GroundVLP: Harnessing Zero-Shot Visual Grounding from Vision-Language Pre-training and Open-Vocabulary Object Detection

Haozhan Shen¹, Tiancheng Zhao^{2*}, Mingwei Zhu¹, Jianwei Yin¹

¹Zhejiang University ²Binjiang Institute of Zhejiang University {hz_shen, zhumw, zjuyjw}@zju.edu.cn, tianchez@zju-bj.com,

Abstract

Visual grounding, a crucial vision-language task involving the understanding of the visual context based on the query expression, necessitates the model to capture the interactions between objects, as well as various spatial and attribute information. However, the annotation data of visual grounding task is limited due to its time-consuming and labor-intensive annotation process, resulting in the trained models being constrained from generalizing its capability to a broader domain. To address this challenge, we propose GroundVLP, a simple yet effective zero-shot method that harnesses visual grounding ability from the existing models trained from imagetext pairs and pure object detection data, both of which are more conveniently obtainable and offer a broader domain compared to visual grounding annotation data. GroundVLP proposes a fusion mechanism that combines the heatmap from GradCAM and the object proposals of open-vocabulary detectors. We demonstrate that the proposed method significantly outperforms other zero-shot methods on RefCO-CO/+/g datasets, surpassing prior zero-shot state-of-the-art by approximately 28% on the test split of RefCOCO and RefCOCO+. Furthermore, GroundVLP performs comparably to or even better than some non-VLP-based supervised models on the Flickr30k entities dataset. Our code is available at https://github.com/om-ai-lab/GroundVLP.

Introduction

Visual grounding seeks to pinpoint the image region described by a linguistic expression containing complex semantic information. It includes two typical tasks, Referring Expression Comprehension (REC) and Phrase Grounding. REC aims to localize an object in an image given a textual referring expression, while phrase grounding seeks to ground every entity in the sentence to objects in the image. Generally, models are trained via task-specific datasets in the supervised setting (Yu et al. 2018; Liu et al. 2019c; Sun et al. 2021; Yang et al. 2019; Liao et al. 2020; Deng et al. 2021) to perform visual grounding.

However, creating these task-specific datasets poses challenges due to their intricate annotation process. Annotating a query demands a detailed examination of object interactions and an understanding of various spatial and attribute



Figure 1: With the combination of existing models trained with image-text matching and object detection, we could conduct zero-shot visual grounding without fine-tuning on any additional supervised dataset.

information within the image. As a result, these datasets are limited in quantity, especially when compared to two other dataset types, as detailed in Table 1. Undoubtedly, this finite amount of data restricts a model's adaptability to broader domain. Compared to the difficulty of obtaining visual grounding data, two alternative data types—image-text pairs and object detection (OD) data—are comparatively easier to obtain, as shown in Table 1.

Recently, Vision-Language Pre-training (VLP) models, when trained on image-text pairs, have demonstrated impressive results in image-text matching (ITM) (Zhang et al. 2021; Kim, Son, and Kim 2021; Li et al. 2021). Similarly, Open-Vocabulary object Detectors (OVD) trained with OD data have excelled in detecting specific categories (Zareian et al. 2021; Gu et al. 2021; Zhou et al. 2022), as referenced in the top of Figure 1. Therefore, a natural question arises: Can we harness the semantic comprehension of VLP and the category-specific detection prowess of OVD to perform visual grounding without any additional training, just

^{*}Corresponding author

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

as shown in the bottom of Figure 1.

Data Type	Name	Total Size
Visual Grounding	RefCOCO/+/g(Yu et al. 2016) CLEVR-Ref+ (Liu et al. 2019b) Cops-Ref (Chen et al. 2020b) Flicker30k (Plummer et al. 2015)	$\sim 220k$
Image-Text Pair	COCO (Lin et al. 2014) Visual Genome (Krishna et al. 2017) CC (Changpinyo et al. 2021) LAION (Schuhmann et al. 2022, 2021)	~ 6.2B
Object Detection	COCO (Lin et al. 2014) LVIS (Gupta, Dollar, and Girshick 2019) OpenImages (Kuznetsova et al. 2020) ImageNet-21K [†] (Deng et al. 2009)	~ 23.3M

Table 1: The list of some widely used datasets for three tasks. Size means the number of images included in the dataset.[†] ImageNet, an image classification dataset, has been proven to be applicable to OD task by (Zhou et al. 2022).

In this paper, we introduce GroundVLP, a novel method for zero-shot visual grounding tasks that encompasses both REC and phrase grounding. GroundVLP comprises three main components: (1) a VLP model that employs GradCAM to identify the image regions that are most semantically relevant to the given $[expression]^1$, (2) an OVD to detect the candidate objects, and (3) a fusion mechanism that combines the aforementioned two parts using a weighted grade to select the answer judiciously. In contrast to its previous usage in the literature (Li et al. 2021; He et al. 2022), we aggregate the GradCAM attention values only for visually recognizable words to optimize the modality mapping from text to image. Another significant difference is that we narrow down the candidate boxes to those belonging to a given object category with an OVD for to reduce noisy candidates compared to previous zero-shot methods (Yao et al. 2021; Subramanian et al. 2022). The object category can be manually defined or predicted from the textual query using offthe-shelf NLP toolboxes such as Spacy (Honnibal and Johnson 2015) or Stanza (Qi et al. 2020).

