

Pic2Word: Mapping Pictures to Words for Zero-shot Composed Image Retrieval

Kuniaki Saito^{1,2*}, Kihyuk Sohn³, Xiang Zhang², Chun-Liang Li²,
Chen-Yu Lee², Kate Saenko^{1,4}, Tomas Pfister²

{keisaito, saenko}@bu.edu

{kihyuks, fancyzhx, chunliang, chenyllee, tpfister}@google.com

¹Boston University, ²Google Cloud AI Research, ³Google Research, ⁴MIT-IBM Watson AI Lab

Abstract

In *Composed Image Retrieval (CIR)*, a user combines a query image with text to describe their intended target. Existing methods rely on supervised learning of CIR models using labeled triplets consisting of the query image, text specification, and the target image. Labeling such triplets is expensive and hinders broad applicability of CIR. In this work, we propose to study an important task, *Zero-Shot Composed Image Retrieval (ZS-CIR)*, whose goal is to build a CIR model without requiring labeled triplets for training. To this end, we propose a novel method, called *Pic2Word*, that requires only weakly labeled image-caption pairs and unlabeled image datasets to train. Unlike existing supervised CIR models, our model trained on weakly labeled or unlabeled datasets shows strong generalization across diverse ZS-CIR tasks, e.g., attribute editing, object composition, and domain conversion. Our approach outperforms several supervised CIR methods on the common CIR benchmark, CIRR and Fashion-IQ. Code will be made publicly available at https://github.com/google-research/composed_image_retrieval

1. Introduction

Composed image retrieval (CIR) [2, 4, 16, 18, 35] aims to retrieve images using a query composed of an image and text. In contrast to the image-based retrieval systems [7], CIR is advantageous as it allows retrieval of images with a higher precision thanks to the text query that incorporates user’s intent, such as a desired modification to the query image. With a surge of image-text models [1, 20, 30], CIR has received attention recently for diverse real-world applications in e-commerce and internet search.

Several approaches [2, 4, 16, 18, 35] have been proposed to solve CIR problems, including attribute manipulation for fashion image search, composing objects, and converting

*Work done during internship at Google Cloud AI Research.

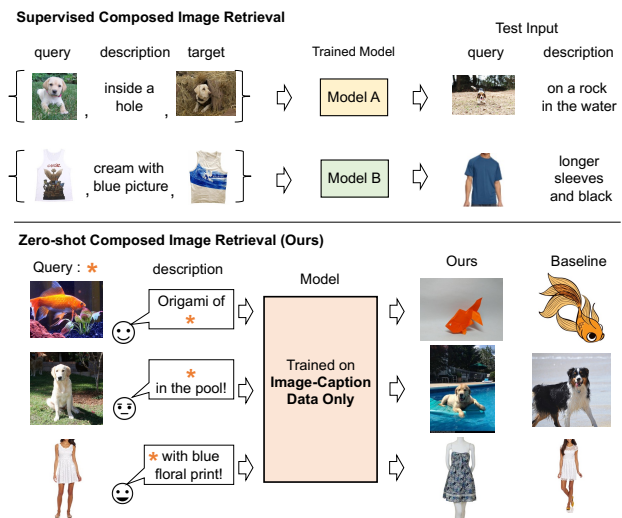


Figure 1. Composed Image Retrieval (CIR) takes a query composed of an image (denoted as \star) plus a text modifier and retrieves image matching images. CIR encompasses diverse tasks, such as domain conversion (origami of \star), scene or object composition (\star in the pool), or fashion-attribute manipulation (\star with blue floral print). **Top:** Existing methods [2, 35] train a separate model for each task, and require strong triplet supervision. **Bottom:** We propose a new task, Zero-shot CIR, and solve diverse CIR sub-tasks using a single model trained only on image-caption pairs and unlabeled image datasets.

the style of images for content creation, as in Fig. 1. At the core of CIR is learning how to compose information from an image and text. We identify two fundamental issues with existing solutions. First, previous methods require a large amount of labeled data, which comes in the form of triplets consisting of a reference image, text, and a target image, to train their retrieval model. The dataset collection involves two processes [25] – collecting pairs of a related reference and the target images as a query-output pair to the CIR system, then giving a description that modifies the reference to the target. Examples of labeled triplets are at

the top of Fig. 1. We note that both steps incur a significant labeling cost. Second, the model trained on labeled data is specialized to specific use-cases and may not generalize to different CIR tasks.

To tackle these challenges, we propose a new task, *zero-shot composed image retrieval* (ZS-CIR). In ZS-CIR, our goal is to build a single CIR model that performs diverse CIR tasks, such as object composition, attribute editing, or domain conversion, as in the bottom of Fig. 1, without requiring an expensive labeled triplet data collection effort. Instead, we propose to train our retrieval model using large-scale image-caption pairs and unlabeled images, which are considerably cheaper to collect than supervised CIR datasets at scale.

To harness these weakly labeled and unlabeled datasets, we propose a two-stage framework for learning ZS-CIR models. The first stage conducts contrastive language-image pretraining (CLIP) [30] on the image-caption dataset, training a two-tower model jointly to maximize the similarity between an image and a caption. We note that some previous works [2] build their models on CLIP followed by a second-stage supervised CIR training. On the contrary, instead of relying on the triplet-labeled training data, we leverage the linguistic capability of the language encoder in CLIP, which excels at composing diverse concepts or attributes to generate embeddings that are close to the corresponding visual representations. The idea is to map a *picture to a word token* such that the language encoder can flexibly and seamlessly compose the query image features and text descriptions. We learn a lightweight mapping network that converts an image embedding of the CLIP vision encoder into a token embedding of its language encoder. The mapping network is trained with a contrastive loss to reconstruct the image embedding, which only requires unlabeled images. We call our method Pic2Word and illustrate it in Fig. 2.

