

# Object-Centric Open-Vocabulary Image Retrieval with Aggregated Features

Hila Levi\*<sup>1</sup>  
hila.levi@gm.com

Guy Heller\*<sup>1</sup>  
guy.heller@gm.com

Dan Levi<sup>1</sup>  
dan.levi@gm.com

Ethan Fetaya<sup>2</sup>  
ethan.fetaya@biu.ac.il

<sup>1</sup> General Motors, RND, Israel

<sup>2</sup> Bar-Ilan University, Israel

---


## Abstract

The task of open-vocabulary object-centric image retrieval involves the retrieval of images containing a specified object of interest, delineated by an open-set text query. As working on large image datasets becomes standard, solving this task efficiently has gained significant practical importance. Applications include targeted performance analysis of retrieved images using ad-hoc queries and hard example mining during training. Recent advancements in contrastive-based open vocabulary systems have yielded remarkable breakthroughs, facilitating large-scale open vocabulary image retrieval. However, these approaches use a single global embedding per image, thereby constraining the system's ability to retrieve images containing relatively small object instances. Alternatively, incorporating local embeddings from detection pipelines faces scalability challenges, making it unsuitable for retrieval from large databases.

In this work, we present a simple yet effective approach to object-centric open-vocabulary image retrieval. Our approach aggregates dense embeddings extracted from CLIP into a compact representation, essentially combining the scalability of image retrieval pipelines with the object identification capabilities of dense detection methods. We show the effectiveness of our scheme to the task by achieving significantly better results than global feature approaches on three datasets, increasing accuracy by up to 15 mAP points. We further integrate our scheme into a large scale retrieval framework and demonstrate our method's advantages in terms of scalability and interpretability.

## 1 Introduction

Retrieving images which include specific objects, according to an on-demand open-set text query, is an important task in computer vision with numerous practical applications. Performing such targeted searches, especially over unlabeled rare concepts, can be used, for example, to analyze the performance of an already trained system, to mine hard examples




**Figure 1: Overall retrieval framework:** Two-stage operation: Offline, we generate per-image global / aggregated embeddings via (a) CLIP or (b) Cluster-CLIP. Subsequently, an online stage enables on-demand retrieval using textual queries. Retrieval is performed by ranking text-image similarity using cosine distance in a shared embeddings space, with potential acceleration using Large Scale Index (details in Section 3.4). This approach accommodates any dual-encoder architecture. We present top retrieved images for the query ‘bird’ using Cluster-CLIP features.

during training, or to guide the process of gathering the data for manual annotations. Relevant use cases vary in scale from web-scale search (e.g., Google Bard, Microsoft Bing) to searching in application specific datasets (e.g., e-commerce, automotive, medical applications). In both cases, scalability and efficiency play critical roles in adopting the technology.

Despite its importance, literature lacks direct references to this task. One possible reason might be the task complexity: common vision-language (VL) representation of open-set objects was hard to achieve until the accelerated evolution of contrastive-based open-vocabulary models. These models (e.g., CLIP [67], Florence [74], Coca [73]), trained on web-scale image-caption data, produce a common embedding space for global image and caption representations, maximized directly via training. Retrieval is then performed by ranking the text-image similarity using cosine distance in the common embeddings space and can be scaled by using frameworks as schematically illustrated in Figure 1. However, empirical experiments with CLIP using open-set object queries on the similar task of object-centric retrieval produce less satisfactory results (see Figure 2 and direct quantitative comparison in Section 4). In particular, performance degrades with the increase in image complexity and the decrease in relative objects sizes. Alternatively, the use of pure detections from SoTA open-vocabulary detection frameworks (e.g., OwViT [47]) is more compatible with cluttered images but is considered ill-suited for retrieval tasks; Running the detection model on each image for each query will require an enormous amount of computational resources, while precomputing and saving its internal dense embeddings will require orders of magnitude more storage.

In this work, we visit the task of object-centric open-vocabulary image retrieval and present a simple approach to tackle it based on the complementary advantages of classification and detection open-vocabulary frameworks. The main challenge in this task is to combine the scalability of image retrieval pipelines, which can operate on huge datasets, with object-level processing of detection systems that commonly operate on a single image at a time. A second challenge is to preserve the good zero-shot accuracy obtained from web-scaled pretrained open vocabulary schemes, which holds significant importance for the task. We address these challenges by exploring the use of aggregated features in two steps.

As a first step towards the solution, we explore the use of local features, extracting dense embeddings from an intermediate feature-map of CLIP vision encoder and manipulating them, keeping CLIP visual-language association as is (abbreviated as Dense-CLIP). Retrieval is performed by ranking according to the maximum similarity in each image. In the experiments, we show that using Dense-CLIP to represent images achieves on-par results to the OwViT baseline on all populations and significantly better results on rare objects queries,



**Figure 2: Retrieval examples** of (a) CLIP vs. (b) Cluster-CLIP with their corresponding text queries. CLIP (single global embedding) top-scored images mainly include large objects, whereas Cluster-CLIP, which aggregates dense embeddings, also retrieves cluttered images with relatively small objects (highlighted in red circles).

with less than half the embeddings per image. With respect to CLIP, which uses one global feature to represent each image, results are significantly increased by up to 12 mAP points. Note that while Dense-CLIP retrieval results are impressive, its use alone is not enough since it enlarges the search space by a large extent, thus impairing potential scalability.

To address scalability requirements, we explore the use of aggregated visual features by introducing Cluster-CLIP (Fig. 1). Cluster-CLIP aggregates Dense-CLIP’s dense embeddings into sparse representatives with distinct local semantics (aggregation module in Fig. 1). We have examined a variety of aggregation methods that require no training to the task. By manipulating hyper-parameters of each aggregation method, we check various points on the accuracy-efficiency tradeoff (see Section 4). Interestingly, we found out that Cluster-CLIP shows, in some cases, higher retrieval rates than Dense-CLIP (and up to 15 mAP points increase with respect to CLIP) for the small number of 10-50 representatives per image. The investigation is meaningful in that, as far as we know, it is the first work to explore and quantify the range between leveraging open vocabulary features as a global image representation and as local dense embeddings.

