | Training Source | Objects365 | | |
|-------------------------|-----------------|-----------------------|------------------|----------------------------------|
| | | COCO AP ₅₀ | AP ₅₀ | AP ₅₀ ^{rare} |
| Target-supervised* | Target dataset | 67.6 | 38.6 | - |
| LVIS-supervised§ | LVIS-base | 55.4 | 18.9 | 16.2 |
| ViLD (Gu et al. 2021) | LVIS-base | 55.6 | 18.2 | 15.4† |
| DetPro (Du et al. 2022) | LVIS-base | 53.8 | 18.8 | 13.9† |
| F-VLM (Kuo et al. 2022) | LVIS-base | 53.1 | 19.2 | 15.4† |
| ProxyDet (ours) | LVIS-base | 57.0 | 19.1 | 16.5 |

Table 4: Cross-dataset evaluation on COCO and Objects365. * denotes fully-supervised with annotations of all the categories in the target dataset. § denotes trained only with base classes of LVIS. † denotes reproduced by our implementation, while all the other comparison numbers are reproduced based on their official implementations. For fair comparison, all the methods employed ResNet50 backbone.

classes (5th row) provides that there is still large headroom for improvement, achieving 25.5 AP_r. Notably, with image-level supervision from IN-L, our model (7th row) outperforms the fully-supervised detector, reaching 26.2 AP_r. Moreover, compared to Detic which also utilizes IN-L, our method outperforms them by a large margin, achieving +1.6 AP_r and +2.9 AP_r improvement when using ResNet50 and Swin-B backbones, respectively.

Open-vocabulary COCO. Table 3 displays the performance comparison results on OV-COCO benchmark. The pioneering work, OVR-CNN, shows moderate performance on novel classes via pretraining with image-caption pairs. With the advent of massively pretrained VLMs (i.e., CLIP), several works (ViLD, RegionCLIP, Detic) greatly enhanced novel class performance. Notably, our proposed method further improves AP_{novel} by +2.6 compared to Detic, achieving the best performance among the comparisons. This result validates the effectiveness of our proposed supervision with regard to proxy-novel classes.

4.4 Cross-Dataset Evaluation

To further investigate the open-vocabulary property, we evaluate the transfer performance of our model on the other detection datasets in Table 4. Simply adding our proposed method (the last row) on the LVIS-supervised detector (2nd row) achieved large performance improvement on COCO (+1.6 AP₅₀). Notably, our method clearly achieves the best

| Proxy-novel | Mixup granularity | Visual prototype | AP _r |
|-------------|-------------------|------------------------|-----------------|
| ✗ | - | - | 24.6 (+0.0) |
| ✓ | Instance-wise | - | 24.8 (+0.2) |
| ✓ | Class-wise | centroid | 25.6 (+1.0) |
| ✓ | Class-wise | weighted by IoU | 26.0 (+1.4) |
| ✓ | Class-wise | weighted by obj. score | 26.2 (+1.6) |

Table 5: Effects of mixup granularity and weighting methods to construct visual prototype embeddings. We sampled the mixing coefficient λ from Beta(1, 1), when synthesizing proxy-novel classes.

| Class-wise Mixup | Visual prototype | λ sampling | AP _r |
|------------------|------------------|--------------------|-----------------|
| ✗ | - | - | 24.6 (+0.0) |
| ✗ | obj. score | Bernoulli(0.5) | 24.8 (+0.2) |
| ✓ | obj. score | Beta(1, 1) | 26.2 (+1.6) |
| ✓ | obj. score | Beta(5, 5) | 26.0 (+1.4) |
| ✓ | obj. score | Beta(10, 10) | 25.5 (+0.9) |
| ✓ | obj. score | Beta(50, 50) | 25.6 (+1.0) |

Table 6: Effect of sampling strategy for mix coefficient λ .

transfer performance among the comparisons on COCO, trailing right behind the target-supervised case (1st row) that provides the upper-bound performance. When transferred on Objects365, we achieve similar performance with the state-of-the-art method, F-VLM, while ours clearly attained the top AP performance with regard to the rare classes, AP₅₀^{rare}. This result indicates that our model generalizes well even for the rare classes from the different data distributions.

4.5 Ablation Studies

We perform ablation studies to figure out the efficacy of each component and sensitivity of the hyper-parameter comprising our proposed method, ProxyDet. Unless specified, we use both LVIS and IN-L as training datasets as in Zhou et al. (2022) and report mask AP_r on LVIS dataset.

Effect of class-wise mixup with prototype construction.