In experiments, we show the strength of our Pic2Word approach on various CIR tasks, including domain conversion, object composition, scene manipulation, and fashion attribute manipulation. We show that our zero-shot approach performs on-par with or better than several recent supervised CIR methods [8, 11, 25] relying on labeled training data. Our contributions are threefold:

- We propose a new task, zero-shot composed image retrieval (ZS-CIR), which aims to solve diverse CIR tasks without requiring expensive triplet-labeled training datasets.
- We propose Pic2Word, a novel method for ZS-CIR that is trainable using only image-caption and unlabeled image datasets. Pic2Word leverages pre-trained vision-language models and transforms an input image to a language token in order to flexibly compose image and text queries.

- Pic2Word improves the ZS-CIR performance, e.g., relative improvement of 10 - 100% on four CIR tasks, which is on-par with several recent CIR methods using labeled training data.

2. Related Work

Composed Image Retrieval. Composed image retrieval (CIR) is proposed to retrieve images with a pair consisting of a reference image and text [35]. The topic has been explored in the field of fashion [37] and scene composition [25]. Current state-of-the-art CIR models rely on late-fusion, i.e., combining visual and language features after extracting them with different encoders [2–4, 8, 18, 25], where a pre-trained CLIP model shows strong performance [2]. Wu *et al.* [37] train a shallow transformer from scratch to fuse image and language features at the input-level. Goenka *et al.* [13] utilize a pre-trained BERT [10] model to fuse image, text, and tag features while Han *et al.* [14] leverage a large-scale fashion dataset to pre-train vision-language model by early fusion, but both still need to train on a CIR dataset. These approaches employ the compositional ability of the pre-trained language model by tuning their transformer model on a CIR dataset. Unlike these methods, our approach does not need a CIR dataset for training, yet can handle diverse scenarios.

Vision-language foundation models. Vision language models such as CLIP [30], ALIGN [20], *etc.*, pre-train an image and language encoder pair on large-scale data containing hundreds of millions of image-caption pair. Following CLIP, many vision-language “foundation” models incorporate more data into training, introduce new architectural designs, or utilize new objectives [1, 17, 20, 32, 38]. Since these models have seen varying text describing the concepts in an image, they can handle diverse tasks, e.g., caption-based image retrieval, zero-shot classification [30], few-shot classification [39], image-captioning [28], and VQA [34], with a small or no extra annotation cost. In this work, we present an approach to adapt the foundation models to CIR tasks without the need for a CIR dataset. We are the first to apply vision-language contrastive learning models to CIR in a zero-shot manner.

Representing image as one word. Several approaches attempt to represent an image region as a token during the pre-training of vision-language models [5, 21, 27, 36]. The typical framework consists of (i) detecting objects in an image using a pre-trained object detector, (ii) feeding the regions and sentence to a text encoder, and (iii) optimizing several multi-modal objectives to obtain a strong vision-language model. These approaches require a high-performing object detector in the pre-training stage, while we present an approach applicable to pre-trained contrastive vision-language models. A concurrent work [6] learns a network with cycle contrastive loss to tackle different prob-

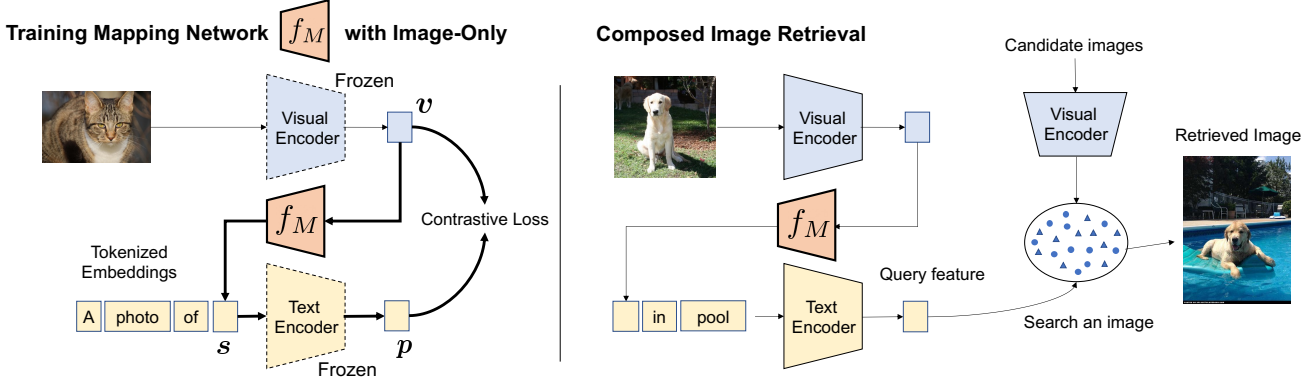


Figure 2. **Left:** Overview of training our Pic2Word mapping network. Given a frozen visual and text encoder, the mapping network, f_M , is optimized to minimize the contrastive loss between the image embedding and the language embedding generated by the pseudo word token s . **Right:** Overview of inference. The estimated token is used to fill in the given prompt sentence.

lems, personalized image retrieval and semantic segmentation. To achieve those, they learn to transform a *set of images* into a concept word token, which requires images with class-wise annotations and caption annotations. In addition, they require fine-tuning the word token given a new concept consisting of a few images. Instead, we focus on composed image retrieval. The proposed Pic2Word does not need such annotations or fine-tuning the word token on few-shot labeled images. Gal *et al.* [12] propose to represent a few images depicting the same object with a word token and use this for a text-to-image generation task. Our approach does not require such a group of images, nor any training at inference time. We believe our approach can potentially be applied to image generation in future work.