To summarize our contributions:

1. We visit the task of object-centric open-vocabulary image retrieval and introduce Dense-CLIP, which uses CLIP’s local features, keeping its original zero-shot properties.
2. We present Cluster-CLIP which enables scalability via a compact representation.
3. We show the effectiveness of our approaches by achieving significantly better results compared with a global feature (CLIP) on three datasets: COCO [40], LVIS [20], and nuImages [9], increasing retrieval accuracy by up to 15 points.
4. We integrate Cluster-CLIP into a retrieval framework, showcasing its scalability and presenting empirical evidence of its efficacy through plausible results.

## 2 Related Literature

Our work is closely related to research on instance retrieval frameworks, cross-modal retrieval, and open-vocabulary VL models. We briefly review related work in these domains.

**Instance Retrieval Frameworks.** Way before the deep learning era, large scale image-to-image retrieval research was dominated by mining for large-sized geographical landmarks applications, presenting methods that use local handcrafted features [6, 43] to query from large databases of descriptors [43, 46, 52], sometimes followed by aggregation techniques such as BoW [60], FV [27], and VLAD [26]. Later, at the beginning of the learning era, the

use of learnable CNN’s global descriptor [2, 4, 18, 52], local descriptors [4, 23, 51], or a combination of both [9, 59, 70] has been proposed, introducing accuracy gains with respect to the use of handcrafted features. Compared to these works, our application focuses on the open-vocabulary text-to-image retrieval task, yet draws inspiration from the use of local descriptors and aggregation methods that lacks in current cross modal retrieval literature.


**Cross-Modal Retrieval.** Aligning vision and language has a long-standing history of research (e.g., [10, 15, 17, 52, 48]). Among early cross-modal retrieval works, one line of research suggests learning a shared VL embedding space via a dual-path architecture [10, 13, 14, 51, 57, 51, 63]. Exemplars include VSRN [57], which employs bottom-up attention and visual semantic reasoning, and PCME [10], which represents samples as probabilistic distributions. Another approach suggests using joint image-text modules [24, 28, 55, 41, 66, 68, 76]. However, this necessitates processing all images per new text query, limiting scalability for cross-modal retrieval. A commonality across these early works is that they learn vision-language (VL) alignment using medium-sized datasets, like MS-COCO [40] or Flickr30K [72], resulting in decreased performance on novel datasets.

Recently, a novel paradigm that performs Visual-Language pre-training on larger datasets has emerged [10, 25, 36, 38, 42, 52]. This paradigm has catalyzed a transformative shift in the field driven by VL models pre-trained on massive web-scale corpora (CLIP [57], ALIGN [29], and others [65, 67, 73, 74, 75]), allowing high zero-shot performance on diverse datasets and enabling open-vocabulary retrieval. As before the VL pre-training era, some of these works employ a joint image-text module [10, 36, 38, 57] that limits their scalability for cross-modal retrieval. In contrast, a more relevant line of research adopts a dual-stream approach that generates a shared embedding space for global image and caption representations, providing a straightforward solution for caption-based image retrieval tasks [29, 39, 57]. This line of work can be easily integrated into the retrieval scheme presented in Fig. 1. However, our study reveals that directly using global image representation for object-centric retrieval yields sub-optimal results, particularly for small-sized objects.

**Open-Vocabulary Region-Level Representation.** This line of research employs VL pre-training for open-vocabulary object-centric tasks, such as object detection, by using region-level representation and supervision [79]; Out of which, most relevant to us are two stream approaches [8, 32, 42, 45, 69, 78]. Some of these works utilize a class-agnostic object detector for region proposals [16, 19, 77], restricting the system’s open-vocabulary capabilities. In contrast, others divide the image into a grid [49, 58, 71]. A representative example of the latter is OwlViT [47], which is a SoTA dual-encoder open-vocabulary object detector. However, using an object detection framework directly is unfeasible for retrieval due to the need to process every image in the dataset for each new query, leading to significant computational demands. Alternatively, storing the internal embeddings of each image for the detection framework would entail a substantial increase in storage requirements, making it impractical. As our method is compatible with any dual-encoder VL open-vocabulary model, it presents a solution for leveraging progress in this line of research. For demonstration, we employ Cluster-CLIP architecture with the OwlViT backbone (see supplementary materials).

### 3 Method

The first parts of this section focus on the creation of the image embeddings and are organized as follows: Section 3.1 briefly reviews the implementation details of CLIP [57], which



**Figure 3: Simplified overview of image encoders:** Encoders receive image  $I \in R^{H \times W \times 3}$  as input. Dense-CLIP uses only the value (V) and output (O) linear layers of CLIP’s Multi-Head Attention module. Cluster-CLIP clusters Dense-CLIP output and transfers a single representative per cluster.

produces one embedding per image. Following, Section 3.2 presents Dense-CLIP<sup>1</sup>, which creates dense embeddings while keeping the same embedding space as CLIP. Finally, Section 3.3 presents Cluster-CLIP (which sets an aggregation module on top of Dense-CLIP) and counts several clustering instantiations to create a compact representation. The last part of the section (3.4) presents the whole retrieval framework, as illustrated in Figure 1.