We conduct main experiments on RefCOCO/+/g datasets for REC and Flickr30k Entities dataset for phrase grounding. GroundVLP outperforms all other zero-shot methods, which obtains an accuracy on average $\sim 18\%$ better than ReCLIP (Subramanian et al. 2022) on all splits of RefCO-CO/+/g and $\sim 36\%$ better than CPT (Yao et al. 2021) on all splits of Flickr30k Entities. Experiment results further show that it performs on par or even better than some non-VLP-based supervised models on most of the test data. This outstanding performance indicates that we can tackle visual grounding tasks, which is traditionally constrained by limited annotations, using easily accessible data such as imagetext pairs and pure object detection data. Additionally, we take ablation studies on each component of GroundVLP, demonstrating their effectiveness.

Our contributions could be summarized as: (1) We propose a simple yet effective zero-shot method supporting both REC and phrase grounding, which achieves performance comparable to some non-VLP-based supervised models, demonstrating that visual grounding could be addressed using easily accessible data. (2) We probe the cause of the decline in performance when not using the ground-truth category and discover inherent noise and bias on RefCOCO/+/g datasets. (3) We conduct detailed ablation studies to verify the effectiveness of each proposed component and demonstrate the weak visual grounding capability of OVD.

Preliminary

There are two widely used attention modules: self-attention and co-attention, where the former employs the query(\mathbf{Q}), key(\mathbf{K}), and value(\mathbf{V}) matrices created by the input sequence itself while the latter collects \mathbf{K} and \mathbf{V} from another sequence(Vaswani et al. 2017). Existing VLP models can be roughly grouped into three types of architecture consisting of the aforementioned two attention modules: one-stream, two-stream, and dual-encoders. We mainly take the first two types into account and depict them in Figure. 2. Our usage of GradCAM (Selvaraju et al. 2017) for these two attention architectures is described in detail here.

Given a text-image pair, we input them into VLP and define *T*, *I* as the number of text and image input tokens respectively. The attention map of a certain layer, denoted as **A**, can be computed by the product of **Q** and **K** with furthermore post-processes, as defined specifically in the equation: $\mathbf{A} = softmax(\frac{\mathbf{Q}\cdot\mathbf{K}^{\top}}{\sqrt{d_h}})$, where $\mathbf{Q} \in \mathbb{R}^{N_h \times s \times d_h}$, $\mathbf{K} \in \mathbb{R}^{N_h \times q \times d_h}$ and $\mathbf{A} \in \mathbb{R}^{N_h \times s \times q}$. N_h is he number of attention heads of multi-head attention, d_h means the dimension of hidden states, and s, q are assigned different values in different architecture:

$$(s,q) = \begin{cases} (T+I, T+I) \text{ one-stream} \\ (T,I) \text{ two-stream} \end{cases}$$
(1)

where for the two-stream architecture, we compute **A** in the co-attention module of the fusion encoder in which **Q** is from the language encoder and **K** from the image. Next, we compute the gradients map via back propagation: $\nabla \mathbf{A} = (\frac{\partial L_{itm}}{\partial \mathbf{A}})^+$, where L_{itm} represents the VLP model's output value of the ITM head, and we remove the negative contributions. Finally, the result map $\mathbf{G} \in \mathbb{R}^{s \times ||}$ is given by:

$$\mathbf{G} = \mathbb{E}_h(\nabla \mathbf{A} \odot \mathbf{A})$$

where \mathbb{E}_h is the average calculation across heads dimension and \odot means element-wise multiplication.

To sum up, for a given layer in the encoder, we can obtain a map $\mathbf{G} \in \mathbb{R}^{s \times q}$ via GradCAM.

The Proposed Method

Figure. 3 demonstrates an overview of GroundVLP. We first obtain ${\bf G}$ via GradCAM. Then we crop its size and apply our

¹In this paper, [query] refers to the query sentence provided by the grounding datasets and [expression] refers to a specific object we need to ground. For REC, [expression] is equivalent to [query]. For phrase grounding, it denotes a certain entity phrase included in [query].



Figure 2: Two types of attention architectures: (a) Onestream. (b) Two-stream. The two-stream model has an additional co-attention module in its cross-modality encoder compared to the one-stream model, which is used to interact with the information from two modalities.

proposed visual-word attention aggression to obtain the heat map of VLP. For the OVD module, we input the category of the referring target into it to receive n instances including boxes and confidence scores. Finally, we fuse two parts to calculate the *weighted grades* of each instance and output the one with the highest grade.

Generating a Heat-map for VLP

To ground the *[expression]*, we prompt the *[query]* and feed it into the model along with the image to obtain **G** by Grad-CAM. Then we can generate a heat map for VLP. First, we crop $\mathbf{G} \in \mathbb{R}^{x \times q}$ to $\mathbf{G}' \in \mathbb{R}^{T \times I}$, where \mathbf{G}' denotes the influence of the image tokens on each text token:

$$\mathbf{G}' = \begin{cases} \mathbf{G}[i,j]_{1 \le i \le T}^{I \le j \le T+I} & \text{one-stream} \\ \mathbf{G} & \text{two-stream} \end{cases}$$
(2)

Next, \mathbf{G}' needs to be further squeezed to $\tilde{\mathbf{G}} \in \mathbb{R}^I$ so that it can represent the connections between the whole *[expression]* and each image token. To this end, we propose *visual-word attention aggregation*, distinguished from previous methods which are implemented by either using the row corresponding to the [CLS] token (Subramanian et al. 2022) directly or averaging the scores of rows across all text tokens (Li et al. 2021; He et al. 2022), as shown in Figure. 4.

Visual-Word Attention Aggregation: An [expression] consists of words with various part-of-speech (POS) tags, where some words can be easily mapped to a specific image region. For instance, for the phrase "black and white cat", it is easy to locate "black", "white" and "cat", but "and" is less clear. Thus, we conjecture that VLP models will also perform well when mapping visually recognizable words to the image.