3. Method

In this section, we introduce our approach for ZS-CIR, Pic2Word. The overview of the training and the inference is described in Fig. 2. We utilize a pre-trained image-text model, CLIP [30], that consists of a language and a vision encoder. The output embeddings of two encoders are aligned with respect to each other’s modality. Given a frozen pre-trained model, we train a mapping network to convert the visual embedding into the corresponding *pseudo language token*. The network is optimized such that the pseudo token represents the visual embedding faithfully. At test time the predicted token is simply inserted into a template together with the query text and the resulting features are compared to candidate images. We detail the training and inference mechanisms in the following subsections.

3.1. Preliminaries

Contrastive Language-Image Pre-training (CLIP). Suppose we have access to an image-caption dataset with image-caption pairs $\mathcal{S} = \{(\mathbf{x}_n, \mathbf{t}_n)\}_{n=1}^N$, where $\mathbf{x} \in \mathcal{X}$ is an image and $\mathbf{t} \in \mathcal{T}$ is a tokenized language description. For

an image \mathbf{x} , an image encoder model f_θ parameterized by θ extracts a visual representation $\tilde{\mathbf{v}} \in \mathbb{R}^{d \times 1}$: $\tilde{\mathbf{v}} = f_\theta(\mathbf{x})$. For each text description $\mathbf{t} \in \mathcal{T}$ (e.g., caption or prompt), a word embedding layer E_w parameterized by w extracts the *token embedding* $E_w(\mathbf{t})$. We will use the terms *token embedding* and *token* interchangeably for concise presentation when there is no ambiguity. Then a text encoder f_ϕ parameterized by ϕ extracts a language representation $\tilde{\mathbf{u}} \in \mathbb{R}^{d \times 1}$ following $E_w(\mathbf{t})$ as $\tilde{\mathbf{u}} = f_\phi(E_w(\mathbf{t}))$.

CLIP [30] is designed to find representations that match an image to its paired caption while separating unpaired ones. For i -th image \mathbf{x}_i and j -th language description \mathbf{t}_j in a batch \mathcal{B} , their features are normalized using $\mathbf{v}_i = \frac{\tilde{\mathbf{v}}_i}{\|\tilde{\mathbf{v}}_i\|}$ and $\mathbf{u}_j = \frac{\tilde{\mathbf{u}}_j}{\|\tilde{\mathbf{u}}_j\|}$. Finally, CLIP optimizes the symmetric multi-class N-pair loss [33]:

$$\min_{\{\theta, \phi, w\}} \mathcal{L}_{con} = \mathcal{L}_{t2i} + \mathcal{L}_{i2t}, \quad (1)$$

which includes two contrastive terms, with a temperature hyper-parameter τ that controls the strength of penalties on hard negative samples, as follows:

$$\mathcal{L}_{t2i} = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \mathbf{u}_i^T \mathbf{v}_i)}{\sum_{j \in \mathcal{B}} \exp(\tau \mathbf{u}_i^T \mathbf{v}_j)}, \quad (2)$$

$$\mathcal{L}_{i2t} = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \mathbf{v}_i^T \mathbf{u}_i)}{\sum_{j \in \mathcal{B}} \exp(\tau \mathbf{v}_i^T \mathbf{u}_j)}. \quad (3)$$

Note that image-caption retrieval is performed on the normalized feature space, i.e., \mathbf{v}_i and \mathbf{u}_i .

3.2. Learning the Pic2Word Mapping Network

We train Pic2Word with contrastive loss to map an input visual embedding $\tilde{\mathbf{v}}$ to a pseudo language token that is compatible with the embedding generated by the frozen CLIP language encoder. In CLIP, the language and vision encoders are optimized to align the language embedding \mathbf{u}

and the image embedding v . We hypothesize that a pseudo token represents the image semantics faithfully if a language embedding from a generic prompt plus the pseudo token embedding is close to the corresponding image embedding. Following this intuition, we propose to train a mapping network to output pseudo language tokens such that the language embedding obtained by a fixed language encoder can minimize the contrastive loss between the visual and language embeddings, i.e., enforcing the network to form a cycle starting from the visual embedding to the final language embedding as shown in the left of Fig. 2.

For an unnormalized visual embedding \tilde{v} , the mapping network f_M with parameters M extracts a pseudo language token embedding $s = f_M(\tilde{v})$. We build with a three-layered fully-connected network for f_M with roughly 0.8M parameters. As shown in the left of Fig. 2, we append s at the end of token embeddings of the prompt sentence, a photo of \tilde{v} , resulting in \hat{s} . Then, we feed \hat{s} into the language encoder f_ϕ to obtain the language embedding $\tilde{p} = f_\phi(\hat{s})$, hoping that \tilde{p} is able to represent an input image embedding \tilde{v} . To achieve this, we propose to minimize the contrastive loss with respect to the mapping network, i.e.,

$$\min_M \mathcal{L} = \mathcal{L}_{t2i}(\mathbf{p}, \mathbf{v}) + \mathcal{L}_{i2t}(\mathbf{p}, \mathbf{v}), \quad (4)$$

which includes two contrastive terms:

$$\mathcal{L}_{t2i}(\mathbf{p}, \mathbf{v}) = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \mathbf{p}_i^T \mathbf{v}_i)}{\sum_{j \in \mathcal{B}} \exp(\tau \mathbf{p}_i^T \mathbf{v}_j)}, \quad (5)$$

$$\mathcal{L}_{i2t}(\mathbf{p}, \mathbf{v}) = -\frac{1}{|\mathcal{B}|} \sum_{i \in \mathcal{B}} \log \frac{\exp(\tau \mathbf{v}_i^T \mathbf{p}_i)}{\sum_{j \in \mathcal{B}} \exp(\tau \mathbf{v}_i^T \mathbf{p}_j)}, \quad (6)$$

where $\mathbf{p}_i = \frac{\tilde{p}_i}{\|\tilde{p}_i\|}$. In this way, we train f_M to generate s_i that makes \mathbf{p}_i close to \mathbf{v}_i . Note that these are different from Eq. (1, 2, 3) in that (i) optimization is done for M while fixing the rest, and (ii) we use \mathbf{p} in replace of \mathbf{u} . That being said, the mapping network is trained in an unsupervised way using unlabeled images only *without* image-text paired or labeled image data.