### 3.1 Preliminaries: CLIP

CLIP [57] is a vision-language model, pretrained on massive amounts of web-scaled image-caption data via contrastive learning. It consists of two separate streams: a text-encoder, implemented as a transformer [64], and a vision-encoder, implemented either by a modified ResNet backbone [22] or by a ViT backbone [17]. In both cases, the last layer is a multi-head attention layer, which sums information from all the pixels in the input tensor weighted by their similarity to a query vector and projects it by an output linear layer (Fig. 3, left). The multi-head attention layer for the modified ResNet backbone is formulated as:

$$y = out \left( concat \left[ y^1, y^2, \dots, y^M \right] \right), \quad y^m = softmax \left( \frac{q^m(\bar{x}) \cdot k^m(X)^T}{\sqrt{C_q}} \right) v^m(X) \quad (1)$$

Here  $X \in R^{K \times C_e}$  is the input tensor and  $y \in R^{1 \times C_o}$  is the output embedding (one global vector of  $C_o$  channels in the output representation of CLIP).  $\bar{x} = \frac{1}{K} \sum_{i=1}^K x_i$  represents the average of all spatial locations,  $\{x_i\}_{i=1}^K$ , of the input tensor  $X$ .  $q^m: R^{C_e} \rightarrow R^{C_q}$ ,  $k^m: R^{C_e} \rightarrow R^{C_q}$ ,  $v^m: R^{C_e} \rightarrow R^{C_v}$  and  $out: R^{M C_v} \rightarrow R^{C_o}$  are respectively the query, key, value and output linear layers, where  $m \in \{1 \dots M\}$  is the index of a specific head in the multi-head architecture and  $\sqrt{C_q}$  is a normalization factor. The described attention layer, which sums information from all spatial locations of the input tensor, is optimized to capture the "average" semantics in the image as forced to match the embedding of the corresponding caption. A reasonable hypothesis, also raised in [17], is that the average semantic is built upon local semantics, already captured by the spatial locations at the input to the attention layer.

### 3.2 Dense-CLIP Module

Inspired by the use of dense embeddings in detection frameworks and based on the above hypothesis, we use the following reformulation of CLIP multi-head attention layer:

$$y_i = out \left( concat \left[ y_i^1, y_i^2, \dots, y_i^M \right] \right), \quad y_i^m = v^m(x_i) \quad (2)$$

<sup>1</sup>we note that this is similar yet distinct from the denseCLIP defined in [23] and similarly used in [58] as they use a fine-tuning process and address different tasks.

Compared to the previous formulation, here the output embedding  $Y \in R^{K \times C_o}$  is a tensor, and  $y_i$  is the representation of its  $i$ 'th spatial pixel:  $\{y_i \in R^{1 \times C_o}\}_{i=1}^K$ . This reformulation, implemented by removing the query and key linear layers and implementing the value and output linear layers as  $1 \times 1$  convolutional layers (with the same weights), essentially creates dense patch embeddings with the same output space as CLIP (Fig. 3, middle). We use it as is, without fine-tuning. Empirical results at Section 4 reveal that Dense-CLIP achieves on-par retrieval accuracy as SoTA detection frameworks (i.e. OwlViT [47]) with fewer representatives. Given this finding, we next explore whether we can further reduce the number of representatives and to which extent.

### 3.3 Cluster-CLIP

To improve scalability, we introduce Cluster-CLIP, which produces aggregated embeddings by an additional aggregation module on top of Dense-CLIP embeddings. The aggregation module first clusters the dense features predicted by Dense-CLIP ( $Y = \{y_i\}_{i=1}^K$ ) within  $N$  clusters, denoted as  $\{C_j\}_{j=1}^N$ , where  $C_j \subset Y$  and  $N \ll K$ . Then, it transfers one representative embedding per cluster (the average of the embeddings within the cluster) for future retrieval use (see Figures 1 and 3 right). Note that the aggregation module is generally defined, and many clustering variants fit into it. We empirically examined a variety of clustering mechanisms and present here the most effective methods (while deferring the complete list and full implementation details to the Supplementary Materials):

**K-Means (Cluster-CLIP, K.M.).** In this method, we perform K-Means clustering on top of each image's dense embeddings. Once clustered, the representatives of an image are the clusters' centroids. Examples of interest that demonstrate grouping by semantic similarity can be seen in Figure 4.

**Agglomerative Clustering (Cluster-CLIP, AG.).** This method applies Agglomerative clustering (hierarchical clustering using a bottom-up approach) where the average embedding from each cluster is declared as the cluster representative. Results are presented with connectivity constraints (AG-T) and without (AG-F).

**Region Proposals (Cluster-CLIP, R.P.).** In this method, we use Segment Anything (SAM) [33] to segment each image. Cluster-CLIP is then provided with both the image and the masks, with the masks serving as guidance for aggregating the dense embeddings.

### 3.4 Overall Framework

A schematic illustration of our overall framework is shown in Fig. 1. Our framework consists of two separate stages. The first stage receives a dataset of images as input and uses an image encoder from a vision-language model to create embeddings through sequential processing, followed by indexing to allow a quick approximate nearest neighbour (ANN) search. In our experiments, we considered embeddings based on global embedding strategy (i.e., from CLIP), local embedding strategies (OwlViT, Ours: Dense-CLIP), and aggregations (Ours: Cluster-CLIP). The second stage receives two inputs: the large scale index created in the first stage and a textual object query wrapped by textual prompt/s. Processing includes applying the corresponding text encoder (of the same vision-language model) to create a search vector, followed by an ANN search to get a final list of ranked images. Any dual-encoder vision-language model can be integrated into this scheme. Notice that whereas the first stage is computationally expensive (i.e., in terms of number of FLOPS), it

is executed only once. The above partitioning into two separate flows essentially allows us to refer to the index as given, enabling on-line interactive retrieval at the second stage. This is a significant advantage compared to object detection pipelines. The performance of the on-line system is dominated by the tradeoff between the quality of representation, the number of representatives per image, and the parameters of the ANN search. In our experiments, we follow common practice and exclude the last factor, as it is not the scope of our work, and report quantitative results based on ranking all images by cosine similarity.