We define \mathcal{V} as a set of POS tags, including nouns, adjectives, verbs, proper nouns, and numerals, which are relatively easy to visualize. An off-the-shelf NLP processing toolbox is utilized to parse the POS tag of each word in the *[expression]* and only those whose tag is included in \mathcal{V} will

remain. We denote \mathcal{W} as the set of *[expression]*'s text tokens. Filtered by \mathcal{V} , original set \mathcal{W} is cut down to \mathcal{W}' , and we further add [CLS] token into \mathcal{W}' when conducting REC because of its general representation of all the tokens. After that $\tilde{\mathbf{G}} \in \mathbb{R}^{I}$ can be calculated as:

$$\tilde{\mathbf{G}} = \mathbb{E}_t(\mathbf{G}'), t \in \mathcal{W}'$$

where \mathbb{E}_t means the average calculation across text dimension and we only calculate the text tokens included in \mathcal{W}' .

We then reshape $\tilde{\mathbf{G}}$ to $\mathbf{H} \in \mathbb{R}^{h \times w}$ having the same size as the input image. VLP can be divided into region-based and end-to-end by whether relying on an external object detector to obtain the visual inputs, which should be applied by different reshaping patterns. We introduce both types into GroundVLP and use different patterns for each type.

Heat-map for Region-Based Models: Region-based VLP models transform the image into a set of region proposals as visual features. Thus, we select a subset of image tokens with high attention values and superimpose their values onto the corresponding image regions to generate the heat-map.

To do this, we sort each element of $\hat{\mathbf{G}}$ in descending order according to their attention values and select the top *m* tokens among them, defined as $\{(v_k, \mathbf{b}_k)\}_1^m$, where v_k and $\mathbf{b}_k = (x_{k1}, y_{k1}, x_{k2}, y_{k2})$ are the corresponding attention value and coordinates of the proposal. The heat map of the region-based model \mathbf{H}_R is computed as:

$$\mathbf{H}_{k}[i,j] = \begin{cases} v_{k} & (i,j) \in \mathbf{b}_{k} \\ 0 & \text{otherwise} \end{cases}$$
(3)

 $\mathbf{H}_{R} = \sum_{k=1}^{m} \mathbf{H}_{k}$, where \mathbf{H}_{k} is a $h \times w$ matrix, R represents Region-based, and the sum calculation for \mathbf{H}_{k} is implemented by element-wise addition.

Heat-map for End-to-End Models: End-to-end VLP models process visual input as a set of patch embeddings with vision-transformer (Dosovitskiy et al. 2020). Its image tokens are a series of image patches. Following (Li et al. 2021; He et al. 2022), we employ a bicubic interpolation on $\tilde{\mathbf{G}}$ to reshape it to $\mathbf{H}_E \in \mathbb{R}^{h \times w}$, where *E* represents End-to-end.

Fusion with Open-Vocabulary Detectors

Having obtained the heat-map $\mathbf{H} \in \mathbb{R}^{h \times w}$, we proceed to generate a set of candidate boxes, calculate the *weighted grades* of regions enclosed by each one, and output the box with the highest grade. We focus on the boxes belonging to the predetermined category merely, simplifying the selection of the final answer box as it reduces the number of candidate boxes. In view of a user should have a specific category in mind when it comes to real-world applications, we employ the ground-truth category during the REC task to mimic the user's input. Meanwhile, we also present an alternative manner to extract the target unit from the *[expression]* as the predicted category when no category is provided.

Category Extraction: Inspired by (Sun et al. 2021), we exploit an NLP toolbox to extract the target unit of the *[expression]* as the predicted category. Specifically, a dependency tree of the *[expression]* is generated by the NLP toolbox, and its rightmost Normal Noun (NN) node of the far



Figure 3: Overview of GroundVLP. The underline words indicate the *[expression]*. The blue arrow lines denote the procedures involving GradCAM, and the pink lines are related to the open-vocabulary object detector. The dashed line indicates that the module may not exist. $\{(s_k, box_k)\}_1^n$ represents an array of confidence scores and bounding boxes detected by OVD based on the given category. The symbol \odot means element-wise multiplication of attention map and gradients, and \oplus means fusing the instances detected by OVD and heat-map. At location ①, we utilize *visual-word attention aggregation*, while at location ②, we employ *weighted grade* to select the answer box.



Figure 4: The difference between the proposed visual word attention aggregation and prior methods

left-bottom Noun Phrase (NP) node is viewed as the predicted category. An example is illustrated in Figure. 5.



Figure 5: An example of category extract. The original query is "*a red and white checkered table with two wooden chairs*". The blue dash rectangle indicates the bottom-left NP node and yellow NN node is the target unit.

Furthermore, we map the extracted predicted categories to the class vocabulary of the evaluation dataset to conform to the ground-truth categories. Let C be the set of classes vocabulary of the evaluation dataset, $c_i \in C$ be one of the classes, and c^p be the extracted predicted category. Next, we employ CLIP (Radford et al. 2021) to embed c_i and c^p as $\mathbf{e}_{c_i} \in \mathbb{R}^D$ and $\mathbf{e}^p \in \mathbb{R}^D$ for projecting them into an uniform embedding space, where D denotes the dimension of CLIP embedding. As per CLIP, the prompt "*a photo of*" is added as a prefix to c_i and c^p before embedding them. We then define c^{map} as the category mapped from c^p to C and sim_i as the similarity between c^p and c_i , which are given as follow:

$$sim_{i} = P(c^{map} = c_{i} \mid c^{p}) = \frac{exp(\mathbf{e}_{c_{i}}^{\top} \mathbf{e}^{p})}{\sum_{j=1}^{|\mathcal{C}|} exp(\mathbf{e}_{c_{j}}^{\top} \mathbf{e}^{p})} \qquad (4)$$
$$c^{map} = c_{i}, \ i = argmax_{i}(sim_{i})$$

Through this equation, each c^p can be mapped to a specific class. For instance, "*table*" in Figure. 5 will be mapped to "*dining table*" when using COCO vocabulary.