In summary, the mapping network is trained to reconstruct the visual representations in the space of the language embedding. We represent each image as one word token, yet using multiple tokens might be optimal considering the richness of image information. As discussed in Sec 4.3, even one word token can represent an image very effectively. We hypothesize that more word tokens may be required to describe very fine details of the image and we can simply extend our idea to obtain multiple word tokens. Also, unlike state-of-the-art supervised CIR methods relying on late-fusion [2–4, 8, 18, 25], i.e., training a network to combine visual and language features, Pic2Word achieves early-fusion in the language token space. This allows us to harness the capacity of pre-trained language model to compose diverse concepts. Additionally, Pic2Word has the flex-

ibility to accept multiple images to create one query feature since each image is represented as a word (See Fig. 8).

3.3. Inference

At inference our goal is to compose the image and text query and compare it to candidate images. We compose the pseudo token embedding of the image, from the mapping network, with the text description as follows. As shown on the right of Fig. 2, the idea is to add a pseudo token into pre-defined prompts as if it were a real word token. The result is embedded by the text encoder and compared to visual features of candidate images. We note that the prompt design plays an important role of the final performance [19, 24]. As our focus is on studying the composition of image embeddings and language descriptions in a zero-shot way, we rely on simple prompts without further tuning. We provide examples of adopted prompts for the different tasks we study. In all examples, [*] indicates the pseudo token from mapping network.

(a) Domain conversion. Under the scenario where we want to modify the domain of the query image, e.g., a real image to a sketch-style image, we compose the domain and image as a [domain] of [*], where [domain] will be replaced by a word specifying the domain.

(b) Object/Scene composition. In object composition, we aim to retrieve an image consisting of an object specified with a query image, and scene/objects described with text. We compose the query by a photo of [*], [obj₁] and [obj₂], ..., and [obj_n], where [obj_i] are text descriptions of objects or scenes.

(c) Sentence specification. The modification to the reference image can be given by a sentence. In such cases, we simply append the sentence after the prompt with a pseudo token such as a photo of [*], [text], where [text] denotes modification text.

4. Experiments

In this section, we describe experiments evaluating Pic2Word on zero-shot composed image retrieval. We leverage four datasets to assess the performance of the model in diverse scenarios, including the standard CIR benchmarks CIRR [25] and Fashion-IQ [37]. This section is organized as follows: (i) an explanation of the experimental setup, (ii) an introduction of the main results, and (iii) a detailed analysis.

4.1. Setup

Training Details. Unless otherwise stated, we use ViT-L/14 CLIP [30] pre-trained with 400M image-text paired data.¹ We train the mapping network on the Conceptual Caption dataset [31], consisting of 3M images. The mapping network consists of three-layered MLP of 512 hid-

¹We employ ViT-L/14 published by OpenAI.

Supervision	Methods	Cartoon		Origami		Toy		Sculpture		Average	
		R10	R50	R10	R50	R10	R50	R10	R50	R10	R50
ZERO-SHOT	Image-only	0.3	4.5	0.2	1.8	0.6	5.7	0.3	4.0	0.4	4.0
	Text-only	0.2	1.1	0.8	3.7	0.8	2.4	0.4	2.0	0.5	2.3
	Image+Text	2.2	13.3	2.0	10.3	1.2	9.7	1.6	11.6	1.7	11.2
	Pic2Word	8.0	21.9	13.5	25.6	8.7	21.6	10.0	23.8	10.1	23.2
CIRR	Combiner [2]	6.1	14.8	10.5	21.3	7.0	17.7	8.5	20.4	8.0	18.5
Fashion-IQ	Combiner [2]	6.0	16.9	7.6	20.2	2.7	10.9	8.0	21.6	6.0	17.4

Table 1. **Results of the domain conversion experiment using ImageNet.** Pic2Word outperforms all baselines by a large margin. In particular, Pic2Word outperforms supervised models trained on other composed image retrieval tasks (bottom two rows).




Figure 3. **Top-3 retrievals in the domain-conversion experiment.** Note that in our approach, the query is composed with the sentence, “a *domain* of *”, where *domain* and * is replaced with domain name and image features respectively. The successful results are highlighted with the green outline. Our approach excels at converting the domain of input image features.

Supervision	Methods	R1	R5	R10
ZERO-SHOT	Image-only	8.6	15.4	18.9
	Text-only	6.1	15.7	23.5
	Image+Text	10.2	20.2	26.6
	Pic2Word	11.5	24.8	33.4
CIRR	Combiner [2]	9.9	22.8	32.2
Fashion-IQ	Combiner [2]	13.2	27.1	35.2

Table 2. **Evaluation on COCO**, i.e., object composition task.

den units with ReLU activation (no activation on the output units). We use AdamW [26] with 10^{-4} learning rate and 0.1 weight decay. The batch size of the contrastive learning is 1024. The mapping network is trained on 8 Tesla V100 GPUs. We report the performance averaged over three trials.