## 4 Experiments

We evaluated our approach on the task of object-centric image retrieval on three publicly available datasets (COCO [40], LVIS [20] and nuImages[2]), using datasets’ semantic categories as queries. We compared the performance, in terms of retrieval accuracy vs. number of embeddings, to global and local features from existing methods. Our approach demonstrates increased performance with a small number of representatives per image, thereby allowing scalability with better retrieval rates.

### 4.1 Datasets and Metrics

**COCO 2017** [40] is a very popular object detection and instance segmentation dataset of common objects in context, consisting of 120K training images and 5K validation images, fully annotated with 80 semantic categories. Categories are varied from large objects (e.g., car, elephant, tv, refrigerator) to much smaller objects (e.g., fork, book, bird, frisbee, donut). **LVIS** [20] is a federated dataset, which includes 20K images in its validation set, intensively used on the long-tail object detection task [6, 19, 47, 77]. LVIS is annotated with 1203 semantic categories, 337 of which are considered rare objects (less than 10 training examples). In our experiments, we separately report a retrieval metric for rare objects as an applicable approximation for the open-set retrieval task.

**nuImages** [2] is a public autonomous driving dataset, significantly different from the former two datasets in terms of resolution, context, RGB distribution and annotated classes. nuImages validation set includes 16.5K images, 1600x900 sized, annotated with 23 diverse semantic categories. In the experiments, we define 7 of them (those that appear in less than 0.3% of the total annotations number) as rare categories, reporting their accuracy separately.

**Evaluation Protocol.** We evaluated our pipeline in two steps: a first processing step used to store the embeddings followed by a ranking step, which uses the datasets’ categories names as queries and sorts all images in descending order of relevance per query, based on the maximal similarity over all patches. For each query, images are declared as *true positive* if they include an object of that category. We report mean average precision (*mAP*), as widely reported in retrieval tasks [9, 60, 61, 65, 70], and *mAP@50* (defined in [21]), previously used in [40], which considers top-k images only, as our main criteria for comparison.

**Implementation Details.** We used the CLIP backbones from the CLIP [66] library. Clustering (K-Means and hierarchical clustering) was performed via the sklearn library [63]. Region proposals were calculated by Segment Anything (SAM) [63] library and tuned with different number point-prompts (64, 256 and 1024 points) which created, approximately 25, 50 and 100 representatives per image. All Dense-CLIP and clustering experiments were conducted on 1 Nvidia GPU machine. Images were resized to a square aspect ratio following hyperparameter search, and positional embeddings were interpolated to match the resolution of the

Backbone	Res.	#rep.	COCO			LVIS		LVIS-rare	
			mAP@50	mAP	mAP@50 <sub>s-m</sub>	mAP@50	mAP	mAP@50	mAP
VSRN, Flicker30K	600	1	44.28	37.23	21.49	31.56	37.35	10.52	11.53
VSRN, COCO-Caption	600	1	70.29	52.33	30.99	42.09	45.60	20.46	21.36
PCME	224	7	69.69	57.98	29.98	47.38	51.06	27.85	28.58
CLIP, RN50	224	1	56.70	50.80	21.91	52.69	55.94	39.84	40.74
CLIP, RN50x4	288	1	64.58	56.39	29.39	57.85	60.35	43.37	44.28
CLIP, RN50x64	448	1	70.62	61.03	36.60	62.60	64.71	53.14	53.92
OwlViT, ViT-B/32	768	576	77.31	<b>70.37</b>	<b>52.22</b>	66.09	67.86	42.60	42.95
OwlViT, ViT-B/16	768	2304	74.96	65.91	47.02	61.28	63.39	34.75	35.82
OwlViT, ViT-L/14	840	3600	76.61	<b>71.06</b>	<b>58.95</b>	66.15	67.86	40.62	41.10
Dense-CLIP, RN50	512	256	58.83 (+2.13)	52.35 (+1.55)	32.58 (+0.67)	55.41 (+2.72)	57.46 (+1.52)	38.80 (-1.04)	39.86 (-0.88)
Dense-CLIP, RN50x4	512	256	69.61 (+5.03)	62.10 (+5.71)	41.18 (+1.79)	63.88 (+6.03)	65.88 (+5.53)	55.32 (+11.95)	56.40 (+12.12)
Dense-CLIP, RN50x64	448	196	77.78 (+7.16)	69.65 (+8.62)	51.47 (+4.87)	70.86 (+8.26)	71.80 (+7.09)	57.97 (+4.83)	58.72 (+4.80)
CLIP + Dense-CLIP, RN50	512	257	67.27 (+10.57)	59.33 (+8.53)	34.99 (+13.08)	61.32 (+8.63)	63.02 (+7.08)	46.96 (+7.12)	47.82 (+7.08)
CLIP + Dense-CLIP, RN50x4	512	257	66.21 (+1.63)	58.08 (+1.69)	29.80 (+0.41)	61.82 (+3.97)	64.12 (+3.77)	49.49 (+6.12)	50.35 (+6.07)
CLIP + Dense-CLIP, RN50x64	448	197	77.48 (+6.86)	69.45 (+8.42)	48.73 (+12.13)	<b>72.24</b> (+9.64)	<b>73.36</b> (+8.65)	<b>61.68</b> (+8.54)	<b>62.33</b> (+8.41)
Cluster-CLIP, K.M., RN50x64	448	10	<b>78.67</b> (+8.05)	63.66 (+2.63)	46.69 (+10.09)	62.72 (+0.12)	64.11 (-0.6)	51.66 (-1.48)	52.45 (-1.47)
Cluster-CLIP, AG-T., RN50x64	448	10	76.72 (+6.10)	64.00 (+2.97)	45.03 (+8.43)	63.70 (+1.1)	65.16 (+0.45)	49.53 (-3.61)	50.38 (-3.54)
Cluster-CLIP, AG-F., RN50x64	448	50	78.35 (+7.73)	69.84 (+8.81)	51.95 (+15.35)	71.63 (+9.03)	72.60 (+7.89)	58.29 (+5.15)	59.07 (+5.15)
Cluster-CLIP, R.P., RN50x64	448	91	<b>79.24</b> (+8.62)	69.43 (+8.4)	50.51 (+13.91)	70.74 (+8.14)	71.92 (+7.21)	58.28 (+5.14)	58.88 (+4.96)
CLIP + Cluster-CLIP, K.M., RN50x64	448	11	75.51 (+4.89)	64.60 (+3.57)	41.61 (+5.01)	66.39 (+3.79)	67.96 (+3.25)	57.02 (+3.88)	57.73 (+3.81)
CLIP + Cluster-CLIP, AG-T., RN50x64	448	11	75.42 (+4.8)	65.19 (+4.16)	43.22 (+6.62)	67.16 (+4.56)	68.64 (+3.93)	57.15 (+4.01)	57.85 (+3.93)
CLIP + Cluster-CLIP, AG-F., RN50x64	448	51	77.06 (+6.44)	69.03 (+8.00)	48.08 (+11.48)	<b>71.79</b> (+9.19)	<b>73.02</b> (+8.31)	60.92 (+7.78)	<b>61.61</b> (+7.69)
CLIP + Cluster-CLIP, R.P., RN50x64	448	92	77.75 (+7.13)	68.59 (+7.56)	46.75 (+10.15)	71.26 (+8.66)	72.59 (+7.88)	<b>60.98</b> (+7.84)	61.21 (+7.29)