Generating Candidate Boxes: Given an image \mathcal{I} , the predetermined category *c* and a score threshold θ , an OVD is allowed to detect *n* instances:

$$\{(s_k, \mathbf{box}_k)\}_1^n = \mathbf{OVD}(c, \theta, \mathcal{I})$$
(5)

$$\mathbf{box}_k = (x_{k1}, y_{k1}, x_{k2}, y_{k2}) \tag{6}$$

$$s_k = P(o_k \in c \mid \mathcal{I}), \ s_k > \theta \tag{7}$$

where o_k represents the object entity included in **box**_k, and s_k denotes the confidence score that o_k belongs to c, which should be greater than θ .

Weighted Grade: A crucial challenge we face now is ascertaining the value of θ . If it is set too low, the superfluous boxes, which do not belong to c, may be included. Conversely, if θ is too high, certain boxes belonging to c may be missed. For this problem, we raise a formula that set a relatively low threshold and let s_k be a weight of r_k , where r_k is defined as the total heat-map value of the region enclosed by \mathbf{box}_k . Through it, we balance both not omitting boxes that pertain to c and preventing low-scoring boxes from disturbing the result. Finally, we calculate g_k which represents a *weighted grade* of (s_k, \mathbf{box}_k) and output \mathbf{box}_{pred} with the highest grade as the prediction of GroundVLP:

$$r_k = \sum_{i=x_{k1}}^{x_{k2}} \sum_{j=y_{k1}}^{y_{k2}} \mathbf{H}[i, j]$$
(8)

$$g_k = \frac{1}{A_k^{\alpha}} \cdot s_k \cdot r_k \tag{9}$$

$$\mathbf{box}_{pred} = \mathbf{box}_k, \ k = argmax_k(g_k) \tag{10}$$

where A_k means the area of \mathbf{box}_k used to avoid the tendency to choose boxes with large areas and α is a hyperparameter.

Experiments

Datasets

Referring Expression Comprehension: We adopt three widely used datasets: RefCOCO, RefCOCO+ (Yu et al. 2016) and RefCOCOg (Mao et al. 2016). RefCOCO and RefCOCO+ are both split into validation, testA, and testB sets, where testA generally contains queries with persons as referring targets and testB contains other types. RefCOCO is described by more spatial information compared to RefCOCO+, whereas RefCOCO+ contains queries using more appearance-related words instead. In contrast, RefCOCOg has longer and more detailed expressions than the other two datasets.

Phrase Grounding: We adopt Flickr30k entities dataset (Plummer et al. 2015) for the task and evaluate the performance in terms of Recall@1, 5. On Flickr30k, a sentence contains several phrases that need to be grounded, each of which may correspond to multiple bounding boxes. Hence, previous researches propose two protocols named ANY-BOX and MERGED-BOX by MDETR (Kamath et al. 2021). In our evaluation, we use the ANY-BOX protocol.

Implementation Details

Selected VLP models and Prompt Templates: \mathcal{I}, \mathcal{T} is defined as the input image and text. We introduce a typical model from both region-based and end-to-end: VinVL (Zhang et al. 2021) and ALBEF (Li et al. 2021). We adopt VinVL-Large and ALBEF-14M as the checkpoints for two models. The input format of VinVL is a triple tuple $\{w, q, v\}$ and can be formed as two types. We adopt the VQA-resemble and prompt the query to adapt it to this format. Specifically, we let $\mathcal{T} =$ "there is a [query]?" on REC task and $\mathcal{T} =$ "[query]?" on phrase grounding, and let q always be "yes". For ALBEF, we prompt $\mathcal{T} =$ "there is a [query]." on phrase grounding, where [query] denotes the query sentence provided by the datasets.

GradCAM Layer: For ALBEF, we use the 3^{rd} layer of the cross-modality encoder for GradCAM. For VinVL, we

use the 20^{th} layer of the cross-modality encoder and select m = 7. All setting is based on tuning on the RefCOCOg validation dataset.

Methods for Category Prediction: We employ Stanza (Qi et al. 2020) to extract predicted category. When testing on RefCOCO/+/g, we map the predicted category to the COCO class via equation 4. For Flickr30k entities, given that its ground-truth category is slightly abstract², we use the predicted category directly. Besides, in order to detect a person entity better, we set $c = \{c^p, person\}$ if the cosine similarity of \mathbf{e}^p and \mathbf{e}^{p*} is greater than 0.9, where \mathbf{e}^{p*} denotes the textual CLIP embedding of "*a photo of person*".

Selected Open-vocabulary Detector: We choose Detic (Zhou et al. 2022) as our open-vocabulary detector (OVD). Other OVDs can also be considered (Li et al. 2022; Zhao et al. 2022). For REC, we set $\alpha = 0.5$, $\theta = 0.15$ when using ground-truth category and $\theta = 0.3$ for predicted category. For phrase grounding, we set $\alpha = 0.25$ and $\theta = 0.15$. If Detic detects no box, we use all proposals as candidate boxes instead. For RefCOCO/+/g, we adopt the proposals from MAttNet (Yu et al. 2018). For Flickr30k entities, we use all proposals detected by Detic.