(a). Domain conversion. In this setup, we evaluate the ability to compose domain information and query image representation. We utilize ImageNet [9] and ImageNet-R [15], which is comprised of 200 classes with diverse domains and has domain annotations. Considering the noise in the annotation, we pick *cartoon*, *origami*, *toy* and *sculpture* as the evaluation target. Then, given the real image from ImageNet and target domain name, we compose the query fol-

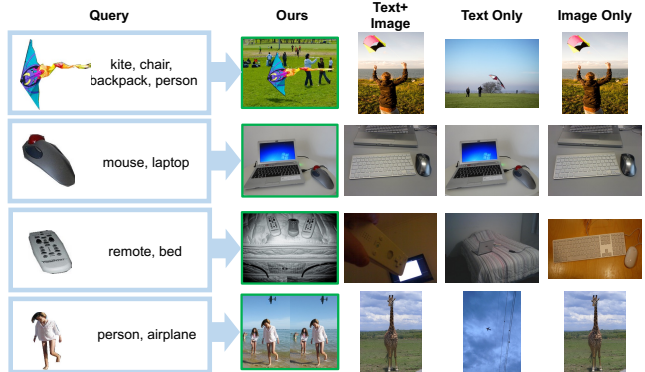


Figure 4. **Qualitative results on COCO.** We visualize top-1 retrievals of zero-shot methods. Correct retrievals are highlighted with green outline.

lowing the procedure in (a) in Sec 3.3. We regard the retrieved image as correct if its class and domain match with the query image and domain. Due to the lack of datasets that can evaluate instance-level cross-domain retrieval on diverse samples, we rely on class- and domain-level evaluation. While not ideal this evaluation can still give us good insights into the model’s ability to accurately translate the domain of the query image compared to baselines.

Supervision	Methods	R1	R5	R10	R50
ZERO-SHOT	Image-only	7.4	23.6	34.0	57.4
	Text-only	20.9	44.8	56.7	79.1
	Image+Text	12.4	36.2	49.1	78.2
	Pic2Word	23.9	51.7	65.3	87.8
CIRR Fashion-IQ	Combiner [2]	30.3	60.4	73.2	92.6
	Combiner [2]	20.1	47.7	61.6	85.9
CIRR	Combiner* [2]	33.6	65.4	77.4	95.2
CIRR	TIRG [35]	14.6	48.4	64.1	90.0
CIRR	ARTEMIS [8]	17.0	46.1	61.3	87.7
CIRR	CIRPLANT [25]	19.6	52.6	68.4	92.4

Table 3. **Evaluation on CIRR test set.** Combiner* is the result reported by the authors using ResNet50x4 as a backbone.

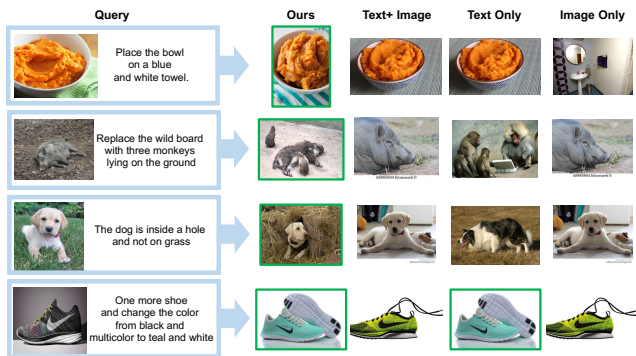


Figure 5. **Qualitative results on CIRR.** We visualize top-1 retrievals of zero-shot methods. Correct retrievals are highlighted with green outline. We can observe that our approach captures both the object in image and the caption well.

The search is performed on 16,983 images from ImageNet-R and ImageNet, i.e., *cartoon, origami, toy, sculpture and real* images are candidates.

(b). Object composition. We evaluate the ability to compose an instance, given as an image, and other scenes or objects, described by text. Following Neculai *et al.* [29], who use COCO [22] for composed image retrieval, COCO validation set (5,000 images) is used for evaluation. To create a query for each image, we randomly crop one object and mask its background using its instance mask. The list of object classes, including the class of the query image, in the image is used as text specification. Given the reference image and class list, we compose query by following (b) in Sec 3.3, e.g., a photo of [*], [car], [cat], and dog.

(c). Scene manipulation by text description. CIRR [25] (Fig. 5) is employed to evaluate image manipulation described by text. The modification from the reference to the target image is given as a sentence. We follow the standard evaluation protocol and compose query text following (c) in Sec 3.3. We report the results on the test split only in Table 3, the validation split is used in other analysis.

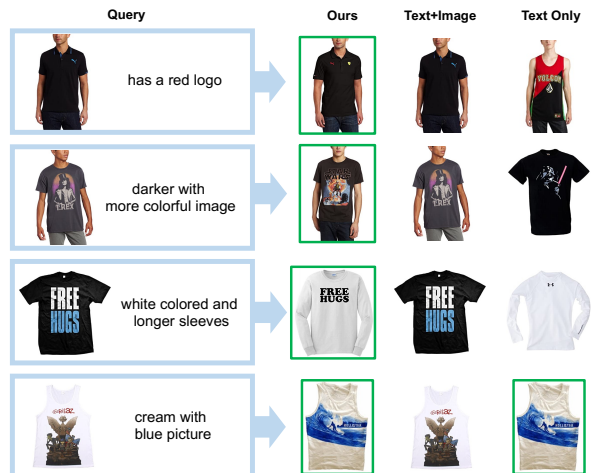


Figure 6. **Retrieval examples in Fashion-IQ.** The search is done in the validation split of Fashion-IQ dataset. Note that Text + Image baseline returns the reference image. The baseline is likely to put more importance on image features although the text and image features are averaged with an equal weight.