In Table 5, we analyze the effect of proxy loss via class-wise mixup and the technique for constructing the visual prototype. Using instance-wise mixup (2nd row) without the construction of class-wise visual prototype embeddings shows a marginal performance difference with the vanilla case (1st row). Since instance-wise mixup blends individual instances of region embeddings for each base class, it comes with the drawback of incorporating the noisy embeddings from low-quality box proposals (Figure 4) into the mixup without any form of filtering, which might impede stable training. In contrast, leveraging the class-wise prototype embeddings (3-5th rows) largely enhances the performance of novel classes. Especially, robustly designing the prototype embedding based on IoU or objectness score further strengthens novel class performance by reducing the influence of such low-quality region embeddings on prototype construction.

Ablation on mixing coefficient. In Table 6, we analyze the effect of sampling strategy for the coefficient of mixing base classes, λ . Only leveraging robust visual prototype without

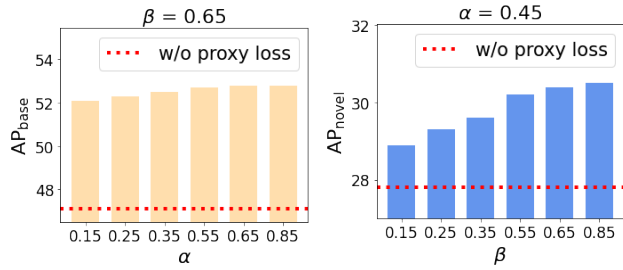


Figure 5: Ablation on fusion parameters α and β . We report AP results on Open-Vocabulary COCO benchmark.

| LVIS Mixup | IN-L Mixup | AP _r |
|------------|------------|-----------------|
| ✓ | ✗ | 26.2 (+0.0) |
| ✓ | ✓ | 23.6 (-2.6) |

Table 7: Effect of class-wise mixup on IN-L.

class-wise mixing (2nd row) by sampling λ from $\{0, 1\}$ following Bernoulli(0.5) shows similar performance with the vanilla case (1st row). However, employing class-wise mixing via Beta distribution sampling (3-6th rows) substantially enhances the performance of novel classes, implying the significance of constructing and utilizing proxy-novel classes through the fusion of base classes to expand the proximity to the embedding space near novel classes.

Effect of score fusion parameters. In Figure 5, we analyze the effect of score fusion parameters applied with regard to the detector head trained with our proxy loss (Eq 5). Generally, all different α or β cases stably outperform the vanilla case (red dotted line) where the detector head trained with our proxy loss does not participate in score fusion (i.e., $\alpha = 0, \beta = 0$). Moreover, as the value of α is raised, the performance of base classes increases gradually while raising β results in a more rapid improvement in terms of the performance for novel classes. This indicates the essential role of training with our proxy loss in terms of improving performance for both base and, more significantly, novel classes.

Scale-up of mixup on ImageNet-LVIS. In Table 7, we validate the effect of additionally applying class-wise mixup on IN-L (Zhou et al. 2022). Since IN-L does not provide the human-labeled bounding box annotations, we simply leveraged max-size proposal from RPN following Zhou et al. (2022) to construct the visual prototype embedding for our proxy loss. We observed that performance of the novel classes largely degrades. We conjecture that since max-size proposal is too coarsely localized, it might contain the other class objects and non-objects which can be served as noise to construct the visual prototype embedding for proxy loss.

Comparison with the other supervision for novel classes. In Table 8, we compared the performance with the other supervision technique which distills the knowledge of novel class objects from the outputs of CLIP visual encoder (Gu et al. 2021). By applying distillation supervision along with the vanilla BCE loss, novel class performance is improved by +0.7 AP_r. Meanwhile, adding our proposed proxy loss shows the best performance, further improving by +1.6 AP_r.

| Supervision | AP _r |
|-------------------------------------|-----------------|
| BCE loss (base classes) | 24.6 (+0.0) |
| BCE + distillation (Gu et al. 2021) | 25.3 (+0.7) |
| BCE + proxy loss (ours, Eq. 4) | 26.2 (+1.6) |

Table 8: Comparison of proxy-novel supervision with knowledge distillation from CLIP.

| Proxy-novel | Novel category prior | Mixing class selection | AP _r |
|-------------|----------------------|------------------------|-----------------|
| ✗ | ✗ | - | 24.6 (+0.0) |
| ✓ | ✗ | random | 26.2 (+1.6) |
| ✓ | ✓ | novel-nearest | 26.9 (+2.3) |

Table 9: The effect of selection strategy for the base classes to be mixed.