Compared Baseline: We choose two previous zero-shot REC methods to compare—ReCLIP (Subramanian et al. 2022) and CPT (Yao et al. 2021). CPT masks the regions of each proposal with different colors and predict the color word in " [query] is in [MASK] color", while ReCLIP scores each proposal by using the contrastive scoring ability of CLIP. Furthermore, we construct a CPT-adapted baseline for phrase grounding to compare. A [query] on Flickr30k contains N phrases, denoted as { $[expression]_i$ }^N. We then copy [query] with N times and add "where [expression]_i is in [MASK] color" after ith duplication. We use all proposals detected by Detic and colored blocks (Yao et al. 2021) for CPT-adapted.

Main Results

Referring Expression Comprehension: Table 2 shows the results on RefCOCO/+/g. GroundVLP outperforms other zero-shot methods, especially in the testA split of RefCOCO and RefCOCO+. When using the ground-truth category, GroundVLP is comparable or superior to some non-VLP-based supervised models.

However, it is also noted that there is an inevitable decline in performance when using the predicted category compared to the case of using the ground truth. It can be attributed to several factors: (1) Unclear referring targets: the word of the target unit in the query may not clearly indicate the referring target. For instance, the query "black hat" indicates a person wearing a black hat, and the word "hat" will be extracted. However, "hat" cannot be mapped to "person" exactly, leading to mistakes. (2) Undisciplined grammar: there is a part of coarse queries on RefCOCO/+/g, where the NLP toolbox cannot extract the target unit accurately. For example, "woman red coat", an undisciplined expression of "woman in red coat" or "woman wearing red coat", will cause the

²They are: *people*, *clothing*, *bodyparts*, *animals*, *vehicles*, *instruments*, *scene*, and *other*.

The Thirty-Eighth AAAI	Conference on Artificial	Intelligence	(AAAI-24
			(

]	RefCOC)		R	RefCOCC)+	RefC	OCOg
Model	val	testA	testB		val	testA	testB	val	test
Supervised SOTA (Yan et al. 2023)	92.64	94.33	91.46		85.24	89.63	79.79	88.73	89.37
Supervised method w/o VLP	76.41	01.01	70.00		(70.00	<i></i>	65.07	(())
TransVG (Deng et al. 2019a)	76.41 81.02	81.21 82.72	70.09 78.35		66.46 64.82	72.02	57.52 56.94	65.87	66.44 67.73
Other zero-shot									
CPT-Seg w/ VinVL (Yao et al. 2021)	32.2	36.1	30.3		31.9	35.2	28.8	36.7	36.5
ReCLIP w/ relations (Subramanian et al. 2022)	45.78	46.10	47.07		47.87	50.10	45.10	59.33	59.01
Ours using predicted category									
GroundVLP w/ ALBEF	52.58	61.30	43.53		56.38	64.77	47.43	64.30	63.54
GroundVLP w/ VinVL	<u>59.05</u>	<u>69.21</u>	48.71		61.80	70.56	50.97	69.08	68.98
Ours using ground-truth category									
GroundVLP w/ ALBEF	58.22	65.48	<u>49.75</u>		<u>63.57</u>	72.02	53.57	69.63	<u>69.15</u>
GroundVLP w/ VinVL	65.01	73.50	55.01		68.87	78.10	57.31	74.73	75.03

Table 2: Accuracy (%) on referral expression comprehension datasets. We show both the results of GroundVLP using the predicted category and that using the ground-truth category. The best zero-shot accuracy in each column is in bold, and the second best is underlined. Supervised SOTA refers to UNINEXT (Yan et al. 2023).

	Flickr	30k val	Flickr.	30k test
Model	R@1	R@5	R@1	R@5
Supervised method with VLP				
MDETR-ENB5 (Kamath et al. 2021)	83.6	93.4	84.3	93.9
GLIP-L (Li et al. 2022)	86.7	96.4	87.1	96.9
Supervised method w/o VLP				
BAN (Kim, Jun, and Zhang 2018)	-	-	69.69	84.22
Zero-shot method				
CPT-adapted (Yao et al. 2021)	27.06	61.78	27.57	61.55
GroundVLP w/ ALBEF(Ours)	63.76	75.02	63.89	74.80
GroundVLP w/ VinVL(Ours)	63.89	74.53	64.19	74.57

Table 3: Accuracy (%) on the Flickr30k entities dataset. We compare GroundVLP with VLP-based, non-VLP-based supervised methods and prior zero-shot method. The best zero-shot accuracy in each column is in bold. Note that we only use the predicted category during this task.

NLP toolbox to regard "*woman*" as a noun adjective used to describe "*red*" and treat "*coat*" as the target instead. (3) No target in query: there are a few queries consisting of pure spatial information, not containing the referring target (e.g. "*left*", "*the closest to you*"). We argue that these datasets inherently include bias and noise, which makes it difficult to accurately map the predicted category to a COCO class, resulting in a decline in performance.

Phrase Grounding: Table 3 shows the results on the Flickr30k entities dataset. GroundVLP is far ahead of CPT for the R@1 score, outperforming it by 38.29% and 37.79% in the val and test split. Moreover, GroundVLP performs comparably to or even better than some non-VLP-based supervised approaches. Finally, it is noticed that we only use predicted category for phrase grounding, demonstrating the effectiveness of the proposed fusion method for grounding

tasks and the disciplined expression on the Flickr30k dataset that is beneficial for our predicted category extraction.

Model	RefCOCO	RefCOCO+	RefCOCOg	Difference
GroundVLP	•			
w/ TCL	59.59	64.96	70.92	5.37
w/ PTP	56.81	61.17	68.53	4.36
w/ Lxmert	63.14	60.99	67.99	-2.15

Table 4: Accuracy (%) of using other VLP models. We report the difference of the score on RefCOCO+ minus that on RefCOCO for each model in column *Difference*. All datasets in the table indicate their val split and all results on RefCO-CO/+/g are obtained by using the ground-truth category.