(d). Fashion attribute manipulation. Fashion-IQ [37] is employed to evaluate the manipulation of fashion images. The attribute manipulation is given as a sentence. Similarly to CIRR, we follow the standard evaluation protocol and compose query text following (c) in Sec 3.3. The validation set is used for evaluation following previous work [2].

Zero-shot baselines. We provide three zero-shot baselines to fairly compare with our approach.

- **Image Only.** This baseline retrieves images by computing similarity between target image features and query image features.
- **Text Only.** This baseline employs text features only to compute similarity with target images.
- **Image and Text Baseline.** This baseline utilizes the mean of image and text features as query features. Before taking the average, both image and text features are normalized to have a unit L_2 -norm.

Supervised baselines. To better understand the performance of zero-shot methods, we compare with baselines trained with labeled triplets of a CIR dataset. We implement Combiner [2] model with a CLIP ViT-L/14 backbone following the author’s code.² We train the model with the training split of CIRR or Fashion-IQ; approximately, CIRR has 28,000 training triplets and Fashion-IQ includes 18,000. We also show the reported performance of several CIR methods for CIRR and Fashion-IQ.

4.2. Main Results

Tables 1-4 present quantitative results, Fig. 3-6 illustrate qualitative results. In the experiment on domain conversion

²<https://github.com/ABaldrati/CLIP4CirDemo>

Supervision	Methods	Dress		Shirt		TopTee		Average	
		R10	R50	R10	R50	R10	R50	R10	R50
ZERO-SHOT	Image-only	5.4	13.9	9.9	20.8	8.3	17.7	7.9	17.5
	Text-only	13.6	29.7	18.9	31.8	19.3	37.0	17.3	32.9
	Image+Text	16.3	33.6	21.0	34.5	22.2	39.0	19.8	35.7
	Pic2Word	20.0	40.2	26.2	43.6	27.9	47.4	24.7	43.7
CIRR	Combiner [2]	17.2	37.9	23.7	39.4	24.1	43.9	21.7	40.4
Fashion-IQ	Combiner [2]	30.3	54.5	37.2	55.8	39.2	61.3	35.6	57.2
Fashion-IQ	Combiner* [2]	31.6	56.7	36.4	58.0	38.2	62.4	35.4	59.0
Fashion-IQ	CIRPLANT [25]	17.5	40.4	17.5	38.8	21.6	45.4	18.9	41.5
Fashion-IQ	ALTEMIS [8]	27.2	52.4	21.8	43.6	29.2	54.8	26.1	50.3
Fashion-IQ	MAAF [11]	23.8	48.6	21.3	44.2	27.9	53.6	24.3	48.8

Table 4. **Results on Fashion-IQ validation set.** Combiner* is the result reported by the authors using ResNet50x4 as a backbone.

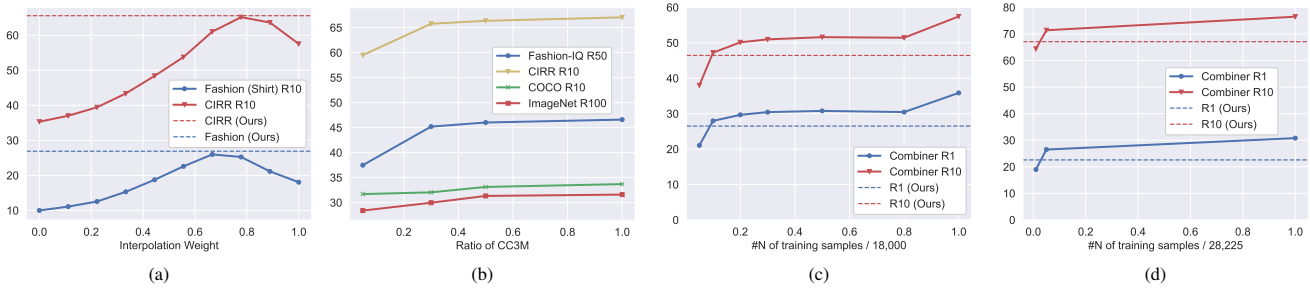


Figure 7. (a): Analysis on the importance weight to average text and image features. (b): Performance change by the number of CC3M samples used to train mapping network. (c): Performance change of the supervised baseline (Combiner) by the number of supervised training samples in Fashion-IQ dataset. (d): Performance change of the supervised baseline (Combiner) by the number of supervised training samples in CIRR dataset.

using ImageNet (Table 1), Pic2Word outperforms all of the baselines with a large margin. This indicates the superiority of Pic2Word in composing domain words and image features. Fig. 3 demonstrates that Pic2Word captures both the domain and image representations well.

In the experiment on object composition using COCO (Table 2), Pic2Word outperforms the zero-shot baselines. Combiner [2] trained on Fashion-IQ performs better than Pic2Word while the model trained on CIRR performs worse than Pic2Word. A model trained on one dataset is not necessarily transferable to other datasets because each composed retrieval dataset can require different dataset bias, e.g., relative importance of image features and caption features.

On scene and fashion attribute manipulation datasets, e.g., CIRR (Table 3) and Fashion-IQ (Table 4), Pic2Word outperforms zero-shot baselines and some of the supervised approaches. For these datasets, we observe that the captions alone can be informative enough to retrieve the correct target images (e.g., last row in Fig. 5), and that some reference images are not very relevant to the target image. Then, the supervised model can learn the dataset-specific bias on how to combine the two modalities. For example, the model can

prioritize learning from text or ignoring the reference image if the image and text are not related well. We hypothesize that this is affecting the transferability of the supervised models. More detailed analysis is provided in the following section.