**Table 1:** Evaluation results on COCO2017 and LVIS val sets. **First** and **second** best scores are marked in red and blue. Using Dense-CLIP improves retrieval accuracy but increases the number of features. Using Cluster-CLIP compensates both, enabling scaling.

image. For a fair comparison, we ensemble over the embeddings space of the 7 best CLIP prompts [47] in all VL pre-trained modules. Full architectures descriptions, design choices, and hyperparameters are specified in the supplementary.

**Baselines.** We compare our work with existing global embeddings from CLIP and local embeddings from OwlViT [47]. As described in Section 2, OwlViT is an open-vocabulary dual-branch VL detection model that achieves SoTA results, making it a strong baseline for our task. It is pre-trained on a web-scale dataset and then fine-tuned for open-vocabulary object detection, which can lead to a forgetting effect. For image representation, we directly used local embeddings from OwlViT’s ViT [47] vision-encoder output, excluding bounding boxes prediction. We additionally compare to dual-stream cross-modal retrieval methods preceding CLIP that are trained on distinct medium-size datasets using caption annotations. Specifically, we compare to PCME and VSRN trained on COCO using the split defined in [47] and to VSRN trained on Flicker30K [47] (VSRN is referred to as either ‘VSRN, COCO-Caption’ or ‘VSRN, Flicker30K’, depending on the dataset).

## 4.2 Results

**Dense-CLIP Results.** Tables 1 and 2 compare the retrieval results of Dense-CLIP to the baselines. The ‘#rep.’ column shows the average number of embeddings for each image. Dense-CLIP achieves on-par results compared to OwlViT with fewer embeddings per image and leads to a significant and consistent improvement of up to 12 points in the retrieval rates over CLIP. Furthermore, Dense-CLIP surpasses VSRN and PCME on COCO and, to a greater extent, on LVIS, even when the latter two were fine-tuned on these datasets’ images. Small object retrieval ( $mAP@50_{s-m}$  in Table 1) benefits from the use of local features, where Dense-CLIP achieves competitive results with respect to OwlViT. Rare object retrieval (right side of the tables) proved to be more difficult (lower accuracy). Interestingly, results show higher retrieval rates for CLIP over OwlViT, maybe because of forgetting effects due to finetuning. Dense-CLIP shows a significant improvement over all baselines, exploiting CLIP open vocabulary representations in a dense manner.



Backbone	res.	#rep.	nuImages		nuImages - rare	
			mAP@50	mAP	mAP@50	mAP
PCME	224	7	25.67	15.08	0.2	0.75
CLIP, RN50	224	1	27.20	17.61	0.84	1.58
CLIP, RN50x4	288	1	28.62	19.07	2.70	3.39
CLIP, RN50x64	448	1	31.93	20.77	3.16	4.28
OWL-ViT, ViT-B/32	768	576	36.93	27.19	1.88	2.48
OWL-ViT, ViT-B/16	768	2304	33.45	26.65	2.25	2.58
OWL-ViT, ViT-L/14	840	3600	30.30	25.10	2.12	3.43
Dense-CLIP, RN50	768	576	31.64 (+4.44)	24.93 (+7.32)	3.33 (+2.49)	4.95 (+3.37)
Dense-CLIP, RN50x4	768	576	32.04 (+3.42)	26.77 (+7.70)	2.80 (+4.10)	4.88 (+1.49)
Dense-CLIP, RN50x64	768	576	34.08 (+2.15)	30.74 (+9.97)	<b>10.11</b> (+6.95)	<b>10.88</b> (+6.60)
CLIP + Dense-CLIP, RN50	768	577	32.60 (+5.40)	25.68 (+8.07)	3.41 (+2.57)	5.13 (+3.55)
CLIP + Dense-CLIP, RN50x4	768	577	33.79 (+5.17)	25.94 (+6.87)	6.32 (+3.62)	7.70 (+4.31)
CLIP + Dense-CLIP, RN50x64	768	577	37.95 (+6.02)	<b>31.26</b> (+10.49)	<b>10.23</b> (+7.07)	<b>11.24</b> (+6.96)
Cluster-CLIP, K.M., RN50x64	768	10	37.87 (+5.94)	25.02 (+4.25)	6.07 (+2.91)	6.93 (+2.65)
Cluster-CLIP, AG-T, RN50x64	768	10	<b>40.55</b> (+8.62)	27.28 (+6.51)	7.35 (+4.19)	8.51 (+4.23)
Cluster-CLIP, AG-F, RN50x64	768	50	37.00 (+5.07)	30.61 (+9.84)	9.32 (+6.16)	10.40 (+6.12)
Cluster-CLIP, R.P., RN50x64	768	97	40.05 (+8.12)	30.05 (+9.28)	9.10 (+5.94)	10.16 (+5.88)
CLIP + Cluster-CLIP, K.M., RN50x64	768	11	39.78 (+7.85)	26.95 (+6.18)	8.13 (+4.97)	9.15 (+4.87)
CLIP + Cluster-CLIP, AG-T, RN50x64	768	11	40.24 (+8.31)	28.06 (+7.29)	8.82 (+5.66)	9.80 (+5.52)
CLIP + Cluster-CLIP, AG-F, RN50x64	768	50	38.55 (+6.62)	<b>30.75</b> (+9.98)	9.10 (+5.94)	10.37 (+6.09)
CLIP + Cluster-CLIP, R.P., RN50x64	768	98	<b>40.54</b> (+8.61)	29.84 (+9.07)	8.84 (+5.68)	10.15 (+5.87)