Extending to Other VLP Models

In order to verify the versatility of our method, we further incorporate more VLP models into GroundVLP. We present the results using TCL (Yang et al. 2022), PTPčitewang2022position and Lxmert (Tan and Bansal 2019) in Table 4, among which TCL and PTP belong to end-to-end and Lxmert belongs to region-based. We found that GroundVLP with all models outperforms other zeroshot methods recorded in Table 2, showing its versatility that could be applied to various VLP models effectively. We also report the difference in accuracy between the models obtained on RefCOCO and RefCOCO+ in the last column. As previously mentioned, RefCOCO includes more spatial information while RefCOCO+ is composed of more appearance-related queries. Thus, the Difference could indicate whether the model is better at position or appearance attributes. It can be observed that Lxmert is better at spatial information while the other two end-to-end models are the opposite. We conjecture that the preliminary modeling of objects by the OD module could facilitate the understanding of the visual context for the region-based models and render it position-sensitive. (Yao et al. 2022; Wang et al. 2022).

Ablation Studies

The evaluation in this section uses the val split of all datasets and ground-truth category on RefCOCO/+/g datasets if there is no supplementary statement.

Different Assembly of θ and s_k : Table 5 investigates the effect of the value of θ and the usage of the *weighted grade*. It can be observed that a low threshold for OVD ($\theta = 0.15$) leads to the detection of superfluous boxes, which will impair the performance if we calculate r_k as grade directly. For this condition, our proposed weighted grade considering both s_k and r_k could effectively mitigate the interference from redundant boxes. A high threshold with the weighted grade ($\theta = 0.5$ and s_k is used), on the other hand, uplifts the quality of detected boxes but probably excludes the answer box. Thus, the optimal assembly is a relatively low threshold with the *weighted grade*. In Figure 6, we illustrate the impact of using the *weighted grade* on the final results. It can be observed that when employing the *weighted grade*, GroundVLP produces the correct answer.

Backbone	θ	Use s_k	RefC	RefC+	RefCg	Flickr30k
ALREE	0.15	(55.13	60.12 62.64	66.52 68.24	62.88
ALDEF	0.30	\checkmark	57.31 58.22	63.57	69.63	63.76
	0.15		63.82	67.10	73.43	63.33
VinVL	0.50	\checkmark	64.17	68.00	72.61	61.01
	0.15	\checkmark	65.01	68.87	74.73	63.89

Table 5: Ablation study on the value setting of θ and whether using s_k as the weight of r_k .

Backbone	Candidate	RefC	RefC+	RefCg	Flickr30k
ALBEF	all pred gt	49.51 52.58 58.22	53.80 56.38 63.57	55.11 64.30 69.63	55.05 63.76
VinVL	all pred gt	58.39 59.05 65.01	60.87 61.80 68.87	63.40 69.08 74.73	53.23 63.89

Table 6: Ablation study on the type of candidate boxes. *all* means using all proposals, *pred* means using predicted category, and *gt* means using ground-truth category.

Type of Candidate Boxes: Table 6 investigates the influence of different candidate boxes. The performance is enhanced after shrinking the number of candidate boxes by a predetermined category. We observe that the improvements on RefCOCO and RefCOCO+ were not as pronounced as those on RefCOCOg and Flickr30k entities when using the predicted category. Combined with the aforementioned brief



GradCAM + Detected Boxes by OVD



Figure 6: An example of the *weighted grade*. If s_k is not used, GroundVLP would achieve an error prediction. This query is from RefCOCOg val set.

about the datasets, it is evident that the extracting for a category is more precise on datasets with concrete and disciplined expressions, such as RefCOCOg and Flickr30k entities that are more practical and common in the real world.

Backbone	Use vw	RefC	RefC+	RefCg	Flickr30k
ALBEF	\checkmark	57.34 58.22	63.02 63.57	69.34 69.63	63.47 63.76
PTP	\checkmark	55.42 56.81	59.43 61.17	67.89 68.53	-

Table 7: Ablation study on the type of aggregating attention. vw denotes visual word attention aggregation.

Type of Aggregating Attention: Table 7 investigates the improvement of using *visual word attention aggregation*. The compared baseline is implemented by averaging the attention scores across all text tokens (i.e. Figure 4 (b)). PTP (Wang et al. 2022) is another VLP we incorporate. It

could be observed that our method achieves better performance, especially on PTP, indicating that filtering based on visually recognizable words could facilitate models' text-toimage mapping, which is beneficial for grounding tasks.

Model	RefCOCO	RefCOCO+	RefCOCOg
Detic	6.65	6.69	8.52
GroundVLP w/ ALBEF	58.22	63.57	69.63

Table 8: Comparison of using Detic solely and GroundVLP with ALBEF.

Applying OVD for visual grounding: In order to figure out whether an OVD could be applied to visual grounding, we exclusively utilized Detic for testing on RefCO-CO/+/g. Initially, we fed the ground-truth category into Detic to produce a set of candidate boxes. Subsequently, we input [query] to Detic, taking into account only these candidate boxes when evaluating the similarity score between each proposal and the text embedding of [query]. The box with the highest score is the final output. The results in Table 8 demonstrate that an OVD struggles to grasp the intricate semantic information of visual grounding without the semantic insight offered by ITM.

Model Type	val	testA	testB
GroundVLP	(0.57	72.02	50.57
W/ ALBEF	63.57	72.02	53.57
w/ ALBEF [†]	68.40	77.07	55.82

Table 9: The comparison of GroundVLP with pre-trained model and fine-tuned model. [†] means the fine-tuned model.