4.3. Analysis

Does the pseudo language token capture the image information? To analyze the information in the estimated token, we conduct evaluation using CC3M validation 13,164 images. Concretely, we evaluate whether the language embedding generated by the pseudo token can retrieve the input image. Recall at top 1 and 5 is 99.8 and 100.0 respectively, indicating that the token captures the unique image features very well.

Insights on existing composed retrieval datasets. The results on CIRR and Fashion-IQ (Table 3 and 4) imply that some target images can be correctly retrieved just by using text. We find that reference images can be irrelevant to text query or text query is informative enough to search target image. This relative importance of reference image and text is a dataset-specific bias. We investigate the importance in

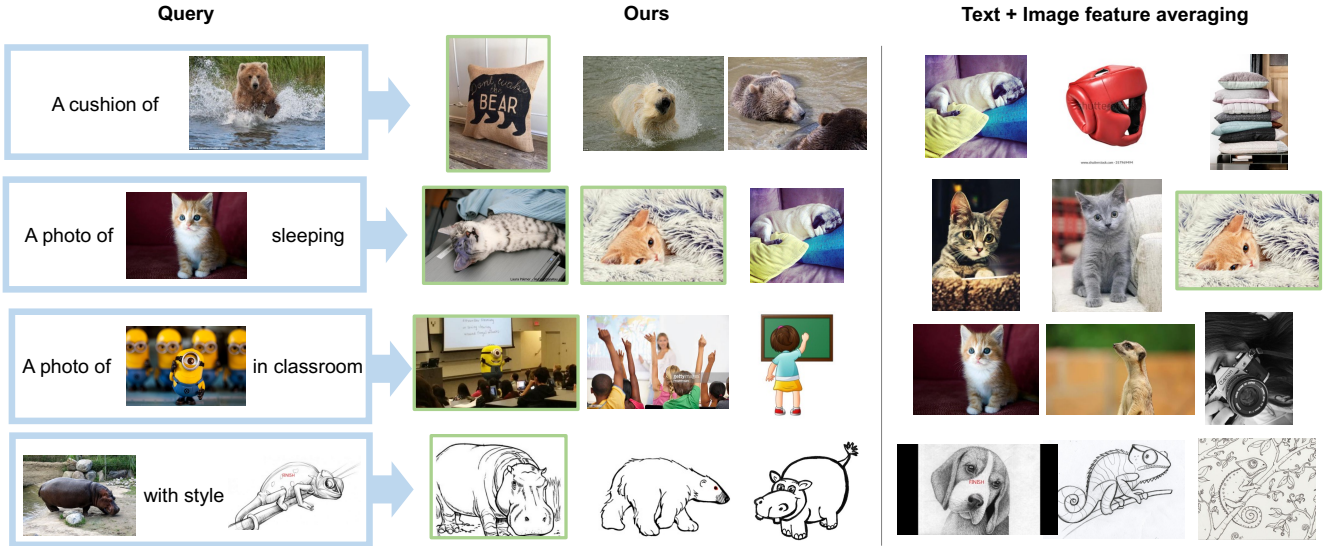


Figure 8. Top-3 examples retrieved from CC3M validation set (13,000 images). Correct retrievals are highlighted with green outline.




Figure 9. Retrieval with sketch-style images. We qualitatively find that it can be hard to handle sketch-to-real image retrieval (Top two).

Fig. 7a, where we produce a query embedding by varying the interpolation weight, w , in computing the average of text t and image features v , i.e., $q = w * t + (1 - w) * v$. The result indicates that the optimal weight can be unique to each dataset and the performance is sensitive to the weight parameter. A supervised method can learn the relative importance from labeled triplets, but it is hard to learn for a zero-shot method.

Comparison to supervised baseline with fewer training samples. Fig. 7c and 7d show the performance of the supervised baseline (Combiner [2]) with fewer training samples. Overall, our approach outperforms the baseline if fewer than 1,000 triplets are given as training samples.

Number of training samples to train the mapping network. In Fig. 7b, we investigate the performance w.r.t the number of data (from CC3M) used to train the mapping net-

work. We observe that using only 10% degrades the performance while the use 50% data is comparable to all data.

More qualitative examples. Fig. 8 shows retrieval results using CC3M validation data. Our model can compose images with adjective (second) or a place (third) and read the style given as an image (bottom).

Failure cases. One of the popular applications is sketch-based image retrieval [23], where the user draws a sketch to retrieve natural images, as shown in Fig. 9. Qualitatively, retrieving an image from a non-natural image domain is likely to be successful (bottom two) while retrieving natural images from sketch images is not trivial (top two). Prompts other than a real image of do not work well for this task either. Natural images are not described with a domain word in pre-training image-caption dataset while other domains, e.g., art, cartoon, sketch and origami, are often described with corresponding domain name. An appropriate domain description for natural images may not exist.

5. Conclusion

In this paper, we present a novel task, zero-shot composed image retrieval, and the first method to approach the problem. We propose to employ a pre-trained CLIP model that treats an image as a text token so that the language encoder can flexibly compose the image features and text description. We perform a comprehensive analysis on four datasets, demonstrating that our approach, Pic2Word, shows strong generalization across diverse CIR tasks, with performance being on-par with or better than several recent CIR methods that require labeled training data.