**Table 2:** Evaluation results on nuImages val set. **First** and **second** best scores are marked in red and blue.




**Figure 4:** Top matching clusters.

**Cluster-CLIP Results.** We empirically studied the use of aggregated features using several clustering mechanisms, described in section 3.3. Fig. 5 displays Cluster-CLIP’s results for different clustering methods and numbers of clusters, while promising working points are presented in the sixth parts of Tab. 1 and 2 (see results for RN50/x4 backbones in the supplementary). To reduce variance, all results are averaged across three experiments.

Notably, K.M. and AG-F (blue and lime) outperform Dense-CLIP when employing a modest count of 10 representatives per image on COCO. With a larger representatives budget of 50, aggregating the dense features according to R.P. (star markers) is made an alternative clustering mechanism. On LVIS, which has almost twice the number of instances per image compared to COCO, AG-F with 50 representatives achieves the highest retrieval rates. Interestingly, while AG-F and K.M. share certain clustering characteristics, it appears that AG-F bottom-up clustering is preferred when fine-grained understanding is required (as in LVIS) and on small objects. On nuImages, AG-T and R.P., which incorporate localization considerations into the clustering process, outperform other mechanisms with merely 5-50 representatives, implying localization is a strong cue for clustering in high-resolution images. As mentioned, Cluster-CLIP outperforms Dense-CLIP in several cases. This can be explained by noticing that applying clustering to the output of Dense-CLIP introduces two conflicting phenomena. On the one hand, clustering reduces the number of representatives per image, thereby decreasing the presence of distractors. On the other hand, averaging across multiple patches may lead to the averaging out of small-sized objects or fine-grained details, potentially hindering performance.

We consider the above results as key contributions, highlighting the potential of aggregated features to offer both efficiency and performance.

**Mixed Architectures.** We found that in many cases, Dense-CLIP and Cluster-CLIP embeddings are in fact complementary to the global embedding produced by CLIP, and those can be effectively combined to boost their performance (referred to as CLIP + Dense/Cluster-CLIP). Results are reported in the fifth and seventh part of Tables 1 and 2. Notably, incorporating CLIP embeddings enhances Dense-CLIP results by up to 3.8 mAP@50 points. Moreover, integrating CLIP embeddings into Cluster-CLIP boosts results by up to 7.6 mAP@50 points while reinforcing its efficacy across tasks.



**Figure 5: Cluster-CLIP accuracy-efficiency scatter plots:** retrieval accuracy (mAP@50) vs. average numbers of embeddings per image for COCO, LVIS and nuImages datasets. Top left is better. Cluster-CLIP achieves high accuracy for the small number of 10-50 embeddings per image.

### 4.3 Qualitative Results with Overall Framework

We build a simple image retrieval framework to study the behavior of Cluster-CLIP aggregated embeddings (Figure 1). For demonstration purposes, we use an index consisting of 120K COCO training set images. Specifically, we used Cluster-CLIP with RN50x64 backbone and K-Means with 10 centroids, then utilized FAISS [50] to map the embeddings into a large-scale index. Interactive querying is enabled using text or image queries (created by CLIP text or image encoder). Qualitative examples of interest are demonstrated in Figures 1, 2 and 6. Please refer to the Supplementary Materials for additional qualitative examples.

## 5 Conclusions

We examine the possible use of local features for the object-centric image retrieval task and introduce Dense-CLIP, which extracts dense embedding, manipulated such that CLIP vision-language association is kept. We compare our method with the use of global features extracted from CLIP with the same backbone and report a significant increase in retrieval rates of up to 12 mAP points. Compared to SoTA detection pipelines, our results are competitive, with a significant increase in retrieval rates for rare categories.

Practically, utilizing local features within a retrieval framework significantly enlarges the search space, impairing scalability. To address this, we introduce Cluster-CLIP, which innovatively represents images using features aggregated from local features. Our approach achieves improved retrieval rates with fewer representatives, practically enabling scaling. From a broader perspective, the potential use of compact representation to efficiently carry useful image information is interesting by itself and might contribute to future work on a wide range of applications (e.g., detection, segmentation, image generation).



**Figure 6: Qualitative Examples:** top-1 retrieved images with Dense-CLIP heat-maps for rare textual and visual queries. Index created from 120K images with K.M. clustered embeddings.