Fine-tuning GroundVLP

We note that GroundVLP can be further fine-tuned with annotation data to enhance the performance. Using the Ref-COCO+ training set, we paired queries with their images as image-text pairs, fine-tuned ALBEF using ITM loss and observed performance improvements, as shown in Table 9.

Related Work

Visual Grounding. The widely used pipelines to resolve visual grounding can be broadly grouped into *two-stage* (Yu et al. 2018; Liu et al. 2019c; Sun et al. 2021) and *one-stage* (Yang et al. 2019; Liao et al. 2020; Deng et al. 2021), where *two-stage* methods exploit a proposal-query matching paradigm while *one-stage* methods generate the answer box with end-to-end. Thanks to the emergence of self-supervised pre-training, the results on visual grounding have been improved substantially by pre-trained models (Chen et al. 2020a; Kamath et al. 2021; Li et al. 2022). Additionally, though pre-trained with objectives not related to grounding task, there exist various pre-trained models having a strong capacity for vision-language alignment (Zhang et al. 2021; Radford et al. 2021). Thus, the works utilizing

their strengths to conduct zero-shot visual grounding were proposed, such as CPT and ReCLIP (Yao et al. 2021; Subramanian et al. 2022). It is noteworthy that there is another definition of zero-shot different from us, which predicates the objects that are unseen during training and still need to be trained on a grounding dataset (Sadhu, Chen, and Nevatia 2019; Shi et al. 2022). CPT, ReCLIP and us were not trained on any grounding dataset while could carry grounding tasks via triggering the capacity of VLP.

GradCAM for Grounding. GradCAM (Selvaraju et al. 2017) is proposed to visualize the regions that the model focuses on for a specific output head. When it is used in the VLP models' ITM head, GradCAM could represent a modality mapping from text to image, which is adapted for visual grounding. Therefore, it was employed in REC with a weakly-supervised setting (Li et al. 2021; He et al. 2022) and in robot 3D-navigation (Ha and Song 2022) by some works. Different from theirs, our method (1) uses the presented visual word attention aggregation to optimize the text-toimage mapping, (2) could generate a heat-map for any VLP models pre-trained with ITM by the approach described in section, which is a universal method, and (3) introduce the weighted grade to improve the matching between the heatmap and candidate boxes instead of calculating the heat-map values enclosed by boxes directly as other methods.

Conclusion

We note the ready availability of image-text pairs and object detection data and then present GroundVLP, a zeroshot method for visual grounding via combining the models trained using these datasets. GroundVLP employ GradCAM for a VLP model to identify the image regions, introduce an open-vocabulary object detector to generate the object proposals that belong to a given category, and fuse these two components via the *weighted grade*. Experiments show the state-of-the-art performance of GroundVLP, which outperforms other zero-shot methods and is comparable to some non-VLP-based models. In the future, we plan to refine the method of category prediction and also continue to introduce diverse VLP models and open-vocabulary object detectors for better performance and application of zero-shot grounding models such as embodied agents.

Limitations

Despite the strong accuracy that GroundVLP achieves, there are still some potential limitations. GroundVLP may inadvertently inherit biases or errors presented in those foundational models. However, we note that both VLP and OVD serve as plug-and-play modules in our implementation. This modular design means that GroundVLP stands to benefit from advancements in both of these areas. Should there be a more robust or improved foundational model in the future, it can be seamlessly integrated into our framework to replace any prior model exhibiting errors or biases. Furthermore, as shown in Table 9, GroundVLP achieves a performance improvement after fine-tuning its VLP backbone, showing that the employed foundation model could be fine-tuned to effectively alleviate these deficiencies.

Acknowledgements

This research is supported by National Key R&D Program of China under grant (2022YFF0902600) and Key R&D Program of Zhejiang under grant (2023C01048).

References

Changpinyo, S.; Sharma, P.; Ding, N.; and Soricut, R. 2021. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3558–3568.

Chen, Y.-C.; Li, L.; Yu, L.; El Kholy, A.; Ahmed, F.; Gan, Z.; Cheng, Y.; and Liu, J. 2020a. Uniter: Universal image-text representation learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXX*, 104–120. Springer.

Chen, Z.; Wang, P.; Ma, L.; Wong, K.-Y. K.; and Wu, Q. 2020b. Cops-ref: A new dataset and task on compositional referring expression comprehension. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 10086–10095.

Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, 248–255. Ieee.

Deng, J.; Yang, Z.; Chen, T.; Zhou, W.; and Li, H. 2021. Transvg: End-to-end visual grounding with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1769–1779.

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. In *International Conference on Learning Representations*.

Gu, X.; Lin, T.-Y.; Kuo, W.; and Cui, Y. 2021. Open-Vocabulary 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.

Ha, H.; and Song, S. 2022. Semantic abstraction: Openworld 3d scene understanding from 2d vision-language models. In *Conference on Robot Learning*.

He, S.; Guo, T.; Dai, T.; Qiao, R.; Wu, C.; Shu, X.; and Ren, B. 2022. VLMAE: Vision-Language Masked Autoencoder. *arXiv preprint arXiv:2208.09374*.

Honnibal, M.; and Johnson, M. 2015. An improved nonmonotonic transition system for dependency parsing. In *Proceedings of the 2015 conference on empirical methods in natural language processing*, 1373–1378.

Kamath, A.; Singh, M.; LeCun, Y.; Synnaeve, G.; Misra, I.; and Carion, N. 2021. Mdetr-modulated detection for endto-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1780–1790. Kim, J.-H.; Jun, J.; and Zhang, B.-T. 2018. Bilinear attention networks. *Advances in neural information processing systems*, 31.