Acknowledgment. We thank Zizhao Zhang for their helpful feedback on the manuscript. This work was supported in part by DARPA LwLL.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *arXiv preprint arXiv:2204.14198*, 2022. 1, 2
- [2] Alberto Baldrati, Marco Bertini, Tiberio Uricchio, and Alberto Del Bimbo. Effective conditioned and composed image retrieval combining clip-based features. In *CVPR*, pages 21466–21474, 2022. 1, 2, 4, 5, 6, 7, 8
- [3] Yanbei Chen and Loris Bazzani. Learning joint visual semantic matching embeddings for language-guided retrieval. In *ECCV*, pages 136–152. Springer, 2020. 2, 4
- [4] Yanbei Chen, Shaogang Gong, and Loris Bazzani. Image search with text feedback by visiolinguistic attention learning. In *CVPR*, pages 3001–3011, 2020. 1, 2, 4
- [5] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Learning universal image-text representations. *arXiv preprint*, 2019. 2
- [6] Niv Cohen, Rinon Gal, Eli A Meir, Gal Chechik, and Yuval Atzmon. "this is my unicorn, fluffy": Personalizing frozen vision-language representations. In *ECCV*, 2022. 2
- [7] Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys (Csur)*, 40(2):1–60, 2008. 1
- [8] Ginger Delmas, Rafael Sampaio de Rezende, Gabriela Csurka, and Diane Larlus. Artemis: Attention-based retrieval with text-explicit matching and implicit similarity. In *ICLR*, 2022. 2, 4, 6, 7
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [11] Eric Dodds, Jack Culpepper, Simao Herdade, Yang Zhang, and Kofi Boakye. Modality-agnostic attention fusion for visual search with text feedback. *arXiv preprint arXiv:2007.00145*, 2020. 2, 7
- [12] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 3
- [13] Sonam Goenka, Zhaoheng Zheng, Ayush Jaiswal, Rakesh Chada, Yue Wu, Varsha Hedau, and Pradeep Natarajan. Fashionvlp: Vision language transformer for fashion retrieval with feedback. In *CVPR*, pages 14105–14115, 2022. 2
- [14] Xiao Han, Licheng Yu, Xiatian Zhu, Li Zhang, Yi-Zhe Song, and Tao Xiang. Fashionvil: Fashion-focused vision-and-language representation learning. In *ECCV*, 2022. 2
- [15] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. 5
- [16] Mehrdad Hosseinzadeh and Yang Wang. Composed query image retrieval using locally bounded features. In *CVPR*, pages 3596–3605, 2020. 1
- [17] 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 *Int. Conf. Machine Learning.*, pages 4904–4916. PMLR, 2021. 2
- [18] Seungmin Lee, Dongwan Kim, and Bohyung Han. Cosmo: Content-style modulation for image retrieval with text feedback. In *CVPR*, pages 802–812, 2021. 1, 2, 4
- [19] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021. 4
- [20] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 34, 2021. 1, 2
- [21] Xiujun Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, et al. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, pages 121–137. Springer, 2020. 2
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [23] Li Liu, Fumin Shen, Yuming Shen, Xianglong Liu, and Ling Shao. Deep sketch hashing: Fast free-hand sketch-based image retrieval. In *CVPR*, pages 2862–2871, 2017. 8
- [24] Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *arXiv preprint arXiv:2107.13586*, 2021. 4
- [25] Zheyuan Liu, Cristian Rodriguez-Opazo, Damien Teney, and Stephen Gould. Image retrieval on real-life images with pre-trained vision-and-language models. In *ICCV*, pages 2125–2134, 2021. 1, 2, 4, 6, 7
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- [27] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 32, 2019. 2
- [28] Ron Mokady, Amir Hertz, and Amit H Bermano. Clip-cap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021. 2
- [29] Andrei Neculai, Yanbei Chen, and Zeynep Akata. Probabilistic compositional embeddings for multimodal image retrieval. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4547–4557, 2022. 6

- [30] 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 *Int. Conf. Machine Learning.*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [3](#), [4](#)
- [31] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018. [4](#)
- [32] Amanpreet Singh, Ronghang Hu, Vedanuj Goswami, Guillaume Couairon, Wojciech Galuba, Marcus Rohrbach, and Douwe Kiela. Flava: A foundational language and vision alignment model. *arXiv preprint arXiv:2112.04482*, 2021. [2](#)
- [33] Kihyuk Sohn. Improved deep metric learning with multi-class n-pair loss objective. *Advances in neural information processing systems*, 29, 2016. [3](#)
- [34] Haoyu Song, Li Dong, Wei-Nan Zhang, Ting Liu, and Furu Wei. Clip models are few-shot learners: Empirical studies on vqa and visual entailment. *arXiv preprint arXiv:2203.07190*, 2022. [2](#)
- [35] Nam Vo, Lu Jiang, Chen Sun, Kevin Murphy, Li-Jia Li, Li Fei-Fei, and James Hays. Composing text and image for image retrieval-an empirical odyssey. In *CVPR*, pages 6439–6448, 2019. [1](#), [2](#), [6](#)
- [36] Teng Wang, Wenhao Jiang, Zhichao Lu, Feng Zheng, Ran Cheng, Chengguo Yin, and Ping Luo. Vlmixer: Unpaired vision-language pre-training via cross-modal cutmix. In *Int. Conf. Machine Learning.*, pages 22680–22690. PMLR, 2022. [2](#)
- [37] Hui Wu, Yupeng Gao, Xiaoxiao Guo, Ziad Al-Halah, Steven Rennie, Kristen Grauman, and Rogerio Feris. Fashion iq: A new dataset towards retrieving images by natural language feedback. In *CVPR*, pages 11307–11317, 2021. [2](#), [4](#), [6](#)
- [38] 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. [2](#)
- [39] Kaiyang Zhou, Jingkang Yang, Chen Change Loy, and Ziwei Liu. Conditional prompt learning for vision-language models. *arXiv preprint arXiv:2203.05557*, 2022. [2](#)