## References

- [1] Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 6077–6086. Computer Vision Foundation / IEEE Computer Society, 2018.
- [2] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016.
- [3] Artem Babenko, Anton Slesarev, Alexandr Chigorin, and Victor Lempitsky. Neural codes for image retrieval. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I 13*, pages 584–599. Springer, 2014.
- [4] Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. Learning local feature descriptors with triplets and shallow convolutional neural networks. In Richard C. Wilson, Edwin R. Hancock, and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference 2016, BMVC 2016, York, UK, September 19-22, 2016*. BMVA Press, 2016.
- [5] Hanoona Bangalath, Muhammad Maaz, Muhammad Uzair Khattak, Salman H Khan, and Fahad Shahbaz Khan. Bridging the gap between object and image-level representations for open-vocabulary detection. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [6] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 110(3):346–359, 2008.
- [7] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11618–11628. Computer Vision Foundation / IEEE, 2020.
- [8] Zhaowei Cai, Gukyeong Kwon, Avinash Ravichandran, Erhan Bas, Zhuowen Tu, Rahul Bhotika, and Stefano Soatto. X-detr: A versatile architecture for instance-wise vision-language tasks. In *European Conference of Computer Vision (ECCV)*, 2022.
- [9] Bingyi Cao, Andre Araujo, and Jack Sim. Unifying deep local and global features for image search. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pages 726–743. Springer, 2020.
- [10] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: universal image-text representation learning. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*, volume 12375 of *Lecture Notes in Computer Science*, pages 104–120. Springer, 2020.

- [11] Sanghyuk Chun, Seong Joon Oh, Rafael Sampaio De Rezende, Yannis Kalantidis, and Diane Larlus. Probabilistic embeddings for cross-modal retrieval. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [13] Martin Engilberge, Louis Chevallier, Patrick Pérez, and Matthieu Cord. Finding beans in burgers: Deep semantic-visual embedding with localization. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3984–3993. Computer Vision Foundation / IEEE Computer Society, 2018.
- [14] Fartash Faghri, David J. Fleet, Jamie Ryan Kiros, and Sanja Fidler. VSE++: improving visual-semantic embeddings with hard negatives. In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*, page 12. BMVA Press, 2018.
- [15] Andrea Frome, Greg S Corrado, Jon Shlens, Samy Bengio, Jeff Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. Devise: A deep visual-semantic embedding model. *Advances in neural information processing systems (NeurIPS)*, 26, 2013.
- [16] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner, editors, *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, Proceedings, Part XXXVI*, volume 13696 of *Lecture Notes in Computer Science*, pages 540–557. Springer, 2022.
- [17] Yunchao Gong, Liwei Wang, Micah Hodosh, Julia Hockenmaier, and Svetlana Lazebnik. Improving image-sentence embeddings using large weakly annotated photo collections. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part IV*, volume 8692 of *Lecture Notes in Computer Science*, pages 529–545. Springer, 2014.
- [18] Albert Gordo, Jon Almazán, Jerome Revaud, and Diane Larlus. Deep image retrieval: Learning global representations for image search. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VI 14*, pages 241–257. Springer, 2016.
- [19] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Open-vocabulary object detection via vision and language knowledge distillation. In *International Conference on Learning Representations (ICLR)*, 2022.
- [20] Agrim Gupta, Piotr Dollár, and Ross B. Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 5356–5364. Computer Vision Foundation / IEEE, 2019.

- [21] Nandeshwar Gupta. Mean average precision map@k metric explained code, 2022. URL <https://www.kaggle.com/code/nandeshwar/mean-average-precision-map-k-metric-explained-code/notebook>.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [23] Kun He, Yan Lu, and Stan Sclaroff. Local descriptors optimized for average precision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 596–605, 2018.
- [24] Yan Huang, Wei Wang, and Liang Wang. Instance-aware image and sentence matching with selective multimodal LSTM. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*, pages 7254–7262. IEEE Computer Society, 2017.
- [25] Zhicheng Huang, Zhaoyang Zeng, Bei Liu, Dongmei Fu, and Jianlong Fu. Pixelbert: Aligning image pixels with text by deep multi-modal transformers. *CoRR*, abs/2004.00849, 2020.
- [26] Hervé Jégou, Matthijs Douze, Cordelia Schmid, and Patrick Pérez. Aggregating local descriptors into a compact image representation. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 3304–3311. IEEE, 2010.
- [27] Hervé Jégou, Florent Perronnin, Matthijs Douze, Jorge Sánchez, Patrick Pérez, and Cordelia Schmid. Aggregating local image descriptors into compact codes. *IEEE transactions on pattern analysis and machine intelligence*, 34(9):1704–1716, 2011.
- [28] Zhong Ji, Haoran Wang, Jungong Han, and Yanwei Pang. Saliency-guided attention network for image-sentence matching. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5753–5762. IEEE, 2019.
- [29] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021.
- [30] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547, 2019.
- [31] Andrej Karpathy, Armand Joulin, and Li Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping. In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pages 1889–1897, 2014.

- [32] Dahun Kim, Anelia Angelova, and Weicheng Kuo. Region-aware pretraining for open-vocabulary object detection with vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11144–11154, 2023.
- [33] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- [34] Ryan Kiros, Ruslan Salakhutdinov, and Richard S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *CoRR*, abs/1411.2539, 2014. URL <http://arxiv.org/abs/1411.2539>.
- [35] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, volume 11208 of *Lecture Notes in Computer Science*, pages 212–228. Springer, 2018.
- [36] Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Gotmare, Shafiq R. Joty, Caiming Xiong, and Steven Chu-Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. In Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan, editors, *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 9694–9705, 2021.
- [37] Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. Visual semantic reasoning for image-text matching. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 4653–4661. IEEE, 2019.
- [38] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXX*, volume 12375 of *Lecture Notes in Computer Science*, pages 121–137. Springer, 2020.
- [39] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [40] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer, 2014.
- [41] Chunxiao Liu, Zhendong Mao, An-An Liu, Tianzhu Zhang, Bin Wang, and Yongdong Zhang. Focus your attention: A bidirectional focal attention network for image-text