Kim, W.; Son, B.; and Kim, I. 2021. Vilt: Vision-andlanguage transformer without convolution or region supervision. In *International Conference on Machine Learning*, 5583–5594. PMLR.

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

Kuznetsova, A.; Rom, H.; Alldrin, N.; Uijlings, J.; Krasin, I.; Pont-Tuset, J.; Kamali, S.; Popov, S.; Malloci, M.; Kolesnikov, A.; et al. 2020. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7): 1956–1981.

Li, J.; Selvaraju, R.; Gotmare, A.; Joty, S.; Xiong, C.; and Hoi, S. C. H. 2021. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems*, 34: 9694–9705.

Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10965–10975.

Liao, Y.; Liu, S.; Li, G.; Wang, F.; Chen, Y.; Qian, C.; and Li, B. 2020. A real-time cross-modality correlation filtering method for referring expression comprehension. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10880–10889.

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, D.; Zhang, H.; Wu, F.; and Zha, Z.-J. 2019a. Learning to assemble neural module tree networks for visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4673–4682.

Liu, R.; Liu, C.; Bai, Y.; and Yuille, A. L. 2019b. Clevr-ref+: Diagnosing visual reasoning with referring expressions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4185–4194.

Liu, X.; Li, L.; Wang, S.; Zha, Z.-J.; Su, L.; and Huang, Q. 2019c. Knowledge-guided pairwise reconstruction network for weakly supervised referring expression grounding. In *Proceedings of the 27th ACM International Conference on Multimedia*, 539–547.

Mao, J.; Huang, J.; Toshev, A.; Camburu, O.; Yuille, A. L.; and Murphy, K. 2016. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 11– 20. Plummer, B. A.; Wang, L.; Cervantes, C. M.; Caicedo, J. C.; Hockenmaier, J.; and Lazebnik, S. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, 2641–2649.

Qi, P.; Zhang, Y.; Zhang, Y.; Bolton, J.; and Manning, C. D. 2020. Stanza: A Python Natural Language Processing Toolkit for Many Human Languages. In *ACL (demo)*.

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.

Sadhu, A.; Chen, K.; and Nevatia, R. 2019. Zero-shot grounding of objects from natural language queries. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4694–4703.

Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35: 25278–25294.

Schuhmann, C.; Vencu, R.; Beaumont, R.; Kaczmarczyk, R.; Mullis, C.; Katta, A.; Coombes, T.; Jitsev, J.; and Komatsuzaki, A. 2021. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*.

Selvaraju, R. R.; Cogswell, M.; Das, A.; Vedantam, R.; Parikh, D.; and Batra, D. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, 618–626.

Shi, Z.; Shen, Y.; Jin, H.; and Zhu, X. 2022. Improving Zero-Shot Phrase Grounding via Reasoning on External Knowledge and Spatial Relations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2253–2261.

Subramanian, S.; Merrill, W.; Darrell, T.; Gardner, M.; Singh, S.; and Rohrbach, A. 2022. ReCLIP: A Strong Zero-Shot Baseline for Referring Expression Comprehension. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5198–5215.

Sun, M.; Xiao, J.; Lim, E. G.; Liu, S.; and Goulermas, J. Y. 2021. Discriminative triad matching and reconstruction for weakly referring expression grounding. *IEEE transactions on pattern analysis and machine intelligence*, 43(11): 4189–4195.

Tan, H.; and Bansal, M. 2019. LXMERT: Learning Cross-Modality Encoder Representations from Transformers. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), 5100–5111.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. At-

tention is all you need. Advances in neural information processing systems, 30.

Wang, A. J.; Zhou, P.; Shou, M. Z.; and Yan, S. 2022. Position-guided Text Prompt for Vision-Language Pre-training. *arXiv preprint arXiv:2212.09737*.

Yan, B.; Jiang, Y.; Wu, J.; Wang, D.; Luo, P.; Yuan, Z.; and Lu, H. 2023. Universal instance perception as object discovery and retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15325–15336.

Yang, J.; Duan, J.; Tran, S.; Xu, Y.; Chanda, S.; Chen, L.; Zeng, B.; Chilimbi, T.; and Huang, J. 2022. Vision-language pre-training with triple contrastive learning. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 15671–15680.

Yang, Z.; Gong, B.; Wang, L.; Huang, W.; Yu, D.; and Luo, J. 2019. A fast and accurate one-stage approach to visual grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4683–4693.

Yao, Y.; Chen, Q.; Zhang, A.; Ji, W.; Liu, Z.; Chua, T.-S.; and Sun, M. 2022. PEVL: Position-enhanced pre-training and prompt tuning for vision-language models. *arXiv* preprint arXiv:2205.11169.

Yao, Y.; Zhang, A.; Zhang, Z.; Liu, Z.; Chua, T.-S.; and Sun, M. 2021. Cpt: Colorful prompt tuning for pre-trained visionlanguage models. *arXiv preprint arXiv:2109.11797*.

Yu, L.; Lin, Z.; Shen, X.; Yang, J.; Lu, X.; Bansal, M.; and Berg, T. L. 2018. Mattnet: Modular attention network for referring expression comprehension. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1307–1315.

Yu, L.; Poirson, P.; Yang, S.; Berg, A. C.; and Berg, T. L. 2016. Modeling context in referring expressions. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II 14*, 69–85. Springer.

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, P.; Li, X.; Hu, X.; Yang, J.; Zhang, L.; Wang, L.; Choi, Y.; and Gao, J. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5579–5588.

Zhao, T.; Liu, P.; Lu, X.; and Lee, K. 2022. OmDet: Language-Aware Object Detection with Large-scale Vision-Language Multi-dataset Pre-training. *arXiv preprint arXiv:2209.05946*.

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.