4.6 Towards Better Approximation of Novel Class

In this subsection, we take a deeper look at the effectiveness of proxy-novel by assuming that a set of vocabularies including the names of the novel categories are given as prior knowledge. With the knowledge of the novel classes and their corresponding text embeddings, we can control which base classes should be mixed to more accurately approximate each of the target novel classes. Specifically, for each novel class, we extract the optimal combination of base classes that exhibit the highest similarity to the novel class in CLIP textual embedding space. Subsequently, we employ the identified combination of base classes to synthesize the proxy-novel and implement our proposed proxy loss.

As shown in Table 9, selecting the optimal combination of base classes nearest to each of novel classes improves AP_r by +2.3 compared to the vanilla case (1st row), achieving better performance compared to the random selection strategy (2nd row). However, in real-world scenarios of OVOD, the knowledge about the names of novel categories is not given during the training phase and hence the random selection is a more practical strategy for generating proxy-novel classes. Under this real-world scenario without explicitly knowing the novel categories, enlarging the set of prior vocabularies for mixing to ideally cover all the novel categories is a practical alternative, which is our future work.

5 Conclusion

Motivated by the observation that novel class representations can be approximated through mixup between base classes, we devise a new learning strategy to effectively detect novel category objects. Our method introduces proxy loss on proxy-novel classes synthesized by well-defined prototypes of base classes, which encourages our models to explore the embedding space close to the embeddings of novel classes. By the extensive experiments on various benchmarks and rigorous analysis with ablations, we demonstrate that our simple add-on technique helps generalization for detecting broad range of novel classes.

Acknowledgments

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant (No.2019-0-00075, Artificial Intelligence Graduate School Program (KAIST), No.2022-0-00713, Meta-learning applicable to real-world problems) and the National Research Foundation of Korea (NRF) grants (No.RS-2023-00209060, A Study on Optimization and Network Interpretation Method for Large-Scale Machine Learning) funded by the Korea government (MSIT). This work was also supported by KAIST-NAVER Hypercreative AI Center.

References

- Bansal, A.; Sikka, K.; Sharma, G.; Chellappa, R.; and Divakaran, A. 2018. Zero-shot object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 384–400.
- Bilen, H.; and Vedaldi, A. 2016. Weakly supervised deep detection networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2846–2854.
- Cai, Z.; and Vasconcelos, N. 2018. Cascade r-cnn: Delving into high quality object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 6154–6162.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, 213–229. Springer.
- Chen, Q.; Wang, J.; Han, C.; Zhang, S.; Li, Z.; Chen, X.; Chen, J.; Wang, X.; Han, S.; Zhang, G.; et al. 2022. Group detr v2: Strong object detector with encoder-decoder pre-training. *arXiv preprint arXiv:2211.03594*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Du, Y.; Wei, F.; Zhang, Z.; Shi, M.; Gao, Y.; and Li, G. 2022. Learning to prompt for open-vocabulary object detection with vision-language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14084–14093.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2009. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88: 303–308.
- 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 *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, 701–717. Springer.
- Gao, M.; Xing, C.; Niebles, J. C.; Li, J.; Xu, R.; Liu, W.; and Xiong, C. 2022. Open vocabulary object detection with pseudo bounding-box labels. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X*, 266–282. Springer.
- Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2021. Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921*.
- Gupta, A.; Dollar, P.; and Girshick, R. 2019. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5356–5364.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Jeong, J.; Cha, S.; Choi, J.; Yun, S.; Moon, T.; and Yoo, Y. 2023. Observations on K-Image Expansion of Image-Mixing Augmentation. *IEEE Access*, 11: 16631–16643.
- Jeong, J.; Kim, B.; Yu, J.; and Yoo, Y. 2022. Rediscovery of the Effectiveness of Standard Convolution for Lightweight Face Detection. *arXiv preprint arXiv:2204.01209*.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 4904–4916. PMLR.
- Kim, J.; Choo, W.; Jeong, H.; and Song, H. O. 2021. Co-Mixup: Saliency Guided Joint Mixup with Supermodular Diversity. In *International Conference on Learning Representations*.
- Kim, J.-H.; Choo, W.; and Song, H. O. 2020. Puzzle mix: Exploiting saliency and local statistics for optimal mixup. In *International Conference on Machine Learning*, 5275–5285. PMLR.
- Kuo, W.; Cui, Y.; Gu, X.; Piergiovanni, A.; and Angelova, A. 2022. F-VLM: Open-Vocabulary Object Detection upon Frozen Vision and Language Models. *arXiv preprint arXiv:2209.15639*.
- Lin, C.; Sun, P.; Jiang, Y.; Luo, P.; Qu, L.; Haffari, G.; Yuan, Z.; and Cai, J. 2022. Learning Object-Language Alignments for Open-Vocabulary Object Detection. *arXiv preprint arXiv:2211.14843*.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017a. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017b. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- 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 *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.