- matching. In *Proceedings of the 27th ACM International Conference on Multimedia, MM 2019, Nice, France, October 21-25, 2019*, pages 3–11.
- [42] Sheng Liu, Kevin Lin, Lijuan Wang, Junsong Yuan, and Zicheng Liu. OVIS: open-vocabulary visual instance search via visual-semantic aligned representation learning. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelfth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 1773–1781. AAAI Press, 2022.
- [43] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60:91–110, 2004.
- [44] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13–23, 2019.
- [45] Huaishao Luo, Junwei Bao, Youzheng Wu, Xiaodong He, and Tianrui Li. Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. *CoRR*, abs/2211.14813, 2022.
- [46] Krystian Mikolajczyk and Cordelia Schmid. An affine invariant interest point detector. In *Computer Vision—ECCV 2002: 7th European Conference on Computer Vision Copenhagen, Denmark, May 28–31, 2002 Proceedings, Part I 7*, pages 128–142. Springer, 2002.
- [47] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, Xiao Wang, Xiaohua Zhai, Thomas Kipf, and Neil Houlsby. Simple open-vocabulary object detection with vision transformers. In *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part X*, page 728–755. Springer-Verlag, 2022.
- [48] Y. Mori, H. Takahashi, and R. Oka. Image-to-word transformation based on dividing and vector quantizing images with words. In *MISRM’99 First International Workshop on Multimedia Intelligent Storage and Retrieval Management*, 1999.
- [49] Jishnu Mukhoti, Tsung-Yu Lin, Omid Poursaeed, Rui Wang, Ashish Shah, Philip H. S. Torr, and Ser-Nam Lim. Open vocabulary semantic segmentation with patch aligned contrastive learning. *CoRR*, abs/2212.04994, 2022.
- [50] Tony Ng, Vassileios Balntas, Yurun Tian, and Krystian Mikolajczyk. Solar: second-order loss and attention for image retrieval. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXV 16*, pages 253–270. Springer, 2020.
- [51] Hyeonwoo Noh, Andre Araujo, Jack Sim, Tobias Weyand, and Bohyung Han. Large-scale image retrieval with attentive deep local features. In *Proceedings of the IEEE international conference on computer vision*, pages 3456–3465, 2017.

- [52] Stepan Obdrzalek and Jiri Matas. Sub-linear indexing for large scale object recognition. In *BMVC*, volume 2, 2005.
- [53] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [54] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Cnn image retrieval learns from bow: Unsupervised fine-tuning with hard examples. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 3–20. Springer, 2016.
- [55] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning cnn image retrieval with no human annotation. *IEEE transactions on pattern analysis and machine intelligence*, 41(7):1655–1668, 2018.
- [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021.
- [57] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [58] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Denseclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022.
- [59] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12716–12725, 2019.
- [60] Josef Sivic and Andrew Zisserman. Video google: A text retrieval approach to object matching in videos. In *Computer Vision, IEEE International Conference on*, volume 3, pages 1470–1470. IEEE Computer Society, 2003.
- [61] Yale Song and Mohammad Soleymani. Polysemous visual-semantic embedding for cross-modal retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019*, pages 1979–1988. Computer Vision Foundation / IEEE, 2019.
- [62] Hao Tan and Mohit Bansal. 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 2019, Hong Kong, China, November 3-7, 2019*, pages 5099–5110. Association for Computational Linguistics, 2019.
- [63] Christopher Thomas and Adriana Kovashka. Preserving semantic neighborhoods for robust cross-modal retrieval. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XVIII*, volume 12363 of *Lecture Notes in Computer Science*, pages 317–335. Springer, 2020.
- [64] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [65] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. GIT: A generative image-to-text transformer for vision and language. *Trans. Mach. Learn. Res.*, 2022, 2022.
- [66] Zihao Wang, Xihui Liu, Hongsheng Li, Lu Sheng, Junjie Yan, Xiaogang Wang, and Jing Shao. CAMP: cross-modal adaptive message passing for text-image retrieval. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pages 5763–5772. IEEE, 2019.
- [67] Zirui Wang, Jiahui Yu, Adams Wei Yu, Zihang Dai, Yulia Tsvetkov, and Yuan Cao. Simvlm: Simple visual language model pretraining with weak supervision. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [68] Xi Wei, Tianzhu Zhang, Yan Li, Yongdong Zhang, and Feng Wu. Multi-modality cross attention network for image and sentence matching. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10938–10947. Computer Vision Foundation / IEEE, 2020.
- [69] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas M. Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 18113–18123. IEEE, 2022.
- [70] Min Yang, Dongliang He, Miao Fan, Baorong Shi, Xuotong Xue, Fu Li, Errui Ding, and Jizhou Huang. Dolg: Single-stage image retrieval with deep orthogonal fusion of local and global features. In *Proceedings of the IEEE/CVF International conference on Computer Vision*, pages 11772–11781, 2021.
- [71] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. FILIP: fine-grained interactive language-image pre-training. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022.
- [72] Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Trans. Assoc. Comput. Linguistics*, 2:67–78, 2014.

- [73] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022.
- [74] Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- [75] Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 18102–18112. IEEE, 2022.
- [76] Qi Zhang, Zhen Lei, Zhaoxiang Zhang, and Stan Z. Li. Context-aware attention network for image-text retrieval. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 3533–3542. Computer Vision Foundation / IEEE, 2020.
- [77] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Lianian Harold Li, Luowei Zhou, Xiyang Dai, Lu Yuan, Yin Li, and Jianfeng Gao. Regionclip: Region-based language-image pretraining. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, 2022.
- [78] Chong Zhou, Chen Change Loy, and Bo Dai. Denseclip: Extract free dense labels from clip. *arXiv preprint arXiv:2112.01071*, 2021.
- [79] Chaoyang Zhu and Long Chen. A survey on open-vocabulary detection and segmentation: Past, present, and future. *arXiv preprint arXiv.2307.09220*, 2023.