- Liu, Z.; Hu, H.; Lin, Y.; Yao, Z.; Xie, Z.; Wei, Y.; Ning, J.; Cao, Y.; Zhang, Z.; Dong, L.; et al. 2022. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12009–12019.
- Loshchilov, I.; and Hutter, F. 2016. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*.
- Loshchilov, I.; and Hutter, F. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Maaz, M.; Rasheed, H.; Khan, S.; Khan, F. S.; Anwer, R. M.; and Yang, M.-H. 2022. Class-agnostic object detection with multi-modal transformer. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X*, 512–531. Springer.
- McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- 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 with vision transformers. *arXiv preprint arXiv:2205.06230*.
- 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 *International conference on machine learning*, 8748–8763. PMLR.
- Rasheed, H.; Maaz, M.; Khattak, M. U.; Khan, S.; and Khan, F. S. 2022. Bridging the gap between object and image-level representations for open-vocabulary detection. *arXiv preprint arXiv:2207.03482*.
- Redmon, J.; Divvala, S.; Girshick, R.; and Farhadi, A. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 779–788.
- Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28.
- Ridnik, T.; Ben-Baruch, E.; Noy, A.; and Zelnik-Manor, L. 2021. Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972*.
- 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 *Proceedings of the IEEE/CVF international conference on computer vision*, 8430–8439.
- Sharma, P.; Ding, N.; Goodman, S.; and Soricut, R. 2018. Conceptual captions: A cleaned, hypernymed, 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)*, 2556–2565.
- So, J.; Oh, C.; Shin, M.; and Song, K. 2022. Multi-modal mixup for robust fine-tuning. *arXiv preprint arXiv:2203.03897*.
- Su, W.; Zhu, X.; Tao, C.; Lu, L.; Li, B.; Huang, G.; Qiao, Y.; Wang, X.; Zhou, J.; and Dai, J. 2022. Towards All-in-one Pre-training via Maximizing Multi-modal Mutual Information. *arXiv preprint arXiv:2211.09807*.
- Verma, V.; Lamb, A.; Beckham, C.; Najafi, A.; Mitliagkas, I.; Lopez-Paz, D.; and Bengio, Y. 2019. Manifold mixup: Better representations by interpolating hidden states. In *International conference on machine learning*, 6438–6447. PMLR.
- Wang, W.; Bao, H.; Dong, L.; Bjorck, J.; Peng, Z.; Liu, Q.; Aggarwal, K.; Mohammed, O. K.; Singhal, S.; Som, S.; et al. 2022. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*.
- Ye, K.; Zhang, M.; Kovashka, A.; Li, W.; Qin, D.; and Berent, J. 2019. Cap2det: Learning to amplify weak caption supervision for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9686–9695.
- Yun, S.; Han, D.; Oh, S. J.; Chun, S.; Choe, J.; and Yoo, Y. 2019. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE International Conference on Computer Vision*, 6023–6032.
- Zareian, A.; Rosa, K. D.; Hu, D. H.; and Chang, S.-F. 2021. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14393–14402.
- Zhang, H.; Cisse, M.; Dauphin, Y. N.; and Lopez-Paz, D. 2017. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*.
- Zhao, S.; Zhang, Z.; Schuster, S.; Zhao, L.; Vijay Kumar, B.; Stathopoulos, A.; Chandraker, M.; and Metaxas, D. N. 2022. Exploiting unlabeled data with vision and language models for object detection. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, 159–175. Springer.
- Zhong, Y.; Yang, J.; Zhang, P.; Li, C.; Codella, N.; Li, L. H.; Zhou, L.; Dai, X.; Yuan, L.; Li, Y.; et al. 2022. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16793–16803.
- Zhou, X.; Girdhar, R.; Joulin, A.; Krähenbühl, P.; and Misra, I. 2022. Detecting twenty-thousand classes using image-level supervision. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, 350–368. Springer.
- Zhou, X.; Koltun, V.; and Krähenbühl, P. 2021. Probabilistic two-stage detection. *arXiv preprint arXiv:2103.07461*.
- Zhu, X.; Su, W.; Lu, L.; Li, B.; Wang, X.; and Dai, J. 2020. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*.
- Zong, Z.; Song, G.; and Liu, Y. 2022. DETRs with Collaborative Hybrid Assignments Training. *arXiv preprint arXiv:2211.12860*.